

Package ‘EnvStats’

August 15, 2013

Type Package

Title Package for Environmental Statistics, including US EPA Guidance

Version 0.0-7

Date 2013-07-23

Author Steven P. Millard <EnvStats@ProbStatInfo.com>

Maintainer Steven P. Millard <EnvStats@ProbStatInfo.com>

Depends R (>= 2.15.2)

Suggests lattice, qcc, sp, MASS

Description *BETA* Version: several help files still need to be added. Functions and datasets to perform analyses of environmental data.

License GPL (>=3)

URL <http://www.probstatinfo.com>

LazyLoad yes

LazyData yes

R topics documented:

EnvStats-package	7
anovaPE	12
aovN	13
aovPower	16
base	19
Benthic.df	20
boxcox	22
boxcox.object	31
boxcoxCensored	33
boxcoxCensored.object	42
boxcoxLm.object	45
boxcoxTransform	47
calibrate	51
CastilloAndHadi1994	54

cdfCompare	56
cdfPlot	60
Chi	63
ciBinomHalfWidth	64
ciBinomN	70
ciNormHalfWidth	77
ciNormN	80
cv	83
detectionLimitCalibrate	86
Distribution.df	91
ebeta	98
ebinom	101
ecdfPlot	107
ecdfPlotCensored	111
eevd	116
eexp	122
egamma	124
egammaAltCensored	132
egammaCensored	140
egeom	148
egevd	150
ehyper	156
elnorm	159
elnorm3	161
elnormAlt	173
elnormAltCensored	185
elnormCensored	200
elogis	209
Empirical	212
enbinom	216
enorm	219
enormCensored	223
enparCensored	241
Environmental	249
EPA.02d.Ex.2.ug.per.L.vec	251
EPA.02d.Ex.4.mg.per.kg.vec	251
EPA.02d.Ex.6.mg.per.kg.vec	252
EPA.02d.Ex.9.mg.per.L.vec	252
EPA.09.Ex.10.1.nickel.df	253
EPA.09.Ex.11.1.arsenic.df	253
EPA.09.Ex.12.1.ccl4.df	254
EPA.09.Ex.12.4.naphthalene.df	255
EPA.09.Ex.13.1.iron.df	255
EPA.09.Ex.14.1.manganese.df	256
EPA.09.Ex.14.3.alkalinity.df	257
EPA.09.Ex.14.4.arsenic.df	257
EPA.09.Ex.14.8.df	258
EPA.09.Ex.15.1.manganese.df	259
EPA.09.Ex.16.1.sulfate.df	259
EPA.09.Ex.16.2.benzene.df	260
EPA.09.Ex.16.4.copper.df	261
EPA.09.Ex.16.5.PCE.df	261

EPA.09.Ex.17.1.loglead.df	262
EPA.09.Ex.17.2.toluene.df	263
EPA.09.Ex.17.3.chrysene.df	263
EPA.09.Ex.17.3.log.chrysene.df	264
EPA.09.Ex.17.4.copper.df	265
EPA.09.Ex.17.5.chloride.df	265
EPA.09.Ex.17.6.sulfate.df	266
EPA.09.Ex.17.7.sodium.df	267
EPA.09.Ex.18.1.arsenic.df	268
EPA.09.Ex.18.2.chrysene.df	268
EPA.09.Ex.18.3.TCE.df	269
EPA.09.Ex.18.4.xylene.df	270
EPA.09.Ex.19.1.sulfate.df	270
EPA.09.Ex.19.2.chloride.df	271
EPA.09.Ex.19.5.mercury.df	272
EPA.09.Ex.20.1.nickel.df	272
EPA.09.Ex.21.1.aldicarb.df	273
EPA.09.Ex.21.2.benzene.df	274
EPA.09.Ex.21.5.beryllium.df	274
EPA.09.Ex.21.6.nitrate.df	275
EPA.09.Ex.21.7.TCE.df	276
EPA.09.Ex.22.1.VC.df	276
EPA.09.Ex.22.2.Specific.Conductance.df	277
EPA.09.Ex.6.3.sulfate.df	278
EPA.09.Ex.7.1.arsenic.df	278
EPA.09.Table.9.1.TCE.df	279
EPA.09.Table.9.3.df	280
EPA.09.Table.9.4.nickel.vec	280
EPA.89b.aldicarb1.df	281
EPA.89b.aldicarb2.df	282
EPA.89b.benzene.df	282
EPA.89b.cadmium.df	283
EPA.89b.chlordane1.df	283
EPA.89b.chlordane2.df	284
EPA.89b.edb.df	284
EPA.89b.lead.df	285
EPA.89b.loglead.df	286
EPA.89b.manganese.df	286
EPA.89b.sulfate.df	287
EPA.89b.t29.df	287
EPA.89b.toc.vec	288
EPA.92c.arsenic1.df	288
EPA.92c.arsenic2.df	289
EPA.92c.arsenic3.df	289
EPA.92c.benzene1.df	290
EPA.92c.benzene2.df	290
EPA.92c.ccl4.df	291
EPA.92c.chrysene.df	292
EPA.92c.copper1.df	292
EPA.92c.copper2.df	293
EPA.92c.lognickel1.df	293
EPA.92c.nickel1.df	294

EPA.92c.nickel2.df	294
EPA.92c.toluene.df	295
EPA.92c.zinc.df	295
EPA.92d.chromium.df	296
EPA.92d.chromium.vec	297
EPA.94b.lead.df	297
EPA.94b.tccb.df	298
EPA.97.cadmium.111.df	298
epareto	299
epdfPlot	301
epois	304
epoisCensored	307
eqbeta	315
eqbinom	317
eqevd	320
eqexp	323
eqgamma	325
eqgeom	330
eqgevd	332
eqhyper	335
eqlnorm	338
eqlnorm3	342
eqlogis	345
eqnbinom	347
eqnorm	349
eqnpar	353
eqpareto	361
eqpois	363
equnif	367
eqweibull	369
eqzmlnorm	371
eqzmnorm	374
estimate.object	376
estimateCensored.object	380
EulersConstant	384
eunif	385
EVD	388
eweibull	390
ezmlnorm	393
ezmnorm	398
FcnsByCat	402
FcnsByCatCalibration	403
FcnsByCatCensoredData	403
FcnsByCatDataTrans	406
FcnsByCatEstDistParams	407
FcnsByCatEstDistQuants	407
FcnsByCatGOFTests	408
FcnsByCatHypothTests	409
FcnsByCatMCandRisk	410
FcnsByCatPlotProbDists	410
FcnsByCatPower	411
FcnsByCatPredInts	414

FcnsByCatPrintPlot	415
FcnsByCatProbDists	417
FcnsByCatSumStats	418
FcnsByCatTolInts	418
FcnsByCatTrend	419
GammaAlt	419
geoMean	422
geoSD	424
GEVD	426
Gibbons.et.al.09.Akilinity.vec	428
Gibbons.et.al.09.Vinyl.Chloride.vec	428
gof.object	429
gofGroup.object	432
gofGroupTest	435
gofTest	441
gofTwoSample.object	455
gpqCiNormCensored	457
gpqTolIntNormCensored	460
Graham.et.al.75.etu.df	463
Helsel.Cohn.88.app.b.df	464
Helsel.Cohn.88.silver.df	464
HoskingEtAl1985	465
inversePredictCalibrate	469
iqr	472
kendallSeasonalTrendTest	474
kendallTrendTest	481
kurtosis	485
Lin.Evans.80.df	489
lMoment	489
Lognormal3	494
LognormalAlt	498
LognormalMix	500
LognormalMixAlt	502
LognormalTrunc	504
LognormalTruncAlt	507
longToWide	510
Millard.Deverel.88.df	512
Modified.TcCB.df	513
NIOSH.89.air.lead.vec	514
NormalMix	514
NormalTrunc	516
Olympic.NH4.df	519
Ozone.NE.df	520
Pareto	521
pdfPlot	523
plot.boxcox	526
plot.boxcoxCensored	529
plot.boxcoxLm	533
plot.gof	536
plot.gofGroup	541
plot.gofTwoSample	545
plotAovDesign	550

plotCiBinomDesign	553
plotCiNormDesign	559
pointwise	563
ppointsCensored	567
predict.lm	578
predIntGamma	583
predIntGammaSimultaneous	589
predIntLnorm	598
predIntLnormSimultaneous	606
predIntNorm	614
predIntNormK	620
predIntNormSimultaneous	626
predIntNormSimultaneousK	636
predIntNpar	644
predIntNparSimultaneous	652
predIntPois	664
print.bboxcox	673
print.bboxcoxCensored	674
print.bboxcoxLm	675
print.gof	676
print.gofGroup	677
print.gofTwoSample	678
propTestN	679
propTestPower	685
pwMoment	691
qqPlot	695
qqPlotCensored	703
qqPlotGestalt	711
Refinery.CO.df	715
simulateMvMatrix	716
simulateVector	723
Skagit.NH3_N.df	727
skewness	728
stripChart	731
summaryFull	736
summaryStats	742
tolIntGamma	750
tolIntLnorm	758
tolIntNorm	763
tolIntNormCensored	768
tolIntNormK	772
tolIntNpar	777
tolIntPois	783
Total.P.df	787
Triangular	788
varGroupTest	790
varTest	794
ZeroModifiedLognormal	796
ZeroModifiedLognormalAlt	799
ZeroModifiedNormal	802

Description

EnvStats is a comprehensive R package for Environmental Statistics and is the successor to the S-PLUS module *EnvironmentalStats for S-PLUS*, the first copy of which was sold in April, 1997. **EnvStats** provides a set of powerful functions for performing graphical and statistical analyses of environmental data, bringing major environmental statistical methods found in the literature and regulatory guidance documents into one statistical package, along with an extensive hypertext help system that explains what these methods do, how to use these methods, and where to find them in the environmental statistics literature. Also included are numerous built-in data sets from regulatory guidance documents and the environmental statistics literature. For a complete list of functions and datasets, you can do any of the following:

- If you are in the on-line help, scroll to the bottom of this help page and click on the **Index** link.
- See the help file [Functions By Category](#) for a listing of functions by category.
- Type `library(help="EnvStats")` at the command prompt.

Note: The names of all **EnvStats** functions start with a lowercase letter, and the names of all **EnvStats** datasets and data objects start an uppercase letter. You can type `newsEnvStats()` at the R command prompt for the latest news for the **EnvStats** package.

Details

Package:	EnvStats
Type:	Package
Version:	0.0-7
Date:	2013-07-29
License:	GPL (>=3)
LazyLoad:	yes

A companion file **EnvStats-manual.pdf** containing a listing of all the current help files is located in the **doc** subdirectory of the directory where the **EnvStats** package was installed. For example, if you installed R under Windows, this file might be located in the directory **C:\Program Files\R-*.**.*\library\EnvStats\doc**, where ***.***** denotes the version of R you are using (e.g., 2.15.1) or in the directory **C:\Users\Name\Documents\R\win-library*.***\EnvStats\doc**, where **Name** denotes your user name on the Windows operating system.

EnvStats comes with companion scripts, located in the **scripts** subdirectory of the directory where the package was installed. One set of scripts lets you reproduce the examples in the User's Manual (currently is still in preparation). There are also scripts that let you reproduce examples from US EPA guidance documents.

See the **References** section below for documentation for the predecessor to **EnvStats**, *EnvironmentalStats for S-PLUS* for Windows.

Features of EnvStats include:

- New functions for computing [summary statistics](#) and creating [summary plots](#) to compare the

distributions of groups side-by-side.

- New [probability distributions](#) have been added to the ones already available in R, including the extreme value distribution and the zero-modified lognormal (delta) distribution. You can compute quantities associated with these probability distributions (probability density functions, cumulative distribution functions, and quantiles), and generate random numbers from these distributions.
- [Plot probability distributions](#) so you can see how they change with the value of the distribution parameter(s).
- [Estimate distribution parameters](#) and [distribution quantiles](#), and compute confidence intervals for commonly used probability distributions, including special methods for the lognormal and gamma distributions.
- Perform and plot the results of [goodness-of-fit tests](#):
 - Observed and Fitted Distributions
 - Quantile-Quantile Plots
 - Results of Shapiro-Wilk test, Kolmogorov-Smirnov test, etc.

Includes a new generalized goodness-of-fit test for any continuous distribution.

- Functions for assessing optimal [Box-Cox data transformations](#).
- Compute parametric and non-parametric [prediction intervals](#), [simultaneous prediction intervals](#), and [tolerance intervals](#).
- New functions for [hypothesis tests](#), including:
 - Nonparametric estimation and tests for seasonal trend
 - Fisher's one-sample randomization (permutation) test for location
 - Quantile test to detect a shift in the tail of one population relative to another
 - Two-sample linear rank tests
 - Test for serial correlation based on von Neumann rank test
- Perform [calibration](#) based on a machine signal to determine decision and detection limits and report estimated concentrations along with confidence intervals.
- Easily perform [power and sample size](#) computations and create companion plots for sampling designs based on confidence intervals, hypothesis tests, prediction intervals, and tolerance intervals.
- Handle singly and multiply [censored \(less-than-detection-limit\) data](#):
 - Empirical CDF and Quantile-Quantile Plots
 - Parameter/Quantile Estimation and Confidence Intervals
 - Prediction and Tolerance Intervals
 - Goodness-of-Fit Tests
 - Optimal Box-Cox Transformations
 - Two-Sample Rank Tests
- Functions for performing [Monte Carlo simulation and probabilistic risk assesement](#).
- Reproduce specific examples in EPA guidance documents by using built-in data sets from these documents and running companion scripts.

Author(s)

Steven P. Millard

Maintainer: Steven P. Millard <EnvStats@ProbStatInfo.com>

References

- Millard, S.P. (In Preparation). *EnvStats: An R Package for Environmental Statistics*. Springer-Verlag, New York.
- Millard, S.P. (2002). *EnvironmentalStats for S-PLUS: User's Manual for Version 2.0*. Second Edition. Springer-Verlag, New York.
- Millard, S.P., and N.K. Neerchal. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, FL.

Examples

```
# Look at plots and summary statistics for the TcCB data given in
# USEPA (1994b), (the data are stored in EPA.94b.tccb.df).
# Arbitrarily set the one censored observation to the censoring level.
# Group by the variable Area.
```

	EPA.94b.tccb.df
#	TcCB.orig TcCB Censored Area
#1	0.22 0.22 FALSE Reference
#2	0.23 0.23 FALSE Reference
#...	
#46	1.20 1.20 FALSE Reference
#47	1.33 1.33 FALSE Reference
#48	<0.09 0.09 TRUE Cleanup
#49	0.09 0.09 FALSE Cleanup
#...	
#123	51.97 51.97 FALSE Cleanup
#124	168.64 168.64 FALSE Cleanup

```

# First plot the data
#-----
windows()
stripChart(TcCB ~ Area, data = EPA.94b.tccb.df,
  xlab = "Area", ylab = "TcCB (ppb)")
mtext("TcCB Concentrations by Area", line = 3, cex = 1.25, font = 2)

windows()
stripChart(log10(TcCB) ~ Area, data = EPA.94b.tccb.df,
  p.value = TRUE,
  xlab = "Area", ylab = expression(paste(log[10], " [ TcCB (ppb) ]")))
mtext(expression(paste(log[10], "(TcCB) Concentrations by Area")),
  line = 3, cex = 1.25, font = 2)

#-----

# Now compute summary statistics
#-----

sum(EPA.94b.tccb.df$Censored)
#[1] 1

EPA.94b.tccb.df$TcCB[EPA.94b.tccb.df$Censored]
#0.09

# Summary statistics will treat the one censored value
```

```

# as assuming the detection limit.

summaryFull(TcCB ~ Area, data = EPA.94b.tccb.df)
#           Cleanup Reference
#N           77      47
#Mean         3.915    0.5985
#Median        0.43     0.54
#10% Trimmed Mean 0.6846 0.5728
#Geometric Mean 0.5784 0.5382
#Skew          7.717    0.9019
#Kurtosis      62.67    0.132
#Min           0.09     0.22
#Max          168.6     1.33
#Range         168.5     1.11
#1st Quartile   0.23     0.39
#3rd Quartile   1.1      0.75
#Standard Deviation 20.02 0.2836
#Geometric Standard Deviation 3.898 1.597
#Interquartile Range 0.87 0.36
#Median Absolute Deviation 0.3558 0.2669
#Coefficient of Variation 5.112 0.4739

summaryStats(TcCB ~ Area, data = EPA.94b.tccb.df, digits = 1)
#           N Mean  SD Median Min  Max
#Cleanup   77  3.9 20.0   0.4 0.1 168.6
#Reference  47  0.6  0.3   0.5 0.2   1.3

#-----

# Compute Shapiro-Wilk Goodness-of-Fit statistic for the
# Reference Area TcCB data assuming a lognormal distribution
#-----

sw.list <- gofTest(TcCB ~ 1, data = EPA.94b.tccb.df,
  subset = Area == "Reference", dist = "lnorm")
sw.list

# Results of Goodness-of-Fit Test
# -----
#
# Test Method:           Shapiro-Wilk GOF
#
# Hypothesized Distribution: Lognormal
#
# Estimated Parameter(s): meanlog = -0.6195712
#                          sdlog   = 0.4679530
#
# Estimation Method:     mvue
#
# Data:                  TcCB
#
# Subset With:           Area == "Reference"
#
# Data Source:           EPA.94b.tccb.df
#
# Sample Size:           47
#

```

```

# Test Statistic:          W = 0.978638
#
# Test Statistic Parameter:  n = 47
#
# P-value:                 0.5371935
#
# Alternative Hypothesis:   True cdf does not equal the
#                           Lognormal Distribution.

#-----

# Plot results of GOF test
windows()
plot(sw.list)

#-----

# Based on the Reference Area data, estimate 90th percentile
# and compute a 95% confidence limit for the 90th percentile
# assuming a lognormal distribution.
#-----

TcCB.ref <- EPA.94b.tccb.df$TcCB[EPA.94b.tccb.df$Area == "Reference"]
eqlnorm(TcCB.ref, p = 0.9, ci = TRUE)

# Results of Distribution Parameter Estimation
# -----
#
# Assumed Distribution:      Lognormal
#
# Estimated Parameter(s):    meanlog = -0.6195712
#                           sdlog   =  0.4679530
#
# Estimation Method:        mvue
#
# Estimated Quantile(s):     90'th %ile = 0.9803307
#
# Quantile Estimation Method: qmle
#
# Data:                     TcCB.ref
#
# Sample Size:              47
#
# Confidence Interval for:    90'th %ile
#
# Confidence Interval Method: Exact
#
# Confidence Interval Type:   two-sided
#
# Confidence Level:          95%
#
# Confidence Interval:        LCL = 0.8358791
#                           UCL = 1.2154977
#-----

# Cleanup
rm(TcCB.ref, sw.list)

```

anovaPE

*Compute Lack-of-Fit and Pure Error Anova Table for a Linear Model***Description**

Compute a lack-of-fit and pure error anova table for a linear model with one predictor variable and replicate observations for at least one value of the predictor variable.

Usage

```
anovaPE(object)
```

Arguments

`object` an object of `class "lm"`. The object can have only one predictor variable in the formula and must have replicate observations for at least one value of the predictor variable.

Details

Produces an anova table with the the sums of squares partitioned by “Lack of Fit” and “Pure Error”. See Draper and Smith (1998, pp.47-53) for details. This function is called by the function `calibrate`.

Value

An object of `class "anova"` inheriting from class `"data.frame"`.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Draper, N., and H. Smith. (1998). *Applied Regression Analysis*. Third Edition. John Wiley and Sons, New York, pp.47-53.

Millard, S.P., and Neerchal, N.K. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, Florida.

See Also

`anova.lm`, `lm`, `calibrate`.

Examples

```
# The data frame EPA.97.cadmium.111.df contains calibration data for
# cadmium at mass 111 (ng/L) that appeared in Gibbons et al. (1997b)
# and were provided to them by the U.S. EPA.
# The Examples section of the help file for calibrate displays a plot
# of these data along with the fitted calibration line and 99%
# non-simultaneous prediction limits.
# Here we will just fit the linear model and produce the anova table
# to check for lack of fit.
```

```
fit <- lm(Cadmium ~ Spike, data = EPA.97.cadmium.111.df)

anova(fit)
#Analysis of Variance Table
#
#Response: Cadmium
#      Df Sum Sq Mean Sq F value    Pr(>F)
#Spike    1  43220    43220  9356.9 < 2.2e-16 ***
#Residuals 33    152         5
#---
#Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
#Analysis of Variance Table
#
#Response: Cadmium
#
#Terms added sequentially (first to last)
#      Df Sum of Sq Mean Sq  F Value Pr(F)
# Spike    1  43220.27  43220.27  9356.879    0
#Residuals 33    152.43     4.62

anovaPE(fit)
#      Df Sum Sq Mean Sq  F value Pr(>F)
#Spike    1  43220    43220 9341.559 <2e-16 ***
#Lack of Fit    3     14         5    0.982 0.4144
#Pure Error   30     139         5
#---
#Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

rm(fit)
```

aovN	<i>Compute Sample Size Necessary to Achieve Specified Power for One-Way Fixed-Effects Analysis of Variance</i>
------	--

Description

Compute the sample sizes necessary to achieve a specified power for a one-way fixed-effects analysis of variance test, given the population means, population standard deviation, and significance level.

Usage

```
aovN(mu.vec, sigma = 1, alpha = 0.05, power = 0.95,
     round.up = TRUE, tol = 1e-07)
```

Arguments

mu.vec required numeric vector of population means. The length of mu.vec must be at least 2. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed.

sigma	optional numeric scalar specifying the population standard deviation (σ) for each group. The default value is sigma=1.
alpha	optional numeric scalar between 0 and 1 indicating the Type I error level associated with the hypothesis test. The default value is alpha=0.05.
power	optional numeric scalar between 0 and 1 indicating the power associated with the hypothesis test. The default value is power=0.95.
round.up	optional logical scalar indicating whether to round up the value of the computed sample size to the next smallest integer. The default value is round.up=TRUE.
tol	optional numeric scalar indicating the tolerance to use in the search algorithm. The default value is tol=1e-7.

Details

The F-statistic to test the equality of k population means assuming each population has a normal distribution with the same standard deviation σ is presented in most basic statistics texts, including Zar (2010), Berthouex and Brown (1994, Chapter 17), and Helsel and Hirsh (1992, pp.164-169). The formula for the power of this test is given in Scheffe (1959, pp.38-39,62-65). The power of the one-way fixed-effects ANOVA depends on the sample sizes for each of the k groups, the value of the population means for each of the k groups, the population standard deviation σ , and the significance level α . The function `aovN` assumes equal sample sizes for each of the k groups and uses a search algorithm to determine the sample size n required to attain a specified power, given the values of the population means and the significance level.

Value

numeric scalar indicating the required sample size for each group. (The number of groups is equal to the length of the argument `mu.vec`.)

Note

The normal and lognormal distribution are probably the two most frequently used distributions to model environmental data. Sometimes it is necessary to compare several means to determine whether any are significantly different from each other (e.g., USEPA, 2009, p.6-38). In this case, assuming normally distributed data, you perform a one-way parametric analysis of variance.

In the course of designing a sampling program, an environmental scientist may wish to determine the relationship between sample size, Type I error level, power, and differences in means if one of the objectives of the sampling program is to determine whether a particular mean differs from a group of means. The functions `aovPower`, `aovN`, and `plotAovDesign` can be used to investigate these relationships for the case of normally-distributed observations.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Berthouex, P.M., and L.C. Brown. (1994). *Statistics for Environmental Engineers*. Lewis Publishers, Boca Raton, FL, Chapter 17.
- Helsel, D.R., and R.M. Hirsch. (1992). *Statistical Methods in Water Resources Research*. Elsevier, New York, NY, Chapter 7.
- Johnson, N. L., S. Kotz, and N. Balakrishnan. (1995). *Continuous Univariate Distributions, Volume 2*. Second Edition. John Wiley and Sons, New York, Chapters 27, 29, 30.

Millard, S.P., and Neerchal, N.K. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, Florida.

Scheffe, H. (1959). *The Analysis of Variance*. John Wiley and Sons, New York, 477pp.

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.6-38.

Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ, Chapter 10.

See Also

[aovPower](#), [plotAovDesign](#), [Normal](#), [aov](#).

Examples

```
# Look at how the required sample size for a one-way ANOVA
# increases with increasing power:
```

```
aovN(mu.vec = c(10, 12, 15), sigma = 5, power = 0.8)
#[1] 21
```

```
aovN(mu.vec = c(10, 12, 15), sigma = 5, power = 0.9)
#[1] 27
```

```
aovN(mu.vec = c(10, 12, 15), sigma = 5, power = 0.95)
#[1] 33
```

```
#-----
```

```
# Look at how the required sample size for a one-way ANOVA,
# given a fixed power, decreases with increasing variability
# in the population means:
```

```
aovN(mu.vec = c(10, 10, 11), sigma=5)
#[1] 581
```

```
aovN(mu.vec = c(10, 10, 15), sigma = 5)
#[1] 25
```

```
aovN(mu.vec = c(10, 13, 15), sigma = 5)
#[1] 33
```

```
aovN(mu.vec = c(10, 15, 20), sigma = 5)
#[1] 10
```

```
#-----
```

```
# Look at how the required sample size for a one-way ANOVA,
# given a fixed power, decreases with increasing values of
# Type I error:
```

```
aovN(mu.vec = c(10, 12, 14), sigma = 5, alpha = 0.001)
#[1] 89
```

```

aovN(mu.vec = c(10, 12, 14), sigma = 5, alpha = 0.01)
#[1] 67

aovN(mu.vec = c(10, 12, 14), sigma = 5, alpha = 0.05)
#[1] 50

aovN(mu.vec = c(10, 12, 14), sigma = 5, alpha = 0.1)
#[1] 42

```

aovPower

*Compute the Power of a One-Way Fixed-Effects Analysis of Variance***Description**

Compute the power of a one-way fixed-effects analysis of variance, given the sample sizes, population means, population standard deviation, and significance level.

Usage

```

aovPower(n.vec, mu.vec = rep(0, length(n.vec)), sigma = 1,
         alpha = 0.05)

```

Arguments

n.vec	numeric vector of sample sizes for each group. The i^{th} element of n.vec denotes the sample size for group i . The length of n.vec must be at least 2, and all elements of n.vec must be greater than or equal to 2. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed.
mu.vec	optional numeric vector of population means. The length of mu.vec must be the same as the length of n.vec. The default value is a vector of zeros. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed.
sigma	optional numeric scalar specifying the population standard deviation (σ) for each group. The default value is sigma=1.
alpha	optional numeric scalar between 0 and 1 indicating the Type I error level associated with the hypothesis test. The default value is alpha=0.05.

Details

The F-statistic to test the equality of k population means assuming each population has a normal distribution with the same standard deviation σ is presented in most basic statistics texts, including Zar(1999, p.183), Berthouex and Brown (1994, Chapter 17), and Helsel and Hirsh (1992, pp.164-169). The formula for the power of this test is given in Scheffe (1959, pp.38-39,62-65). The power of the one-way fixed-effects ANOVA depends on the sample sizes for each of the k groups, the value of the population means for each of the k groups, the population standard deviation σ , and the significance level α .

Value

a numeric scalar indicating the power of the one-way fixed-effects ANOVA for the given sample sizes, population means, population standard deviation, and significance level.

Note

The normal and lognormal distribution are probably the two most frequently used distributions to model environmental data. Sometimes it is necessary to compare several means to determine whether any are significantly different from each other (e.g., USEPA, 2009, p.6-38). In this case, assuming normally distributed data, you perform a one-way parametric analysis of variance.

In the course of designing a sampling program, an environmental scientist may wish to determine the relationship between sample size, Type I error level, power, and differences in means if one of the objectives of the sampling program is to determine whether a particular mean differs from a group of means. The functions `aovPower`, `aovN`, and `plotAovDesign` can be used to investigate these relationships for the case of normally-distributed observations.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Berthouex, P.M., and L.C. Brown. (1994). *Statistics for Environmental Engineers*. Lewis Publishers, Boca Raton, FL, Chapter 17.
- Helsel, D.R., and R.M. Hirsch. (1992). *Statistical Methods in Water Resources Research*. Elsevier, New York, NY, Chapter 7.
- Johnson, N. L., S. Kotz, and N. Balakrishnan. (1995). *Continuous Univariate Distributions, Volume 2*. Second Edition. John Wiley and Sons, New York, Chapters 27, 29, 30.
- Millard, S.P., and Neerchal, N.K. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, Florida.
- Scheffe, H. (1959). *The Analysis of Variance*. John Wiley and Sons, New York, 477pp.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.6-38.
- Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ, Chapter 10.

See Also

`aovN`, `plotAovDesign`, `Normal`, `aov`.

Examples

```
# Look at how the power of a one-way ANOVA increases
# with increasing sample size:
```

```
aovPower(n.vec = rep(5, 3), mu.vec = c(10, 15, 20), sigma = 5)
#[1] 0.7015083
```

```
aovPower(n.vec = rep(10, 3), mu.vec = c(10, 15, 20), sigma = 5)
#[1] 0.9732551
```

```
#-----
```

```
# Look at how the power of a one-way ANOVA increases
# with increasing variability in the population means:
```

```

aovPower(n.vec = rep(5,3), mu.vec = c(10, 10, 11), sigma=5)
#[1] 0.05795739

aovPower(n.vec = rep(5, 3), mu.vec = c(10, 10, 15), sigma = 5)
#[1] 0.2831863

aovPower(n.vec = rep(5, 3), mu.vec = c(10, 13, 15), sigma = 5)
#[1] 0.2236093

aovPower(n.vec = rep(5, 3), mu.vec = c(10, 15, 20), sigma = 5)
#[1] 0.7015083

#-----

# Look at how the power of a one-way ANOVA increases
# with increasing values of Type I error:

aovPower(n.vec = rep(10,3), mu.vec = c(10, 12, 14),
  sigma = 5, alpha = 0.001)
#[1] 0.02655785

aovPower(n.vec = rep(10,3), mu.vec = c(10, 12, 14),
  sigma = 5, alpha = 0.01)
#[1] 0.1223527

aovPower(n.vec = rep(10,3), mu.vec = c(10, 12, 14),
  sigma = 5, alpha = 0.05)
#[1] 0.3085313

aovPower(n.vec = rep(10,3), mu.vec = c(10, 12, 14),
  sigma = 5, alpha = 0.1)
#[1] 0.4373292

#-----

# The example on pages 5-11 to 5-14 of USEPA (1989b) shows
# log-transformed concentrations of lead (mg/L) at two
# background wells and four compliance wells, where observations
# were taken once per month over four months (the data are
# stored in EPA.89b.loglead.df.) Assume the true mean levels
# at each well are 3.9, 3.9, 4.5, 4.5, 4.5, and 5, respectively.
# Compute the power of a one-way ANOVA to test for mean
# differences between wells. Use alpha=0.05, and assume the
# true standard deviation is equal to the one estimated from
# the data in this example.

# First look at the data
names(EPA.89b.loglead.df)
#[1] "LogLead" "Month" "Well" "Well.type"

windows()
stripChart(LogLead ~ Well, data = EPA.89b.loglead.df,
  show.ci = FALSE, xlab = "Well Number",
  ylab="Log [ Lead (ug/L) ]",
  main="Lead Concentrations at Six Wells")

```

```
# Note: The assumption of a constant variance across
# all wells is suspect.

# Now perform the ANOVA and get the estimated sd
aov.list <- aov(LogLead ~ Well, data=EPA.89b.loglead.df)

summary(aov.list)
#           Df Sum Sq Mean Sq F value    Pr(>F)    
#Well         5  5.7447   1.14895    3.3469  0.02599 * 
#Residuals    18  6.1791   0.34328                      
#---
#Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

# Now call the function aovPower
aovPower(n.vec = rep(4, 6),
  mu.vec = c(3.9,3.9,4.5,4.5,4.5,5), sigma=sqrt(0.34))
#[1] 0.5523148

# Clean up
rm(aov.list)
```

base

*Base b Representation of a Number***Description**

For any number represented in base 10, compute the representation in any user-specified base.

Usage

```
base(n, base = 10, num.digits = max(0, floor(log(n, base))) + 1)
```

Arguments

n	a non-negative integer (base 10).
base	a positive integer greater than 1 indicating what base to represent n in.
num.digits	a positive integer indicating how many digits to use to represent n in base base. By default, num.digits is equal to just the number of required digits (i.e., $\max(0, \text{floor}(\log(n, \text{base}))) + 1$). Setting num.digits to a larger number than this will result in 0's padding the left.

Details

If b is a positive integer greater than 1, and n is a positive integer, then n can be expressed uniquely in the form

$$n = a_k b^k + a_{k-1} b^{k-1} + \dots + a_1 b + a_0$$

where k is a non-negative integer, the coefficients a_0, a_1, \dots, a_k are non-negative integers less than b , and $a_k > 0$ (Rosen, 1988, p.105). The function base computes the coefficients a_0, a_1, \dots, a_k .

Value

A numeric vector of length `num.digits` showing the representation of `n` in base `base`.

Note

The function `base` is included in **EnvStats** because it is called by the function [oneSamplePermutationTest](#).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Rosen, K.H. (1988). *Discrete Mathematics and Its Applications*. Random House, New York, pp.105-107.

See Also

[oneSamplePermutationTest](#).

Examples

```
# Compute the value of 7 in base 2.

base(7, 2)
#[1] 1 1 1

base(7, 2, num.digits=5)
#[1] 0 0 1 1 1
```

Benthic.df

Benthic Data from Monitoring Program in Chesapeake Bay

Description

Benthic data from a monitoring program in the Chesapeake Bay, Maryland, covering July 1994 - December 1991.

Usage

```
Benthic.df
```

Format

A data frame with 585 observations on the following 7 variables.

Site.ID Site ID

Stratum Stratum Number (101-131)

Latitude Latitude (degrees North)

Longitude Longitude (negative values; degrees West)

Index Benthic Index (between 1 and 5)

Salinity Salinity (ppt)

Silt Silt Content (% clay in soil)

Details

Data from the Long Term Benthic Monitoring Program of the Chesapeake Bay. The data consist of measurements of benthic characteristics and a computed index of benthic health for several locations in the bay. Sampling methods and designs of the program are discussed in Ranasinghe et al. (1992).

The data represent observations collected at 585 separate point locations (sites). The sites are divided into 31 different strata, numbered 101 through 131, each strata consisting of geographically close sites of similar degradation conditions. The benthic index values range from 1 to 5 on a continuous scale, where high values correspond to healthier benthos. Salinity was measured in parts per thousand (ppt), and silt content is expressed as a percentage of clay in the soil with high numbers corresponding to muddy areas.

The United States Environmental Protection Agency (USEPA) established an initiative for the Chesapeake Bay in partnership with the states bordering the bay in 1984. The goal of the initiative is the restoration (abundance, health, and diversity) of living resources to the bay by reducing nutrient loadings, reducing toxic chemical impacts, and enhancing habitats. USEPA's Chesapeake Bay Program Office is responsible for implementing this initiative and has established an extensive monitoring program that includes traditional water chemistry sampling, as well as collecting data on living resources to measure progress towards meeting the restoration goals.

Sampling benthic invertebrate assemblages has been an integral part of the Chesapeake Bay monitoring program due to their ecological importance and their value as biological indicators. The condition of benthic assemblages is a measure of the ecological health of the bay, including the effects of multiple types of environmental stresses. Nevertheless, regional-scale assessment of ecological status and trends using benthic assemblages are limited by the fact that benthic assemblages are strongly influenced by naturally variable habitat elements, such as salinity, sediment type, and depth. Also, different state agencies and USEPA programs use different sampling methodologies, limiting the ability to integrate data into a unified assessment. To circumvent these limitations, USEPA has standardized benthic data from several different monitoring programs into a single database, and from that database developed a Restoration Goals Benthic Index that identifies whether benthic restoration goals are being met.

Source

Ranasinghe, J.A., L.C. Scott, and R. Newport. (1992). *Long-term Benthic Monitoring and Assessment Program for the Maryland Portion of the Bay*, Jul 1984-Dec 1991. Report prepared for the Maryland Department of the Environment and the Maryland Department of Natural Resources by Versar, Inc., Columbia, MD.

Examples

```
attach(Benthic.df)

# Show station locations
#-----
windows()
plot(Longitude, Latitude,
      xlab = "-Longitude (Degrees West)",
      ylab = "Latitude",
      main = "Sampling Station Locations")
```

```

# Scatterplot matrix of benthic index, salinity, and silt
#-----
windows()
pairs(~ Index + Salinity + Silt, data = Benthic.df)

# Contour and perspective plots based on loess fit
# showing only predicted values within the convex hull
# of station locations
#-----
library(sp)

loess.fit <- loess(Index ~ Longitude * Latitude,
  data=Benthic.df, normalize=FALSE, span=0.25)
lat <- Benthic.df$Latitude
lon <- Benthic.df$Longitude
Latitude <- seq(min(lat), max(lat), length=50)
Longitude <- seq(min(lon), max(lon), length=50)
predict.list <- list(Longitude=Longitude,
  Latitude=Latitude)
predict.grid <- expand.grid(predict.list)
predict.fit <- predict(loess.fit, predict.grid)
index.chull <- chull(lon, lat)
inside <- point.in.polygon(point.x = predict.grid$Longitude,
  point.y = predict.grid$Latitude,
  pol.x = lon[index.chull],
  pol.y = lat[index.chull])
predict.fit[inside == 0] <- NA

windows()
contour(Longitude, Latitude, predict.fit,
  levels=seq(1, 5, by=0.5), labcex=0.75,
  xlab="-Longitude (degrees West)",
  ylab="Latitude (degrees North)")
title(main=paste("Contour Plot of Benthic Index",
  "Based on Loess Smooth", sep="\n"))

windows()
persp(Longitude, Latitude, predict.fit,
  xlim = c(-77.3, -75.9), ylim = c(38.1, 39.5), zlim = c(0, 6),
  theta = -45, phi = 30, d = 0.5,
  xlab="-Longitude (degrees West)",
  ylab="Latitude (degrees North)",
  zlab="Benthic Index", ticktype = "detailed")
title(main=paste("Surface Plot of Benthic Index",
  "Based on Loess Smooth", sep="\n"))

detach("Benthic.df")

rm(loess.fit, lat, lon, Latitude, Longitude, predict.list,
  predict.grid, predict.fit, index.chull, inside)

```

Description

boxcox is a generic function used to compute the value(s) of an objective for one or more Box-Cox power transformations, or to compute an optimal power transformation based on a specified objective. The function invokes particular [methods](#) which depend on the [class](#) of the first argument.

Currently, there is a default method and a method for objects of class "lm".

Usage

```
boxcox(x, ...)

## Default S3 method:
boxcox(x,
  lambda = {if (optimize) c(-2, 2) else seq(-2, 2, by = 0.5)},
  optimize = FALSE, objective.name = "PPCC",
  eps = .Machine$double.eps, include.x = TRUE, ...)

## S3 method for class 'lm'
boxcox(x,
  lambda = {if (optimize) c(-2, 2) else seq(-2, 2, by = 0.5)},
  optimize = FALSE, objective.name = "PPCC",
  eps = .Machine$double.eps, include.x = TRUE, ...)
```

Arguments

x	an object of class "lm" for which the response variable is all positive numbers, or else a numeric vector of positive numbers. When x is an object of class "lm", the object must have been created with a call to the function lm that includes the data argument. When x is a numeric vector of positive observations, missing (NA), undefined (NaN), and infinite (-Inf, Inf) values are allowed but will be removed.
lambda	numeric vector of finite values indicating what powers to use for the Box-Cox transformation. When optimize=FALSE, the default value is lambda=seq(-2, 2, by=0.5). When optimize=TRUE, lambda must be a vector with two values indicating the range over which the optimization will occur and the range of these two values must include 1. In this case, the default value is lambda=c(-2, 2).
optimize	logical scalar indicating whether to simply evaluate the objective function at the given values of lambda (optimize=FALSE; the default), or to compute the optimal power transformation within the bounds specified by lambda (optimize=TRUE).
objective.name	character string indicating what objective to use. The possible values are "PPCC" (probability plot correlation coefficient; the default), "Shapiro-Wilk" (the Shapiro-Wilk goodness-of-fit statistic), and "Log-Likelihood" (the log-likelihood function).
eps	finite, positive numeric scalar. When the absolute value of lambda is less than eps, lambda is assumed to be 0 for the Box-Cox transformation. The default value is eps=.Machine\$double.eps.
include.x	logical scalar indicating whether to include the finite, non-missing values of the argument x with the returned object. The default value is include.x=TRUE.
...	optional arguments for possible future methods. Currently not used.

Details

Two common assumptions for several standard parametric hypothesis tests are:

1. The observations all come from a normal distribution.
2. The observations all come from distributions with the same variance.

For example, the standard one-sample t-test assumes all the observations come from the same normal distribution, and the standard two-sample t-test assumes that all the observations come from a normal distribution with the same variance, although the mean may differ between the two groups.

When the original data do not satisfy the above assumptions, data transformations are often used to attempt to satisfy these assumptions. The rest of this section is divided into two parts: one that discusses Box-Cox transformations in the context of the original observations, and one that discusses Box-Cox transformations in the context of linear models.

Box-Cox Transformations Based on the Original Observations

Box and Cox (1964) presented a formalized method for deciding on a data transformation. Given a random variable X from some distribution with only positive values, the Box-Cox family of power transformations is defined as:

$$\begin{aligned} Y &= \frac{X^\lambda - 1}{\lambda} & \lambda \neq 0 \\ \log(X) & & \lambda = 0 \end{aligned} \quad (1)$$

where Y is assumed to come from a normal distribution. This transformation is continuous in λ . Note that this transformation also preserves ordering. See the help file for [boxcoxTransform](#) for more information on data transformations.

Let $\underline{x} = x_1, x_2, \dots, x_n$ denote a random sample of n observations from some distribution and assume that there exists some value of λ such that the transformed observations

$$\begin{aligned} y_i &= \frac{x_i^\lambda - 1}{\lambda} & \lambda \neq 0 \\ \log(x_i) & & \lambda = 0 \end{aligned} \quad (2)$$

($i = 1, 2, \dots, n$) form a random sample from a normal distribution.

Box and Cox (1964) proposed choosing the appropriate value of λ based on maximizing the likelihood function. Alternatively, an appropriate value of λ can be chosen based on another objective, such as maximizing the probability plot correlation coefficient or the Shapiro-Wilk goodness-of-fit statistic.

In the case when `optimize=TRUE`, the function `boxcox` calls the R function `nlminb` to minimize the negative value of the objective (i.e., maximize the objective) over the range of possible values of λ specified in the argument `lambda`. The starting value for the optimization is always $\lambda = 1$ (i.e., no transformation).

The rest of this sub-section explains how the objective is computed for the various options for `objective.name`.

Objective Based on Probability Plot Correlation Coefficient (`objective.name="PPCC"`)

When `objective.name="PPCC"`, the objective is computed as the value of the normal probability plot correlation coefficient based on the transformed data (see the description of the Probability Plot Correlation Coefficient (PPCC) goodness-of-fit test in the help file for [gofTest](#)). That is, the

objective is the correlation coefficient for the normal [quantile-quantile plot](#) for the transformed data. Large values of the PPCC tend to indicate a good fit to a normal distribution.

Objective Based on Shapiro-Wilk Goodness-of-Fit Statistic (`objective.name="Shapiro-Wilk"`)
When `objective.name="Shapiro-Wilk"`, the objective is computed as the value of the Shapiro-Wilk goodness-of-fit statistic based on the transformed data (see the description of the Shapiro-Wilk test in the help file for [gofTest](#)). Large values of the Shapiro-Wilk statistic tend to indicate a good fit to a normal distribution.

Objective Based on Log-Likelihood Function (`objective.name="Log-Likelihood"`)
When `objective.name="Log-Likelihood"`, the objective is computed as the value of the log-likelihood function. Assuming the transformed observations in Equation (2) above come from a normal distribution with mean μ and standard deviation σ , we can use the change of variable formula to write the log-likelihood function as:

$$\log[L(\lambda, \mu, \sigma)] = \frac{-n}{2} \log(2\pi) - \frac{n}{2} \log(\sigma^2) - \frac{1}{2\sigma^2} \sum_{i=1}^n (y_i - \mu)^2 + (\lambda - 1) \sum_{i=1}^n \log(x_i) \quad (3)$$

where y_i is defined in Equation (2) above (Box and Cox, 1964). For a fixed value of λ , the log-likelihood function is maximized by replacing μ and σ with their maximum likelihood estimators:

$$\hat{\mu} = \frac{1}{n} \sum_{i=1}^n y_i \quad (4)$$

$$\hat{\sigma} = \left[\frac{1}{n} \sum_{i=1}^n (y_i - \bar{y})^2 \right]^{1/2} \quad (5)$$

Thus, when `optimize=TRUE`, Equation (3) is maximized by iteratively solving for λ using the values for μ and σ given in Equations (4) and (5). When `optimize=FALSE`, the value of the objective is computed by using Equation (3), using the values of λ specified in the argument `lambda`, and using the values for μ and σ given in Equations (4) and (5).

Box-Cox Transformation for Linear Models

In the case of a standard linear regression model with n observations and p predictors:

$$Y_i = \beta_0 + \beta_1 X_{i1} + \dots + \beta_p X_{ip} + \epsilon_i, \quad i = 1, 2, \dots, n \quad (6)$$

the standard assumptions are:

1. The error terms ϵ_i come from a normal distribution with mean 0.
2. The variance is the same for all of the error terms and does not depend on the predictor variables.

Assuming Y is a random variable from some distribution that may depend on the predictor variables and Y takes on only positive values, the Box-Cox family of power transformations is defined as:

$$\begin{aligned} Y^* &= \frac{Y^\lambda - 1}{\lambda} & \lambda \neq 0 \\ \log(Y) & & \lambda = 0 \end{aligned} \quad (7)$$

where Y^* becomes the new response variable and the errors are now assumed to come from a normal distribution with a mean of 0 and a constant variance.

In this case, the objective is computed as described above, but it is based on the residuals from the fitted linear model in which the response variable is now Y^* instead of Y .

Value

When x is an object of class "lm", boxcox returns a list of class "boxcoxLm" containing the results. See the help file for [boxcoxLm.object](#) for details.

When x is simply a numeric vector of positive numbers, boxcox returns a list of class "boxcox" containing the results. See the help file for [boxcox.object](#) for details.

Note

Data transformations are often used to induce normality, homoscedasticity, and/or linearity, common assumptions of parametric statistical tests and estimation procedures. Transformations are not "tricks" used by the data analyst to hide what is going on, but rather useful tools for understanding and dealing with data (Berthouex and Brown, 2002, p.61). Hoaglin (1988) discusses "hidden" transformations that are used everyday, such as the pH scale for measuring acidity. Johnson and Wichern (2007, p.192) note that "Transformations are nothing more than a reexpression of the data in different units."

In the case of a linear model, there are at least two approaches to improving a model fit: transform the Y and/or X variable(s), and/or use more predictor variables. Often in environmental data analysis, we assume the observations come from a lognormal distribution and automatically take logarithms of the data. For a simple linear regression (i.e., one predictor variable), if regression diagnostic plots indicate that a straight line fit is not adequate, but that the variance of the errors appears to be fairly constant, you may only need to transform the predictor variable X or perhaps use a quadratic or cubic model in X . On the other hand, if the diagnostic plots indicate that the constant variance and/or normality assumptions are suspect, you probably need to consider transforming the response variable Y . Data transformations for linear regression models are discussed in Draper and Smith (1998, Chapter 13) and Helsel and Hirsch (1992, pp. 228-229).

One problem with data transformations is that translating results on the transformed scale back to the original scale is not always straightforward. Estimating quantities such as means, variances, and confidence limits in the transformed scale and then transforming them back to the original scale usually leads to biased and inconsistent estimates (Gilbert, 1987, p.149; van Belle et al., 2004, p.400). For example, exponentiating the confidence limits for a mean based on log-transformed data does not yield a confidence interval for the mean on the original scale. Instead, this yields a confidence interval for the median (see the help file for [elnormAlt](#)). It should be noted, however, that quantiles (percentiles) and rank-based procedures are invariant to monotonic transformations (Helsel and Hirsch, 1992, p.12).

Finally, there is no guarantee that a Box-Cox transformation based on the "optimal" value of λ will provide an adequate transformation to allow the assumption of approximate normality and constant variance. Any set of transformed data should be inspected relative to the assumptions you want to make about it (Johnson and Wichern, 2007, p.194).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Berthouex, P.M., and L.C. Brown. (2002). *Statistics for Environmental Engineers, Second Edition*. Lewis Publishers, Boca Raton, FL.

- Box, G.E.P., and D.R. Cox. (1964). An Analysis of Transformations (with Discussion). *Journal of the Royal Statistical Society, Series B* **26**(2), 211–252.
- Draper, N., and H. Smith. (1998). *Applied Regression Analysis*. Third Edition. John Wiley and Sons, New York, pp.47–53.
- Gilbert, R.O. (1987). *Statistical Methods for Environmental Pollution Monitoring*. Van Nostrand Reinhold, NY.
- Helsel, D.R., and R.M. Hirsch. (1992). *Statistical Methods in Water Resources Research*. Elsevier, New York, NY.
- Hinkley, D.V., and G. Runger. (1984). The Analysis of Transformed Data (with Discussion). *Journal of the American Statistical Association* **79**, 302–320.
- Hoaglin, D.C., F.M. Mosteller, and J.W. Tukey, eds. (1983). *Understanding Robust and Exploratory Data Analysis*. John Wiley and Sons, New York, Chapter 4.
- Hoaglin, D.C. (1988). Transformations in Everyday Experience. *Chance* **1**, 40–45.
- Johnson, N. L., S. Kotz, and A.W. Kemp. (1992). *Univariate Discrete Distributions, Second Edition*. John Wiley and Sons, New York, p.163.
- Johnson, R.A., and D.W. Wichern. (2007). *Applied Multivariate Statistical Analysis, Sixth Edition*. Pearson Prentice Hall, Upper Saddle River, NJ, pp.192–195.
- Shumway, R.H., A.S. Azari, and P. Johnson. (1989). Estimating Mean Concentrations Under Transformations for Environmental Data With Detection Limits. *Technometrics* **31**(3), 347–356.
- Stoline, M.R. (1991). An Examination of the Lognormal and Box and Cox Family of Transformations in Fitting Environmental Data. *Environmetrics* **2**(1), 85–106.
- van Belle, G., L.D. Fisher, Heagerty, P.J., and Lumley, T. (2004). *Biostatistics: A Methodology for the Health Sciences, 2nd Edition*. John Wiley & Sons, New York.
- Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ, Chapter 13.

See Also

[boxcox.object](#), [plot.boxcox](#), [print.boxcox](#), [boxcoxLm.object](#), [plot.boxcoxLm](#), [print.boxcoxLm](#), [boxcoxTransform](#), [Data Transformations](#), [Goodness-of-Fit Tests](#).

Examples

```
# Generate 30 observations from a lognormal distribution with
# mean=10 and cv=2. Look at some values of various objectives
# for various transformations. Note that for both the PPCC and
# the Log-Likelihood objective, the optimal value of lambda is
# about 0, indicating that a log transformation is appropriate.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
x <- rlnormAlt(30, mean = 10, cv = 2)

# Using the PPCC objective:
#-----

boxcox(x)
#Results of Box-Cox Transformation
#-----
#
#Objective Name:                PPCC
```

```

#
#Data:                                x
#
#Sample Size:                         30
#
# lambda      PPCC
#   -2.0 0.5423739
#   -1.5 0.6402782
#   -1.0 0.7818160
#   -0.5 0.9272219
#    0.0 0.9921702
#    0.5 0.9581178
#    1.0 0.8749611
#    1.5 0.7827009
#    2.0 0.7004547

boxcox(x, optimize = TRUE)
#Results of Box-Cox Transformation
#-----
#
#Objective Name:                      PPCC
#
#Data:                                x
#
#Sample Size:                         30
#
#Bounds for Optimization:             lower = -2
#                                     upper = 2
#
#Optimal Value:                      lambda = 0.04530789
#
#Value of Objective:                 PPCC = 0.9925919

# Using the Log-Likelihood objective
#-----

boxcox(x, objective.name = "Log-Likelihood")
#Results of Box-Cox Transformation
#-----
#
#Objective Name:                      Log-Likelihood
#
#Data:                                x
#
#Sample Size:                         30
#
# lambda Log-Likelihood
#   -2.0 -154.94255
#   -1.5 -128.59988
#   -1.0 -106.23882
#   -0.5 -90.84800
#    0.0 -85.10204
#    0.5 -88.69825
#    1.0 -99.42630
#    1.5 -115.23701
#    2.0 -134.54125

```

```

boxcox(x, objective.name = "Log-Likelihood", optimize = TRUE)
#Results of Box-Cox Transformation
#-----
#
#Objective Name:          Log-Likelihood
#
#Data:                   x
#
#Sample Size:            30
#
#Bounds for Optimization: lower = -2
#                          upper = 2
#
#Optimal Value:          lambda = 0.0405156
#
#Value of Objective:      Log-Likelihood = -85.07123
#-----

# Plot the results based on the PPCC objective
#-----
boxcox.list <- boxcox(x)
windows()
plot(boxcox.list)

#Look at QQ-Plots for the candidate values of lambda
#-----
plot(boxcox.list, plot.type = "Q-Q Plots", same.window = FALSE)

#=====

# The data frame Environmental.df contains daily measurements of
# ozone concentration, wind speed, temperature, and solar radiation
# in New York City for 153 consecutive days between May 1 and
# September 30, 1973. In this example, we'll plot ozone vs.
# temperature and look at the Q-Q plot of the residuals. Then
# we'll look at possible Box-Cox transformations. The "optimal" one
# based on the PPCC looks close to a log-transformation
# (i.e., lambda=0). The power that produces the largest PPCC is
# about 0.2, so a cube root (lambda=1/3) transformation might work too.

# Fit the model with the raw Ozone data
#-----
ozone.fit <- lm(ozone ~ temperature, data = Environmental.df)

# Plot Ozone vs. Temperature, with fitted line
#-----
windows()
with(Environmental.df,
  plot(temperature, ozone, xlab = "Temperature (degrees F)",
    ylab = "Ozone (ppb)", main = "Ozone vs. Temperature"))
abline(ozone.fit)

# Look at the Q-Q Plot for the residuals
#-----
windows()

```

```

qqPlot(ozone.fit$residuals, add.line = TRUE)

# Look at Box-Cox transformations of Ozone
#-----
boxcox.list <- boxcox(ozone.fit)
boxcox.list
#Results of Box-Cox Transformation
#-----
#
#Objective Name:                PPCC
#
#Linear Model:                  ozone.fit
#
#Sample Size:                   116
#
# lambda      PPCC
#   -2.0 0.4286781
#   -1.5 0.4673544
#   -1.0 0.5896132
#   -0.5 0.8301458
#    0.0 0.9871519
#    0.5 0.9819825
#    1.0 0.9408694
#    1.5 0.8840770
#    2.0 0.8213675

# Plot PPCC vs. lambda based on Q-Q plots of residuals
#-----
windows()
plot(boxcox.list)

# Look at Q-Q plots of residuals for the various transformation
#-----
plot(boxcox.list, plot.type = "Q-Q Plots", same.window = FALSE)

# Compute the "optimal" transformation
#-----
boxcox(ozone.fit, optimize = TRUE)
#Results of Box-Cox Transformation
#-----
#
#Objective Name:                PPCC
#
#Linear Model:                  ozone.fit
#
#Sample Size:                   116
#
#Bounds for Optimization:       lower = -2
#                                upper = 2
#
#Optimal Value:                  lambda = 0.2004305
#
#Value of Objective:             PPCC = 0.9940222

#=====

# Clean up

```

```
#-----
rm(x, boxcox.list, ozone.fit)
graphics.off()
```

boxcox.object

S3 Class "boxcox"

Description

Objects of S3 class "boxcox" are returned by the **EnvStats** function [boxcox](#), which computes objective values for user-specified powers, or computes the optimal power for the specified objective.

Details

Objects of class "boxcox" are lists that contain information about the powers that were used, the objective that was used, the values of the objective for the given powers, and whether an optimization was specified.

Value

Required Components

The following components must be included in a legitimate list of class "boxcox".

lambda	Numeric vector containing the powers used in the Box-Cox transformations. If the value of the optimize component is FALSE, then lambda contains the values of all of the powers at which the objective was evaluated. If the value of the optimize component is TRUE, then lambda is a scalar containing the value of the power that maximizes the objective.
objective	Numeric vector containing the value(s) of the objective for the given value(s) of λ that are stored in the component lambda.
objective.name	character string indicating the objective that was used. The possible values are "PPCC" (probability plot correlation coefficient; the default), "Shapiro-Wilk" (the Shapiro-Wilk goodness-of-fit statistic), and "Log-Likelihood" (the log-likelihood function).
optimize	logical scalar indicating whether the objective was simply evaluated at the given values of lambda (optimize=FALSE), or instead the optimal power transformation was computed within the bounds specified by lambda (optimize=TRUE).
optimize.bounds	Numeric vector of length 2 with a names attribute indicating the bounds within which the optimization took place. When optimize=FALSE, this contains missing values.
eps	finite, positive numeric scalar indicating what value of eps was used. When the absolute value of lambda is less than eps, lambda is assumed to be 0 for the Box-Cox transformation.
sample.size	Numeric scalar indicating the number of finite, non-missing observations.
data.name	The name of the data object used for the Box-Cox computations.
bad.obs	The number of missing (NA), undefined (NaN) and/or infinite (Inf, -Inf) values that were removed from the data object prior to performing the Box-Cox computations.

Optional Component

The following component may optionally be included in a legitimate list of class "boxcox". It must be included if you want to call the function `plot.boxcox` and specify Q-Q plots or Tukey Mean-Difference Q-Q plots.

`data` Numeric vector containing the data actually used for the Box-Cox computations (i.e., the original data without any missing or infinite values).

Methods

Generic functions that have methods for objects of class "boxcox" include:

`link{plot}`, `print`.

Note

Since objects of class "boxcox" are lists, you may extract their components with the `$` and `[[` operators.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

See Also

`boxcox`, `plot.boxcox`, `print.boxcox`, `boxcoxLm.object`.

Examples

```
# Create an object of class "boxcox", then print it out.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
x <- rlnormAlt(30, mean = 10, cv = 2)

boxcox.list <- boxcox(x)

data.class(boxcox.list)
#[1] "boxcox"

names(boxcox.list)
# [1] "lambda"           "objective"         "objective.name"
# [4] "optimize"         "optimize.bounds"   "eps"
# [7] "data"             "sample.size"       "data.name"
#[10] "bad.obs"

boxcox.list
#Results of Box-Cox Transformation
#-----
#
#Objective Name:                PPCC
#
#Data:                          x
#
#Sample Size:                   30
#
# lambda      PPCC
```



```

# -2.0 0.5423739
# -1.5 0.6402782
# -1.0 0.7818160
# -0.5 0.9272219
# 0.0 0.9921702
# 0.5 0.9581178
# 1.0 0.8749611
# 1.5 0.7827009
# 2.0 0.7004547

boxcox(x, optimize = TRUE)
#Results of Box-Cox Transformation
#-----
#
#Objective Name:          PPCC
#
#Data:                   x
#
#Sample Size:            30
#
#Bounds for Optimization: lower = -2
#                          upper = 2
#
#Optimal Value:           lambda = 0.04530789
#
#Value of Objective:      PPCC = 0.9925919

#-----

# Clean up
#-----
rm(x, boxcox.list)

```

boxcoxCensored

*Boxcox Power Transformation for Type I Censored Data***Description**

Compute the value(s) of an objective for one or more Box-Cox power transformations, or to compute an optimal power transformation based on a specified objective, based on Type I censored data.

Usage

```

boxcoxCensored(x, censored, censoring.side = "left",
  lambda = {if (optimize) c(-2, 2) else seq(-2, 2, by = 0.5)}, optimize = FALSE,
  objective.name = "PPCC", eps = .Machine$double.eps,
  include.x.and.censored = TRUE, prob.method = "michael-schucany",
  plot.pos.con = 0.375)

```

Arguments

x a numeric vector of positive numbers. Missing (NA), undefined (NaN), and infinite (-Inf, Inf) values are allowed but will be removed.

censored	numeric or logical vector indicating which values of <code>x</code> are censored. This must be the same length as <code>x</code> . If the mode of <code>censored</code> is "logical", TRUE values correspond to elements of <code>x</code> that are censored, and FALSE values correspond to elements of <code>x</code> that are not censored. If the mode of <code>censored</code> is "numeric", it must contain only 1's and 0's; 1 corresponds to TRUE and 0 corresponds to FALSE. Missing (NA) values are allowed but will be removed.
censoring.side	character string indicating on which side the censoring occurs. The possible values are "left" (the default) and "right".
lambda	numeric vector of finite values indicating what powers to use for the Box-Cox transformation. When <code>optimize=FALSE</code> , the default value is <code>lambda=seq(-2, 2, by=0.5)</code> . When <code>optimize=TRUE</code> , <code>lambda</code> must be a vector with two values indicating the range over which the optimization will occur and the range of these two values must include 1. In this case, the default value is <code>lambda=c(-2, 2)</code> .
optimize	logical scalar indicating whether to simply evaluate the objective function at the given values of <code>lambda</code> (<code>optimize=FALSE</code> ; the default), or to compute the optimal power transformation within the bounds specified by <code>lambda</code> (<code>optimize=TRUE</code>).
objective.name	character string indicating what objective to use. The possible values are "PPCC" (probability plot correlation coefficient; the default), "Shapiro-Wilk" (the Shapiro-Wilk goodness-of-fit statistic), and "Log-Likelihood" (the log-likelihood function).
eps	finite, positive numeric scalar. When the absolute value of <code>lambda</code> is less than <code>eps</code> , <code>lambda</code> is assumed to be 0 for the Box-Cox transformation. The default value is <code>eps=.Machine\$double.eps</code> .
include.x.and.censored	logical scalar indicating whether to include the finite, non-missing values of the argument <code>x</code> and the corresponding values of <code>censored</code> with the returned object. The default value is <code>include.x.and.censored=TRUE</code> .
prob.method	for multiply censored data, character string indicating what method to use to compute the plotting positions (empirical probabilities) when <code>objective.name="PPCC"</code> . Possible values are "kaplan-meier" (product-limit method of Kaplan and Meier (1958)), "kaplan-meier with max" (product-limit method of Kaplan and Meier (1958) with the maximum value included), "nelson" (hazard plotting method of Nelson (1972)), "michael-schucany" (generalization of the product-limit method due to Michael and Schucany (1986)), and "hirsch-stedinger" (generalization of the product-limit method due to Hirsch and Stedinger (1987)). The default value is <code>prob.method="michael-schucany"</code> . The "nelson" method is only available for <code>censoring.side="right"</code> . See the DETAILS section for more explanation. This argument is ignored if <code>objective.name</code> is not equal to "PPCC" and/or the data are singly censored.
plot.pos.con	for multiply censored data, numeric scalar between 0 and 1 containing the value of the plotting position constant when <code>objective.name="PPCC"</code> . The default value is <code>plot.pos.con=0.375</code> . See the DETAILS section for more information. This argument is used only if <code>prob.method</code> is equal to "michael-schucany" or "hirsch-stedinger". This argument is ignored if <code>objective.name</code> is not equal to "PPCC" and/or the data are singly censored.

Details

Two common assumptions for several standard parametric hypothesis tests are:

1. The observations all come from a normal distribution.
2. The observations all come from distributions with the same variance.

For example, the standard one-sample t-test assumes all the observations come from the same normal distribution, and the standard two-sample t-test assumes that all the observations come from a normal distribution with the same variance, although the mean may differ between the two groups.

When the original data do not satisfy the above assumptions, data transformations are often used to attempt to satisfy these assumptions. Box and Cox (1964) presented a formalized method for deciding on a data transformation. Given a random variable X from some distribution with only positive values, the Box-Cox family of power transformations is defined as:

$$Y = \begin{cases} \frac{X^\lambda - 1}{\lambda} & \lambda \neq 0 \\ \log(X) & \lambda = 0 \end{cases} \quad (1)$$

where Y is assumed to come from a normal distribution. This transformation is continuous in λ . Note that this transformation also preserves ordering. See the help file for `boxcoxTransform` for more information on data transformations.

Box and Cox (1964) proposed choosing the appropriate value of λ based on maximizing the likelihood function. Alternatively, an appropriate value of λ can be chosen based on another objective, such as maximizing the probability plot correlation coefficient or the Shapiro-Wilk goodness-of-fit statistic.

Shumway et al. (1989) investigated extending the method of Box and Cox (1964) to the case of Type I censored data, motivated by the desire to produce estimated means and confidence intervals for air monitoring data that included censored values.

In the case when `optimize=TRUE`, the function `boxcoxCensored` calls the R function `nlminb` to minimize the negative value of the objective (i.e., maximize the objective) over the range of possible values of λ specified in the argument `lambda`. The starting value for the optimization is always $\lambda = 1$ (i.e., no transformation).

The next section explains assumptions and notation, and the section after that explains how the objective is computed for the various options for `objective.name`.

Assumptions and Notation

Let \underline{x} denote a random sample of N observations from some continuous distribution. Assume n ($0 < n < N$) of these observations are known and c ($c = N - n$) of these observations are all censored below (left-censored) or all censored above (right-censored) at k fixed censoring levels

$$T_1, T_2, \dots, T_K; \quad K \geq 1 \quad (2)$$

For the case when $K \geq 2$, the data are said to be Type I **multiply censored**. For the case when $K = 1$, set $T = T_1$. If the data are left-censored and all n known observations are greater than or equal to T , or if the data are right-censored and all n known observations are less than or equal to T , then the data are said to be Type I **singly censored** (Nelson, 1982, p.7), otherwise they are considered to be Type I multiply censored.

Let c_j denote the number of observations censored below or above censoring level T_j for $j = 1, 2, \dots, K$, so that

$$\sum_{j=1}^K c_j = c \quad (3)$$

Let $x_{(1)}, x_{(2)}, \dots, x_{(N)}$ denote the “ordered” observations, where now “observation” means either the actual observation (for uncensored observations) or the censoring level (for censored observations). For right-censored data, if a censored observation has the same value as an uncensored one,

the uncensored observation should be placed first. For left-censored data, if a censored observation has the same value as an uncensored one, the censored observation should be placed first.

Note that in this case the quantity $x_{(i)}$ does not necessarily represent the i 'th "largest" observation from the (unknown) complete sample.

Finally, let Ω (omega) denote the set of n subscripts in the "ordered" sample that correspond to uncensored observations, and let Ω_j denote the set of c_j subscripts in the "ordered" sample that correspond to the censored observations censored at censoring level T_j for $j = 1, 2, \dots, k$.

We assume that there exists some value of λ such that the transformed observations

$$\begin{aligned} y_i &= \frac{x_i^\lambda - 1}{\lambda} & \lambda \neq 0 \\ \log(x_i) & & \lambda = 0 \end{aligned} \quad (4)$$

($i = 1, 2, \dots, n$) form a random sample of Type I censored data from a normal distribution.

Note that for the censored observations, Equation (4) becomes:

$$\begin{aligned} y_{(i)} = T_j^* &= \frac{T_j^\lambda - 1}{\lambda} & \lambda \neq 0 \\ \log(T_j) & & \lambda = 0 \end{aligned} \quad (5)$$

where $i \in \Omega_j$.

Computing the Objective

Objective Based on Probability Plot Correlation Coefficient (objective.name="PPCC")

When objective.name="PPCC", the objective is computed as the value of the normal probability plot correlation coefficient based on the transformed data (see the description of the Probability Plot Correlation Coefficient (PPCC) goodness-of-fit test in the help file for [gofTestCensored](#)). That is, the objective is the correlation coefficient for the normal [quantile-quantile plot](#) for the transformed data. Large values of the PPCC tend to indicate a good fit to a normal distribution.

Objective Based on Shapiro-Wilk Goodness-of-Fit Statistic (objective.name="Shapiro-Wilk")

When objective.name="Shapiro-Wilk", the objective is computed as the value of the Shapiro-Wilk goodness-of-fit statistic based on the transformed data (see the description of the Shapiro-Wilk test in the help file for [gofTestCensored](#)). Large values of the Shapiro-Wilk statistic tend to indicate a good fit to a normal distribution.

Objective Based on Log-Likelihood Function (objective.name="Log-Likelihood")

When objective.name="Log-Likelihood", the objective is computed as the value of the log-likelihood function. Assuming the transformed observations in Equation (4) above come from a normal distribution with mean μ and standard deviation σ , we can use the change of variable formula to write the log-likelihood function as follows.

For Type I left censored data, the likelihood function is given by:

$$\log[L(\lambda, \mu, \sigma)] = \log\left[\binom{N}{c_1 c_2 \dots c_k n}\right] + \sum_{j=1}^k c_j \log[F(T_j^*)] + \sum_{i \in \Omega} \log\{f[y_{(i)}]\} + (\lambda - 1) \sum_{i \in \Omega} \log[x_{(i)}] \quad (6)$$

where f and F denote the probability density function (pdf) and cumulative distribution function (cdf) of the population. That is,

$$f(t) = \phi\left(\frac{t - \mu}{\sigma}\right) \quad (7)$$

$$F(t) = \Phi\left(\frac{t - \mu}{\sigma}\right) \quad (8)$$

where ϕ and Φ denote the pdf and cdf of the standard normal distribution, respectively (Shumway et al., 1989). For left singly censored data, Equation (6) simplifies to:

$$\log[L(\lambda, \mu, \sigma)] = \log\left[\binom{N}{c}\right] + c \log[F(T^*)] + \sum_{i=c+1}^N \log\{f[y_{(i)}]\} + (\lambda - 1) \sum_{i=c+1}^N \log[x_{(i)}] \quad (9)$$

Similarly, for Type I right censored data, the likelihood function is given by:

$$\log[L(\lambda, \mu, \sigma)] = \log\left[\binom{N}{c_1 c_2 \dots c_k n}\right] + \sum_{j=1}^k c_j \log[1 - F(T_j^*)] + \sum_{i \in \Omega} \log\{f[y_{(i)}]\} + (\lambda - 1) \sum_{i \in \Omega} \log[x_{(i)}] \quad (10)$$

and for right singly censored data this simplifies to:

$$\log[L(\lambda, \mu, \sigma)] = \log\left[\binom{N}{c}\right] + c \log[1 - F(T^*)] + \sum_{i=1}^n \log\{f[y_{(i)}]\} + (\lambda - 1) \sum_{i=1}^n \log[x_{(i)}] \quad (11)$$

For a fixed value of λ , the log-likelihood function is maximized by replacing μ and σ with their maximum likelihood estimators (see the section *Maximum Likelihood Estimation* in the help file for [enormCensored](#)).

Thus, when `optimize=TRUE`, Equation (6) or (10) is maximized by iteratively solving for λ using the MLEs for μ and σ . When `optimize=FALSE`, the value of the objective is computed by using Equation (6) or (10), using the values of λ specified in the argument `lambda`, and using the MLEs of μ and σ .

Value

`boxcoxCensored` returns a list of class "boxcoxCensored" containing the results. See the help file for [boxcoxCensored.object](#) for details.

Note

Data transformations are often used to induce normality, homoscedasticity, and/or linearity, common assumptions of parametric statistical tests and estimation procedures. Transformations are not "tricks" used by the data analyst to hide what is going on, but rather useful tools for understanding and dealing with data (Berthouex and Brown, 2002, p.61). Hoaglin (1988) discusses "hidden" transformations that are used everyday, such as the pH scale for measuring acidity. Johnson and Wichern (2007, p.192) note that "Transformations are nothing more than a reexpression of the data in different units."

Shumway et al. (1989) investigated extending the method of Box and Cox (1964) to the case of Type I censored data, motivated by the desire to produce estimated means and confidence intervals for air monitoring data that included censored values.

Stoline (1991) compared the goodness-of-fit of Box-Cox transformed data (based on using the "optimal" power transformation from a finite set of values between -1.5 and 1.5) with log-transformed data for 17 groundwater chemistry variables. Using the Probability Plot Correlation Coefficient statistic for censored data as a measure of goodness-of-fit (see [gofTest](#)), Stoline (1991) found that

only 6 of the variables were adequately modeled by a Box-Cox transformation ($p > 0.10$ for these 6 variables). Of these variables, five were adequately modeled by a log transformation. Ten of variables were “marginally” fit by an optimal Box-Cox transformation, and of these 10 only 6 were marginally fit by a log transformation. Based on these results, Stoline (1991) recommends checking the assumption of lognormality before automatically assuming environmental data fit a lognormal distribution.

One problem with data transformations is that translating results on the transformed scale back to the original scale is not always straightforward. Estimating quantities such as means, variances, and confidence limits in the transformed scale and then transforming them back to the original scale usually leads to biased and inconsistent estimates (Gilbert, 1987, p.149; van Belle et al., 2004, p.400). For example, exponentiating the confidence limits for a mean based on log-transformed data does not yield a confidence interval for the mean on the original scale. Instead, this yields a confidence interval for the median (see the help file for [elnormAltCensored](#)). It should be noted, however, that quantiles (percentiles) and rank-based procedures are invariant to monotonic transformations (Helsel and Hirsch, 1992, p.12).

Finally, there is no guarantee that a Box-Cox transformation based on the “optimal” value of λ will provide an adequate transformation to allow the assumption of approximate normality and constant variance. Any set of transformed data should be inspected relative to the assumptions you want to make about it (Johnson and Wichern, 2007, p.194).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Berthouex, P.M., and L.C. Brown. (2002). *Statistics for Environmental Engineers, Second Edition*. Lewis Publishers, Boca Raton, FL.
- Box, G.E.P., and D.R. Cox. (1964). An Analysis of Transformations (with Discussion). *Journal of the Royal Statistical Society, Series B* **26**(2), 211–252.
- Cohen, A.C. (1991). Truncated and Censored Samples. *Marcel Dekker*, New York, New York, pp.50–59.
- Draper, N., and H. Smith. (1998). *Applied Regression Analysis*. Third Edition. John Wiley and Sons, New York, pp.47–53.
- Gilbert, R.O. (1987). *Statistical Methods for Environmental Pollution Monitoring*. Van Nostrand Reinhold, NY.
- Helsel, D.R., and R.M. Hirsch. (1992). *Statistical Methods in Water Resources Research*. Elsevier, New York, NY.
- Hinkley, D.V., and G. Runger. (1984). The Analysis of Transformed Data (with Discussion). *Journal of the American Statistical Association* **79**, 302–320.
- Hoaglin, D.C., F.M. Mosteller, and J.W. Tukey, eds. (1983). *Understanding Robust and Exploratory Data Analysis*. John Wiley and Sons, New York, Chapter 4.
- Hoaglin, D.C. (1988). Transformations in Everyday Experience. *Chance* **1**, 40–45.
- Johnson, N. L., S. Kotz, and A.W. Kemp. (1992). *Univariate Discrete Distributions, Second Edition*. John Wiley and Sons, New York, p.163.
- Johnson, R.A., and D.W. Wichern. (2007). *Applied Multivariate Statistical Analysis, Sixth Edition*. Pearson Prentice Hall, Upper Saddle River, NJ, pp.192–195.
- Shumway, R.H., A.S. Azari, and P. Johnson. (1989). Estimating Mean Concentrations Under Transformations for Environmental Data With Detection Limits. *Technometrics* **31**(3), 347–356.

Stoline, M.R. (1991). An Examination of the Lognormal and Box and Cox Family of Transformations in Fitting Environmental Data. *Environmetrics* 2(1), 85–106.

van Belle, G., L.D. Fisher, Heagerty, P.J., and Lumley, T. (2004). *Biostatistics: A Methodology for the Health Sciences, 2nd Edition*. John Wiley & Sons, New York.

Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ, Chapter 13.

See Also

[boxcoxCensored.object](#), [plot.boxcoxCensored](#), [print.boxcoxCensored](#), [boxcox](#), [Data Transformations](#), [Goodness-of-Fit Tests](#).

Examples

```
# Generate 15 observations from a lognormal distribution with
# mean=10 and cv=2 and censor the observations less than 2.
# Then generate 15 more observations from this distribution and
# censor the observations less than 4.
# Then Look at some values of various objectives for various transformations.
# Note that for both the PPCC objective the optimal value is about -0.3,
# whereas for the Log-Likelihood objective it is about 0.3.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)

x.1 <- rlnormAlt(15, mean = 10, cv = 2)
censored.1 <- x.1 < 2
x.1[censored.1] <- 2

x.2 <- rlnormAlt(15, mean = 10, cv = 2)
censored.2 <- x.2 < 4
x.2[censored.2] <- 4

x <- c(x.1, x.2)
censored <- c(censored.1, censored.2)

#-----
# Using the PPCC objective:
#-----

boxcoxCensored(x, censored)

#Results of Box-Cox Transformation
#Based on Type I Censored Data
#-----
#
#Objective Name:          PPCC
#
#Data:                   x
#
#Censoring Variable:      censored
#
#Censoring Side:         left
#
#Censoring Level(s):      2 4
#
```

```

#Sample Size:          30
#
#Percent Censored:     26.7%
#
# lambda      PPCC
#   -2.0 0.8954683
#   -1.5 0.9338467
#   -1.0 0.9643680
#   -0.5 0.9812969
#    0.0 0.9776834
#    0.5 0.9471025
#    1.0 0.8901990
#    1.5 0.8187488
#    2.0 0.7480494

```

```
boxcoxCensored(x, censored, optimize = TRUE)
```

```

#Results of Box-Cox Transformation
#Based on Type I Censored Data
#-----
#
#Objective Name:          PPCC
#
#Data:                   x
#
#Censoring Variable:     censored
#
#Censoring Side:         left
#
#Censoring Level(s):     2 4
#
#Sample Size:            30
#
#Percent Censored:       26.7%
#
#Bounds for Optimization: lower = -2
#                          upper = 2
#
#Optimal Value:          lambda = -0.3194799
#
#Value of Objective:     PPCC = 0.9827546

```

```

#-----
# Using the Log-Likelihood objective
#-----

```

```
boxcoxCensored(x, censored, objective.name = "Log-Likelihood")
```

```

#Results of Box-Cox Transformation
#Based on Type I Censored Data
#-----
#
#Objective Name:          Log-Likelihood
#
#Data:                   x

```



```

#
#Censoring Variable:      censored
#
#Censoring Side:         left
#
#Censoring Level(s):     2 4
#
#Sample Size:            30
#
#Percent Censored:       26.7%
#
# lambda Log-Likelihood
#   -2.0      -95.38785
#   -1.5      -84.76697
#   -1.0      -75.36204
#   -0.5      -68.12058
#    0.0      -63.98902
#    0.5      -63.56701
#    1.0      -66.92599
#    1.5      -73.61638
#    2.0      -82.87970

boxcoxCensored(x, censored, objective.name = "Log-Likelihood",
  optimize = TRUE)

#Results of Box-Cox Transformation
#Based on Type I Censored Data
#-----
#
#Objective Name:          Log-Likelihood
#
#Data:                    x
#
#Censoring Variable:      censored
#
#Censoring Side:         left
#
#Censoring Level(s):     2 4
#
#Sample Size:            30
#
#Percent Censored:       26.7%
#
#Bounds for Optimization: lower = -2
#                          upper = 2
#
#Optimal Value:           lambda = 0.3049744
#
#Value of Objective:      Log-Likelihood = -63.2733

#-----

# Plot the results based on the PPCC objective
#-----
boxcox.list <- boxcoxCensored(x, censored)
windows()
```

```

plot(boxcox.list)

#Look at QQ-Plots for the candidate values of lambda
#-----
plot(boxcox.list, plot.type = "Q-Q Plots", same.window = FALSE)

#=====

# Clean up
#-----
rm(x.1, censored.1, x.2, censored.2, x, censored, boxcox.list)
graphics.off()

```

boxcoxCensored.object *S3 Class "boxcoxCensored"*

Description

Objects of S3 class "boxcoxCensored" are returned by the **EnvStats** function [boxcoxCensored](#), which computes objective values for user-specified powers, or computes the optimal power for the specified objective, based on Type I censored data.

Details

Objects of class "boxcoxCensored" are lists that contain information about the powers that were used, the objective that was used, the values of the objective for the given powers, and whether an optimization was specified.

Value

Required Components

The following components must be included in a legitimate list of class "boxcoxCensored".

lambda	Numeric vector containing the powers used in the Box-Cox transformations. If the value of the optimize component is FALSE, then lambda contains the values of all of the powers at which the objective was evaluated. If the value of the optimize component is TRUE, then lambda is a scalar containing the value of the power that maximizes the objective.
objective	Numeric vector containing the value(s) of the objective for the given value(s) of λ that are stored in the component lambda.
objective.name	Character string indicating the objective that was used. The possible values are "PPCC" (probability plot correlation coefficient; the default), "Shapiro-Wilk" (the Shapiro-Wilk goodness-of-fit statistic), and "Log-Likelihood" (the log-likelihood function).
optimize	Logical scalar indicating whether the objective was simply evaluated at the given values of lambda (optimize=FALSE), or instead the optimal power transformation was computed within the bounds specified by lambda (optimize=TRUE).
optimize.bounds	Numeric vector of length 2 with a names attribute indicating the bounds within which the optimization took place. When optimize=FALSE, this contains missing values.

eps	Finite, positive numeric scalar indicating what value of eps was used. When the absolute value of lambda is less than eps, lambda is assumed to be 0 for the Box-Cox transformation.
sample.size	Numeric scalar indicating the number of finite, non-missing observations.
censoring.side	Character string indicating the censoring side. Possible values are "left" and "right".
censoring.levels	Numeric vector containing the censoring levels.
percent.censored	Numeric scalar indicating the percent of observations that are censored.
data.name	The name of the data object used for the Box-Cox computations.
censoring.name	The name of the data object indicating which observations are censored.
bad.obs	The number of missing (NA), undefined (NaN) and/or infinite (Inf, -Inf) values that were removed from the data object prior to performing the Box-Cox computations.

Optional Component

The following components may optionally be included in a legitimate list of class "boxcoxCensored". They must be included if you want to call the function [plot.boxcoxCensored](#) and specify Q-Q plots or Tukey Mean-Difference Q-Q plots.

data	Numeric vector containing the data actually used for the Box-Cox computations (i.e., the original data without any missing or infinite values).
censored	Logical vector indicating which of the vales in the component data are censored.

Methods

Generic functions that have methods for objects of class "boxcoxCensored" include: [link{plot}](#), [print](#).

Note

Since objects of class "boxcoxCensored" are lists, you may extract their components with the `$` and `[]` operators.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

See Also

[boxcoxCensored](#), [plot.boxcoxCensored](#), [print.boxcoxCensored](#).

Examples

```
# Create an object of class "boxcoxCensored", then print it out.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
```

```

x.1 <- rlnormAlt(15, mean = 10, cv = 2)
censored.1 <- x.1 < 2
x.1[censored.1] <- 2

x.2 <- rlnormAlt(15, mean = 10, cv = 2)
censored.2 <- x.2 < 4
x.2[censored.2] <- 4

x <- c(x.1, x.2)
censored <- c(censored.1, censored.2)

boxcox.list <- boxcoxCensored(x, censored)

data.class(boxcox.list)
#[1] "boxcoxCensored"

names(boxcox.list)
# [1] "lambda"           "objective"         "objective.name"
# [4] "optimize"         "optimize.bounds"   "eps"
# [7] "data"             "censored"          "sample.size"
#[10] "censoring.side"    "censoring.levels"  "percent.censored"
#[13] "data.name"        "censoring.name"    "bad.obs"

boxcox.list

#Results of Box-Cox Transformation
#Based on Type I Censored Data
#-----
#
#Objective Name:                PPCC
#
#Data:                          x
#
#Censoring Variable:            censored
#
#Censoring Side:                left
#
#Censoring Level(s):            2 4
#
#Sample Size:                   30
#
#Percent Censored:              26.7%
#
# lambda      PPCC
# -2.0 0.8954683
# -1.5 0.9338467
# -1.0 0.9643680
# -0.5 0.9812969
# 0.0 0.9776834
# 0.5 0.9471025
# 1.0 0.8901990
# 1.5 0.8187488
# 2.0 0.7480494

boxcox.list2 <- boxcox(x, optimize = TRUE)
names(boxcox.list2)
# [1] "lambda"           "objective"         "objective.name"

```

```

# [4] "optimize"          "optimize.bounds" "eps"
# [7] "data"              "sample.size"     "data.name"
#[10] "bad.obs"

boxcox.list2
#Results of Box-Cox Transformation
#-----
#
#Objective Name:          PPCC
#
#Data:                    x
#
#Sample Size:             30
#
#Bounds for Optimization: lower = -2
#                          upper = 2
#
#Optimal Value:           lambda = -0.5826431
#
#Value of Objective:      PPCC = 0.9755402

#=====

# Clean up
#-----
rm(x.1, censored.1, x.2, censored.2, x, censored, boxcox.list, boxcox.list2)

```

boxcoxLm.object	<i>S3 Class "boxcoxLm"</i>
-----------------	----------------------------

Description

Objects of S3 class "boxcoxLm" are returned by the **EnvStats** function `boxcox` when the argument `x` is an object of class "lm". In this case, `boxcox` computes values of an objective function for user-specified powers, or computes the optimal power for the specified objective, based on residuals from the linear model.

Details

Objects of class "boxcoxLm" are lists that contain information about the "lm" object that was supplied, the powers that were used, the objective that was used, the values of the objective for the given powers, and whether an optimization was specified.

Value

The following components must be included in a legitimate list of class "boxcoxLm".

lambda	Numeric vector containing the powers used in the Box-Cox transformations. If the value of the optimize component is FALSE, then lambda contains the values of all of the powers at which the objective was evaluated. If the value of the optimize component is TRUE, then lambda is a scalar containing the value of the power that maximizes the objective.
objective	Numeric vector containing the value(s) of the objective for the given value(s) of λ that are stored in the component lambda.

objective.name	character string indicating the objective that was used. The possible values are "PPCC" (probability plot correlation coefficient; the default), "Shapiro-Wilk" (the Shapiro-Wilk goodness-of-fit statistic), and "Log-Likelihood" (the log-likelihood function).
optimize	logical scalar indicating whether the objective was simply evaluated at the given values of lambda (optimize=FALSE), or instead the optimal power transformation was computed within the bounds specified by lambda (optimize=TRUE).
optimize.bounds	Numeric vector of length 2 with a names attribute indicating the bounds within which the optimization took place. When optimize=FALSE, this contains missing values.
eps	finite, positive numeric scalar indicating what value of eps was used. When the absolute value of lambda is less than eps, lambda is assumed to be 0 for the Box-Cox transformation.
lm.obj	the value of the argument x provided to boxcox (an object that must inherit from class "lm").
sample.size	Numeric scalar indicating the number of finite, non-missing observations.
data.name	The name of the data object used for the Box-Cox computations.

Methods

Generic functions that have methods for objects of class "boxcoxLm" include: [link{plot}](#), [print](#).

Note

Since objects of class "boxcoxLm" are lists, you may extract their components with the \$ and [[operators.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

See Also

[boxcox](#), [plot.boxcoxLm](#), [print.boxcoxLm](#), [boxcox.object](#).

Examples

```
# Create an object of class "boxcoxLm", then print it out.

# The data frame Environmental.df contains daily measurements of
# ozone concentration, wind speed, temperature, and solar radiation
# in New York City for 153 consecutive days between May 1 and
# September 30, 1973. In this example, we'll plot ozone vs.
# temperature and look at the Q-Q plot of the residuals. Then
# we'll look at possible Box-Cox transformations. The "optimal" one
# based on the PPCC looks close to a log-transformation
# (i.e., lambda=0). The power that produces the largest PPCC is
# about 0.2, so a cube root (lambda=1/3) transformation might work too.

# Fit the model with the raw Ozone data
#-----
```

```

ozone.fit <- lm(ozone ~ temperature, data = Environmental.df)

# Plot Ozone vs. Temperature, with fitted line
#-----
windows()
with(Environmental.df,
     plot(temperature, ozone, xlab = "Temperature (degrees F)",
          ylab = "Ozone (ppb)", main = "Ozone vs. Temperature"))
abline(ozone.fit)

# Look at the Q-Q Plot for the residuals
#-----
windows()
qqPlot(ozone.fit$residuals, add.line = TRUE)

# Look at Box-Cox transformations of Ozone
#-----
boxcox.list <- boxcox(ozone.fit)
boxcox.list
#Results of Box-Cox Transformation
#-----
#
#Objective Name:                PPCC
#
#Linear Model:                  ozone.fit
#
#Sample Size:                   116
#
# lambda      PPCC
#   -2.0 0.4286781
#   -1.5 0.4673544
#   -1.0 0.5896132
#   -0.5 0.8301458
#    0.0 0.9871519
#    0.5 0.9819825
#    1.0 0.9408694
#    1.5 0.8840770
#    2.0 0.8213675

#-----

# Clean up
#-----
rm(ozone.fit, boxcox.list)

```

boxcoxTransform

Apply a Box-Cox Power Transformation to a Set of Data

Description

Apply a Box-Cox power transformation to a set of data to attempt to induce normality and homogeneity of variance.

Usage

```
boxcoxTransform(x, lambda, eps = .Machine$double.eps)
```

Arguments

x	a numeric vector of positive numbers.
lambda	finite numeric scalar indicating what power to use for the Box-Cox transformation.
eps	finite, positive numeric scalar. When the absolute value of lambda is less than eps, lambda is assumed to be 0 for the Box-Cox transformation. The default value is <code>eps=.Machine\$double.eps</code> .

Details

Two common assumptions for several standard parametric hypothesis tests are:

1. The observations all come from a normal distribution.
2. The observations all come from distributions with the same variance.

For example, the standard one-sample t-test assumes all the observations come from the same normal distribution, and the standard two-sample t-test assumes that all the observations come from a normal distribution with the same variance, although the mean may differ between the two groups. For standard linear regression models, these assumptions can be stated as: the error terms all come from a normal distribution with mean 0 and a constant variance.

Often, especially with environmental data, the above assumptions do not hold because the original data are skewed and/or they follow a distribution that is not really shaped like a normal distribution. It is sometimes possible, however, to transform the original data so that the transformed observations in fact come from a normal distribution or close to a normal distribution. The transformation may also induce homogeneity of variance and, for the case of a linear regression model, a linear relationship between the response and predictor variable(s).

Sometimes, theoretical considerations indicate an appropriate transformation. For example, count data often follow a [Poisson distribution](#), and it can be shown that taking the square root of observations from a Poisson distribution tends to make these data look more bell-shaped (Johnson et al., 1992, p.163; Johnson and Wichern, 2007, p.192; Zar, 2010, p.291). A common example in the environmental field is that chemical concentration data often appear to come from a [lognormal distribution](#) or some other positively-skewed distribution (e.g., [gamma](#)). In this case, taking the logarithm of the observations often appears to yield normally distributed data.

Ideally, a data transformation is chosen based on knowledge of the process generating the data, as well as graphical tools such as [quantile-quantile plots](#) and histograms.

Box and Cox (1964) presented a formalized method for deciding on a data transformation. Given a random variable X from some distribution with only positive values, the Box-Cox family of power transformations is defined as:

$$Y = \begin{cases} \frac{X^\lambda - 1}{\lambda} & \lambda \neq 0 \\ \log(X) & \lambda = 0 \end{cases} \quad (1)$$

where Y is assumed to come from a normal distribution. This transformation is continuous in λ . Note that this transformation also preserves ordering; that is, if $X_1 < X_2$ then $Y_1 < Y_2$.

Box and Cox (1964) proposed choosing the appropriate value of λ based on maximizing a likelihood function. See the help file for [boxcox](#) for details.

Note that for non-zero values of λ , instead of using the formula of Box and Cox in Equation (1), you may simply use the power transformation:

$$Y = X^\lambda \quad (2)$$

since these two equations differ only by a scale difference and origin shift, and the essential character of the transformed distribution remains unchanged.

The value $\lambda = 1$ corresponds to no transformation. Values of λ less than 1 shrink large values of X , and are therefore useful for transforming positively-skewed (right-skewed) data. Values of λ larger than 1 inflate large values of X , and are therefore useful for transforming negatively-skewed (left-skewed) data (Helsel and Hirsch, 1992, pp.13-14; Johnson and Wichern, 2007, p.193). Commonly used values of λ include 0 (log transformation), 0.5 (square-root transformation), -1 (reciprocal), and -0.5 (reciprocal root).

It is often recommend that when dealing with several similar data sets, it is best to find a common transformation that works reasonably well for all the data sets, rather than using slightly different transformations for each data set (Helsel and Hirsch, 1992, p.14; Shumway et al., 1989).

Value

numeric vector of transformed observations.

Note

Data transformations are often used to induce normality, homoscedasticity, and/or linearity, common assumptions of parametric statistical tests and estimation procedures. Transformations are not “tricks” used by the data analyst to hide what is going on, but rather useful tools for understanding and dealing with data (Berthouex and Brown, 2002, p.61). Hoaglin (1988) discusses “hidden” transformations that are used everyday, such as the pH scale for measuring acidity.

In the case of a linear model, there are at least two approaches to improving a model fit: transform the Y and/or X variable(s), and/or use more predictor variables. Often in environmental data analysis, we assume the observations come from a lognormal distribution and automatically take logarithms of the data. For a simple linear regression (i.e., one predictor variable), if regression diagnostic plots indicate that a straight line fit is not adequate, but that the variance of the errors appears to be fairly constant, you may only need to transform the predictor variable X or perhaps use a quadratic or cubic model in X . On the other hand, if the diagnostic plots indicate that the constant variance and/or normality assumptions are suspect, you probably need to consider transforming the response variable Y . Data transformations for linear regression models are discussed in Draper and Smith (1998, Chapter 13) and Helsel and Hirsch (1992, pp. 228-229).

One problem with data transformations is that translating results on the transformed scale back to the original scale is not always straightforward. Estimating quantities such as means, variances, and confidence limits in the transformed scale and then transforming them back to the original scale usually leads to biased and inconsistent estimates (Gilbert, 1987, p.149; van Belle et al., 2004, p.400). For example, exponentiating the confidence limits for a mean based on log-transformed data does not yield a confidence interval for the mean on the original scale. Instead, this yields a confidence interval for the median (see the help file for [elnormAlt](#)). It should be noted, however, that quantiles (percentiles) and rank-based procedures are invariant to monotonic transformations (Helsel and Hirsch, 1992, p.12).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Berthouex, P.M., and L.C. Brown. (2002). *Statistics for Environmental Engineers, Second Edition*. Lewis Publishers, Boca Raton, FL.

- Box, G.E.P., and D.R. Cox. (1964). An Analysis of Transformations (with Discussion). *Journal of the Royal Statistical Society, Series B* **26**(2), 211–252.
- Draper, N., and H. Smith. (1998). *Applied Regression Analysis*. Third Edition. John Wiley and Sons, New York, pp.47-53.
- Gilbert, R.O. (1987). *Statistical Methods for Environmental Pollution Monitoring*. Van Nostrand Reinhold, NY.
- Helsel, D.R., and R.M. Hirsch. (1992). *Statistical Methods in Water Resources Research*. Elsevier, New York, NY.
- Hinkley, D.V., and G. Runger. (1984). The Analysis of Transformed Data (with Discussion). *Journal of the American Statistical Association* **79**, 302–320.
- Hoaglin, D.C., F.M. Mosteller, and J.W. Tukey, eds. (1983). *Understanding Robust and Exploratory Data Analysis*. John Wiley and Sons, New York, Chapter 4.
- Hoaglin, D.C. (1988). Transformations in Everyday Experience. *Chance* **1**, 40–45.
- Johnson, N. L., S. Kotz, and A.W. Kemp. (1992). *Univariate Discrete Distributions, Second Edition*. John Wiley and Sons, New York, p.163.
- Johnson, R.A., and D.W. Wichern. (2007). *Applied Multivariate Statistical Analysis, Sixth Edition*. Pearson Prentice Hall, Upper Saddle River, NJ, pp.192–195.
- Shumway, R.H., A.S. Azari, and P. Johnson. (1989). Estimating Mean Concentrations Under Transformations for Environmental Data With Detection Limits. *Technometrics* **31**(3), 347–356.
- Stoline, M.R. (1991). An Examination of the Lognormal and Box and Cox Family of Transformations in Fitting Environmental Data. *Environmetrics* **2**(1), 85–106.
- van Belle, G., L.D. Fisher, Heagerty, P.J., and Lumley, T. (2004). *Biostatistics: A Methodology for the Health Sciences, 2nd Edition*. John Wiley & Sons, New York.
- Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ, Chapter 13.

See Also

[boxcox](#), [Data Transformations](#), [Goodness-of-Fit Tests](#).

Examples

```
# Generate 30 observations from a lognormal distribution with
# mean=10 and cv=2, then look at some normal quantile-quantile
# plots for various transformations.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
x <- rlnormAlt(30, mean = 10, cv = 2)

windows()
qqPlot(x, add.line = TRUE)

windows()
qqPlot(boxcoxTransform(x, lambda = 0.5), add.line = TRUE)

windows()
qqPlot(boxcoxTransform(x, lambda = 0), add.line = TRUE)

# Clean up
```

```
#-----
rm(x)
```

calibrate

Fit a Calibration Line or Curve

Description

Fit a calibration line or curve based on linear regression.

Usage

```
calibrate(formula, data, max.order = 4, p.crit = 0.05, weights,
  subset, na.action, method = "qr", model = FALSE, x = FALSE,
  y = FALSE, contrasts = NULL, ...)
```

Arguments

- | | |
|-----------------|--|
| formula | a formula object, with the response on the left of a ~ operator, and the single predictor variable on the right. For example, Cadmium ~ Spike. |
| data | an optional data frame, list or environment (or object coercible by as.data.frame to a data frame) containing the variables in the model. If not found in data, the variables are taken from environment(formula), typically the environment from which calibrate is called. |
| max.order | optional integer indicating the maximum order of the polynomial to consider for the calibration curve. The default value is max.order=4. |
| p.crit | optional numeric scaler between 0 and 1 indicating the p-value to use for the stepwise regression when determining which polynomial model to use. The default value is p.crit=0.05. |
| weights | optional vector of observation weights; if supplied, the algorithm fits to minimize the sum of the weights multiplied into the squared residuals. The length of weights must be the same as the number of observations. The weights must be nonnegative and it is strongly recommended that they be strictly positive, since zero weights are ambiguous, compared to use of the subset argument. |
| subset | optional expression saying which subset of the rows of the data should be used in the fit. This can be a logical vector (which is replicated to have length equal to the number of observations), or a numeric vector indicating which observation numbers are to be included, or a character vector of the row names to be included. All observations are included by default. |
| na.action | optional function which indicates what should happen when the data contain NAs. The default is set by the na.action setting of options , and is na.fail if that is unset. The 'factory-fresh' default is na.omit . Another possible value is NULL, no action. Value na.exclude can be useful. |
| method | optional method to be used; for fitting, currently only method = "qr" is supported; method = "model.frame" returns the model frame (the same as with model = TRUE, see below). |
| model, x, y, qr | optional logicals. If TRUE the corresponding components of the fit (the model frame, the model matrix, the response, the QR decomposition) are returned. |

`contrasts` an optional list. See the argument `contrasts.arg` of `model.matrix`.

`...` additional arguments to be passed to the low level regression fitting functions (see [lm](#)).

Details

A simple and frequently used calibration model is a straight line where the response variable S denotes the signal of the machine and the predictor variable C denotes the true concentration in the physical sample. The error term is assumed to follow a normal distribution with mean 0. Note that the average value of the signal for a blank ($C = 0$) is the intercept. Other possible calibration models include higher order polynomial models such as a quadratic or cubic model.

In a typical setup, a small number of samples (e.g., $n = 6$) with known concentrations are measured and the signal is recorded. A sample with no chemical in it, called a blank, is also measured. (You have to be careful to define exactly what you mean by a “blank.” A blank could mean a container from the lab that has nothing in it but is prepared in a similar fashion to containers with actual samples in them. Or it could mean a field blank: the container was taken out to the field and subjected to the same process that all other containers were subjected to, except a physical sample of soil or water was not placed in the container.) Usually, replicate measures at the same known concentrations are taken. (The term “replicate” must be well defined to distinguish between for example the same physical samples that are measured more than once vs. two different physical samples of the same known concentration.)

The function `calibrate` initially fits a linear calibration model. If the argument `max.order` is greater than 1, `calibrate` then performs forward stepwise linear regression to determine the “best” polynomial model.

In the case where replicates are available, for each model `calibrate` computes the p-value of the ANOVA for lack-of-fit vs. pure error (Draper and Smith, 1998, Chapters 2). If the p-value is greater than or equal to `p.crit`, then this is the final model; otherwise the next higher-order term is added to the polynomial and the model is re-fit.

In the case where replicates are not available, `calibrate` uses standard stepwise ANOVA to compare models (Draper and Smith, 1998, p.335). In this case, if the p-value for the partial F-test to compare models is greater than or equal to `p.crit`, then the model with fewer terms is used as the final model.

The stepwise algorithm terminates when either the p-value is greater than or equal to `p.crit`, or the currently selected model in the algorithm is of order `max.order`.

Value

An object of `class` “`lm`” that includes a component called `x` that stores the model matrix (the values of the predictor variables for the final calibration model).

Note

Almost always the process of determining the concentration of a chemical in a soil, water, or air sample involves using some kind of machine that produces a signal, and this signal is related to the concentration of the chemical in the physical sample. The process of relating the machine signal to the concentration of the chemical is called **calibration**. Once calibration has been performed, estimated concentrations in physical samples with unknown concentrations are computed using inverse regression (see [inversePredictCalibrate](#)). The uncertainty in the process used to estimate the concentration may be quantified with decision, detection, and quantitation limits.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Draper, N., and H. Smith. (1998). *Applied Regression Analysis*. Third Edition. John Wiley and Sons, New York, Chapter 3 and p.335.

Millard, S.P., and N.K. Neerchal. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, FL, pp.562-575.

See Also

[inversePredictCalibrate](#), [detectionLimitCalibrate](#), [lm](#).

Examples

```
# The data frame EPA.97.cadmium.111.df contains calibration data for
# cadmium at mass 111 (ng/L) that appeared in Gibbons et al. (1997b)
# and were provided to them by the U.S. EPA.
# Display a plot of these data along with the fitted calibration line
# and 99% non-simultaneous prediction limits. See
# Millard and Neerchal (2001, pp.566-569) for more details on this
# example.
```

```
Cadmium <- EPA.97.cadmium.111.df$Cadmium
```

```
Spike <- EPA.97.cadmium.111.df$Spike
```

```
calibrate.list <- calibrate(Cadmium ~ Spike,
  data=EPA.97.cadmium.111.df)
```

```
newdata <- data.frame(Spike = seq(min(Spike), max(Spike), len=100))
```

```
pred.list <- predict(calibrate.list, newdata=newdata, se.fit=TRUE)
```

```
pointwise.list <- pointwise(pred.list, coverage=0.99,
  individual=TRUE)
```

```
plot(Spike, Cadmium, ylim=c(min(pointwise.list$lower),
  max(pointwise.list$upper)), xlab="True Concentration (ng/L)",
  ylab="Observed Concentration (ng/L)")
```

```
abline(calibrate.list, lwd=2)
```

```
lines(newdata$Spike, pointwise.list$lower, lty=8, lwd=2)
```

```
lines(newdata$Spike, pointwise.list$upper, lty=8, lwd=2)
```

```
title(paste("Calibration Line and 99% Prediction Limits",
  "for US EPA Cadmium 111 Data", sep="\n"))
```

```
rm(Cadmium, Spike, newdata, calibrate.list, pred.list,
  pointwise.list)
```

CastilloAndHadi1994 *Abstract: Castillo and Hadi (1994)*

Description

Detailed abstract of the manuscript:

Castillo, E., and A. Hadi. (1994). Parameter and Quantile Estimation for the Generalized Extreme-Value Distribution. *Environmetrics* **5**, 417–432.

Details

Abstract

Castillo and Hadi (1994) introduce a new way to estimate the parameters and quantiles of the [generalized extreme value distribution](#) (GEVD) with parameters $\text{location}=\eta$, $\text{scale}=\theta$, and $\text{shape}=\kappa$. The estimator is based on a two-stage procedure using order statistics, denoted here by “TSOE”, which stands for two-stage order-statistics estimator. Castillo and Hadi (1994) compare the TSOE to the maximum likelihood estimator (MLE; Jenkinson, 1969; Prescott and Walden, 1983) and probability-weighted moments estimator (PWME; [Hosking et al., 1985](#)).

Castillo and Hadi (1994) note that for some samples the likelihood may not have a local maximum, and also when $\kappa > 1$ the likelihood can be made infinite so the MLE does not exist. They also note, as do [Hosking et al., 1985](#)), that when $\kappa \leq -1$, the moments and probability-weighted moments of the GEVD do not exist, hence neither does the PWME. (Hosking et al., however, claim that in practice the shape parameter usually lies between $-1/2$ and $1/2$.) On the other hand, the TSOE exists for all values of κ .

Based on computer simulations, Castillo and Hadi (1994) found that the performance (bias and root mean squared error) of the TSOE is comparable to the PWME for values of κ in the range $-1/2 \leq \kappa \leq 1/2$. They also found that the TSOE is superior to the PWME for large values of κ . Their results, however, are based on using the PWME computed using the approximation given in equation (14) of [Hosking et al. \(1985, p.253\)](#). The true PWME is computed using equation (12) of [Hosking et al. \(1985, p.253\)](#). [Hosking et al. \(1985\)](#) introduced the approximation as a matter of computational convenience, and noted that it is valid in the range $-1/2 \leq \kappa \leq 1/2$. If Castillo and Hadi (1994) had used the true PWME for values of κ larger than $1/2$, they probably would have gotten very different results for the PWME. (Note: the function [egevd](#) with `method="pwme"` uses the exact equation (12) of [Hosking et al. \(1985\)](#), not the approximation (14)).

Castillo and Hadi (1994) suggest using the bootstrap or jackknife to obtain variance estimates and confidence intervals for the distribution parameters based on the TSOE.

More Details Let $\underline{x} = (x_1, x_2, \dots, x_n)$ be a vector of n observations from a [generalized extreme value distribution](#) with parameters $\text{location}=\eta$, $\text{scale}=\theta$, and $\text{shape}=\kappa$ with cumulative distribution function F . Also, let $x(1), x(2), \dots, x(n)$ denote the ordered values of \underline{x} .

First Stage

Castillo and Hadi (1994) propose as initial estimates of the distribution parameters the solutions to the following set of simultaneous equations based on just three observations from the total sample of size n :

$$F[x(1); \eta, \theta, \kappa] = p_{1,n}$$

$$F[x(j); \eta, \theta, \kappa] = p_{j,n}$$

$$F[x(n); \eta, \theta, \kappa] = p_{n,n} \quad (1)$$

where $2 \leq j \leq n - 1$, and

$$p_{i,n} = \hat{F}[x(i); \eta, \theta, \kappa]$$

denotes the i 'th plotting position for a sample of size n ; that is, a nonparametric estimate of the value of F at $x(i)$. Typically, plotting positions have the form:

$$p_{i,n} = \frac{i - a}{n + b} \quad (2)$$

where $b > -a > -1$. In their simulation studies, Castillo and Hadi (1994) used $a=0.35$, $b=0$.

Since j is arbitrary in the above set of equations (1), denote the solutions to these equations by:

$$\hat{\eta}_j, \hat{\theta}_j, \hat{\kappa}_j$$

There are thus $n - 2$ sets of estimates.

Castillo and Hadi (1994) show that the estimate of the shape parameter, κ , is the solution to the equation:

$$\frac{x(j) - x(n)}{x(1) - x(n)} = \frac{1 - A_{jn}^{\kappa}}{1 - A_{1n}^{\kappa}} \quad (3)$$

where

$$A_{ik} = C_i / C_k \quad (4)$$

$$C_i = -\log(p_{i,n}) \quad (5)$$

Castillo and Hadi (1994) show how to easily solve equation (3) using the method of bisection.

Once the estimate of the shape parameter is obtained, the other estimates are given by:

$$\hat{\theta}_j = \frac{\hat{\kappa}_j [x(1) - x(n)]}{(C_n)^{\hat{\kappa}_j} - (C_1)^{\hat{\kappa}_j}} \quad (6)$$

$$\hat{\eta}_j = x(1) - \frac{\hat{\theta}_j [1 - (C_1)^{\hat{\kappa}_j}]}{\hat{\kappa}_j} \quad (7)$$

Second Stage

Apply a robust function to the $n - 2$ sets of estimates obtained in the first stage. Castillo and Hadi (1994) suggest using either the median or the least median of squares (using a column of 1's as the predictor variable; see the help file for [lmsreg](#) in the package **MASS**). Using the median, for example, the final distribution parameter estimates are given by:

$$\hat{\eta} = \text{Median}(\hat{\eta}_2, \hat{\eta}_3, \dots, \hat{\eta}_{n-1})$$

$$\hat{\theta} = \text{Median}(\hat{\theta}_2, \hat{\theta}_3, \dots, \hat{\theta}_{n-1})$$

$$\hat{\kappa} = \text{Median}(\hat{\kappa}_2, \hat{\kappa}_3, \dots, \hat{\kappa}_{n-1})$$

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.
- Hosking, J.R.M. (1985). Algorithm AS 215: Maximum-Likelihood Estimation of the Parameters of the Generalized Extreme-Value Distribution. *Applied Statistics* **34**(3), 301–310.
- Jenkinson, A.F. (1969). *Statistics of Extremes*. *Technical Note 98*, World Meteorological Office, Geneva.
- Johnson, N. L., S. Kotz, and N. Balakrishnan. (1995). *Continuous Univariate Distributions, Volume 2*. Second Edition. John Wiley and Sons, New York.
- Prescott, P., and A.T. Walden. (1983). Maximum Likelihood Estimation of the Three-Parameter Generalized Extreme-Value Distribution from Censored Samples. *Journal of Statistical Computing and Simulation* **16**, 241–250.

See Also

[Generalized Extreme Value Distribution](#), [egevd](#), [Hosking et al., 1985](#)).

cdfCompare

Plot Two Cumulative Distribution Functions

Description

For one sample, plots the empirical cumulative distribution function (ecdf) along with a theoretical cumulative distribution function (cdf). For two samples, plots the two ecdf's. These plots are used to graphically assess goodness of fit.

Usage

```
cdfCompare(x, y = NULL, discrete = FALSE,
  prob.method = ifelse(discrete, "emp.probs", "plot.pos"),
  plot.pos.con = NULL, distribution = "norm", param.list = NULL,
  estimate.params = is.null(param.list), est.arg.list = NULL, x.col = "blue",
  y.or.fitted.col = "black", x.lwd = 3 * par("cex"), y.or.fitted.lwd = 3 * par("cex"),
  x.lty = 1, y.or.fitted.lty = 2, digits = .Options$digits, ...,
  type = ifelse(discrete, "s", "l"), main = NULL, xlab = NULL, ylab = NULL,
  xlim = NULL, ylim = NULL)
```

Arguments

- | | |
|----------|--|
| x | numeric vector of observations. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed. |
| y | a numeric vector (not necessarily of the same length as x). Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed. The default value is y=NULL, in which case the empirical cdf of x will be plotted along with the theoretical cdf specified by the argument distribution. |
| discrete | logical scalar indicating whether the assumed parent distribution of x is discrete (discrete=TRUE) or continuous (discrete=FALSE; the default). |

<code>prob.method</code>	character string indicating what method to use to compute the plotting positions (empirical probabilities). Possible values are <code>plot.pos</code> (plotting positions, the default if <code>discrete=FALSE</code>) and <code>emp.probs</code> (empirical probabilities, the default if <code>discrete=TRUE</code>). See the help file for ecdfPlot for more explanation.
<code>plot.pos.con</code>	numeric scalar between 0 and 1 containing the value of the plotting position constant. When <code>y</code> is supplied, the default value is <code>plot.pos.con=0.375</code> . When <code>y</code> is not supplied, for the normal, lognormal, three-parameter lognormal, zero-modified normal, and zero-modified lognormal distributions, the default value is <code>plot.pos.con=0.375</code> . For the Type I extreme value (Gumbel) distribution (<code>distribution="evd"</code>), the default value is <code>plot.pos.con=0.44</code> . For all other distributions, the default value is <code>plot.pos.con=0.4</code> . See the help files for ecdfPlot and qqPlot for more information. This argument is ignored if <code>prob.method="emp.probs"</code> .
<code>distribution</code>	when <code>y</code> is not supplied, a character string denoting the distribution abbreviation. The default value is <code>distribution="norm"</code> . See the help file for Distribution.df for a list of possible distribution abbreviations. This argument is ignored if <code>y</code> is supplied.
<code>param.list</code>	when <code>y</code> is not supplied, a list with values for the parameters of the distribution. The default value is <code>param.list=list(mean=0, sd=1)</code> . See the help file for Distribution.df for the names and possible values of the parameters associated with each distribution. This argument is ignored if <code>y</code> is supplied or <code>estimate.params=TRUE</code> .
<code>estimate.params</code>	when <code>y</code> is not supplied, a logical scalar indicating whether to compute the cdf for <code>x</code> based on estimating the distribution parameters (<code>estimate.params=TRUE</code>) or using the known distribution parameters specified in <code>param.list</code> (<code>estimate.params=FALSE</code>). The default value is <code>TRUE</code> unless the argument <code>param.list</code> is supplied. The argument <code>estimate.params</code> is ignored if <code>y</code> is supplied.
<code>est.arg.list</code>	when <code>y</code> is not supplied and <code>estimate.params=TRUE</code> , a list whose components are optional arguments associated with the function used to estimate the parameters of the assumed distribution (see the help file Estimating Distribution Parameters). For example, all functions used to estimate distribution parameters have an optional argument called <code>method</code> that specifies the method to use to estimate the parameters. (See the help file for Distribution.df for a list of available estimation methods for each distribution.) To override the default estimation method, supply the argument <code>est.arg.list</code> with a component called <code>method</code> ; for example <code>est.arg.list=list(method="mle")</code> . The default value is <code>est.arg.list=NULL</code> so that all default values for the estimating function are used. This argument is ignored if <code>estimate.params=FALSE</code> or <code>y</code> is supplied.
<code>x.col</code>	a numeric scalar or character string determining the color of the empirical cdf (based on <code>x</code>) line or points. The default value is <code>x.col="blue"</code> . See the entry for <code>col</code> in the help file for par for more information.
<code>y.or.fitted.col</code>	a numeric scalar or character string determining the color of the empirical cdf (based on <code>y</code>) or the theoretical cdf line or points. The default value is <code>y.or.fitted.col="black"</code> . See the entry for <code>col</code> in the help file for par for more information.
<code>x.lwd</code>	a numeric scalar determining the width of the empirical cdf (based on <code>x</code>) line. The default value is <code>x.lwd=3*par("cex")</code> . See the entry for <code>lwd</code> in the help file for par for more information.

<code>y.or.fitted.lwd</code>	a numeric scalar determining the width of the empirical cdf (based on <code>y</code>) or theoretical cdf line. The default value is <code>y.or.fitted.lwd=3*par("cex")</code> . See the entry for <code>lwd</code> in the help file for par for more information.
<code>x.lty</code>	a numeric scalar determining the line type of the empirical cdf (based on <code>x</code>) line. The default value is <code>x.lty=1</code> . See the entry for <code>lty</code> in the help file for par for more information.
<code>y.or.fitted.lty</code>	a numeric scalar determining the line type of the empirical cdf (based on <code>y</code>) or theoretical cdf line. The default value is <code>y.or.fitted.lty=2</code> . See the entry for <code>lty</code> in the help file for par for more information.
<code>digits</code>	when <code>y</code> is not supplied, a scalar indicating how many significant digits to print for the distribution parameters. The default value is <code>digits=Options\$digits</code> .
<code>type, main, xlab, ylab, xlim, ylim, ...</code>	additional graphical parameters (see lines and par). In particular, the argument <code>type</code> specifies the kind of line type. By default, the function <code>cdfCompare</code> plots a step function (<code>type="s"</code>) when <code>discrete=TRUE</code> , and plots a straight line between points (<code>type="l"</code>) when <code>discrete=FALSE</code> . The user may override these defaults by supplying the graphics parameter <code>type</code> (<code>type="s"</code> for a step function, <code>type="l"</code> for linear interpolation, <code>type="p"</code> for points only, etc.).

Details

When both `x` and `y` are supplied, the function `cdfCompare` creates the empirical cdf plot of `x` and `y` on the same plot by calling the function [ecdfPlot](#).

When `y` is not supplied, the function `cdfCompare` creates the empirical cdf plot of `x` (by calling [ecdfPlot](#)) and the theoretical cdf plot (by calling [cdfPlot](#) and using the argument `distribution`) on the same plot.

Value

When `y` is supplied, `cdfCompare` invisibly returns a list with components `x.ecdf.list` and `y.ecdf.list`. Each of these components is itself a list, with the components `Order.Statistics` and `Cumulative.Probabilities`, giving coordinates of the points that have been plotted.

When `y` is not supplied, `cdfCompare` invisibly returns a list with components `x.ecdf.list` and `fitted.cdf.list`. The component `x.ecdf.list` is itself a list with the components `Order.Statistics` and `Cumulative.Probabilities`, giving coordinates of the points that have been plotted for the `x` values. The component `fitted.cdf.list` is itself a list with the components `Quantiles` and `Cumulative.Probabilities`, giving coordinates of the points that have been plotted for the fitted cdf.

Note

An empirical cumulative distribution function (ecdf) plot is a graphical tool that can be used in conjunction with other graphical tools such as histograms, strip charts, and boxplots to assess the characteristics of a set of data. It is easy to determine quartiles and the minimum and maximum values from such a plot. Also, ecdf plots allow you to assess local density: a higher density of observations occurs where the slope is steep.

Chambers et al. (1983, pp.11-16) plot the observed order statistics on the *y*-axis vs. the ecdf on the *x*-axis and call this a quantile plot.

Empirical cumulative distribution function (ecdf) plots are often plotted with theoretical cdf plots (see [cdfPlot](#) and [cdfCompare](#)) to graphically assess whether a sample of observations comes from

a particular distribution. The Kolmogorov-Smirnov goodness-of-fit test (see [gofTest](#)) is the statistical companion of this kind of comparison; it is based on the maximum vertical distance between the empirical cdf plot and the theoretical cdf plot. More often, however, quantile-quantile (Q-Q) plots are used instead of ecdf plots to graphically assess departures from an assumed distribution (see [qqPlot](#)).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Chambers, J.M., W.S. Cleveland, B. Kleiner, and P.A. Tukey. (1983). *Graphical Methods for Data Analysis*. Duxbury Press, Boston, MA, pp.11-16.

Cleveland, W.S. (1993). *Visualizing Data*. Hobart Press, Summit, New Jersey, 360pp.

D'Agostino, R.B. (1986a). Graphical Analysis. In: D'Agostino, R.B., and M.A. Stephens, eds. *Goodness-of Fit Techniques*. Marcel Dekker, New York, Chapter 2, pp.7-62.

See Also

[cdfPlot](#), [ecdfPlot](#), [qqPlot](#).

Examples

```
# Generate 20 observations from a normal (Gaussian) distribution
# with mean=10 and sd=2 and compare the empirical cdf with a
# theoretical normal cdf that is based on estimating the parameters.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
x <- rnorm(20, mean = 10, sd = 2)
windows()
cdfCompare(x)

#-----

# Generate 30 observations from an exponential distribution with parameter
# rate=0.1 (see the R help file for Exponential) and compare the empirical
# cdf with the empirical cdf of the normal observations generated in the
# previous example:

set.seed(432)
y <- rexp(30, rate = 0.1)
windows()
cdfCompare(x, y)

#=====

# Generate 20 observations from a Poisson distribution with parameter lambda=10
# (see the R help file for Poisson) and compare the empirical cdf with a
# theoretical Poisson cdf based on estimating the distribution parameters.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
x <- rpois(20, lambda = 10)
```

```

windows()
cdfCompare(x, dist = "pois")

#=====

# Clean up
#-----
rm(x, y)
graphics.off()

```

cdfPlot

*Plot Cumulative Distribution Function***Description**

Produce a cumulative distribution function (cdf) plot for a user-specified distribution.

Usage

```

cdfPlot(distribution = "norm", param.list = list(mean = 0, sd = 1),
  left.tail.cutoff = ifelse(is.finite(supp.min), 0, 0.001),
  right.tail.cutoff = ifelse(is.finite(supp.max), 0, 0.001), plot.it = TRUE,
  add = FALSE, n.points = 1000, cdf.col = "black", cdf.lwd = 3 * par("cex"),
  cdf.lty = 1, curve.fill = FALSE, curve.fill.col = "cyan",
  digits = .Options$digits, ..., type = ifelse(discrete, "s", "l"),
  main = NULL, xlab = NULL, ylab = NULL, xlim = NULL, ylim = NULL)

```

Arguments

- | | |
|-------------------|--|
| distribution | a character string denoting the distribution abbreviation. The default value is <code>distribution="norm"</code> . See the help file for Distribution.df for a list of possible distribution abbreviations. |
| param.list | a list with values for the parameters of the distribution. The default value is <code>param.list=list(mean=0, sd=1)</code> . See the help file for Distribution.df for the names and possible values of the parameters associated with each distribution. |
| left.tail.cutoff | a numeric scalar indicating what proportion of the left-tail of the probability distribution to omit from the plot. For densities with a finite support minimum (e.g., Lognormal) the default value is 0; for all other densities the default value is 0.001. |
| right.tail.cutoff | a scalar indicating what proportion of the right-tail of the probability distribution to omit from the plot. For densities with a finite support maximum (e.g., Binomial) the default value is 0; for all other densities the default value is 0.001. |
| plot.it | a logical scalar indicating whether to create a plot or add to the existing plot (see <code>add</code>) on the current graphics device. If <code>plot.it=FALSE</code> , no plot is produced, but a list of (x, y) values is returned (see the section <code>VALUE</code> below). The default value is <code>plot.it=TRUE</code> . |
| add | a logical scalar indicating whether to add the cumulative distribution function curve to the existing plot (<code>add=TRUE</code>), or to create a new plot (<code>add=FALSE</code> ; the default). This argument is ignored if <code>plot.it=FALSE</code> . |

<code>n.points</code>	a numeric scalar specifying at how many evenly-spaced points the cumulative distribution function will be evaluated. The default value is <code>n.points=1000</code> .
<code>cdf.col</code>	a numeric scalar or character string determining the color of the cdf line in the plot. The default value is <code>pdf.col="black"</code> . See the entry for <code>col</code> in the help file for par for more information.
<code>cdf.lwd</code>	a numeric scalar determining the width of the cdf line in the plot. The default value is <code>pdf.lwd=3*par("cex")</code> . See the entry for <code>lwd</code> in the help file for par for more information.
<code>cdf.lty</code>	a numeric scalar determining the line type of the cdf line in the plot. The default value is <code>pdf.lty=1</code> . See the entry for <code>lty</code> in the help file for par for more information.
<code>curve.fill</code>	a logical value indicating whether to fill in the area below the cumulative distribution function curve with the color specified by <code>curve.fill.col</code> . The default value is <code>curve.fill=FALSE</code> .
<code>curve.fill.col</code>	when <code>curve.fill=TRUE</code> , a numeric scalar or character string indicating what color to use to fill in the area below the cumulative distribution function curve. The default value is <code>curve.fill.col="cyan"</code> . See the entry for <code>col</code> in the help file for par for more information.
<code>digits</code>	a scalar indicating how many significant digits to print for the distribution parameters. The default value is <code>digits=Options\$digits</code> .
<code>type, main, xlab, ylab, xlim, ylim, ...</code>	additional graphical parameters (see lines and par). In particular, the argument <code>type</code> specifies the kind of line type. By default, the function <code>cdfPlot</code> plots a step function (<code>type="s"</code>) for discrete distributions, and plots a straight line between points (<code>type="l"</code>) otherwise. The user may override these defaults by supplying the graphics parameter <code>type</code> (<code>type="s"</code> for a step function, <code>type="l"</code> for linear interpolation, <code>type="p"</code> for points only, etc.).

Details

The **cumulative distribution function (cdf)** of a random variable X , usually denoted F , is defined as:

$$F(x) = Pr(X \leq x) \quad (1)$$

That is, $F(x)$ is the probability that X is less than or equal to x . This is the probability that the random variable X takes on a value in the interval $(-\infty, x]$ and is simply the (Lebesgue) integral of the pdf evaluated between $-\infty$ and x . That is,

$$F(x) = Pr(X \leq x) = \int_{-\infty}^x f(t)dt \quad (2)$$

where $f(t)$ denotes the probability density function of X evaluated at t . For discrete distributions, Equation (2) translates to summing up the probabilities of all values in this interval:

$$F(x) = Pr(X \leq x) = \sum_{t \in (-\infty, x]} f(t) = \sum_{t \in (-\infty, x]} Pr(X = t) \quad (3)$$

A **cumulative distribution function (cdf) plot** plots the values of the cdf against quantiles of the specified distribution. Theoretical cdf plots are sometimes plotted along with [empirical cdf plots](#) to visually assess whether data have a particular distribution.

Value

cdfPlot invisibly returns a list giving coordinates of the points that have been or would have been plotted:

Quantiles	The quantiles used for the plot.
Cumulative.Probabilities	The values of the cdf associated with the quantiles.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.

Johnson, N. L., S. Kotz, and A.W. Kemp. (1992). *Univariate Discrete Distributions, Second Edition*. John Wiley and Sons, New York.

Johnson, N. L., S. Kotz, and N. Balakrishnan. (1994). *Continuous Univariate Distributions, Volume 1*. Second Edition. John Wiley and Sons, New York.

Johnson, N. L., S. Kotz, and N. Balakrishnan. (1995). *Continuous Univariate Distributions, Volume 2*. Second Edition. John Wiley and Sons, New York.

See Also

[Distribution.df](#), [ecdfPlot](#), [cdfCompare](#), [pdfPlot](#).

Examples

```
# Plot the cdf of the standard normal distribution
#-----
windows()
cdfPlot()

#=====

# Plot the cdf of the standard normal distribution
# and a N(2, 2) distribution on the sample plot.
#-----
windows()
cdfPlot(param.list = list(mean=2, sd=2), main = "")

cdfPlot(add = TRUE, cdf.col = "red")

legend("topleft", legend = c("N(2,2)", "N(0,1)"),
      col = c("black", "red"), lwd = 3 * par("cex"))

title("CDF Plots for Two Normal Distributions")

#=====

# Clean up
#-----
graphics.off()
```

Description

Density, distribution function, quantile function, and random generation for the chi distribution.

Usage

```
dchi(x, df)
pchi(q, df)
qchi(p, df)
rchi(n, df)
```

Arguments

x	vector of (positive) quantiles.
q	vector of (positive) quantiles.
p	vector of probabilities between 0 and 1.
n	sample size. If length(n) is larger than 1, then length(n) random values are returned.
df	vector of (positive) degrees of freedom (> 0). Non-integer values are allowed.

Details

Elements of x, q, p, or df that are missing will cause the corresponding elements of the result to be missing.

The chi distribution with n degrees of freedom is the distribution of the positive square root of a random variable having a [chi-squared](#) distribution with n degrees of freedom.

The chi density function is given by:

$$f(x, \nu) = g(x^2, \nu)2x, x > 0$$

where $g(x, \nu)$ denotes the density function of a chi-square random variable with n degrees of freedom.

Value

density (dchi), probability (pchi), quantile (qchi), or random sample (rchi) for the chi distribution with df degrees of freedom.

Note

The chi distribution takes on positive real values. It is important because for a sample of n observations from a [normal](#) distribution, the sample standard deviation multiplied by the square root of the degrees of freedom ν and divided by the true standard deviation follows a chi distribution with ν degrees of freedom. The chi distribution is also used in computing exact prediction intervals for the next k observations from a normal distribution (see [predIntNorm](#)).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.

Johnson, N. L., S. Kotz, and N. Balakrishnan. (1995). *Continuous Univariate Distributions, Volume 1*. Second Edition. John Wiley and Sons, New York.

See Also

[Chisquare](#), [Normal](#), [predIntNorm](#), [Probability Distributions and Random Numbers](#).

Examples

```
# Density of a chi distribution with 4 degrees of freedom, evaluated at 3:

dchi(3, 4)
#[1] 0.1499715

#-----

# The 95'th percentile of a chi distribution with 10 degrees of freedom:

qchi(.95, 10)
#[1] 4.278672

#-----

# The cumulative distribution function of a chi distribution with
# 5 degrees of freedom evaluated at 3:

pchi(3, 5)
#[1] 0.8909358

#-----

# A random sample of 2 numbers from a chi distribution with 7 degrees of freedom.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(20)
rchi(2, 7)
#[1] 3.271632 2.035179
```

ciBinomHalfWidth

Compute Half-Width of a Confidence Interval for a Binomial Proportion or the Difference Between Two Proportions

Description

Compute the half-width of a confidence interval for a binomial proportion or the difference between two proportions, given the sample size(s), estimated proportion(s), and confidence level.

Usage

```
ciBinomHalfWidth(n.or.n1, p.hat.or.p1.hat = 0.5,
  n2 = n.or.n1, p2.hat = 0.4, conf = 0.95,
  sample.type = "one.sample", ci.method = "score",
  correct = TRUE, warn = TRUE)
```

Arguments

n.or.n1	<p>numeric vector of sample sizes.</p> <p>When <code>sample.type="one.sample"</code>, <code>n.or.n1</code> denotes n, the number of observations in the single sample.</p> <p>When <code>sample.type="two.sample"</code>, <code>n.or.n1</code> denotes n_1, the number of observations from group 1.</p> <p>Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed.</p>
p.hat.or.p1.hat	<p>numeric vector of estimated proportions.</p> <p>When <code>sample.type="one.sample"</code>, <code>p.hat.or.p1.hat</code> denotes the estimated value of p, the probability of "success".</p> <p>When <code>sample.type="two.sample"</code>, <code>p.hat.or.p1.hat</code> denotes the estimated value of p_1, the probability of "success" in group 1.</p> <p>Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed.</p>
n2	<p>numeric vector of sample sizes for group 2. The default value is the value of <code>n.or.n1</code>. This argument is ignored when <code>sample.type="one.sample"</code>. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed.</p>
p2.hat	<p>numeric vector of estimated proportions for group 2. This argument is ignored when <code>sample.type="one.sample"</code>. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed.</p>
conf	<p>numeric vector of numbers between 0 and 1 indicating the confidence level associated with the confidence interval(s). The default value is <code>conf=0.95</code>.</p>
sample.type	<p>character string indicating whether this is a one-sample or two-sample confidence interval. When <code>sample.type="one.sample"</code>, the computed half-width is based on a confidence interval for a single proportion. When <code>sample.type="two.sample"</code>, the computed half-width is based on a confidence interval for the difference between two proportions. The default value is <code>sample.type="one.sample"</code> unless the argument <code>n2</code> or <code>p2.hat</code> is supplied.</p>
ci.method	<p>character string indicating which method to use to construct the confidence interval. Possible values are "score" (the default), "exact", "adjusted Wald", and "Wald" (the "Wald" method is never recommended but is included for historical purposes). The exact method is only available for the one-sample case, i.e., when <code>sample.type="one.sample"</code>.</p>
correct	<p>logical scalar indicating whether to use the continuity correction when <code>ci.method="score"</code> or <code>ci.method="Wald"</code>.</p> <p>The default value is <code>correct=TRUE</code>.</p>
warn	<p>logical scalar indicating whether to issue a warning when <code>ci.method="Wald"</code> for cases when the normal approximation to the binomial distribution probably is not accurate. The default value is <code>warn=TRUE</code>.</p>

Details

If the arguments `n.or.n1`, `p.hat.or.p1.hat`, `n2`, `p2.hat`, and `conf` are not all the same length, they are replicated to be the same length as the length of the longest argument.

The values of `p.hat.or.p1.hat` and `p2.hat` are automatically adjusted to the closest legitimate values, given the user-supplied values of `n.or.n1` and `n2`. For example, if `n.or.n1=5`, legitimate values for `p.hat.or.p1.hat` are 0, 0.2, 0.4, 0.6, 0.8 and 1. In this case, if the user supplies `p.hat.or.p1.hat=0.45`, then `p.hat.or.p1.hat` is reset to `p.hat.or.p1.hat=0.4`, and if the user supplies `p.hat.or.p1.hat=0.55`, then `p.hat.or.p1.hat` is reset to `p.hat.or.p1.hat=0.6`. In cases where the two closest legitimate values are equal distance from the user-supplied value of `p.hat.or.p1.hat` or `p2.hat`, the value closest to 0.5 is chosen since that will tend to yield the wider confidence interval.

One-Sample Case (`sample.type="one.sample"`).

`ci.method="score"` The confidence interval for p based on the score method was developed by Wilson (1927) and is discussed by Newcombe (1998a), Agresti and Coull (1998), and Agresti and Caffo (2000). When `ci=TRUE` and `ci.method="score"`, the function `ebinom` calls the R function `prop.test` to compute the confidence interval. This method has been shown to provide the best performance (in terms of actual coverage matching assumed coverage) of all the methods provided here, although unlike the exact method, the actual coverage can fall below the assumed coverage.

`ci.method="exact"` The confidence interval for p based on the exact (Clopper-Pearson) method is discussed by Newcombe (1998a), Agresti and Coull (1998), and Zar (2010, pp.543-547). This is the method used in the R function `binom.test`. This method ensures the actual coverage is greater than or equal to the assumed coverage.

`ci.method="Wald"` The confidence interval for p based on the Wald method (with or without a correction for continuity) is the usual “normal approximation” method and is discussed by Newcombe (1998a), Agresti and Coull (1998), Agresti and Caffo (2000), and Zar (2010, pp.543-547). This method is **never** recommended but is included for historical purposes.

`ci.method="adjusted Wald"` The confidence interval for p based on the adjusted Wald method is discussed by Agresti and Coull (1998), Agresti and Caffo (2000), and Zar (2010, pp.543-547). This is a simple modification of the Wald method and performs surprisingly well.

Two-Sample Case (`sample.type="two.sample"`).

`ci.method="score"` This method is presented in Newcombe (1998b) and is based on the score method developed by Wilson (1927) for the one-sample case. This is the method used by the R function `prop.test`. In a comparison of 11 methods, Newcombe (1998b) showed this method performs remarkably well.

`ci.method="Wald"` The confidence interval for the difference between two proportions based on the Wald method (with or without a correction for continuity) is the usual “normal approximation” method and is discussed by Newcombe (1998b), Agresti and Caffo (2000), and Zar (2010, pp.549-552). This method is **not** recommended but is included for historical purposes.

`ci.method="adjusted Wald"` This method is discussed by Agresti and Caffo (2000), and Zar (2010, pp.549-552). This is a simple modification of the Wald method and performs surprisingly well.

Value

a list with information about the half-widths, sample sizes, and estimated proportions.

One-Sample Case (`sample.type="one.sample"`).

When `sample.type="one.sample"`, the function `ciBinomHalfWidth` returns a list with these components:

`half.width` the half-width(s) of the confidence interval(s)

n	the sample size(s) associated with the confidence interval(s)
p.hat	the estimated proportion(s)
method	the method used to construct the confidence interval(s)

Two-Sample Case (sample.type="two.sample").

When sample.type="two.sample", the function ciBinomHalfWidth returns a list with these components:

half.width	the half-width(s) of the confidence interval(s)
n1	the sample size(s) for group 1 associated with the confidence interval(s)
p1.hat	the estimated proportion(s) for group 1
n2	the sample size(s) for group 2 associated with the confidence interval(s)
p2.hat	the estimated proportion(s) for group 2
method	the method used to construct the confidence interval(s)

Note

The binomial distribution is used to model processes with binary (Yes-No, Success-Failure, Heads-Tails, etc.) outcomes. It is assumed that the outcome of any one trial is independent of any other trial, and that the probability of "success", p , is the same on each trial. A binomial discrete random variable X is the number of "successes" in n independent trials. A special case of the binomial distribution occurs when $n = 1$, in which case X is also called a Bernoulli random variable.

In the context of environmental statistics, the binomial distribution is sometimes used to model the proportion of times a chemical concentration exceeds a set standard in a given period of time (e.g., Gilbert, 1987, p.143), or to compare the proportion of detects in a compliance well vs. a background well (e.g., USEPA, 1989b, Chapter 8, p.3-7). (However, USEPA 2009, p.8-27 recommends using the Wilcoxon rank sum test ([wilcox.test](#)) instead of comparing proportions.)

In the course of designing a sampling program, an environmental scientist may wish to determine the relationship between sample size, confidence level, and half-width if one of the objectives of the sampling program is to produce confidence intervals. The functions ciBinomHalfWidth, ciBinomN, and plotCiBinomDesign can be used to investigate these relationships for the case of binomial proportions.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Agresti, A., and B.A. Coull. (1998). Approximate is Better than "Exact" for Interval Estimation of Binomial Proportions. *The American Statistician*, **52**(2), 119–126.
- Agresti, A., and B. Caffo. (2000). Simple and Effective Confidence Intervals for Proportions and Differences of Proportions Result from Adding Two Successes and Two Failures. *The American Statistician*, **54**(4), 280–288.
- Berthouex, P.M., and L.C. Brown. (1994). *Statistics for Environmental Engineers*. Lewis Publishers, Boca Raton, FL, Chapters 2 and 15.
- Cochran, W.G. (1977). *Sampling Techniques*. John Wiley and Sons, New York, Chapter 3.
- Fisher, R.A., and F. Yates. (1963). *Statistical Tables for Biological, Agricultural, and Medical Research*. 6th edition. Hafner, New York, 146pp.

- Fleiss, J. L. (1981). *Statistical Methods for Rates and Proportions*. Second Edition. John Wiley and Sons, New York, Chapters 1-2.
- Gilbert, R.O. (1987). *Statistical Methods for Environmental Pollution Monitoring*. Van Nostrand Reinhold, New York, NY, Chapter 11.
- Millard, S.P., and Neerchal, N.K. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, Florida.
- Newcombe, R.G. (1998a). Two-Sided Confidence Intervals for the Single Proportion: Comparison of Seven Methods. *Statistics in Medicine*, **17**, 857–872.
- Newcombe, R.G. (1998b). Interval Estimation for the Difference Between Independent Proportions: Comparison of Eleven Methods. *Statistics in Medicine*, **17**, 873–890.
- Ott, W.R. (1995). *Environmental Statistics and Data Analysis*. Lewis Publishers, Boca Raton, FL, Chapter 4.
- USEPA. (1989b). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities, Interim Final Guidance*. EPA/530-SW-89-026. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.6-38.
- Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ, Chapter 24.

See Also

[ciBinomN](#), [plotCiBinomDesign](#), [ebinom](#), [binom.test](#), [prop.test](#).

Examples

```
# Look at how the half-width of a one-sample confidence interval
# decreases with sample size:
```

```
ciBinomHalfWidth(n.or.n1 = c(10, 50, 100, 500))
#$half.width
#[1] 0.26340691 0.13355486 0.09616847 0.04365873
#
#$n
#[1] 10 50 100 500
#
#$p.hat
#[1] 0.5 0.5 0.5 0.5
#
#$method
#[1] "Score normal approximation, with continuity correction"
```

```
#-----
```

```
# Look at how the half-width of a one-sample confidence interval
# tends to decrease as the estimated value of p decreases below
# 0.5 or increases above 0.5:
```

```
seq(0.2, 0.8, by = 0.1)
#[1] 0.2 0.3 0.4 0.5 0.6 0.7 0.8
```

```

ciBinomHalfWidth(n.or.n1 = 30, p.hat = seq(0.2, 0.8, by = 0.1))
#$half.width
#[1] 0.1536299 0.1707256 0.1801322 0.1684587 0.1801322 0.1707256
#[7] 0.1536299
#
#$n
#[1] 30 30 30 30 30 30 30
#
#$p.hat
#[1] 0.2 0.3 0.4 0.5 0.6 0.7 0.8
#
#$method
#[1] "Score normal approximation, with continuity correction"

#-----

# Look at how the half-width of a one-sample confidence interval
# increases with increasing confidence level:

ciBinomHalfWidth(n.or.n1 = 20, conf = c(0.8, 0.9, 0.95, 0.99))
#$half.width
#[1] 0.1377380 0.1725962 0.2007020 0.2495523
#
#$n
#[1] 20 20 20 20
#
#$p.hat
#[1] 0.5 0.5 0.5 0.5
#
#$method
#[1] "Score normal approximation, with continuity correction"

#-----

# Compare the half-widths for a one-sample
# confidence interval based on the different methods:

ciBinomHalfWidth(n.or.n1 = 30, ci.method = "score")$half.width
#[1] 0.1684587

ciBinomHalfWidth(n.or.n1 = 30, ci.method = "exact")$half.width
#[1] 0.1870297

ciBinomHalfWidth(n.or.n1 = 30, ci.method = "adjusted Wald")$half.width
#[1] 0.1684587

ciBinomHalfWidth(n.or.n1 = 30, ci.method = "Wald")$half.width
#[1] 0.1955861

#-----

# Look at how the half-width of a two-sample
# confidence interval decreases with increasing
# sample sizes:

ciBinomHalfWidth(n.or.n1 = c(10, 50, 100, 500), sample.type = "two")

```

```

#$half.width
#[1] 0.53385652 0.21402654 0.14719748 0.06335658
#
#$n1
#[1] 10 50 100 500
#
#$p1.hat
#[1] 0.5 0.5 0.5 0.5
#
#$n2
#[1] 10 50 100 500
#
#$p2.hat
#[1] 0.4 0.4 0.4 0.4
#
#$method
#[1] "Score normal approximation, with continuity correction"

```

ciBinomN

Compute Sample Size Necessary to Achieve Specified Half-Width of a Confidence Interval for a Binomial Proportion or the Difference Between Two Proportions

Description

Compute the sample size necessary to achieve a specified half-width of a confidence interval for a binomial proportion or the difference between two proportions, given the estimated proportion(s), and confidence level.

Usage

```

ciBinomN(half.width, p.hat.or.p1.hat = 0.5, p2.hat = 0.4,
  conf = 0.95, sample.type = "one.sample", ratio = 1,
  ci.method = "score", correct = TRUE, warn = TRUE,
  n.or.n1.min = 2, n.or.n1.max = 10000,
  tol.half.width = 5e-04, tol.p.hat = 5e-04,
  tol = 1e-7, maxiter = 1000)

```

Arguments

half.width	numeric vector of (positive) half-widths. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed.
p.hat.or.p1.hat	<p>numeric vector of estimated proportions.</p> <p>When sample.type="one.sample", p.hat.or.p1.hat denotes the estimated value of p, the probability of "success".</p> <p>When sample.type="two.sample", p.hat.or.p1.hat denotes the estimated value of p_1, the probability of "success" in group 1.</p> <p>Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed.</p>
p2.hat	<p>numeric vector of estimated proportions for group 2. This argument is ignored when sample.type="one.sample". Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed.</p>

conf	numeric vector of numbers between 0 and 1 indicating the confidence level associated with the confidence interval(s). The default value is conf=0.95.
sample.type	character string indicating whether this is a one-sample or two-sample confidence interval. When sample.type="one.sample", the computed half-width is based on a confidence interval for a single proportion. When sample.type="two.sample", the computed half-width is based on a confidence interval for the difference between two proportions. The default value is sample.type="one.sample" unless the argument p2.hat or ratio is supplied.
ratio	numeric vector indicating the ratio of sample size in group 2 to sample size in group 1 (n_2/n_1). The default value is ratio=1. All values of ratio must be greater than or equal to 1. This argument is ignored if sample.type="one.sample".
ci.method	character string indicating which method to use to construct the confidence interval. Possible values are: <ul style="list-style-type: none"> • "score" (the default), • "exact", • "adjusted Wald" and, • "Wald" (the "Wald" method is never recommended but is included for historical purposes). The exact method is only available for the one-sample case, i.e., when sample.type="one.sample".
correct	logical scalar indicating whether to use the continuity correction when ci.method="score" or ci.method="Wald". The default value is correct=TRUE.
warn	logical scalar indicating whether to issue a warning when ci.method="Wald" for cases when the normal approximation to the binomial distribution probably is not accurate. The default value is warn=TRUE.
n.or.n1.min	integer indicating the minimum allowed value for n (sample.type="one.sample") or n_1 (sample.type="two.sample"). The default value is n.or.n1.min=2.
n.or.n1.max	integer indicating the maximum allowed value for n (sample.type="one.sample") or n_1 (sample.type="two.sample"). The default value is n.or.n1.max=10000.
tol.half.width	numeric scalar indicating the tolerance to use for the half width for the search algorithm. The sample sizes are computed so that the actual half width is less than or equal to half.width + tol.half.width. The default value is tol.half.width=5e-04.
tol.p.hat	numeric scalar indicating the tolerance to use for the estimated proportion(s) for the search algorithm. For the one-sample case, the sample sizes are computed so that the absolute value of the difference between the user supplied value of p.hat.or.p1.hat and the actual estimated proportion is less than or equal to tol.p.hat. For the two-sample case, the sample sizes are computed so that the absolute value of the difference between the user supplied value of p.hat.or.p1.hat and the actual estimated proportion for group 1 is less than or equal to tol.p.hat, and the absolute value of the difference between the user

	supplied value of <code>p2.hat</code> and the actual estimated proportion for group 2 is less than or equal to <code>tol.p.hat</code> . The default value is <code>tol.p.hat=0.005</code> .
<code>tol</code>	positive scalar indicating the tolerance to use for the search algorithm (passed to <code>uniroot</code>). The default value is <code>tol=1e-7</code> .
<code>maxiter</code>	integer indicating the maximum number of iterations to use for the search algorithm (passed to <code>uniroot</code>). The default value is <code>maxiter=1000</code> .

Details

If the arguments `half.width`, `p.hat.or.p1.hat`, `p2.hat`, `conf` and `ratio` are not all the same length, they are replicated to be the same length as the length of the longest argument.

For the one-sample case, the arguments `p.hat.or.p1.hat`, `tol.p.hat`, `half.width`, and `tol.half.width` must satisfy:

$(p.hat.or.p1.hat + tol.p.hat + half.width + tol.half.width) \leq 1$,
and
 $(p.hat.or.p1.hat - tol.p.hat - half.width - tol.half.width) \geq 0$.

For the two-sample case, the arguments `p.hat.or.p1.hat`, `p2.hat`, `tol.p.hat`, `half.width`, and `tol.half.width` must satisfy:

$((p.hat.or.p1.hat + tol.p.hat) - (p2.hat - tol.p.hat) + half.width + tol.half.width) \leq 1$,
and
 $((p.hat.or.p1.hat - tol.p.hat) - (p2.hat + tol.p.hat) - half.width - tol.half.width) \geq -1$.

The function `ciBinomN` uses the search algorithm in the function `uniroot` to call the function `ciBinomHalfWidth` to find the values of n (`sample.type="one.sample"`) or n_1 and n_2 (`sample.type="two.sample"`) that satisfy the requirements for the half-width, estimated proportions, and confidence level. See the Details section of the help file for `ciBinomHalfWidth` for more information.

Value

a list with information about the sample sizes, estimated proportions, and half-widths.

One-Sample Case (`sample.type="one.sample"`).

When `sample.type="one.sample"`, the function `ciBinomN` returns a list with these components:

<code>n</code>	the sample size(s) associated with the confidence interval(s)
<code>p.hat</code>	the estimated proportion(s)
<code>half.width</code>	the half-width(s) of the confidence interval(s)
<code>method</code>	the method used to construct the confidence interval(s)

Two-Sample Case (`sample.type="two.sample"`).

When `sample.type="two.sample"`, the function `ciBinomN` returns a list with these components:

<code>n1</code>	the sample size(s) for group 1 associated with the confidence interval(s)
<code>n2</code>	the sample size(s) for group 2 associated with the confidence interval(s)
<code>p1.hat</code>	the estimated proportion(s) for group 1
<code>p2.hat</code>	the estimated proportion(s) for group 2
<code>half.width</code>	the half-width(s) of the confidence interval(s)
<code>method</code>	the method used to construct the confidence interval(s)

Note

The binomial distribution is used to model processes with binary (Yes-No, Success-Failure, Heads-Tails, etc.) outcomes. It is assumed that the outcome of any one trial is independent of any other trial, and that the probability of “success”, p , is the same on each trial. A binomial discrete random variable X is the number of “successes” in n independent trials. A special case of the binomial distribution occurs when $n = 1$, in which case X is also called a Bernoulli random variable.

In the context of environmental statistics, the binomial distribution is sometimes used to model the proportion of times a chemical concentration exceeds a set standard in a given period of time (e.g., Gilbert, 1987, p.143), or to compare the proportion of detects in a compliance well vs. a background well (e.g., USEPA, 1989b, Chapter 8, p.3-7). (However, USEPA 2009, p.8-27 recommends using the Wilcoxon rank sum test ([wilcox.test](#)) instead of comparing proportions.)

In the course of designing a sampling program, an environmental scientist may wish to determine the relationship between sample size, confidence level, and half-width if one of the objectives of the sampling program is to produce confidence intervals. The functions `ciBinomHalfWidth`, `ciBinomN`, and `plotCiBinomDesign` can be used to investigate these relationships for the case of binomial proportions.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Agresti, A., and B.A. Coull. (1998). Approximate is Better than "Exact" for Interval Estimation of Binomial Proportions. *The American Statistician*, **52**(2), 119–126.
- Agresti, A., and B. Caffo. (2000). Simple and Effective Confidence Intervals for Proportions and Differences of Proportions Result from Adding Two Successes and Two Failures. *The American Statistician*, **54**(4), 280–288.
- Berthouex, P.M., and L.C. Brown. (1994). *Statistics for Environmental Engineers*. Lewis Publishers, Boca Raton, FL, Chapters 2 and 15.
- Cochran, W.G. (1977). *Sampling Techniques*. John Wiley and Sons, New York, Chapter 3.
- Fisher, R.A., and F. Yates. (1963). *Statistical Tables for Biological, Agricultural, and Medical Research*. 6th edition. Hafner, New York, 146pp.
- Fleiss, J. L. (1981). *Statistical Methods for Rates and Proportions*. Second Edition. John Wiley and Sons, New York, Chapters 1-2.
- Gilbert, R.O. (1987). *Statistical Methods for Environmental Pollution Monitoring*. Van Nostrand Reinhold, New York, NY, Chapter 11.
- Millard, S.P., and Neerchal, N.K. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, Florida.
- Newcombe, R.G. (1998a). Two-Sided Confidence Intervals for the Single Proportion: Comparison of Seven Methods. *Statistics in Medicine*, **17**, 857–872.
- Newcombe, R.G. (1998b). Interval Estimation for the Difference Between Independent Proportions: Comparison of Eleven Methods. *Statistics in Medicine*, **17**, 873–890.
- Ott, W.R. (1995). *Environmental Statistics and Data Analysis*. Lewis Publishers, Boca Raton, FL, Chapter 4.
- USEPA. (1989b). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities, Interim Final Guidance*. EPA/530-SW-89-026. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C.

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.6-38.

Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ, Chapter 24.

See Also

[ciBinomHalfWidth](#), [uniroot](#), [plotCiBinomDesign](#), [ebinom](#), [binom.test](#), [prop.test](#).

Examples

```
# Look at how the required sample size of a one-sample
# confidence interval increases with decreasing
# required half-width:

ciBinomN(half.width = c(0.1, 0.05, 0.03))
#$n
#[1]  92 374 1030
#
#$p.hat
#[1] 0.5 0.5 0.5
#
#$half.width
#[1] 0.10010168 0.05041541 0.03047833
#
#$method
#[1] "Score normal approximation, with continuity correction"

#-----

# Note that the required sample size decreases if we are less
# stringent about how much the confidence interval width can
# deviate from the supplied value of the 'half.width' argument:

ciBinomN(half.width = c(0.1, 0.05, 0.03), tol.half.width = 0.005)
#$n
#[1]  84 314 782
#
#$p.hat
#[1] 0.5 0.5 0.5
#
#$half.width
#[1] 0.10456066 0.05496837 0.03495833
#
#$method
#[1] "Score normal approximation, with continuity correction"

#-----

# Look at how the required sample size for a one-sample
# confidence interval tends to decrease as the estimated
# value of p decreases below 0.5 or increases above 0.5:
```

```

seq(0.2, 0.8, by = 0.1)
#[1] 0.2 0.3 0.4 0.5 0.6 0.7 0.8

ciBinomN(half.width = 0.1, p.hat = seq(0.2, 0.8, by = 0.1))
# $n
#[1] 70 90 100 92 100 90 70
#
# $p.hat
#[1] 0.2 0.3 0.4 0.5 0.6 0.7 0.8
#
# $half.width
#[1] 0.09931015 0.09839843 0.09910818 0.10010168 0.09910818 0.09839843
#[7] 0.09931015
#
# $method
#[1] "Score normal approximation, with continuity correction"

#-----

# Look at how the required sample size for a one-sample
# confidence interval increases with increasing confidence level:

ciBinomN(half.width = 0.05, conf = c(0.8, 0.9, 0.95, 0.99))
# $n
#[1] 160 264 374 644
#
# $p.hat
#[1] 0.5 0.5 0.5 0.5
#
# $half.width
#[1] 0.05039976 0.05035948 0.05041541 0.05049152
#
# $method
#[1] "Score normal approximation, with continuity correction"

#-----

# Compare required sample size for a one-sample
# confidence interval based on the different methods:

ciBinomN(half.width = 0.05, ci.method = "score")
# $n
#[1] 374
#
# $p.hat
#[1] 0.5
#
# $half.width
#[1] 0.05041541
#
# $method
#[1] "Score normal approximation, with continuity correction"

ciBinomN(half.width = 0.05, ci.method = "exact")
# $n
#[1] 394
#

```

```

#$p.hat
#[1] 0.5
#
#$half.width
#[1] 0.05047916
#
#$method
#[1] "Exact"

ciBinomN(half.width = 0.05, ci.method = "adjusted Wald")
#$n
#[1] 374
#
#$p.hat
#[1] 0.5
#
#$half.width
#[1] 0.05041541
#
#$method
#[1] "Adjusted Wald normal approximation"

ciBinomN(half.width = 0.05, ci.method = "Wald")
#$n
#[1] 398
#
#$p.hat
#[1] 0.5
#
#$half.width
#[1] 0.05037834
#
#$method
#[1] "Wald normal approximation, with continuity correction"

#-----

# Look at how the required sample size of a two-sample
# confidence interval increases with decreasing
# required half-width:

ciBinomN(half.width = c(0.1, 0.05, 0.03), sample.type = "two")
#$n1
#[1] 210 778 2089
#
#$n2
#[1] 210 778 2089
#
#$p1.hat
#[1] 0.5000000 0.5000000 0.4997607
#
#$p2.hat
#[1] 0.4000000 0.3997429 0.4001915
#
#$half.width
#[1] 0.09943716 0.05047044 0.03049753
#

```

```
##$method
#[1] "Score normal approximation, with continuity correction"
```

ciNormHalfWidth	<i>Compute Half-Width of Confidence Interval for Mean of a Normal Distribution or Difference Between Two Means</i>
-----------------	--

Description

Compute the half-width of a confidence interval for the mean of a normal distribution or the difference between two means, given the sample size(s), estimated standard deviation, and confidence level.

Usage

```
ciNormHalfWidth(n.or.n1, n2 = n.or.n1,
  sigma.hat = 1, conf = 0.95,
  sample.type = ifelse(missing(n2), "one.sample", "two.sample"))
```

Arguments

n.or.n1	numeric vector of sample sizes. When sample.type="one.sample", this argument denotes n , the number of observations in the single sample. When sample.type="two.sample", this argument denotes n_1 , the number of observations from group 1. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed.
n2	numeric vector of sample sizes for group 2. The default value is the value of n.or.n1. This argument is ignored when sample.type="one.sample". Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed.
sigma.hat	numeric vector specifying the value(s) of the estimated standard deviation(s).
conf	numeric vector of numbers between 0 and 1 indicating the confidence level associated with the confidence interval(s). The default value is conf=0.95.
sample.type	character string indicating whether this is a one-sample (sample.type="one.sample") or two-sample (sample.type="two.sample") confidence interval. When sample.type="one.sample", the computed half-width is based on a confidence interval for a single mean. When sample.type="two.sample", the computed half-width is based on a confidence interval for the difference between two means. The default value is sample.type="one.sample" unless the argument n2 is supplied.

Details

If the arguments n.or.n1, n2, sigma.hat, and conf are not all the same length, they are replicated to be the same length as the length of the longest argument.

The formulas for the confidence interval of the mean of a normal distribution or the difference between two means are given in basic statistics texts, including Zar(2010) and Millard and Neerchal (2001). For the two-sample case, the function ciNormHalfWidth assumes the two populations have the same standard deviation.

Value

a numeric vector of half-widths.

Note

The normal distribution and lognormal distribution are probably the two most frequently used distributions to model environmental data. In order to make any kind of probability statement about a normally-distributed population (of chemical concentrations for example), you have to first estimate the mean and standard deviation (the population parameters) of the distribution. Once you estimate these parameters, it is often useful to characterize the uncertainty in the estimate of the mean. This is done with confidence intervals.

In the course of designing a sampling program, an environmental scientist may wish to determine the relationship between sample size, confidence level, and half-width if one of the objectives of the sampling program is to produce confidence intervals. The functions `ciNormHalfWidth`, `ciNormN`, and `plotCiNormDesign` can be used to investigate these relationships for the case of normally-distributed observations.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Berthouex, P.M., and L.C. Brown. (2002). *Statistics for Environmental Engineers*. Second Edition. Lewis Publishers, Boca Raton, FL.
- Gilbert, R.O. (1987). *Statistical Methods for Environmental Pollution Monitoring*. Van Nostrand Reinhold, New York, NY.
- Helsel, D.R., and R.M. Hirsch. (1992). *Statistical Methods in Water Resources Research*. Elsevier, New York, NY, Chapter 7.
- Millard, S.P., and N. Neerchal. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, FL.
- Ott, W.R. (1995). *Environmental Statistics and Data Analysis*. Lewis Publishers, Boca Raton, FL.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.21-3.
- Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ, Chapters 7 and 8.

See Also

`ciNormN`, `plotCiNormDesign`, `Normal`, `enorm`, `t.test`
[Estimating Distribution Parameters](#).

Examples

```
# Look at how the half-width of a one-sample confidence interval
# decreases with increasing sample size:

seq(5, 30, by = 5)
#[1] 5 10 15 20 25 30
```

```

hw <- ciNormHalfWidth(n.or.n1 = seq(5, 30, by = 5))

round(hw, 2)
#[1] 1.24 0.72 0.55 0.47 0.41 0.37

#-----

# Look at how the half-width of a one-sample confidence interval
# increases with increasing estimated standard deviation:

seq(0.5, 2, by = 0.5)
#[1] 0.5 1.0 1.5 2.0

hw <- ciNormHalfWidth(n.or.n1 = 20, sigma.hat = seq(0.5, 2, by = 0.5))

round(hw, 2)
#[1] 0.23 0.47 0.70 0.94

#-----

# Look at how the half-width of a one-sample confidence interval
# increases with increasing confidence level:

seq(0.5, 0.9, by = 0.1)
#[1] 0.5 0.6 0.7 0.8 0.9

hw <- ciNormHalfWidth(n.or.n1 = 20, conf = seq(0.5, 0.9, by = 0.1))

round(hw, 2)
#[1] 0.15 0.19 0.24 0.30 0.39

#-----

# Modifying the example on pages 21-4 to 21-5 of USEPA (2009),
# determine how adding another four months of observations to
# increase the sample size from 4 to 8 will affect the half-width
# of a two-sided 95% confidence interval for the Aldicarb level at
# the first compliance well.
#
# Use the estimated standard deviation from the first four months
# of data. (The data are stored in EPA.09.Ex.21.1.aldicarb.df.)
# Note that the half-width changes from 34% of the observed mean to
# 18% of the observed mean by increasing the sample size from
# 4 to 8.

EPA.09.Ex.21.1.aldicarb.df
#   Month   Well Aldicarb.ppb
#1      1 Well.1      19.9
#2      2 Well.1      29.6
#3      3 Well.1      18.7
#4      4 Well.1      24.2
#...

attach(EPA.09.Ex.21.1.aldicarb.df)

mu.hat <- mean(Aldicarb.ppb[Well=="Well.1"])

```

```

mu.hat
#[1] 23.1

sigma.hat <- sd(Aldicarb.ppb[Well=="Well.1"])

sigma.hat
#[1] 4.93491

hw.4 <- ciNormHalfWidth(n.or.n1 = 4, sigma.hat = sigma.hat)

hw.4
#[1] 7.852543

hw.8 <- ciNormHalfWidth(n.or.n1 = 8, sigma.hat = sigma.hat)

hw.8
#[1] 4.125688

100 * hw.4/mu.hat
#[1] 33.99369

100 * hw.8/mu.hat
#[1] 17.86012

rm(mu.hat, sigma.hat, hw.4, hw.8)
detach("EPA.09.Ex.21.1.aldicarb.df")

```

ciNormN

Compute Sample Size Necessary to Achieve Specified Half-Width of Confidence Interval for Mean of a Normal Distribution or Difference Between Two Means

Description

Compute the sample size necessary to achieve a specified half-width of a confidence interval for the mean of a normal distribution or the difference between two means, given the estimated standard deviation and confidence level.

Usage

```

ciNormN(half.width, sigma.hat = 1, conf = 0.95,
        sample.type = ifelse(is.null(n2), "one.sample", "two.sample"),
        n2 = NULL, round.up = TRUE, tol = 1e-07)

```

Arguments

half.width	numeric vector of (positive) half-widths. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed.
sigma.hat	numeric vector specifying the value(s) of the estimated standard deviation(s).
conf	numeric vector of numbers between 0 and 1 indicating the confidence level associated with the confidence interval(s). The default value is conf=0.95.

<code>sample.type</code>	character string indicating whether this is a one-sample (<code>sample.type="one.sample"</code>) or two-sample (<code>sample.type="two.sample"</code>) confidence interval. When <code>sample.type="one.sample"</code> , the computed sample size is based on a confidence interval for a single mean. When <code>sample.type="two.sample"</code> , the computed sample size is based on a confidence interval for the difference between two means. The default value is <code>sample.type="one.sample"</code> unless the argument <code>n2</code> is supplied.
<code>n2</code>	numeric vector of sample sizes for group 2. The default value is <code>NULL</code> , in which case it is assumed that the sample sizes for groups 1 and 2 are equal. This argument is ignored when <code>sample.type="one.sample"</code> . Missing (<code>NA</code>), undefined (<code>NaN</code>), and infinite (<code>Inf</code> , <code>-Inf</code>) values are not allowed.
<code>round.up</code>	logical scalar indicating whether to round up the values of the computed sample size(s) to the next smallest integer. The default value is <code>round.up=TRUE</code> .
<code>tol</code>	numeric scalar indicating the tolerance to use in the search algorithm. The default value is <code>tol=1e-7</code> .

Details

If the arguments `half.width`, `n2`, `sigma.hat`, and `conf` are not all the same length, they are replicated to be the same length as the length of the longest argument.

The formulas for the confidence interval of the mean of a normal distribution or the difference between two means are given in basic statistics texts, including Zar(2010) and Millard and Neerchal (2001). For the two-sample case, the function `ciNormN` assumes the two populations have the same standard deviation.

For the one-sample case, the function `ciNormN` uses formula (7.9) shown in Zar (2010, p.115) to iteratively solve for the sample size.

For the two-sample case, the function `ciNormN` uses formulas (8.20) and (8.21) shown in Zar (2010, p.146) to iteratively solve for the sample size.

Value

When `sample.type="one.sample"`, or `sample.type="two.sample"` and `n2` is not supplied (so equal sample sizes for each group is assumed), the function `ciNormN` returns a numeric vector of sample sizes. When `sample.type="two.sample"` and `n2` is supplied, the function `ciNormN` returns a list with two components called `n1` and `n2`, specifying the sample sizes for each group.

Note

The normal distribution and lognormal distribution are probably the two most frequently used distributions to model environmental data. In order to make any kind of probability statement about a normally-distributed population (of chemical concentrations for example), you have to first estimate the mean and standard deviation (the population parameters) of the distribution. Once you estimate these parameters, it is often useful to characterize the uncertainty in the estimate of the mean. This is done with confidence intervals.

In the course of designing a sampling program, an environmental scientist may wish to determine the relationship between sample size, confidence level, and half-width if one of the objectives of the sampling program is to produce confidence intervals. The functions `ciNormHalfWidth`, `ciNormN`, and `plotCiNormDesign` can be used to investigate these relationships for the case of normally-distributed observations.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Berthouex, P.M., and L.C. Brown. (2002). *Statistics for Environmental Engineers*. Second Edition. Lewis Publishers, Boca Raton, FL.
- Gilbert, R.O. (1987). *Statistical Methods for Environmental Pollution Monitoring*. Van Nostrand Reinhold, New York, NY.
- Helsel, D.R., and R.M. Hirsch. (1992). *Statistical Methods in Water Resources Research*. Elsevier, New York, NY, Chapter 7.
- Millard, S.P., and N. Neerchal. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, FL.
- Ott, W.R. (1995). *Environmental Statistics and Data Analysis*. Lewis Publishers, Boca Raton, FL.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.21-3.
- Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ, Chapters 7 and 8.

See Also

[ciNormHalfWidth](#), [plotCiNormDesign](#), [Normal](#), [enorm](#), [t.test](#), [Estimating Distribution Parameters](#).

Examples

```
# Look at how the required sample size for a one-sample
# confidence interval decreases with increasing half-width:
```

```
seq(0.25, 1, by = 0.25)
#[1] 0.25 0.50 0.75 1.00
```

```
ciNormN(half.width = seq(0.25, 1, by = 0.25))
#[1] 64 18 10 7
```

```
ciNormN(seq(0.25, 1, by=0.25), round = FALSE)
#[1] 63.897899 17.832337 9.325967 6.352717
```

```
#-----
```

```
# Look at how the required sample size for a one-sample
# confidence interval increases with increasing estimated
# standard deviation for a fixed half-width:
```

```
seq(0.5, 2, by = 0.5)
#[1] 0.5 1.0 1.5 2.0
```

```
ciNormN(half.width = 0.5, sigma.hat = seq(0.5, 2, by = 0.5))
#[1] 7 18 38 64
```

```
#-----
```

```

# Look at how the required sample size for a one-sample
# confidence interval increases with increasing confidence
# level for a fixed half-width:

seq(0.5, 0.9, by = 0.1)
#[1] 0.5 0.6 0.7 0.8 0.9

ciNormN(half.width = 0.25, conf = seq(0.5, 0.9, by = 0.1))
#[1] 9 13 19 28 46

#-----

# Modifying the example on pages 21-4 to 21-5 of USEPA (2009),
# determine the required sample size in order to achieve a
# half-width that is 10% of the observed mean (based on the first
# four months of observations) for the Aldicarb level at the first
# compliance well. Assume a 95% confidence level and use the
# estimated standard deviation from the first four months of data.
# (The data are stored in EPA.09.Ex.21.1.aldicarb.df.)
#
# The required sample size is 20, so almost two years of data are
# required assuming observations are taken once per month.

EPA.09.Ex.21.1.aldicarb.df
#   Month   Well Aldicarb.ppb
#1      1 Well.1      19.9
#2      2 Well.1      29.6
#3      3 Well.1      18.7
#4      4 Well.1      24.2
#...

attach(EPA.09.Ex.21.1.aldicarb.df)

mu.hat <- mean(Aldicarb.ppb[Well=="Well.1"])

mu.hat
#[1] 23.1

sigma.hat <- sd(Aldicarb.ppb[Well=="Well.1"])

sigma.hat
#[1] 4.93491

ciNormN(half.width = 0.1 * mu.hat, sigma.hat = sigma.hat)
#[1] 20

#-----
# Clean up

rm(mu.hat, sigma.hat)
detach("EPA.09.Ex.21.1.aldicarb.df")

```

Description

Compute the sample coefficient of variation.

Usage

```
cv(x, method = "moments", sd.method = "sqrt.unbiased",
   l.moment.method = "unbiased", plot.pos.cons = c(a = 0.35, b = 0),
   na.rm = FALSE)
```

Arguments

<code>x</code>	numeric vector of observations.
<code>method</code>	character string specifying what method to use to compute the sample coefficient of variation. The possible values are "moments" (product moment ratio estimator; the default), or "l.moments" (L-moment ratio estimator).
<code>sd.method</code>	character string specifying what method to use to compute the sample standard deviation when <code>method="moments"</code> . The possible values are "sqrt.ubaised" (the square root of the unbiased estimate of variance; the default), or "moments" (the method of moments estimator).
<code>l.moment.method</code>	character string specifying what method to use to compute the <i>L</i> -moments when <code>method="l.moments"</code> . The possible values are "ubaised" (method based on the <i>U</i> -statistic; the default), or "plotting.position" (method based on the plotting position formula).
<code>plot.pos.cons</code>	numeric vector of length 2 specifying the constants used in the formula for the plotting positions when <code>method="l.moments"</code> and <code>l.moment.method="plotting.position"</code> . The default value is <code>plot.pos.cons=c(a=0.35, b=0)</code> . If this vector has a names attribute with the value <code>c("a", "b")</code> or <code>c("b", "a")</code> , then the elements will be matched by name in the formula for computing the plotting positions. Otherwise, the first element is mapped to the name "a" and the second element to the name "b".
<code>na.rm</code>	logical scalar indicating whether to remove missing values from <code>x</code> . If <code>na.rm=FALSE</code> (the default) and <code>x</code> contains missing values, then a missing value (NA) is returned. If <code>na.rm=TRUE</code> , missing values are removed from <code>x</code> prior to computing the coefficient of variation.

Details

Let \underline{x} denote a random sample of n observations from some distribution with mean μ and standard deviation σ .

Product Moment Coefficient of Variation (`method="moments"`)

The coefficient of variation (sometimes denoted CV) of a distribution is defined as the ratio of the standard deviation to the mean. That is:

$$CV = \frac{\sigma}{\mu} \quad (1)$$

The coefficient of variation measures how spread out the distribution is relative to the size of the mean. It is usually used to characterize positive, right-skewed distributions such as the lognormal distribution.

When `sd.method="sqrt.unbiased"`, the coefficient of variation is estimated using the sample mean and the square root of the unbiased estimator of variance:

$$\widehat{CV} = \frac{s}{\bar{x}} \quad (2)$$

where

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i \quad (3)$$

$$s = \left[\frac{1}{n-1} \sum_{i=1}^n (x_i - \bar{x})^2 \right]^{1/2} \quad (4)$$

Note that the estimator of standard deviation in equation (4) is not unbiased.

When `sd.method="moments"`, the coefficient of variation is estimated using the sample mean and the square root of the method of moments estimator of variance:

$$\widehat{CV} = \frac{s_m}{\bar{x}} \quad (5)$$

$$s = \left[\frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2 \right]^{1/2} \quad (6)$$

L-Moment Coefficient of Variation (`method="l.moments"`)

Hosking (1990) defines an *L*-moment analog of the coefficient of variation (denoted the *L*-CV) as:

$$\tau = \frac{l_2}{l_1} \quad (7)$$

that is, the second *L*-moment divided by the first *L*-moment. He shows that for a positive-valued random variable, the *L*-CV lies in the interval (0, 1).

When `l.moment.method="unbiased"`, the *L*-CV is estimated by:

$$t = \frac{l_2}{l_1} \quad (8)$$

that is, the unbiased estimator of the second *L*-moment divided by the unbiased estimator of the first *L*-moment.

When `l.moment.method="plotting.position"`, the *L*-CV is estimated by:

$$\tilde{t} = \frac{\tilde{l}_2}{\tilde{l}_1} \quad (9)$$

that is, the plotting-position estimator of the second *L*-moment divided by the plotting-position estimator of the first *L*-moment.

See the help file for [lMoment](#) for more information on estimating *L*-moments.

Value

A numeric scalar – the sample coefficient of variation.

Note

Traditionally, the coefficient of variation has been estimated using product moment estimators. Hosking (1990) introduced the idea of *L*-moments and the *L*-CV. Vogel and Fennessey (1993) argue that *L*-moment ratios should replace product moment ratios because of their superior performance (they are nearly unbiased and better for discriminating between distributions).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Berthouex, P.M., and L.C. Brown. (2002). *Statistics for Environmental Engineers, Second Edition*. Lewis Publishers, Boca Raton, FL.

Gilbert, R.O. (1987). *Statistical Methods for Environmental Pollution Monitoring*. Van Nostrand Reinhold, NY.

Ott, W.R. (1995). *Environmental Statistics and Data Analysis*. Lewis Publishers, Boca Raton, FL.

Taylor, J.K. (1990). *Statistical Techniques for Data Analysis*. Lewis Publishers, Boca Raton, FL.

Vogel, R.M., and N.M. Fennessey. (1993). *L Moment Diagrams Should Replace Product Moment Diagrams*. *Water Resources Research* **29**(6), 1745–1752.

Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ.

See Also

[Summary Statistics](#), [summaryFull](#), [var](#), [sd](#), [skewness](#), [kurtosis](#).

Examples

```
# Generate 20 observations from a lognormal distribution with
# parameters mean=10 and cv=1, and estimate the coefficient of variation.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
dat <- rlnormAlt(20, mean = 10, cv = 1)

cv(dat)
#[1] 0.5077981

cv(dat, sd.method = "moments")
#[1] 0.4949403

cv(dat, method = "l.moments")
#[1] 0.2804148

#-----
# Clean up
rm(dat)
```

detectionLimitCalibrate

Determine Detection Limit

Description

Determine the detection limit based on using a calibration line (or curve) and inverse regression.

Usage

```
detectionLimitCalibrate(object, coverage = 0.99,
  simultaneous = FALSE)
```

Arguments

object	an object that is the result of calling the function <code>calibrate</code> .
coverage	optional numeric scalar between 0 and 1 indicating the confidence level associated with the prediction intervals used in determining the detection limit. The default value is <code>coverage=0.99</code> .
simultaneous	optional logical scalar indicating whether to base the prediction intervals on simultaneous or non-simultaneous prediction limits. The default value is <code>simultaneous=FALSE</code> .

Details

The idea of a decision limit and detection limit is directly related to calibration and can be framed in terms of a hypothesis test, as shown in the table below. The null hypothesis is that the chemical is not present in the physical sample, i.e., $H_0 : C = 0$, where C denotes the concentration.

Your Decision	H_0 True ($C = 0$)	H_0 False ($C > 0$)
Reject H_0 (Declare Chemical Present)	Type I Error (Probability = α)	
Do Not Reject H_0 (Declare Chemical Absent)		Type II Error (Probability = β)

Ideally, you would like to minimize both the Type I and Type II error rates. Just as we use critical values to compare against the test statistic for a hypothesis test, we need to use a critical signal level S_D called the **decision limit** to decide whether the chemical is present or absent. If the signal is less than or equal to S_D we will declare the chemical is absent, and if the signal is greater than S_D we will declare the chemical is present.

First, suppose no chemical is present (i.e., the null hypothesis is true). If we want to guard against the mistake of declaring that the chemical is present when in fact it is absent (Type I error), then we should choose S_D so that the probability of this happening is some small value α . Thus, the value of S_D depends on what we want to use for α (the Type I error rate), and the true (but unknown) value of σ (the standard deviation of the errors assuming a constant standard deviation) (Massart et al., 1988, p. 111).

When the true concentration is 0, the decision limit is the $(1-\alpha)$ 100th percentile of the distribution of the signal S . Note that the decision limit is on the scale of and in units of the signal S .

Now suppose that in fact the chemical is present in some concentration C (i.e., the null hypothesis is false). If we want to guard against the mistake of declaring that the chemical is absent when in fact it is present (Type II error), then we need to determine a minimal concentration C_{DL} called the **detection limit (DL)** that we know will yield a signal less than the decision limit S_D only a small fraction of the time (β).

In practice we do not know the true value of the standard deviation of the errors (σ), so we cannot compute the true decision limit. Also, we do not know the true values of the intercept and slope of the calibration line, so we cannot compute the true detection limit. Instead, we usually set $\alpha = \beta$

and estimate the decision and detection limits by computing prediction limits for the calibration line and using inverse regression.

The estimated detection limit corresponds to the upper confidence bound on concentration given that the signal is equal to the estimated decision limit. Currie (1997) discusses other ways to define the detection limit, and Glaser et al. (1981) define a quantity called the method detection limit.

Value

A numeric scalar indicating the detection limit. This scalar has two attributes called coverage and simultaneous indicating the values of these arguments that were used in the call to detectionLimitCalibrate.

Note

Perhaps no other topic in environmental statistics has generated as much confusion or controversy as the topic of detection limits. After decades of disparate terminology, ISO and IUPAC provided harmonized guidance on the topic in 1995 (Currie, 1997). Intuitively, the idea of a detection limit is simple to grasp: the **detection limit** is “the smallest amount or concentration of a particular substance that can be reliably detected in a given type of sample or medium by a specific measurement process” (Currie, 1997, p. 152). Unfortunately, because of the exceedingly complex nature of measuring chemical concentrations, this simple idea is difficult to apply in practice.

Detection and quantification capabilities are fundamental performance characteristics of the **Chemical Measurement Process (CMP)** (Currie, 1996, 1997). In this help file we discuss some currently accepted definitions of the terms decision, detection, and quantification limits. For more details, the reader should consult the references listed in this help file.

The **quantification limit** is defined as the concentration C at which the coefficient of variation (also called relative standard deviation or RSD) for the distribution of the signal S is some small value, usually taken to be 10% (Currie, 1968, 1997). In practice the quantification limit is difficult to estimate because we have to estimate both the mean and the standard deviation of the signal S for any particular concentration, and usually the standard deviation varies with concentration. Variations of the quantification limit include the quantitation limit (Keith, 1991, p. 109), minimum level (USEPA, 1993), and alternative minimum level (Gibbons et al., 1997a).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Clark, M.J.R., and P.H. Whitfield. (1994). Conflicting Perspectives About Detection Limits and About the Censoring of Environmental Data. *Water Resources Bulletin* **30**(6), 1063–1079.
- Clayton, C.A., J.W. Hines, and P.D. Elkins. (1987). Detection Limits with Specified Assurance Probabilities. *Analytical Chemistry* **59**, 2506–2514.
- Code of Federal Regulations. (1996). Definition and Procedure for the Determination of the Method Detection Limit–Revision 1.11. Title 40, Part 136, Appendix B, 7-1-96 Edition, pp.265–267.
- Currie, L.A. (1968). Limits for Qualitative Detection and Quantitative Determination: Application to Radiochemistry. *Annals of Chemistry* **40**, 586–593.
- Currie, L.A. (1988). *Detection in Analytical Chemistry: Importance, Theory, and Practice*. American Chemical Society, Washington, D.C.
- Currie, L.A. (1995). Nomenclature in Evaluation of Analytical Methods Including Detection and Quantification Capabilities. *Pure & Applied Chemistry* **67**(10), 1699–1723.

- Currie, L.A. (1996). Foundations and Future of Detection and Quantification Limits. *Proceedings of the Section on Statistics and the Environment*, American Statistical Association, Alexandria, VA.
- Currie, L.A. (1997). Detection: International Update, and Some Emerging Di-Lemmas Involving Calibration, the Blank, and Multiple Detection Decisions. *Chemometrics and Intelligent Laboratory Systems* **37**, 151–181.
- Davis, C.B. (1994). Environmental Regulatory Statistics. In Patil, G.P., and C.R. Rao, eds., *Handbook of Statistics, Vol. 12: Environmental Statistics*. North-Holland, Amsterdam, a division of Elsevier, New York, NY, Chapter 26, 817–865.
- Davis, C.B. (1997). Challenges in Regulatory Environmetrics. *Chemometrics and Intelligent Laboratory Systems* **37**, 43–53.
- Gibbons, R.D. (1995). Some Statistical and Conceptual Issues in the Detection of Low-Level Environmental Pollutants (with Discussion). *Environmetrics* **2**, 125–167.
- Gibbons, R.D., D.E. Coleman, and R.F. Maddalone. (1997a). An Alternative Minimum Level Definition for Analytical Quantification. *Environmental Science & Technology* **31**(7), 2071–2077. Comments and Discussion in Volume **31**(12), 3727–3731, and Volume **32**(15), 2346–2353.
- Gibbons, R.D., D.E. Coleman, and R.F. Maddalone. (1997b). Response to Comment on “An Alternative Minimum Level Definition for Analytical Quantification”. *Environmental Science and Technology* **31**(12), 3729–3731.
- Gibbons, R.D., D.E. Coleman, and R.F. Maddalone. (1998). Response to Comment on “An Alternative Minimum Level Definition for Analytical Quantification”. *Environmental Science and Technology* **32**(15), 2349–2353.
- Gibbons, R.D., N.E. Grams, F.H. Jarke, and K.P. Stoub. (1992). Practical Quantitation Limits. *Chemometrics Intelligent Laboratory Systems* **12**, 225–235.
- Gibbons, R.D., F.H. Jarke, and K.P. Stoub. (1991). Detection Limits: For Linear Calibration Curves with Increasing Variance and Multiple Future Detection Decisions. In Tatsch, D.E., editor. *Waste Testing and Quality Assurance: Volume 3*. American Society for Testing and Materials, Philadelphia, PA.
- Glasser, J.A., D.L. Foerst, G.D. McKee, S.A. Quave, and W.L. Budde. (1981). Trace Analyses for Wastewaters. *Environmental Science and Technology* **15**, 1426–1435.
- Hubaux, A., and G. Vos. (1970). Decision and Detection Limits for Linear Calibration Curves. *Annals of Chemistry* **42**, 849–855.
- Kahn, H.D., C.E. White, K. Stralka, and R. Kuznetsovski. (1997). Alternative Estimates of Detection. *Proceedings of the Twentieth Annual EPA Conference on Analysis of Pollutants in the Environment, May 7-8, Norfolk, VA*. U.S. Environmental Protection Agency, Washington, D.C.
- Kahn, H.D., W.A. Telliard, and C.E. White. (1998). Comment on “An Alternative Minimum Level Definition for Analytical Quantification” (with Response). *Environmental Science & Technology* **32**(5), 2346–2353.
- Kaiser, H. (1965). Zum Problem der Nachweisgrenze. *Fresenius’ Z. Anal. Chem.* **209**, 1.
- Keith, L.H. (1991). *Environmental Sampling and Analysis: A Practical Guide*. Lewis Publishers, Boca Raton, FL, Chapter 10.
- Kimbrough, D.E. (1997). Comment on “An Alternative Minimum Level Definition for Analytical Quantification” (with Response). *Environmental Science & Technology* **31**(12), 3727–3731.
- Lambert, D., B. Peterson, and I. Terpenning. (1991). Nondetects, Detection Limits, and the Probability of Detection. *Journal of the American Statistical Association* **86**(414), 266–277.
- Massart, D.L., B.G.M. Vandeginste, S.N. Deming, Y. Michotte, and L. Kaufman. (1988). *Chemometrics: A Textbook*. Elsevier, New York, Chapter 7.

- Millard, S.P., and Neerchal, N.K. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, Florida.
- Porter, P.S., R.C. Ward, and H.F. Bell. (1988). The Detection Limit. *Environmental Science & Technology* **22**(8), 856–861.
- Rocke, D.M., and S. Lorenzato. (1995). A Two-Component Model for Measurement Error in Analytical Chemistry. *Technometrics* **37**(2), 176–184.
- Singh, A. (1993). Multivariate Decision and Detection Limits. *Analytica Chimica Acta* **277**, 205–214.
- Spiegelman, C.H. (1997). A Discussion of Issues Raised by Lloyd Currie and a Cross Disciplinary View of Detection Limits and Estimating Parameters That Are Often At or Near Zero. *Chemometrics and Intelligent Laboratory Systems* **37**, 183–188.
- USEPA. (1987c). List (Phase 1) of Hazardous Constituents for Ground-Water Monitoring; Final Rule. *Federal Register* **52**(131), 25942–25953 (July 9, 1987).
- Zorn, M.E., R.D. Gibbons, and W.C. Sonzogni. (1997). Weighted Least-Squares Approach to Calculating Limits of Detection and Quantification by Modeling Variability as a Function of Concentration. *Analytical Chemistry* **69**, 3069–3075.

See Also

[calibrate](#), [inversePredictCalibrate](#), [pointwise](#).

Examples

```
# The data frame EPA.97.cadmium.111.df contains calibration
# data for cadmium at mass 111 (ng/L) that appeared in
# Gibbons et al. (1997b) and were provided to them by the U.S. EPA.
#
# The Example section in the help file for calibrate shows how to
# plot these data along with the fitted calibration line and 99%
# non-simultaneous prediction limits.
#
# For the current example, we will compute the decision limit (7.68)
# and detection limit (12.36 ng/L) based on using alpha = beta = 0.01
# and a linear calibration line with constant variance. See
# Millard and Neerchal (2001, pp.566–575) for more details on this
# example.

calibrate.list <- calibrate(Cadmium ~ Spike,
  data = EPA.97.cadmium.111.df)

# Compute decision limit

pred.list <- predict(calibrate.list, newdata = data.frame(Spike=0),
  se.fit = TRUE)

decision.limit <- pointwise(pred.list, coverage = 0.99,
  individual = TRUE)$upper

decision.limit
#      1
#7.677842

# Compute detection limit
```

```
detection.limit <- detectionLimitCalibrate(calibrate.list)

detection.limit
#[1] 12.36467
#attr(,"coverage"):
#[1] 0.99
#attr(,"simultaneous"):
#[1] FALSE

rm(calibrate.list, pred.list, decision.limit, detection.limit)
```

Distribution.df	<i>Data Frame Summarizing Available Probability Distributions and Estimation Methods</i>
-----------------	--

Description

Data frame summarizing information about available probability distributions in **R** and the **EnvStats** package, and which distributions have associated functions for estimating distribution parameters.

Usage

```
Distribution.df
```

Format

A data frame with 35 rows corresponding to 35 different available probability distributions, and 25 columns containing information associated with these probability distributions.

Name a character vector containing the name of the probability distribution (see the column labeled **Name** in the table below).

Type a character vector indicating the type of distribution (see the column labeled **Type** in the table below). Possible values are "Finite Discrete", "Discrete", "Continuous", and "Mixed".

Support.Min a character vector indicating the minimum value the random variable can assume (see the column labeled **Range** in the table below). The reason this is a character vector instead of a numeric vector is because some distributions have a lower bound that depends on the value of a distribution parameter. For example, the minimum value for a [Uniform](#) distribution is given by the value of the parameter min.

Support.Max a character vector indicating the maximum value the random variable can assume (see the column labeled **Range** in the table below). The reason this is a character vector instead of a numeric vector is because some distributions have an upper bound that depends on the value of a distribution parameter. For example, the maximum value for a [Uniform](#) distribution is given by the value of the parameter max.

Estimation.Method(s) a character vector indicating the names of the methods available to estimate the distribution parameter(s) (see the column labeled **Estimation Method(s)** in the table below). Possible values include "mle" (maximum likelihood), "mme" (method of moments), "mmue" (method of moments based on the unbiased estimate of variance), "mvue" (minimum variance unbiased), "qmle" (quasi-mle), etc., or some combination of these. In cases where an estimator is more than one kind, a slash (/) is used to denote all methods covered by the

single estimator. For example, for the Binomial distribution, the sample proportion is the maximum likelihood, method of moments, and minimum variance unbiased estimator, so this method is denoted as "mle/mme/mvue". See the help files for the specific function listed under [Estimating Distribution Parameters](#) for an explanation of each of these estimation methods.

Quantile.Estimation.Method(s) a character vector indicating the names of the methods available to estimate the distribution quantiles. For many distributions, these are the same as Estimation.Method(s). See the help files for the specific function listed under [Estimating Distribution Quantiles](#) for an explanation of each of these estimation methods.

Prediction.Interval.Method(s) a character vector indicating the names of the methods available to create prediction intervals. See the help files for the specific function listed under [Prediction Intervals](#) for an explanation of each of these estimation methods.

Singly.Censored.Estimation.Method(s) a character vector indicating the names of the methods available to estimate the distribution parameter(s) for Type I singly-censored data. See the help files for the specific function listed under *Estimating Distribution Parameters* in the help file for [Censored Data](#) for an explanation of each of these estimation methods.

Multiply.Censored.Estimation.Method(s) a character vector indicating the names of the methods available to estimate the distribution parameter(s) for Type I multiply-censored data. See the help files for the specific function listed under *Estimating Distribution Parameters* in the help file for [Censored Data](#) for an explanation of each of these estimation methods.

Number.parameters a numeric vector indicating the number of parameters associated with the distribution (see the column labeled **Parameters** in the table below).

Parameter.1 the columns labeled Parameter.1, Parameter.2, ..., Parameter.5 are character vectors containing the names of the distribution parameters (see the column labeled **Parameters** in the table below). If a distribution has n parameters and $n < 5$, then the columns labeled Parameter.n+1, ..., Parameter.5 are empty. For example, the [Normal](#) distribution has only two parameters associated with it (mean and sd), so the fields in Parameter.3, Parameter.4, and Parameter.5 are empty.

Parameter.2 see Parameter.1

Parameter.3 see Parameter.1

Parameter.4 see Parameter.1

Parameter.5 see Parameter.1

Parameter.1.Min the columns labeled Parameter.1.Min, Parameter.2.Min, ..., Parameter.5.Min are character vectors containing the minimum values that can be assumed by the distribution parameters (see the column labeled **Parameter Range(s)** in the table below).

The reason these are character vectors instead of numeric vectors is because some parameters have a lower bound of 0 but must be strictly bigger than 0 (e.g., the parameter sd for the [Normal](#) distribution), in which case the lower bound is `.Machine$double.eps`, which may vary from machine to machine. Also, some parameters have a lower bound that depends on the value of another parameter. For example, the parameter max for a [Uniform](#) distribution is bounded below by the value of the parameter min.

If a distribution has n parameters and $n < 5$, then the columns labeled Parameter.n+1.Min, ..., Parameter.5.Min have the missing value code (NA). For example, the [Normal](#) distribution has only two parameters associated with it (mean and sd) so the fields in Parameter.3.Min, Parameter.4.Min, and Parameter.5.Min have NAs in them.

Parameter.2.Min see Parameter.1.Min

Parameter.3.Min see Parameter.1.Min

Parameter.4.Min see Parameter.1.Min

Parameter.5.Min see Parameter.1.Min

Parameter.1.Max the columns labeled Parameter.1.Max, Parameter.2.Max, ..., Parameter.5.Max are character vectors containing the maximum values that can be assumed by the distribution parameters (see the column labeled **Parameter Range(s)** in the table below).

The reason these are character vectors instead of numeric vectors is because some parameters have an upper bound that depends on the value of another parameter. For example, the parameter min for a [Uniform](#) distribution is bounded above by the value of the parameter max. If a distribution has n parameters and $n < 5$, then the columns labeled Parameter.n+1.Max, ..., Parameter.5.Max have the missing value code (NA). For example, the [Normal](#) distribution has only two parameters associated with it (mean and sd) so the fields in Parameter.3.Max, Parameter.4.Max, and Parameter.5.Max have NAs in them.

Parameter.2.Max see Parameter.1.Max

Parameter.3.Max see Parameter.1.Max

Parameter.4.Max see Parameter.1.Max

Parameter.5.Max see Parameter.1.Max

Details

The table below summarizes the probability distributions available in R and **EnvStats**. For each distribution, there are four associated functions for computing density values, percentiles, quantiles, and random numbers. The form of the names of these functions are *dabb*, *pabb*, *qabb*, and *rabb*, where *abb* is the abbreviated name of the distribution (see table below). These functions are described in the help file with the name of the distribution (see the first column of the table below). For example, the help file for [Beta](#) describes the behavior of [dbeta](#), [pbeta](#), [qbeta](#), and [rbeta](#).

For most distributions, there is also an associated function for estimating the distribution parameters, and the form of the names of these functions is *eabb*, where *abb* is the abbreviated name of the distribution (see table below). All of these functions are listed in the help file [Estimating Distribution Parameters](#). For example, the function [ebeta](#) estimates the shape parameters of a Beta distribution based on a random sample of observations from this distribution.

For some distributions, there are functions to estimate distribution parameters based on Type I censored data. The form of the names of these functions is *eabbSinglyCensored* for singly censored data and *eabbMultiplyCensored* for multiply censored data. All of these functions are listed under the heading *Estimating Distribution Parameters* in the help file [Censored Data](#).

Table 1a. Available Distributions: Name, Abbreviation, Type, and Range

Name	Abbreviation	Type	Range
Beta	beta	Continuous	$[0, 1]$
Binomial	binom	Finite Discrete	$[0, size]$ (integer)
Cauchy	cauchy	Continuous	$(-\infty, \infty)$
Chi	chi	Continuous	$[0, \infty)$
Chi-square	chisq	Continuous	$[0, \infty)$
Exponential	exp	Continuous	$[0, \infty)$

Extreme Value	evd	Continuous	$(-\infty, \infty)$
F	f	Continuous	$[0, \infty)$
Gamma	gamma	Continuous	$[0, \infty)$
Gamma (Alternative)	gammaAlt	Continuous	$[0, \infty)$
Generalized Extreme Value	gevd	Continuous	$(-\infty, \infty)$ for $shape = 0$ $(-\infty, location + \frac{scale}{shape}]$ for $shape > 0$ $[location + \frac{scale}{shape}, \infty)$ for $shape < 0$
Geometric	geom	Discrete	$[0, \infty)$ (integer)
Hypergeometric	hyper	Finite Discrete	$[0, \min(k, m)]$ (integer)
Logistic	logis	Continuous	$(-\infty, \infty)$
Lognormal	lnorm	Continuous	$[0, \infty)$
Lognormal (Alternative)	lnormAlt	Continuous	$[0, \infty)$
Lognormal Mixture	lnormMix	Continuous	$[0, \infty)$
Lognormal Mixture (Alternative)	lnormMixAlt	Continuous	$[0, \infty)$
Three-Parameter Lognormal	lnorm3	Continuous	$[threshold, \infty)$
Truncated Lognormal	lnormTrunc	Continuous	$[min, max]$
Truncated Lognormal (Alternative)	lnormTruncAlt	Continuous	$[min, max]$
Negative Binomial	nbinom	Discrete	$[0, \infty)$ (integer)

Normal	norm	Continuous	$(-\infty, \infty)$
Normal Mixture	normMix	Continuous	$(-\infty, \infty)$
Truncated Normal	normTrunc	Continuous	$[min, max]$
Pareto	pareto	Continuous	$[location, \infty)$
Poisson	pois	Discrete	$[0, \infty)$ (integer)
Student's t	t	Continuous	$(-\infty, \infty)$
Triangular	tri	Continuous	$[min, max]$
Uniform	unif	Continuous	$[min, max]$
Weibull	weibull	Continuous	$[0, \infty)$
Wilcoxon Rank Sum	wilcox	Finite Discrete	$[0, mn]$ (integer)
Zero-Modified Lognormal (Delta)	zmlnorm	Mixed	$[0, \infty)$
Zero-Modified Lognormal (Delta) (Alternative)	zmlnormAlt	Mixed	$[0, \infty)$
Zero-Modified Normal	zmnorm	Mixed	$(-\infty, \infty)$

Table 1b. Available Distributions: Name, Parameters, Parameter Default Values, Parameter Ranges, Estimation Method(s)

Name	Parameter(s)	Default Value(s)	Parameter Range(s)	Estimation Method(s)
Beta	shape1 shape2 ncp	0	$(0, \infty)$ $(0, \infty)$ $(0, \infty)$	mle, mme, mmue
Binomial	size prob		$[0, \infty)$ $[0, 1]$	mle/mme/mvue
Cauchy	location scale	0 1	$(-\infty, \infty)$ $(0, \infty)$	

Chi	df		$(0, \infty)$	
Chi-square	df		$(0, \infty)$	
	ncp	0	$(-\infty, \infty)$	
Exponential	rate	1	$(0, \infty)$	mle/mme
Extreme Value	location	0	$(-\infty, \infty)$	mle, mme, mmue, pwme
	scale	1	$(0, \infty)$	
F	df1		$(0, \infty)$	
	df2		$(0, \infty)$	
	ncp	0	$(0, \infty)$	
Gamma	shape		$(0, \infty)$	mle, bcmle, mme, mmue
	scale	1	$(0, \infty)$	
Gamma (Alternative)	mean		$(0, \infty)$	mle, bcmle, mme, mmue
	cv	1	$(0, \infty)$	
Generalized Extreme Value	location	0	$(-\infty, \infty)$	mle, pwme, tsoe
	scale	1	$(0, \infty)$	
	shape	0	$(-\infty, \infty)$	
Geometric	prob		$(0, 1)$	mle/mme, mvue
Hypergeometric	m		$[0, \infty)$	mle, mvue
	n		$[0, \infty)$	
	k		$[1, m + n]$	
Logistic	location	0	$(-\infty, \infty)$	mle, mme, mmue
	scale	1	$(0, \infty)$	
Lognormal	meanlog	0	$(-\infty, \infty)$	mle/mme, mvue
	sdlog	1	$(0, \infty)$	
Lognormal (Alternative)	mean	$\exp(1/2)$	$(0, \infty)$	mle, mme, mmue, mvue, qmle
	cv	$\sqrt{\exp(1)-1}$	$(0, \infty)$	
Lognormal Mixture	meanlog1	0	$(-\infty, \infty)$	
	sdlog1	1	$(0, \infty)$	
	meanlog2	0	$(-\infty, \infty)$	
	sdlog2	1	$(0, \infty)$	
	p.mix	0.5	$[0, 1]$	
Lognormal Mixture (Alternative)	mean1	$\exp(1/2)$	$(0, \infty)$	
	cv1	$\sqrt{\exp(1)-1}$	$(0, \infty)$	
	mean2	$\exp(1/2)$	$(0, \infty)$	
	cv2	$\sqrt{\exp(1)-1}$	$(0, \infty)$	
	p.mix	0.5	$[0, 1]$	

Three- Parameter Lognormal	meanlog	0	$(-\infty, \infty)$	lmle, mme, mmue, mmme, royston.skew, zero.skew
	sdlog	1	$(0, \infty)$	
	threshold	0	$(-\infty, \infty)$	
Truncated Lognormal	meanlog	0	$(-\infty, \infty)$	
	sdlog	1	$(0, \infty)$	
	min	0	$[0, max)$	
	max	Inf	(min, ∞)	
Truncated Lognormal (Alternative)	mean	$\exp(1/2)$	$(0, \infty)$	
	cv	$\sqrt{\exp(1)-1}$	$(0, \infty)$	
	min	0	$[0, max)$	
	max	Inf	(min, ∞)	
Negative Binomial	size		$[1, \infty)$	mle/mme, mvue
	prob		$(0, 1]$	
	mu		$(0, \infty)$	
Normal	mean	0	$(-\infty, \infty)$	mle/mme, mvue
	sd	1	$(0, \infty)$	
Normal Mixture	mean1	0	$(-\infty, \infty)$	
	sd1	1	$(0, \infty)$	
	mean2	0	$(-\infty, \infty)$	
	sd2	1	$(0, \infty)$	
	p.mix	0.5	$[0, 1]$	
Truncated Normal	mean	0	$(-\infty, \infty)$	
	sd	1	$(0, \infty)$	
	min	-Inf	$(-\infty, max)$	
	max	Inf	(min, ∞)	
Pareto	location		$(0, \infty)$	lse, mle
	shape	1	$(0, \infty)$	
Poisson	lambda		$(0, \infty)$	mle/mme/mvue
Student's t	df		$(0, \infty)$	
	ncp	0	$(-\infty, \infty)$	
Triangular	min	0	$(-\infty, max)$	
	max	1	(min, ∞)	
	mode	0.5	(min, max)	
Uniform	min	0	$(-\infty, max)$	mle, mme, mmue
	max	1	(min, ∞)	
Weibull	shape		$(0, \infty)$	mle, mme, mmue
	scale	1	$(0, \infty)$	
Wilcoxon	m		$[1, \infty)$	

Rank Sum	n		$[1, \infty)$	
Zero-Modified Lognormal (Delta)	meanlog	0	$(-\infty, \infty)$	mvue
	sdlog	1	$(0, \infty)$	
	p.zero	0.5	$[0, 1]$	
Zero-Modified Lognormal (Delta) (Alternative)	mean	$\exp(1/2)$	$(0, \infty)$	mvue
	cv	$\sqrt{\exp(1)-1}$	$(0, \infty)$	
	p.zero	0.5	$[0, 1]$	
Zero-Modified Normal	mean	0	$(-\infty, \infty)$	mvue
	sd	1	$(0, \infty)$	
	p.zero	0.5	$[0, 1]$	

Source

The **EnvStats** package.

References

Millard, S.P. (In Preparation). *EnvStats: An R Package for Environmental Statistics*. Springer-Verlag, New York.

 ebeta

Estimate Parameters of a Beta Distribution

Description

Estimate the shape parameters of a [beta distribution](#).

Usage

```
ebeta(x, method = "mle")
```

Arguments

x numeric vector of observations. All observations must be between greater than 0 and less than 1.

method character string specifying the method of estimation. The possible values are "mle" (maximum likelihood; the default), "mme" (method of moments), and "mmue" (method of moments based on the unbiased estimator of variance). See the DETAILS section for more information on these estimation methods.

Details

If x contains any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

Let $\underline{x} = (x_1, x_2, \dots, x_n)$ be a vector of n observations from a [beta distribution](#) with parameters $\text{shape1}=\nu$ and $\text{shape2}=\omega$.

Maximum Likelihood Estimation (method="mle")

The maximum likelihood estimators (mle's) of the shape parameters ν and ω are the solutions of the simultaneous equations:

$$\Psi(\hat{\nu}) - \Psi(\hat{\nu} + \hat{\omega}) = (1/n) \sum_{i=1}^n \log(x_i)$$

$$\Psi(\hat{\nu}) - \Psi(\hat{\nu} + \hat{\omega}) = (1/n) \sum_{i=1}^n \log(1 - x_i)$$

where $\Psi()$ is the [digamma function](#) (Forbes et al., 2011).

Method of Moments Estimators (method="mme")

The method of moments estimators (mme's) of the shape parameters ν and ω are given by (Forbes et al., 2011):

$$\hat{\nu} = \bar{x} \{ [\bar{x}(1 - \bar{x})/s_m^2] - 1 \}$$

$$\hat{\omega} = (1 - \bar{x}) \{ [\bar{x}(1 - \bar{x})/s_m^2] - 1 \}$$

where

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i; s_m^2 = \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2$$

Method of Moments Estimators Based on the Unbiased Estimator of Variance (method="mmue")

These estimators are the same as the method of moments estimators except that the method of moments estimator of variance is replaced with the unbiased estimator of variance:

$$s^2 = \frac{1}{n-1} \sum_{i=1}^n (x_i - \bar{x})^2$$

Value

a list of class "estimate" containing the estimated parameters and other information. See [estimate.object](#) for details.

Note

The beta distribution takes real values between 0 and 1. Special cases of the beta are the [Uniform](#)[0,1] when $\text{shape1}=1$ and $\text{shape2}=1$, and the arcsin distribution when $\text{shape1}=0.5$ and $\text{shape2}=0.5$. The arcsin distribution appears in the theory of random walks. The beta distribution is used in Bayesian analyses as a conjugate to the binomial distribution.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.

Johnson, N. L., S. Kotz, and N. Balakrishnan. (1995). *Continuous Univariate Distributions, Volume 2*. Second Edition. John Wiley and Sons, New York.

See Also

[Beta](#).

Examples

```
# Generate 20 observations from a beta distribution with parameters
# shape1=2 and shape2=4, then estimate the parameters via
# maximum likelihood.
# (Note: the call to set.seed simply allows you to reproduce this example.)
```

```
set.seed(250)
dat <- rbeta(20, shape1 = 2, shape2 = 4)
ebeta(dat)
```

```
#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Beta
#
#Estimated Parameter(s):      shape1 =  5.392221
#                               shape2 = 11.823233
#
#Estimation Method:           mle
#
#Data:                         dat
#
#Sample Size:                  20
```

```
#=====
```

```
# Repeat the above, but use the method of moments estimators:
```

```
ebeta(dat, method = "mme")
```

```
#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Beta
#
#Estimated Parameter(s):      shape1 =  5.216311
#                               shape2 = 11.461341
#
#Estimation Method:           mme
#
#Data:                         dat
#
#Sample Size:                  20
```

```
#=====
```

```
# Clean up
#-----
rm(dat)
```

ebinom

Estimate Parameter of a Binomial Distribution

Description

Estimate p (the probability of “success”) for a binomial distribution, and optionally construct a confidence interval for p .

Usage

```
ebinom(x, size = NULL, method = "mle/mme/mvue", ci = FALSE,
       ci.type = "two-sided", ci.method = "score", correct = TRUE,
       var.denom = "n", conf.level = 0.95, warn = TRUE)
```

Arguments

<code>x</code>	numeric or logical vector of observations. When <code>size</code> is not supplied, <code>x</code> must be a numeric vector of 0s (“failures”) and 1s (“successes”), or else a logical vector of FALSE values (“failures”) and TRUE values (“successes”). When <code>size</code> is supplied, <code>x</code> must be a non-negative integer containing the number of “successes” out of the number of trials indicated by <code>size</code> . Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
<code>size</code>	positive integer indicating the of number of trials; <code>size</code> must be at least as large as the value of <code>x</code> .
<code>method</code>	character string specifying the method of estimation. The only possible value is “mle/mme/mvue” (maximum likelihood, method of moments, and minimum variance unbiased). See the DETAILS section for more information.
<code>ci</code>	logical scalar indicating whether to compute a confidence interval for the mean. The default value is <code>ci=FALSE</code> .
<code>ci.type</code>	character string indicating what kind of confidence interval to compute. The possible values are “two-sided” (the default), “lower”, and “upper”. This argument is ignored if <code>ci=FALSE</code> .
<code>ci.method</code>	character string indicating which method to use to construct the confidence interval. Possible values are “score” (the default), “exact”, “adjusted Wald”, and “Wald”. This argument is ignored if <code>ci=FALSE</code> .
<code>correct</code>	logical scalar indicating whether to use the continuity correction when <code>ci.method=“score”</code> or <code>ci.method=“Wald”</code> . The default value is <code>correct=TRUE</code> .
<code>var.denom</code>	character string indicating what value to use in the denominator of the variance estimator when <code>ci.method=“Wald”</code> . Possible values are “n” (the default) and “n-1”. This argument is ignored if <code>ci=FALSE</code> .
<code>conf.level</code>	a scalar between 0 and 1 indicating the confidence level of the confidence interval. The default value is <code>conf.level=0.95</code> . This argument is ignored if <code>ci=FALSE</code> .

warn a logical scalar indicating whether to issue a warning in the case when `ci=TRUE`, `ci.method="Wald"`, and any of the following conditions is true: the estimated proportion is less than 0.2, the estimated proportion is greater than 0.8, the number of successes or failures is less than 5. The default value is `warn=TRUE`.

Details

If `x` contains any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

If `x` is a vector of n observations from a binomial distribution with parameters `size=1` and `prob=p`, then the sum of all the values in `x` is an observation from a binomial distribution with parameters `size=n` and `prob=p`.

If `x` is an observation from a binomial distribution with parameters `size=n` and `prob=p`, the maximum likelihood estimator (mle), method of moments estimator (mme), and minimum variance unbiased estimator (mvue) of p is simply x/n .

Confidence Intervals.

`ci.method="score"` The confidence interval for p based on the score method was developed by Wilson (1927) and is discussed by Newcombe (1998a), Agresti and Coull (1998), and Agresti and Caffo (2000). When `ci=TRUE` and `ci.method="score"`, the function `ebinom` calls the R function `prop.test` to compute the confidence interval. This method has been shown to provide the best performance (in terms of actual coverage matching assumed coverage) of all the methods provided here, although unlike the exact method, the actual coverage can fall below the assumed coverage.

`ci.method="exact"` The confidence interval for p based on the exact (Clopper-Pearson) method is discussed by Newcombe (1998a), Agresti and Coull (1998), and Zar (2010, pp.543-547). This is the method used in the R function `binom.test`. This method ensures the actual coverage is greater than or equal to the assumed coverage.

`ci.method="Wald"` The confidence interval for p based on the Wald method (with or without a correction for continuity) is the usual "normal approximation" method and is discussed by Newcombe (1998a), Agresti and Coull (1998), Agresti and Caffo (2000), and Zar (2010, pp.543-547). This method is **never** recommended but is included for historical purposes.

`ci.method="adjusted Wald"` The confidence interval for p based on the adjusted Wald method is discussed by Agresti and Coull (1998), Agresti and Caffo (2000), and Zar (2010, pp.543-547). This is a simple modification of the Wald method and performs surprisingly well.

Value

a list of class "estimate" containing the estimated parameters and other information. See `estimate.object` for details.

Note

The binomial distribution is used to model processes with binary (Yes-No, Success-Failure, Heads-Tails, etc.) outcomes. It is assumed that the outcome of any one trial is independent of any other trial, and that the probability of "success", p , is the same on each trial. A binomial discrete random variable X is the number of "successes" in n independent trials. A special case of the binomial distribution occurs when $n = 1$, in which case X is also called a Bernoulli random variable.

In the context of environmental statistics, the binomial distribution is sometimes used to model the proportion of times a chemical concentration exceeds a set standard in a given period of time (e.g., Gilbert, 1987, p.143). The binomial distribution is also used to compute an upper bound on the

overall Type I error rate for deciding whether a facility or location is in compliance with some set standard. Assume the null hypothesis is that the facility is in compliance. If a test of hypothesis is conducted periodically over time to test compliance and/or several tests are performed during each time period, and the facility or location is always in compliance, and each single test has a Type I error rate of α , and the result of each test is independent of the result of any other test (usually not a reasonable assumption), then the number of times the facility is declared out of compliance when in fact it is in compliance is a binomial random variable with probability of “success” $p = \alpha$ being the probability of being declared out of compliance (see USEPA, 2009).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Agresti, A., and B.A. Coull. (1998). Approximate is Better than "Exact" for Interval Estimation of Binomial Proportions. *The American Statistician*, **52**(2), 119–126.
- Agresti, A., and B. Caffo. (2000). Simple and Effective Confidence Intervals for Proportions and Differences of Proportions Result from Adding Two Successes and Two Failures. *The American Statistician*, **54**(4), 280–288.
- Berthouex, P.M., and L.C. Brown. (1994). *Statistics for Environmental Engineers*. Lewis Publishers, Boca Raton, FL, Chapters 2 and 15.
- Cochran, W.G. (1977). *Sampling Techniques*. John Wiley and Sons, New York, Chapter 3.
- Fisher, R.A., and F. Yates. (1963). *Statistical Tables for Biological, Agricultural, and Medical Research*. 6th edition. Hafner, New York, 146pp.
- Fleiss, J. L. (1981). *Statistical Methods for Rates and Proportions*. Second Edition. John Wiley and Sons, New York, Chapters 1-2.
- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.
- Gilbert, R.O. (1987). *Statistical Methods for Environmental Pollution Monitoring*. Van Nostrand Reinhold, New York, NY, Chapter 11.
- Johnson, N. L., S. Kotz, and A.W. Kemp. (1992). *Univariate Discrete Distributions*. Second Edition. John Wiley and Sons, New York, Chapter 3.
- Millard, S.P., and Neerchal, N.K. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, Florida.
- Newcombe, R.G. (1998a). Two-Sided Confidence Intervals for the Single Proportion: Comparison of Seven Methods. *Statistics in Medicine*, **17**, 857–872.
- Ott, W.R. (1995). *Environmental Statistics and Data Analysis*. Lewis Publishers, Boca Raton, FL, Chapter 4.
- USEPA. (1989b). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities, Interim Final Guidance*. EPA/530-SW-89-026. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.6-38.
- Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ, Chapter 24.

See Also

[Binomial](#), [prop.test](#), [binom.test](#), [ciBinomHalfWidth](#), [ciBinomN](#), [plotCiBinomDesign](#).

Examples

```
# Generate 20 observations from a binomial distribution with
# parameters size=1 and prob=0.2, then estimate the 'prob' parameter.
# (Note: the call to set.seed simply allows you to reproduce this
# example. Also, the only parameter estimated is 'prob'; 'size' is
# specified in the call to ebinom. The parameter 'size' is printed
# in order to show all of the parameters associated with the
# distribution.)
```

```
set.seed(251)
dat <- rbinom(20, size = 1, prob = 0.2)
ebinom(dat)
```

```
#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Binomial
#
#Estimated Parameter(s):      size = 20.0
#                              prob =  0.1
#
#Estimation Method:           mle/mme/mvue for 'prob'
#
#Data:                        dat
#
#Sample Size:                 20
```

```
#-----
```

```
# Generate one observation from a binomial distribution with
# parameters size=20 and prob=0.2, then estimate the "prob"
# parameter and compute a confidence interval:
```

```
set.seed(763)
dat <- rbinom(1, size=20, prob=0.2)
ebinom(dat, size = 20, ci = TRUE)
```

```
#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Binomial
#
#Estimated Parameter(s):      size = 20.00
#                              prob =  0.35
#
#Estimation Method:           mle/mme/mvue for 'prob'
#
#Data:                        dat
#
#Sample Size:                 20
#
#Confidence Interval for:      prob
```



```
#
#Confidence Interval Method:      Score normal approximation
#                                (With continuity correction)
#
#Confidence Interval Type:       two-sided
#
#Confidence Level:               95%
#
#Confidence Interval:           LCL = 0.1630867
#                                UCL = 0.5905104
```

```
#-----
```

```
# Using the data from the last example, compare confidence
# intervals based on the various methods
```

```
ebinom(dat, size = 20, ci = TRUE,
       ci.method = "score", correct = TRUE)$interval$limits
#      LCL      UCL
#0.1630867 0.5905104
```

```
ebinom(dat, size = 20, ci = TRUE,
       ci.method = "score", correct = FALSE)$interval$limits
#      LCL      UCL
#0.1811918 0.5671457
```

```
ebinom(dat, size = 20, ci = TRUE,
       ci.method = "exact")$interval$limits
#      LCL      UCL
#0.1539092 0.5921885
```

```
ebinom(dat, size = 20, ci = TRUE,
       ci.method = "adjusted Wald")$interval$limits
#      LCL      UCL
#0.1799264 0.5684112
```

```
ebinom(dat, size = 20, ci = TRUE,
       ci.method = "Wald", correct = TRUE)$interval$limits
#      LCL      UCL
#0.1159627 0.5840373
```

```
ebinom(dat, size = 20, ci = TRUE,
       ci.method = "Wald", correct = FALSE)$interval$limits
#      LCL      UCL
#0.1409627 0.5590373
```

```
#-----
```

```
# Use the cadmium data on page 8-6 of USEPA (1989b) to compute
# two-sided 95% confidence intervals for the probability of
# detection at background and compliance wells. The data are
# stored in EPA.89b.cadmium.df.
```

```

EPA.89b.cadmium.df
#   Cadmium.orig Cadmium Censored Well.type
#1      0.1    0.100    FALSE Background
#2      0.12   0.120    FALSE Background
#3      BDL    0.000     TRUE Background
#...
#86      BDL    0.000     TRUE Compliance
#87      BDL    0.000     TRUE Compliance
#88      BDL    0.000     TRUE Compliance

attach(EPA.89b.cadmium.df)

# Probability of detection at Background well:
#-----

ebinom(!Censored[Well.type=="Background"], ci=TRUE)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:      Binomial
#
#Estimated Parameter(s):    size = 24.0000000
#                           prob =  0.3333333
#
#Estimation Method:        mle/mme/mvue for 'prob'
#
#Data:                      !Censored[Well.type == "Background"]
#
#Sample Size:              24
#
#Confidence Interval for:   prob
#
#Confidence Interval Method: Score normal approximation
#                           (With continuity correction)
#
#Confidence Interval Type:  two-sided
#
#Confidence Level:         95%
#
#Confidence Interval:      LCL = 0.1642654
#                           UCL = 0.5530745

# Probability of detection at Compliance well:
#-----

ebinom(!Censored[Well.type=="Compliance"], ci=TRUE)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:      Binomial
#
#Estimated Parameter(s):    size = 64.000
#                           prob =  0.375

```

```

#
#Estimation Method:          mle/mme/mvue for 'prob'
#
#Data:                      !Censored[Well.type == "Compliance"]
#
#Sample Size:               64
#
#Confidence Interval for:   prob
#
#Confidence Interval Method: Score normal approximation
#                           (With continuity correction)
#
#Confidence Interval Type:  two-sided
#
#Confidence Level:         95%
#
#Confidence Interval:      LCL = 0.2597567
#                           UCL = 0.5053034

#-----

# Clean up
rm(dat)
detach("EPA.89b.cadmium.df")

```

ecdfPlot

Empirical Cumulative Distribution Function Plot

Description

Produce an empirical cumulative distribution function plot.

Usage

```

ecdfPlot(x, discrete = FALSE,
  prob.method = ifelse(discrete, "emp.probs", "plot.pos"),
  plot.pos.con = 0.375, plot.it = TRUE, add = FALSE, ecdf.col = "black",
  ecdf.lwd = 3 * par("cex"), ecdf.lty = 1, curve.fill = FALSE,
  curve.fill.col = "cyan", ..., type = ifelse(discrete, "s", "l"),
  main = NULL, xlab = NULL, ylab = NULL, xlim = NULL, ylim = NULL)

```

Arguments

x	numeric vector of observations. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
discrete	logical scalar indicating whether the assumed parent distribution of x is discrete (discrete=TRUE) or continuous (discrete=FALSE; the default).
prob.method	character string indicating what method to use to compute the plotting positions (empirical probabilities). Possible values are plot.pos (plotting positions, the default if discrete=FALSE) and emp.probs (empirical probabilities, the default if discrete=TRUE). See the DETAILS section for more explanation.

<code>plot.pos.con</code>	numeric scalar between 0 and 1 containing the value of the plotting position constant. The default value is <code>plot.pos.con=0.375</code> . See the DETAILS section for more information. This argument is ignored if <code>prob.method="emp.probs"</code> .
<code>plot.it</code>	logical scalar indicating whether to produce a plot or add to the current plot (see <code>add</code>) on the current graphics device. The default value is <code>plot.it=TRUE</code> .
<code>add</code>	logical scalar indicating whether to add the empirical cdf to the current plot (<code>add=TRUE</code>) or generate a new plot (<code>add=FALSE</code> ; the default). This argument is ignored if <code>plot.it=FALSE</code> .
<code>ecdf.col</code>	a numeric scalar or character string determining the color of the empirical cdf line or points. The default value is <code>ecdf.col=1</code> . See the entry for <code>col</code> in the help file for <code>par</code> for more information.
<code>ecdf.lwd</code>	a numeric scalar determining the width of the empirical cdf line. The default value is <code>ecdf.lwd=3*par("cex")</code> . See the entry for <code>lwd</code> in the help file for <code>par</code> for more information.
<code>ecdf.lty</code>	a numeric scalar determining the line type of the empirical cdf line. The default value is <code>ecdf.lty=1</code> . See the entry for <code>lty</code> in the help file for <code>par</code> for more information.
<code>curve.fill</code>	a logical scalar indicating whether to fill in the area below the empirical cdf curve with the color specified by <code>curve.fill.col</code> . The default value is <code>curve.fill=FALSE</code> .
<code>curve.fill.col</code>	a numeric scalar or character string indicating what color to use to fill in the area below the empirical cdf curve. The default value is <code>curve.fill.col=5</code> . This argument is ignored if <code>curve.fill=FALSE</code> .
<code>type, main, xlab, ylab, xlim, ylim, ...</code>	additional graphical parameters (see <code>lines</code> and <code>par</code>). In particular, the argument <code>type</code> specifies the kind of line type. By default, the function <code>ecdfPlot</code> plots a step function (<code>type="s"</code>) when <code>discrete=TRUE</code> , and plots a straight line between points (<code>type="l"</code>) when <code>discrete=FALSE</code> . The user may override these defaults by supplying the graphics parameter <code>type</code> (<code>type="s"</code> for a step function, <code>type="l"</code> for linear interpolation, <code>type="p"</code> for points only, etc.).

Details

The *cumulative distribution function (cdf)* of a random variable X is the function F such that

$$F(x) = Pr(X \leq x) \quad (1)$$

for all values of x . That is, if $p = F(x)$, then p is the proportion of the population that is less than or equal to x , and x is called the p 'th *quantile*, or the $100p$ 'th percentile. A plot of quantiles on the x -axis (i.e., the possible value for the random variable X) vs. the fraction of the population less than or equal to that number on the y -axis is called the *cumulative distribution function plot*, and the y -axis is usually labeled as the "cumulative probability" or "cumulative frequency".

When we have a sample of data from some population, we usually do not know what percentiles our observations correspond to because we do not know the form of the cumulative distribution function F , so we have to use the sample data to estimate the cdf F . An *empirical cumulative distribution function (ecdf) plot*, also called a *quantile plot*, is a plot of the observed quantiles (i.e., the ordered observations) on the x -axis vs. the estimated cumulative probabilities on the y -axis (Chambers et al., 1983, pp. 11-19; Cleveland, 1993, pp. 17-20; Cleveland, 1994, pp. 136-139; Helsel and Hirsch, 1992, pp. 21-24).

(Note: Some authors (e.g., Chambers et al., 1983, pp.11-16; Cleveland, 1993, pp.17-20) reverse the axes on a quantile plot, i.e., the observed order statistics from the random sample are on the y -axis and the estimated cumulative probabilities are on the x -axis.)

The **empirical cumulative distribution function (ecdf)** is an estimate of the cdf based on a random sample of n observations from the distribution. Let x_1, x_2, \dots, x_n denote the n observations, and let $x_{(1)}, x_{(2)}, \dots, x_{(n)}$ denote the ordered observations (i.e., the order statistics). The cdf is usually estimated by either the **empirical probabilities estimator** or the **plotting-position estimator**. The empirical probabilities estimator is given by:

$$\hat{F}[x_{(i)}] = \hat{p}_i = \frac{\#[x_j \leq x_{(i)}]}{n} \quad (2)$$

where $\#[x_j \leq x_{(i)}]$ denotes the number of observations less than or equal to $x_{(i)}$. The plotting-position estimator is given by:

$$\hat{F}[x_{(i)}] = \hat{p}_i = \frac{i - a}{n - 2a + 1} \quad (3)$$

where $0 \leq a \leq 1$ (Cleveland, 1993, p. 18; D'Agostino, 1986a, pp. 8,25).

For any value x such that $x_{(1)} < x < x_{(n)}$, the ecdf is usually defined as either a step function:

$$\hat{F}(x) = \hat{F}[x_{(i)}], \quad x_{(i)} \leq x < x_{(i+1)} \quad (4)$$

(e.g., D'Agostino, 1986a), or linear interpolation between order statistics is used:

$$\hat{F}(x) = (1 - r)\hat{F}[x_{(i)}] + r\hat{F}[x_{(i+1)}], \quad x_{(i)} \leq x < x_{(i+1)} \quad (5)$$

where

$$r = \frac{x - x_{(i)}}{x_{(i+1)} - x_{(i)}} \quad (6)$$

(e.g., Chambers et al., 1983). For the step function version, the ecdf stays flat until it hits a value on the x -axis corresponding to one of the order statistics, then it makes a jump. For the linear interpolation version, the ecdf plot looks like lines connecting the points. By default, the function `ecdfPlot` uses the step function version when `discrete=TRUE`, and the linear interpolation version when `discrete=FALSE`. The user may override these defaults by supplying the graphics parameter `type="s"` for a step function, `type="l"` for linear interpolation, `type="p"` for points only, etc.).

The empirical probabilities estimator is intuitively appealing. This is the estimator used when `prob.method="emp.probs"`. The disadvantage of this estimator is that it implies the largest observed value is the maximum possible value of the distribution (i.e., the 100'th percentile). This may be satisfactory if the underlying distribution is known to be discrete, but it is usually not satisfactory if the underlying distribution is known to be continuous.

The plotting-position estimator with various values of a is often used when the goal is to produce a probability plot (see `qqPlot`) rather than an empirical cdf plot. It is used to compute the estimated expected values or medians of the order statistics for a probability plot. This is the estimator used when `prob.method="plot.pos"`. The argument `plot.pos.con` refers to the variable a . Based on certain principles from statistical theory, certain values of the constant a make sense for specific underlying distributions (see the help file for `qqPlot` for more information).

Because x is a random sample, the empirical cdf changes from sample to sample and the variability in these estimates can be dramatic for small sample sizes.

Value

ecdfPlot invisibly returns a list with the following components:

`Order.Statistics`

numeric vector of the ordered observations.

`Cumulative.Probabilities`

numeric vector of the associated plotting positions.

Note

An empirical cumulative distribution function (ecdf) plot is a graphical tool that can be used in conjunction with other graphical tools such as histograms, strip charts, and boxplots to assess the characteristics of a set of data. It is easy to determine quartiles and the minimum and maximum values from such a plot. Also, ecdf plots allow you to assess local density: a higher density of observations occurs where the slope is steep.

Chambers et al. (1983, pp.11-16) plot the observed order statistics on the *y*-axis vs. the ecdf on the *x*-axis and call this a quantile plot.

Empirical cumulative distribution function (ecdf) plots are often plotted with theoretical cdf plots (see [cdfPlot](#) and [cdfCompare](#)) to graphically assess whether a sample of observations comes from a particular distribution. The Kolmogorov-Smirnov goodness-of-fit test (see [gofTest](#)) is the statistical companion of this kind of comparison; it is based on the maximum vertical distance between the empirical cdf plot and the theoretical cdf plot. More often, however, quantile-quantile (Q-Q) plots are used instead of ecdf plots to graphically assess departures from an assumed distribution (see [qqPlot](#)).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Chambers, J.M., W.S. Cleveland, B. Kleiner, and P.A. Tukey. (1983). *Graphical Methods for Data Analysis*. Duxbury Press, Boston, MA, pp.11-16.

Cleveland, W.S. (1993). *Visualizing Data*. Hobart Press, Summit, New Jersey, 360pp.

D'Agostino, R.B. (1986a). Graphical Analysis. In: D'Agostino, R.B., and M.A. Stephens, eds. *Goodness-of Fit Techniques*. Marcel Dekker, New York, Chapter 2, pp.7-62.

See Also

[ppoints](#), [cdfPlot](#), [cdfCompare](#), [qqPlot](#), [ecdfPlotCensored](#).

Examples

```
# Generate 20 observations from a normal distribution with
# mean=0 and sd=1 and create an ecdf plot.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
x <- rnorm(20)
windows()
ecdfPlot(x)

#-----
```

```

# Repeat the above example, but fill in the area under the
# empirical cdf curve.

windows()
ecdfPlot(x, curve.fill = TRUE)

#-----

# Repeat the above example, but plot only the points.

windows()
ecdfPlot(x, type = "p")

#-----

# Repeat the above example, but force a step function.

windows()
ecdfPlot(x, type = "s")

#-----

# Clean up
rm(x)

#-----

# The guidance document USEPA (1994b, pp. 6.22--6.25)
# contains measures of 1,2,3,4-Tetrachlorobenzene (TcCB)
# concentrations (in parts per billion) from soil samples
# at a Reference area and a Cleanup area. These data are stored
# in the data frame EPA.94b.tccb.df.
#
# Create an empirical CDF plot for the reference area data.

windows()
with(EPA.94b.tccb.df,
     ecdfPlot(TcCB[Area == "Reference"], xlab = "TcCB (ppb)"))

#=====

# Clean up
#-----
graphics.off()

```

ecdfPlotCensored

Empirical Cumulative Distribution Function Plot Based on Type I Censored Data

Description

Produce an empirical cumulative distribution function plot for Type I left-censored or right-censored data.

Usage

```
ecdfPlotCensored(x, censored, censoring.side = "left", discrete = FALSE,
  prob.method = "michael-schucany", plot.pos.con = 0.375, plot.it = TRUE,
  add = FALSE, ecdf.col = 1, ecdf.lwd = 3 * par("cex"), ecdf.lty = 1,
  include.cen = FALSE, cen.pch = ifelse(censoring.side == "left", 6, 2),
  cen.cex = par("cex"), cen.col = 4, ...,
  type = ifelse(discrete, "s", "l"), main = NULL, xlab = NULL, ylab = NULL,
  xlim = NULL, ylim = NULL)
```

Arguments

<code>x</code>	numeric vector of observations. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
<code>censored</code>	numeric or logical vector indicating which values of <code>x</code> are censored. This must be the same length as <code>x</code> . If the mode of <code>censored</code> is "logical", TRUE values correspond to elements of <code>x</code> that are censored, and FALSE values correspond to elements of <code>x</code> that are not censored. If the mode of <code>censored</code> is "numeric", it must contain only 1's and 0's; 1 corresponds to TRUE and 0 corresponds to FALSE. Missing (NA) values are allowed but will be removed.
<code>censoring.side</code>	character string indicating on which side the censoring occurs. The possible values are "left" (the default) and "right".
<code>discrete</code>	logical scalar indicating whether the assumed parent distribution of <code>x</code> is discrete (<code>discrete=TRUE</code>) or continuous (<code>discrete=FALSE</code> ; the default).
<code>prob.method</code>	character string indicating what method to use to compute the plotting positions (empirical probabilities). Possible values are "kaplan-meier" (product-limit method of Kaplan and Meier (1958)), "nelson" (hazard plotting method of Nelson (1972)), "michael-schucany" (generalization of the product-limit method due to Michael and Schucany (1986)), and "hirsch-stedinger" (generalization of the product-limit method due to Hirsch and Stedinger (1987)). The default value is <code>prob.method="michael-schucany"</code> . The "nelson" method is only available for <code>censoring.side="right"</code> . See the DETAILS section for more explanation.
<code>plot.pos.con</code>	numeric scalar between 0 and 1 containing the value of the plotting position constant. The default value is <code>plot.pos.con=0.375</code> . See the DETAILS section for more information. This argument is used only if <code>prob.method</code> is equal to "michael-schucany" or "hirsch-stedinger".
<code>plot.it</code>	logical scalar indicating whether to produce a plot or add to the current plot (see <code>add</code>) on the current graphics device. The default value is <code>plot.it=TRUE</code> .
<code>add</code>	logical scalar indicating whether to add the empirical cdf to the current plot (<code>add=TRUE</code>) or generate a new plot (<code>add=FALSE</code> ; the default). This argument is ignored if <code>plot.it=FALSE</code> .
<code>ecdf.col</code>	a numeric scalar or character string determining the color of the empirical cdf line or points. The default value is <code>ecdf.col=1</code> . See the entry for <code>col</code> in the help file for par for more information.
<code>ecdf.lwd</code>	a numeric scalar determining the width of the empirical cdf line. The default value is <code>ecdf.lwd=3*par("cex")</code> . See the entry for <code>lwd</code> in the help file for par for more information.
<code>ecdf.lty</code>	a numeric scalar determining the line type of the empirical cdf line. The default value is <code>ecdf.lty=1</code> . See the entry for <code>lty</code> in the help file for par for more information.

<code>include.cen</code>	logical scalar indicating whether to include censored values in the plot. The default value is <code>include.cen=FALSE</code> . If <code>include.cen=TRUE</code> , censored values are plotted using the plotting character indicated by the argument <code>cen.pch</code> (see below).
<code>cen.pch</code>	numeric scalar or character string indicating the plotting character to use to plot censored values. The default value is <code>cen.pch=2</code> (hollow triangle pointing up) when <code>censoring.side="right"</code> , and <code>cen.pch=6</code> (hollow triangle pointing down) when <code>censoring.side="left"</code> . See the help file for points for a list of other possible plotting characters. This argument is ignored if <code>include.cen=FALSE</code> .
<code>cen.cex</code>	numeric scalar that determines the size of the plotting character used to plot censored values. The default value is the current value of the <code>cex</code> graphics parameter. See the entry for <code>cex</code> in the help file for par for more information. This argument is ignored if <code>include.cen=FALSE</code> .
<code>cen.col</code>	numeric scalar or character string that determines the color of the plotting character used to plot censored values. The default value is <code>cen.col=4</code> . See the entry for <code>col</code> in the help file for par for more information. This argument is ignored if <code>include.cen=FALSE</code> .
<code>type, main, xlab, ylab, xlim, ylim, ...</code>	additional graphical parameters (see lines and par). In particular, the argument <code>type</code> specifies the kind of line type. By default, the function <code>ecdfPlotCensored</code> plots a step function (<code>type="s"</code>) when <code>discrete=TRUE</code> , and plots a straight line between points (<code>type="l"</code>) when <code>discrete=FALSE</code> . The user may override these defaults by supplying the graphics parameter <code>type</code> (<code>type="s"</code> for a step function, <code>type="l"</code> for linear interpolation, <code>type="p"</code> for points only, etc.).

Details

The function `ecdfPlotCensored` does exactly the same thing as [ecdfPlot](#), except it calls the function [ppointsCensored](#) to compute the plotting positions (estimated cumulative probabilities) for the uncensored observations.

If `plot.it=TRUE`, the estimated cumulative probabilities for the uncensored observations are plotted against the uncensored observations. By default, the function `ecdfPlotCensored` plots a step function when `discrete=TRUE`, and plots a straight line between points when `discrete=FALSE`. The user may override these defaults by supplying the graphics parameter `type` (`type="s"` for a step function, `type="l"` for linear interpolation, `type="p"` for points only, etc.).

If `include.cen=TRUE`, censored observations are included on the plot as points. The arguments `cen.pch`, `cen.cex`, and `cen.col` control the appearance of these points.

In cases where `x` is a random sample, the empirical cdf will change from sample to sample and the variability in these estimates can be dramatic for small sample sizes. Caution must be used in interpreting the empirical cdf when a large percentage of the observations are censored.

Value

`ecdfPlotCensored` returns a list with the following components:

<code>Order.Statistics</code>	numeric vector of the “ordered” observations.
<code>Cumulative.Probabilities</code>	numeric vector of the associated plotting positions.
<code>Censored</code>	logical vector indicating which of the ordered observations are censored.

Censoring.Side	character string indicating whether the data are left- or right-censored. This is same value as the argument <code>censoring.side</code> .
Prob.Method	character string indicating what method was used to compute the plotting positions. This is the same value as the argument <code>prob.method</code> .
Optional Component (only present when <code>prob.method="michael-schucany"</code> or <code>prob.method="hirsch-stedinger"</code>):	
Plot.Pos.Con	numeric scalar containing the value of the plotting position constant that was used. This is the same as the argument <code>plot.pos.con</code> .

Note

An empirical cumulative distribution function (ecdf) plot is a graphical tool that can be used in conjunction with other graphical tools such as histograms, strip charts, and boxplots to assess the characteristics of a set of data.

Censored observations complicate the procedures used to graphically explore data. Techniques from survival analysis and life testing have been developed to generalize the procedures for constructing plotting positions, empirical cdf plots, and q-q plots to data sets with censored observations (see [ppointsCensored](#)).

Empirical cumulative distribution function (ecdf) plots are often plotted with theoretical cdf plots (see [cdfPlot](#) and [cdfCompareCensored](#)) to graphically assess whether a sample of observations comes from a particular distribution. More often, however, quantile-quantile (Q-Q) plots are used instead (see [qqPlot](#) and [qqPlotCensored](#)).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Chambers, J.M., W.S. Cleveland, B. Kleiner, and P.A. Tukey. (1983). *Graphical Methods for Data Analysis*. Duxbury Press, Boston, MA, pp.11-16.
- Cleveland, W.S. (1993). *Visualizing Data*. Hobart Press, Summit, New Jersey, 360pp.
- D'Agostino, R.B. (1986a). Graphical Analysis. In: D'Agostino, R.B., and M.A. Stephens, eds. *Goodness-of Fit Techniques*. Marcel Dekker, New York, Chapter 2, pp.7-62.
- Gillespie, B.W., Q. Chen, H. Reichert, A. Franzblau, E. Hedgeman, J. Lepkowski, P. Adriaens, A. Demond, W. Luksemburg, and D.H. Garabrant. (2010). Estimating Population Distributions When Some Data Are Below a Limit of Detection by Using a Reverse Kaplan-Meier Estimator. *Epidemiology* **21**(4), S64–S70.
- Helsel, D.R. (2012). *Statistics for Censored Environmental Data Using Minitab and R, Second Edition*. John Wiley & Sons, Hoboken, New Jersey.
- Helsel, D.R., and T.A. Cohn. (1988). Estimation of Descriptive Statistics for Multiply Censored Water Quality Data. *Water Resources Research* **24**(12), 1997-2004.
- Hirsch, R.M., and J.R. Stedinger. (1987). Plotting Positions for Historical Floods and Their Precision. *Water Resources Research* **23**(4), 715-727.
- Kaplan, E.L., and P. Meier. (1958). Nonparametric Estimation From Incomplete Observations. *Journal of the American Statistical Association* **53**, 457-481.
- Lee, E.T., and J.W. Wang. (2003). *Statistical Methods for Survival Data Analysis, Third Edition*. John Wiley & Sons, Hoboken, New Jersey, 513pp.

Michael, J.R., and W.R. Schucany. (1986). Analysis of Data from Censored Samples. In D'Agostino, R.B., and M.A. Stephens, eds. *Goodness-of Fit Techniques*. Marcel Dekker, New York, 560pp, Chapter 11, 461-496.

Nelson, W. (1972). Theory and Applications of Hazard Plotting for Censored Failure Data. *Technometrics* **14**, 945-966.

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. Chapter 15.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

See Also

[ppoints](#), [ppointsCensored](#), [ecdfPlot](#), [qqPlot](#), [qqPlotCensored](#), [cdfPlot](#), [cdfCompareCensored](#).

Examples

```
# Generate 20 observations from a normal distribution with mean=20 and sd=5,
# censor all observations less than 18, then generate an empirical cdf plot
# for the complete data set and the censored data set. Note that the empirical
# cdf plot for the censored data set starts at the first ordered uncensored
# observation, and that for values of x > 18 the two empirical cdf plots are
# exactly the same. This is because there is only one censoring level and
# no uncensored observations fall below the censored observations.
# (Note: the call to set.seed simply allows you to reproduce this example.)
```

```
set.seed(333)
x <- rnorm(20, mean=20, sd=5)
censored <- x < 18

sum(censored)
#[1] 7

new.x <- x
new.x[censored] <- 18

windows()
ecdfPlot(x, xlim = range(pretty(x)),
  main = "Empirical CDF Plot for\nComplete Data Set")
```

```
windows()
ecdfPlotCensored(new.x, censored, xlim = range(pretty(x)),
  main="Empirical CDF Plot for\nCensored Data Set")
```

```
# Clean up
#-----
rm(x, censored, new.x)
```

```
#-----
```

```
# Example 15-1 of USEPA (2009, page 15-10) gives an example of
# computing plotting positions based on censored manganese
# concentrations (ppb) in groundwater collected at 5 monitoring
```

```

# wells. The data for this example are stored in
# EPA.09.Ex.15.1.manganese.df. Here we will create an empirical
# CDF plot based on the Kaplan-Meier method.

EPA.09.Ex.15.1.manganese.df
#   Sample  Well Manganese.Orig.ppb Manganese.ppb Censored
#1      1 Well.1          <5          5.0      TRUE
#2      2 Well.1          12.1         12.1     FALSE
#3      3 Well.1          16.9         16.9     FALSE
#4      4 Well.1          21.6         21.6     FALSE
#5      5 Well.1          <2          2.0      TRUE
#...
#21     1 Well.5          17.9         17.9     FALSE
#22     2 Well.5          22.7         22.7     FALSE
#23     3 Well.5           3.3           3.3     FALSE
#24     4 Well.5           8.4           8.4     FALSE
#25     5 Well.5          <2          2.0      TRUE

windows()
with(EPA.09.Ex.15.1.manganese.df,
     ecdfPlotCensored(Manganese.ppb, Censored,
                      prob.method = "kaplan-meier", ecdf.col = "blue",
                      main = "Empirical CDF of Manganese Data\nBased on Kaplan-Meier"))

#=====

# Clean up
#-----
graphics.off()

```

eevd

Estimate Parameters of an Extreme Value (Gumbel) Distribution

Description

Estimate the location and scale parameters of an [extreme value distribution](#), and optionally construct a confidence interval for one of the parameters.

Usage

```

eevd(x, method = "mle", pwme.method = "unbiased",
     plot.pos.cons = c(a = 0.35, b = 0), ci = FALSE,
     ci.parameter = "location", ci.type = "two-sided",
     ci.method = "normal.approx", conf.level = 0.95)

```

Arguments

x	numeric vector of observations.
method	character string specifying the method of estimation. Possible values are "mle" (maximum likelihood; the default), "mme" (methods of moments), "mmue" (method of moments based on the unbiased estimator of variance), and "pwme" (probability-weighted moments). See the DETAILS section for more information on these estimation methods.

<code>pwme.method</code>	character string specifying what method to use to compute the probability-weighted moments when <code>method="pwme"</code> . The possible values are "ubiased" (method based on the U-statistic; the default), or "plotting.position" (method based on the plotting position formula). See the DETAILS section in this help file and the help file for pwMoment for more information. This argument is ignored if <code>method</code> is not equal to "pwme".
<code>plot.pos.cons</code>	numeric vector of length 2 specifying the constants used in the formula for the plotting positions when <code>method="pwme"</code> and <code>pwme.method="plotting.position"</code> . The default value is <code>plot.pos.cons=c(a=0.35, b=0)</code> . If this vector has a <code>names</code> attribute with the value <code>c("a", "b")</code> or <code>c("b", "a")</code> , then the elements will be matched by name in the formula for computing the plotting positions. Otherwise, the first element is mapped to the name "a" and the second element to the name "b". See the DETAILS section in this help file and the help file for pwMoment for more information. This argument is ignored if <code>method</code> is not equal to "pwme" or if <code>pwme.method="ubiased"</code> .
<code>ci</code>	logical scalar indicating whether to compute a confidence interval for the location or scale parameter. The default value is FALSE.
<code>ci.parameter</code>	character string indicating the parameter for which the confidence interval is desired. The possible values are "location" (the default) and "scale". This argument is ignored if <code>ci=FALSE</code> .
<code>ci.type</code>	character string indicating what kind of confidence interval to compute. The possible values are "two-sided" (the default), "lower", and "upper". This argument is ignored if <code>ci=FALSE</code> .
<code>ci.method</code>	character string indicating what method to use to construct the confidence interval for the location or scale parameter. Currently, the only possible value is "normal.approx" (the default). See the DETAILS section for more information. This argument is ignored if <code>ci=FALSE</code> .
<code>conf.level</code>	a scalar between 0 and 1 indicating the confidence level of the confidence interval. The default value is <code>conf.level=0.95</code> . This argument is ignored if <code>ci=FALSE</code> .

Details

If `x` contains any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

Let $\underline{x} = (x_1, x_2, \dots, x_n)$ be a vector of n observations from an [extreme value distribution](#) with parameters `location= η` and `scale= θ` .

Estimation

Maximum Likelihood Estimation (`method="mle"`)

The maximum likelihood estimators (mle's) of η and θ are the solutions of the simultaneous equations (Forbes et al., 2011):

$$\hat{\eta}_{mle} = \hat{\theta}_{mle} \log \left[\frac{1}{n} \sum_{i=1}^n \exp \left(\frac{-x_i}{\hat{\theta}_{mle}} \right) \right]$$

$$\hat{\theta}_{mle} = \bar{x} - \frac{\sum_{i=1}^n x_i \exp \left(\frac{-x_i}{\hat{\theta}_{mle}} \right)}{\sum_{i=1}^n \exp \left(\frac{-x_i}{\hat{\theta}_{mle}} \right)}$$

where

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i$$

Method of Moments Estimation (method="mme")

The method of moments estimators (mme's) of η and θ are given by (Johnson et al., 1995, p.27):

$$\hat{\eta}_{mme} = \bar{x} - \epsilon \hat{\theta}_{mme}$$

$$\hat{\theta}_{mme} = \frac{\sqrt{6}}{\pi} s_m$$

where ϵ denotes [Euler's constant](#) and s_m denotes the square root of the method of moments estimator of variance:

$$s_m^2 = \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2$$

Method of Moments Estimators Based on the Unbiased Estimator of Variance (method="mmue")

These estimators are the same as the method of moments estimators except that the method of moments estimator of variance is replaced with the unbiased estimator of variance:

$$s^2 = \frac{1}{n-1} \sum_{i=1}^n (x_i - \bar{x})^2$$

Probability-Weighted Moments Estimation (method="pwme")

Greenwood et al. (1979) show that the relationship between the distribution parameters η and θ and the probability-weighted moments is given by:

$$\eta = M(1, 0, 0) - \epsilon \theta$$

$$\theta = \frac{M(1, 0, 0) - 2M(1, 0, 1)}{\log(2)}$$

where $M(i, j, k)$ denotes the ijk 'th probability-weighted moment and ϵ denotes [Euler's constant](#). The probability-weighted moment estimators (pwme's) of η and θ are computed by simply replacing the $M(i, j, k)$'s in the above two equations with estimates of the $M(i, j, k)$'s (and for the estimate of η , replacing θ with its estimated value). See the help file for [pwMoment](#) for more information on how to estimate the $M(i, j, k)$'s. Also, see Landwehr et al. (1979) for an example of this method of estimation using the unbiased (U-statistic type) probability-weighted moment estimators. Hosking et al. (1985) note that this method of estimation using the U-statistic type probability-weighted moments is equivalent to Downton's (1966) linear estimates with linear coefficients.

Confidence Intervals

When ci=TRUE, an approximate $(1 - \alpha)100\%$ confidence intervals for η can be constructed assuming the distribution of the estimator of η is approximately normally distributed. A two-sided confidence interval is constructed as:

$$[\hat{\eta} - t(n-1, 1 - \alpha/2) \hat{\sigma}_{\hat{\eta}}, \hat{\eta} + t(n-1, 1 - \alpha/2) \hat{\sigma}_{\hat{\eta}}]$$

where $t(\nu, p)$ is the p 'th quantile of [Student's t-distribution](#) with ν degrees of freedom, and the quantity

$$\hat{\sigma}_{\hat{\eta}}$$

denotes the estimated asymptotic standard deviation of the estimator of η .

Similarly, a two-sided confidence interval for θ is constructed as:

$$[\hat{\theta} - t(n-1, 1-\alpha/2)\hat{\sigma}_{\hat{\theta}}, \hat{\theta} + t(n-1, 1-\alpha/2)\hat{\sigma}_{\hat{\theta}}]$$

One-sided confidence intervals for η and θ are computed in a similar fashion.

Maximum Likelihood (method="mle")

Downton (1966) shows that the estimated asymptotic variances of the mle's of η and θ are given by:

$$\hat{\sigma}_{\hat{\eta}_{mle}}^2 = \frac{\hat{\theta}_m l e^2}{n} \left[1 + \frac{6(1-\epsilon)^2}{\pi^2} \right] = \frac{1.10867 \hat{\theta}_m l e^2}{n}$$

$$\hat{\sigma}_{\hat{\theta}_{mle}}^2 = \frac{6}{\pi^2} \frac{\hat{\theta}_m l e^2}{n} = \frac{0.60793 \hat{\theta}_m l e^2}{n}$$

where ϵ denotes [Euler's constant](#).

Method of Moments (method="mme" or method="mmue")

Tiago de Oliveira (1963) and Johnson et al. (1995, p.27) show that the estimated asymptotic variance of the mme's of η and θ are given by:

$$\hat{\sigma}_{\hat{\eta}_{mme}}^2 = \frac{\hat{\theta}_m m e^2}{n} \left[\frac{\pi^2}{6} + \frac{\epsilon^2}{4} (\beta_2 - 1) - \frac{\pi \epsilon \sqrt{\beta_1}}{\sqrt{6}} \right] = \frac{1.1678 \hat{\theta}_m m e^2}{n}$$

$$\hat{\sigma}_{\hat{\theta}_{mme}}^2 = \frac{\hat{\theta}_m l e^2}{n} \frac{(\beta_2 - 1)}{4} = \frac{1.1 \hat{\theta}_m m e^2}{n}$$

where the quantities

$$\sqrt{\beta_1}, \beta_2$$

denote the skew and kurtosis of the distribution, and ϵ denotes [Euler's constant](#).

The estimated asymptotic variances of the mmue's of η and θ are the same, except replace the mme of θ in the above equations with the mmue of θ .

Probability-Weighted Moments (method="pwme")

As stated above, Hosking et al. (1985) note that this method of estimation using the U-statistic type probability-weighted moments is equivalent to Downton's (1966) linear estimates with linear coefficients. Downton (1966) provides exact values of the variances of the estimates of location and scale parameters for the smallest extreme value distribution. For the largest extreme value distribution, the formula for the estimate of scale is the same, but the formula for the estimate of location must be modified. Thus, Downton's (1966) equation (3.4) is modified to:

$$\hat{\eta}_{pwme} = \frac{(n-1)\log(2) + (n+1)\epsilon}{n(n-1)\log(2)} v - \frac{2\epsilon}{n(n-1)\log(2)} w$$

where ϵ denotes [Euler's constant](#), and v and w are defined in Downton (1966, p.8). Using Downton's (1966) equations (3.9)-(3.12), the exact variance of the pwme of η can be derived. Note that when method="pwme" and pwme.method="plotting.position", these are only the asymptotically correct variances.

Value

a list of class "estimate" containing the estimated parameters and other information. See [estimate.object](#) for details.

Note

There are three families of extreme value distributions. The one described here is the [Type I, also called the Gumbel extreme value distribution or simply Gumbel distribution](#). The name “extreme value” comes from the fact that this distribution is the limiting distribution (as n approaches infinity) of the greatest value among n independent random variables each having the same continuous distribution.

The Gumbel extreme value distribution is related to the [exponential distribution](#) as follows. Let Y be an [exponential random variable](#) with parameter $\text{rate}=\lambda$. Then $X = \eta - \log(Y)$ has an extreme value distribution with parameters $\text{location}=\eta$ and $\text{scale}=1/\lambda$.

The distribution described above and assumed by eeve is the *largest* extreme value distribution. The smallest extreme value distribution is the limiting distribution (as n approaches infinity) of the smallest value among n independent random variables each having the same continuous distribution. If X has a largest extreme value distribution with parameters $\text{location}=\eta$ and $\text{scale}=\theta$, then $Y = -X$ has a smallest extreme value distribution with parameters $\text{location}=-\eta$ and $\text{scale}=\theta$. The smallest extreme value distribution is related to the [Weibull distribution](#) as follows. Let Y be a [Weibull random variable](#) with parameters $\text{shape}=\beta$ and $\text{scale}=\alpha$. Then $X = \log(Y)$ has a smallest extreme value distribution with parameters $\text{location}=\log(\alpha)$ and $\text{scale}=1/\beta$.

The extreme value distribution has been used extensively to model the distribution of streamflow, flooding, rainfall, temperature, wind speed, and other meteorological variables, as well as material strength and life data.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Castillo, E. (1988). *Extreme Value Theory in Engineering*. Academic Press, New York, pp.184–198.
- Downton, F. (1966). Linear Estimates of Parameters in the Extreme Value Distribution. *Technometrics* **8**(1), 3–17.
- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.
- Greenwood, J.A., J.M. Landwehr, N.C. Matalas, and J.R. Wallis. (1979). Probability Weighted Moments: Definition and Relation to Parameters of Several Distributions Expressible in Inverse Form. *Water Resources Research* **15**(5), 1049–1054.
- Hosking, J.R.M., J.R. Wallis, and E.F. Wood. (1985). Estimation of the Generalized Extreme-Value Distribution by the Method of Probability-Weighted Moments. *Technometrics* **27**(3), 251–261.
- Johnson, N. L., S. Kotz, and N. Balakrishnan. (1995). *Continuous Univariate Distributions, Volume 2*. Second Edition. John Wiley and Sons, New York.
- Landwehr, J.M., N.C. Matalas, and J.R. Wallis. (1979). Probability Weighted Moments Compared With Some Traditional Techniques in Estimating Gumbel Parameters and Quantiles. *Water Resources Research* **15**(5), 1055–1064.
- Tiago de Oliveira, J. (1963). Decision Results for the Parameters of the Extreme Value (Gumbel) Distribution Based on the Mean and Standard Deviation. *Trabajos de Estadística* **14**, 61–81.

See Also

[Extreme Value Distribution, Euler’s Constant](#).

Examples

```
# Generate 20 observations from an extreme value distribution with
# parameters location=2 and scale=1, then estimate the parameters
# and construct a 90% confidence interval for the location parameter.
# (Note: the call to set.seed simply allows you to reproduce this example.)
```

```
set.seed(250)
dat <- revd(20, location = 2)
eevd(dat, ci = TRUE, conf.level = 0.9)
```

```
#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:      Extreme Value
#
#Estimated Parameter(s):   location = 1.9684093
#                           scale    = 0.7481955
#
#Estimation Method:        mle
#
#Data:                     dat
#
#Sample Size:              20
#
#Confidence Interval for:   location
#
#Confidence Interval Method: Normal Approximation
#                           (t Distribution)
#
#Confidence Interval Type:  two-sided
#
#Confidence Level:         90%
#
#Confidence Interval:      LCL = 1.663809
#                           UCL = 2.273009
```

```
#-----
```

```
#Compare the values of the different types of estimators:
```

```
eevd(dat, method = "mle")$parameters
# location      scale
#1.9684093 0.7481955
```

```
eevd(dat, method = "mme")$parameters
# location      scale
#1.9575980 0.8339256
```

```
eevd(dat, method = "mmue")$parameters
# location      scale
#1.9450932 0.8555896
```

```
eevd(dat, method = "pwme")$parameters
# location      scale
#1.9434922 0.8583633
```

```
#-----
# Clean up
#-----
rm(dat)
```

eexp

Estimate Rate Parameter of an Exponential Distribution

Description

Estimate the rate parameter of an [exponential distribution](#), and optionally construct a confidence interval for the rate parameter.

Usage

```
eexp(x, method = "mle/mme", ci = FALSE, ci.type = "two-sided",
     ci.method = "exact", conf.level = 0.95)
```

Arguments

x	numeric vector of observations.
method	character string specifying the method of estimation. Currently the only possible value is "mle/mme" (maximum likelihood/method of moments; the default). See the DETAILS section for more information.
ci	logical scalar indicating whether to compute a confidence interval for the location or scale parameter. The default value is FALSE.
ci.type	character string indicating what kind of confidence interval to compute. The possible values are "two-sided" (the default), "lower", and "upper". This argument is ignored if ci=FALSE.
ci.method	character string indicating what method to use to construct the confidence interval for the location or scale parameter. Currently, the only possible value is "exact" (the default). See the DETAILS section for more information. This argument is ignored if ci=FALSE.
conf.level	a scalar between 0 and 1 indicating the confidence level of the confidence interval. The default value is conf.level=0.95. This argument is ignored if ci=FALSE.

Details

If x contains any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

Let $\underline{x} = (x_1, x_2, \dots, x_n)$ be a vector of n observations from an [exponential distribution](#) with parameter $\text{rate}=\lambda$.

Estimation

The maximum likelihood estimator (mle) of λ is given by:

$$\hat{\lambda}_{mle} = \frac{1}{\bar{x}}$$

where

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i$$

(Forbes et al., 2011). That is, the mle is the reciprocal of the sample mean.

Sometimes the exponential distribution is parameterized with a scale parameter instead of a rate parameter. The scale parameter is the reciprocal of the rate parameter, and the sample mean is both the mle and the minimum variance unbiased estimator (mvue) of the scale parameter.

Confidence Interval

When `ci=TRUE`, an exact $(1 - \alpha)100\%$ confidence intervals for λ can be constructed based on the relationship between the exponential distribution, the [gamma distribution](#), and the [chi-square distribution](#). An exponential distribution with parameter `rate= λ` is equivalent to a gamma distribution with parameters `shape=1` and `scale= $1/\lambda$` . The sum of n iid gamma random variables with parameters `shape=1` and `scale= $1/\lambda$` is a gamma random variable with parameters `shape= n` and `scale= $1/\lambda$` . Finally, a gamma distribution with parameters `shape= n` and `scale= $1/\lambda$` is equivalent to 0.5 times a chi-square distribution with degrees of freedom `df= $2n$` . Thus, the quantity $2n\bar{x}$ has a chi-square distribution with degrees of freedom `df= $2n$` .

A two-sided $(1 - \alpha)100\%$ confidence interval for λ is therefore constructed as:

$$\left[\frac{\chi^2(2n, \alpha/2)}{2n\bar{x}}, \frac{\chi^2(2n, 1 - \alpha/2)}{2n\bar{x}} \right]$$

where $\chi^2(\nu, p)$ is the p 'th quantile of a [chi-square distribution](#) with ν degrees of freedom.

One-sided confidence intervals are computed in a similar fashion.

Value

a list of class "estimate" containing the estimated parameters and other information. See [estimate.object](#) for details.

Note

The [exponential distribution](#) is a special case of the [gamma distribution](#), and takes on positive real values. A major use of the exponential distribution is in life testing where it is used to model the lifetime of a product, part, person, etc.

The exponential distribution is the only continuous distribution with a "lack of memory" property. That is, if the lifetime of a part follows the exponential distribution, then the distribution of the time until failure is the same as the distribution of the time until failure given that the part has survived to time t .

The exponential distribution is related to the double exponential (also called Laplace) distribution, and to the [extreme value distribution](#).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). Statistical Distributions. Fourth Edition. John Wiley and Sons, Hoboken, NJ.
- Johnson, N. L., S. Kotz, and N. Balakrishnan. (1994). *Continuous Univariate Distributions, Volume 1*. Second Edition. John Wiley and Sons, New York.

See Also

[Exponential](#).

Examples

```
# Generate 20 observations from an exponential distribution with parameter
# rate=2, then estimate the parameter and construct a 90% confidence interval.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
dat <- rexp(20, rate = 2)
eexp(dat, ci=TRUE, conf = 0.9)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Exponential
#
#Estimated Parameter(s):       rate = 2.260587
#
#Estimation Method:            mle/mme
#
#Data:                          dat
#
#Sample Size:                  20
#
#Confidence Interval for:      rate
#
#Confidence Interval Method:    Exact
#
#Confidence Interval Type:      two-sided
#
#Confidence Level:              90%
#
#Confidence Interval:          LCL = 1.498165
#                               UCL = 3.151173
#-----

# Clean up
#-----
rm(dat)
```

egamma

Estimate Parameters of Gamma Distribution

Description

Estimate the shape and scale parameters (or the mean and coefficient of variation) of a [Gamma](#) distribution.

Usage

```
egamma(x, method = "mle", ci = FALSE,
       ci.type = "two-sided", ci.method = "normal.approx",
       normal.approx.transform = "kulkarni.powar", conf.level = 0.95)

egammaAlt(x, method = "mle", ci = FALSE,
          ci.type = "two-sided", ci.method = "normal.approx",
          normal.approx.transform = "kulkarni.powar", conf.level = 0.95)
```

Arguments

<code>x</code>	numeric vector of non-negative observations. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
<code>method</code>	character string specifying the method of estimation. The possible values are: "mle" (maximum likelihood; the default), "bcmle" (bias-corrected mle), "mme" (method of moments), and "mmue" (method of moments based on the unbiased estimator of variance). See the DETAILS section for more information.
<code>ci</code>	logical scalar indicating whether to compute a confidence interval for the mean. The default value is <code>ci=FALSE</code> .
<code>ci.type</code>	character string indicating what kind of confidence interval to compute. The possible values are "two-sided" (the default), "lower", and "upper". This argument is ignored if <code>ci=FALSE</code> .
<code>ci.method</code>	character string indicating which method to use to construct the confidence interval. Possible values are "normal.approx" (the default), "profile.likelihood", and "chisq.approx". This argument is ignored if <code>ci=FALSE</code> .
<code>normal.approx.transform</code>	character string indicating which power transformation to use when <code>ci.method="normal.approx"</code> . Possible values are "kulkarni.powar" (the default), "cube.root", and "fourth.root". See the DETAILS section for more information. This argument is ignored if <code>ci=FALSE</code> or <code>ci.method="chisq.approx"</code> .
<code>conf.level</code>	a scalar between 0 and 1 indicating the confidence level of the confidence interval. The default value is <code>conf.level=0.95</code> . This argument is ignored if <code>ci=FALSE</code> .

Details

If `x` contains any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

Let $\underline{x} = x_1, x_2, \dots, x_n$ denote a random sample of n observations from a [gamma distribution](#) with parameters $\text{shape}=\alpha$ and $\text{scale}=\beta$. The relationship between these parameters and the mean ($\text{mean}=\mu$) and coefficient of variation ($\text{cv}=\tau$) of this distribution is given by:

$$\alpha = \tau^{-2} \quad (1)$$

$$\beta = \mu/\alpha \quad (2)$$

$$\mu = \alpha \beta \quad (3)$$

$$\tau = \alpha^{-1/2} \quad (4)$$

The function `egamma` returns estimates of the shape and scale parameters. The function `egammaAlt` returns estimates of the mean (μ) and coefficient of variation (cv) based on the estimates of the shape and scale parameters.

Estimation

Maximum Likelihood Estimation (method="mle")

The maximum likelihood estimators (mle's) of the shape and scale parameters α and β are solutions of the simultaneous equations:

$$\hat{\alpha}_{mle} = \frac{1}{n} \sum_{i=1}^n \log(x_i) - \log(\bar{x}) = \psi(\hat{\alpha}_{mle}) - \log(\hat{\alpha}_{mle}) \quad (5)$$

$$\hat{\beta}_{mle} = \bar{x} / \hat{\alpha} \quad (6)$$

where ψ denotes the [digamma function](#), and \bar{x} denotes the sample mean:

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i \quad (7)$$

(Forbes et al., 2011, chapter 22; Johnson et al., 1994, chapter 17).

Bias-Corrected Maximum Likelihood Estimation (method="bcmle")

The "bias-corrected" maximum likelihood estimator of the shape parameter is based on the suggestion of Anderson and Ray (1975; see also Johnson et al., 1994, p.366 and Singh et al., 2010b, p.48), who noted that the bias of the maximum likelihood estimator of the shape parameter can be considerable when the sample size is small. This estimator is given by:

$$\hat{\alpha}_{bcmle} = \frac{n-3}{n} \hat{\alpha}_{mle} + \frac{2}{3n} \quad (8)$$

The estimate of the scale paramter is not modified (i.e., the mle of β is returned).

Method of Moments Estimation (method="mme")

The method of moments estimators (mme's) of the shape and scale parameters α and β are:

$$\hat{\alpha}_{mme} = (\bar{x} / s_m)^2 \quad (9)$$

$$\hat{\beta}_{mme} = s_m^2 / \bar{x} \quad (10)$$

where s_m^2 denotes the method of moments estimator of variance:

$$s_m^2 = \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2 \quad (11)$$

Method of Moments Estimation Based on the Unbiased Estimator of Variance (method="mmue")

The method of moments estimators based on the unbiased estimator of variance are exactly the same as the method of moments estimators, except that the method of moments estimator of variance is replaced with the unbiased estimator of variance:

$$\hat{\alpha}_{mmue} = (\bar{x} / s)^2 \quad (12)$$

$$\hat{\beta}_{mmue} = s^2 / \bar{x} \quad (13)$$

where s^2 denotes the unbiased estimator of variance:

$$s^2 = \frac{1}{n-1} \sum_{i=1}^n (x_i - \bar{x})^2 \quad (14)$$

Confidence Intervals

This section discusses how confidence intervals for the mean μ are computed.

Normal Approximation (ci.method="normal.approx")

The normal approximation method is based on the method of Kulkarni and Powar (2010), who use a power transformation of the the original data to approximate a sample from a normal distributon, compute the confidence interval for the mean on the transformed scale using the usual formula for a confidence interval for the mean of a normal distributon, and then tranform the limits back to the original space using equations based on the expected value of a gamma random variable raised to a power.

The particular power used for the normal approximation is defined by the argument normal.approx.transform. The value normal.approx.transform="cube.root" uses the cube root transformation suggested by Wilson and Hilferty (1931), and the value "fourth.root" uses the fourth root transformation suggested by Hawkins and Wixley (1986). The default value "kulkarni.powar" uses the "Optimum Power Normal Approximation Method" of Kulkarni and Powar (2010), who show this method performs the best in terms of maintining coverage and minimizing confidence interval width compared to eight other methods. The "optimum" power p is determined by:

$$\begin{aligned} p &= -0.0705 - 0.178\hat{\alpha} + 0.475\sqrt{\hat{\alpha}} & \text{if } \hat{\alpha} \leq 1.5 \\ p &= 0.246 & \text{if } \hat{\alpha} > 1.5 \end{aligned} \quad (15)$$

where $\hat{\alpha}$ denotes the estimate of the shape parameter. Kulkarni and Powar (2010) derived this equation by determining what power transformation yields a skew closest to 0 and a kurtosis closest to 3 for a gamma random variable with a given shape parameter. Although Kulkarni and Powar (2010) use the maximum likelihood estimate of shape to determine the power to use to induce approximate normality, for the functions egamma and egammaAlt the power is based on whatever estimate of shape is used (e.g., method="mle", method="bcmle", etc.).

Likelihood Profile (ci.method="profile.likelihood")

This method was proposed by Cox (1970, p.88), and Venzon and Moolgavkar (1988) introduced an efficient method of computation. This method is also discussed by Stryhn and Christensen (2003) and Royston (2007). The idea behind this method is to invert the likelihood-ratio test to obtain a confidence interval for the mean μ while treating the coefficient of variation τ as a nuisance parameter.

The likelihood function is given by:

$$L(\theta, \tau | \underline{x}) = \prod_{i=1}^n \frac{x_i^{\alpha-1} e^{-x_i/\beta}}{\beta^\alpha \Gamma(\alpha)} \quad (16)$$

where α and β are defined in Equations (1) and (2) above, and $\Gamma(t)$ denotes the [Gamma function](#) evaluated at t .

Following Stryhn and Christensen (2003), denote the maximum likelihood estimates of the mean and coefficient of variation by (θ^*, τ^*) . The likelihood ratio test statistic (G^2) of the hypothesis $H_0 : \theta = \theta_0$ (where θ_0 is a fixed value) equals the drop in $2\log(L)$ between the "full" model and

the reduced model with θ fixed at θ_0 , i.e.,

$$G^2 = 2\{\log[L(\theta^*, \tau^*)] - \log[L(\theta_0, \tau_0^*)]\} \quad (17)$$

where τ_0^* is the maximum likelihood estimate of τ for the reduced model (i.e., when $\theta = \theta_0$). Under the null hypothesis, the test statistic G^2 follows a [chi-squared distribution](#) with 1 degree of freedom.

Alternatively, we may express the test statistic in terms of the profile likelihood function L_1 for the mean θ , which is obtained from the usual likelihood function by maximizing over the parameter τ , i.e.,

$$L_1(\theta) = \max_{\tau} L(\theta, \tau) \quad (18)$$

Then we have

$$G^2 = 2\{\log[L_1(\theta^*)] - \log[L_1(\theta_0)]\} \quad (19)$$

A two-sided $(1 - \alpha)100\%$ confidence interval for the mean θ consists of all values of θ_0 for which the test is not significant at level *alpha*:

$$\theta_0 : G^2 \leq \chi_{1,1-\alpha}^2 \quad (20)$$

where $\chi_{\nu,p}^2$ denotes the p 'th quantile of the [chi-squared distribution](#) with ν degrees of freedom.

Chi-Square Approximation (ci.method="chisq.approx")

This method is based on the relationship between the sample mean of the gamma distribution and the chi-squared distribution (Grice and Bain, 1980). Because this method is exact only when the shape parameter α is known, the method used here is called the "chi-square approximation" method because the estimate of the shape parameter is used. This method is **not** the method proposed by Grice and Bain (1980) in which the confidence interval is adjusted based on adjusting for the fact that the shape parameter is estimated. The chi-square approximation method used by `egamma` and `egammaAlt` is equivalent to the "Approximate gamma" method of Singh et al. (2010b, equation (2-34), p.51).

Value

a list of class "estimate" containing the estimated parameters and other information. See [estimate.object](#) for details.

Warning

When `ci=TRUE` and `ci.method="normal.approx"`, it is possible for the lower confidence limit based on the transformed data to be less than 0. In this case, the lower confidence limit on the original scale is set to 0 and a warning is issued stating that the normal approximation is not accurate in this case.

Note

The gamma distribution takes values on the positive real line. Special cases of the gamma are the [exponential](#) distribution and the [chi-square](#) distributions. Applications of the gamma include life testing, statistical ecology, queuing theory, inventory control, and precipitation processes. A gamma distribution starts to resemble a normal distribution as the shape parameter α tends to infinity.

Some EPA guidance documents (e.g., Singh et al., 2002; Singh et al., 2010a,b) strongly recommend against using a lognormal model for environmental data and recommend trying a gamma distribution instead.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Anderson, C.W., and W.D. Ray. (1975). Improved Maximum Likelihood Estimators for the Gamma Distribution. *Communications in Statistics*, **4**, 437–448.
- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions, Fourth Edition*. John Wiley and Sons, Hoboken, NJ.
- Grice, J.V., and L.J. Bain. (1980). Inferences Concerning the Mean of the Gamma Distribution. *Journal of the American Statistician*, **75**, 929–933.
- Hawkins, D. M., and R.A.J. Wixley. (1986). A Note on the Transformation of Chi-Squared Variables to Normality. *The American Statistician*, **40**, 296–298.
- Johnson, N.L., S. Kotz, and N. Balakrishnan. (1994). *Continuous Univariate Distributions, Volume I*. Second Edition. John Wiley and Sons, New York, Chapter 17.
- Kulkarni, H.V., and S.K. Powar. (2010). A New Method for Interval Estimation of the Mean of the Gamma Distribution. *Lifetime Data Analysis*, **16**, 431–447.
- Singh, A., A.K. Singh, and R.J. Iaci. (2002). *Estimation of the Exposure Point Concentration Term Using a Gamma Distribution*. EPA/600/R-02/084. October 2002. Technology Support Center for Monitoring and Site Characterization, Office of Research and Development, Office of Solid Waste and Emergency Response, U.S. Environmental Protection Agency, Washington, D.C.
- Singh, A., R. Maichle, and N. Armbya. (2010a). *ProUCL Version 4.1.00 User Guide (Draft)*. EPA/600/R-07/041, May 2010. Office of Research and Development, U.S. Environmental Protection Agency, Washington, D.C.
- Singh, A., N. Armbya, and A. Singh. (2010b). *ProUCL Version 4.1.00 Technical Guide (Draft)*. EPA/600/R-07/041, May 2010. Office of Research and Development, U.S. Environmental Protection Agency, Washington, D.C.
- Wilson, E.B., and M.M. Hilferty. (1931). The Distribution of Chi-Squares. *Proceedings of the National Academy of Sciences*, **17**, 684–688.

See Also

[GammaDist](#), [estimate.object](#), [eqgamma](#), [predIntGamma](#), [tolIntGamma](#).

Examples

```
# Generate 20 observations from a gamma distribution with parameters
# shape=3 and scale=2, then estimate the parameters.
# (Note: the call to set.seed simply allows you to reproduce this
# example.)

set.seed(250)
dat <- rgamma(20, shape = 3, scale = 2)
egamma(dat, ci = TRUE)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Gamma
#
#Estimated Parameter(s):      shape = 2.203862
```



```

#Sample Size:          47
#
#Test Statistic:       W = 0.9703827
#
#Test Statistic Parameter:  n = 47
#
#P-value:             0.2739512
#
#Alternative Hypothesis: True cdf does not equal the
#                       Gamma Distribution.

#-----
# Now estimate the paramters and compute the upper confidence
# limit.

egamma(TcCB[Area == "Reference"], method = "bcmle", ci = TRUE,
       ci.type = "upper", conf.level = 0.9)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:   Gamma
#
#Estimated Parameter(s): shape = 4.5695247
#                       scale = 0.1309788
#
#Estimation Method:     bcmle
#
#Data:                  TcCB[Area == "Reference"]
#
#Sample Size:           47
#
#Confidence Interval for: mean
#
#Confidence Interval Method: Optimum Power Normal Approximation
#                           of Kulkarni & Powar (2010)
#                           using bcmle of 'shape'
#
#Normal Transform Power: 0.246
#
#Confidence Interval Type: upper
#
#Confidence Level:      90%
#
#Confidence Interval:    LCL = 0.0000000
#                       UCL = 0.6561838
#-----

# Repeat the above example but use the alternative
# parameterization.

egammaAlt(TcCB[Area == "Reference"], method = "bcmle", ci = TRUE,
          ci.type = "upper", conf.level = 0.9)

#Results of Distribution Parameter Estimation
#-----

```

```
#
#Assumed Distribution:      Gamma
#
#Estimated Parameter(s):  mean = 0.5985106
#                          cv   = 0.4678046
#
#Estimation Method:       bcmle of 'shape'
#
#Data:                    TcCB[Area == "Reference"]
#
#Sample Size:             47
#
#Confidence Interval for:  mean
#
#Confidence Interval Method: Optimum Power Normal Approximation
#                             of Kulkarni & Powar (2010)
#                             using bcmle of 'shape'
#
#Normal Transform Power:   0.246
#
#Confidence Interval Type: upper
#
#Confidence Level:         90%
#
#Confidence Interval:      LCL = 0.0000000
#                          UCL = 0.6561838

#-----

# Clean up
#-----
detach("EPA.94b.tccb.df")
```

egammaAltCensored	<i>Estimate Mean and Coefficient of Variation for a Gamma Distribution Based on Type I Censored Data</i>
-------------------	--

Description

Estimate the mean and coefficient of variation of a [gamma distribution](#) given a sample of data that has been subjected to Type I censoring, and optionally construct a confidence interval for the mean.

Usage

```
egammaAltCensored(x, censored, method = "mle", censoring.side = "left",
  ci = FALSE, ci.method = "profile.likelihood", ci.type = "two-sided",
  conf.level = 0.95, n.bootstraps = 1000, use.acc.con = FALSE,
  pivot.statistic = "z", ci.sample.size = sum(!censored))
```

Arguments

x numeric vector of observations. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.

censored	numeric or logical vector indicating which values of x are censored. This must be the same length as x . If the mode of censored is "logical", TRUE values correspond to elements of x that are censored, and FALSE values correspond to elements of x that are not censored. If the mode of censored is "numeric", it must contain only 1's and 0's; 1 corresponds to TRUE and 0 corresponds to FALSE. Missing (NA) values are allowed but will be removed.
method	character string specifying the method of estimation. Currently, the only available method is maximum likelihood (method="mle").
censoring.side	character string indicating on which side the censoring occurs. The possible values are "left" (the default) and "right".
ci	logical scalar indicating whether to compute a confidence interval for the mean or variance. The default value is ci=FALSE.
ci.method	character string indicating what method to use to construct the confidence interval for the mean. The possible values are "profile.likelihood" (profile likelihood; the default), "normal.approx" (normal approximation), and "bootstrap" (based on bootstrapping). See the DETAILS section for more information. This argument is ignored if ci=FALSE.
ci.type	character string indicating what kind of confidence interval to compute. The possible values are "two-sided" (the default), "lower", and "upper". This argument is ignored if ci=FALSE.
conf.level	a scalar between 0 and 1 indicating the confidence level of the confidence interval. The default value is conf.level=0.95. This argument is ignored if ci=FALSE.
n.bootstraps	numeric scalar indicating how many bootstraps to use to construct the confidence interval for the mean when ci.type="bootstrap". This argument is ignored if ci=FALSE and/or ci.method does not equal "bootstrap".
use.acc.con	logical scalar indicating whether to use the acceleration constant when computing the bias-corrected bootstrap confidence interval (see the DETAILS section). The default value is FALSE. This argument is ignored if ci=FALSE and/or ci.method does not equal "bootstrap".
pivot.statistic	character string indicating which pivot statistic to use in the construction of the confidence interval for the mean when ci.method="normal.approx" or ci.method="normal.approx.w.cov" (see the DETAILS section). The possible values are pivot.statistic="z" (the default) and pivot.statistic="t". When pivot.statistic="t" you may supply the argument ci.sample.size (see below). The argument pivot.statistic is ignored if ci=FALSE.
ci.sample.size	numeric scalar indicating what sample size to assume to construct the confidence interval for the mean if pivot.statistic="t" and ci.method="normal.approx". The default value is the number of uncensored observations.

Details

If x or censored contain any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

Let \underline{x} denote a vector of N observations from a [gamma distribution](#) with parameters shape= α and scale= β . The relationship between these parameters and the mean μ and coefficient of variation τ of this distribution is given by:

$$\alpha = \tau^{-2} \quad (1)$$

$$\beta = \mu/\alpha \quad (2)$$

$$\mu = \alpha \beta \quad (3)$$

$$\tau = \alpha^{-1/2} \quad (4)$$

Assume n ($0 < n < N$) of these observations are known and c ($c = N - n$) of these observations are all censored below (left-censored) or all censored above (right-censored) at k fixed censoring levels

$$T_1, T_2, \dots, T_k; \quad k \geq 1 \quad (5)$$

For the case when $k \geq 2$, the data are said to be Type I **multiply censored**. For the case when $k = 1$, set $T = T_1$. If the data are left-censored and all n known observations are greater than or equal to T , or if the data are right-censored and all n known observations are less than or equal to T , then the data are said to be Type I **singly censored** (Nelson, 1982, p.7), otherwise they are considered to be Type I multiply censored.

Let c_j denote the number of observations censored below or above censoring level T_j for $j = 1, 2, \dots, k$, so that

$$\sum_{j=1}^k c_j = c \quad (6)$$

Let $x_{(1)}, x_{(2)}, \dots, x_{(N)}$ denote the “ordered” observations, where now “observation” means either the actual observation (for uncensored observations) or the censoring level (for censored observations). For right-censored data, if a censored observation has the same value as an uncensored one, the uncensored observation should be placed first. For left-censored data, if a censored observation has the same value as an uncensored one, the censored observation should be placed first.

Note that in this case the quantity $x_{(i)}$ does not necessarily represent the i ’th “largest” observation from the (unknown) complete sample.

Finally, let Ω (omega) denote the set of n subscripts in the “ordered” sample that correspond to uncensored observations.

Estimation

Maximum Likelihood Estimation (method="mle")

For Type I left censored data, the likelihood function is given by:

$$L(\mu, \tau | \underline{x}) = \binom{N}{c_1 c_2 \dots c_k n} \prod_{j=1}^k [F(T_j)]^{c_j} \prod_{i \in \Omega} f[x_{(i)}] \quad (7)$$

where f and F denote the probability density function (pdf) and cumulative distribution function (cdf) of the population (Cohen, 1963; Cohen, 1991, pp.6, 50). That is,

$$f(t) = \frac{t^{\alpha-1} e^{-t/\beta}}{\beta^\alpha \Gamma(\alpha)} \quad (8)$$

(Johnson et al., 1994, p.343), where α and β are defined in terms of μ and τ by Equations (1) and (2) above.

For left singly censored data, equation (7) simplifies to:

$$L(\mu, \tau | \underline{x}) = \binom{N}{c} [F(T)]^c \prod_{i=c+1}^n f[x_{(i)}] \quad (9)$$

Similarly, for Type I right censored data, the likelihood function is given by:

$$L(\mu, \tau | \underline{x}) = \binom{N}{c_1 c_2 \dots c_k n} \prod_{j=1}^k [1 - F(T_j)]^{c_j} \prod_{i \in \Omega} f[x_{(i)}] \quad (10)$$

and for right singly censored data this simplifies to:

$$L(\alpha, \beta | \underline{x}) = \binom{N}{c} [1 - F(T)]^c \prod_{i=1}^n f[x_{(i)}] \quad (11)$$

The maximum likelihood estimators are computed by minimizing the negative log-likelihood function.

Confidence Intervals

This section explains how confidence intervals for the mean μ are computed. In this section, do not confuse the parameter α used to define the confidence level of the confidence interval with the parameter α that was used earlier to denote the shape parameter of the gamma distribution.

Likelihood Profile (ci.method="profile.likelihood")

This method was proposed by Cox (1970, p.88), and Venzon and Moolgavkar (1988) introduced an efficient method of computation. This method is also discussed by Stryhn and Christensen (2003) and Royston (2007). The idea behind this method is to invert the likelihood-ratio test to obtain a confidence interval for the mean μ while treating the coefficient of variation τ as a nuisance parameter. Equation (7) above shows the form of the likelihood function $L(\mu, \tau | \underline{x})$ for multiply left-censored data, where μ and τ are defined by Equations (3) and (4), and Equation (10) shows the function for multiply right-censored data.

Following Stryhn and Christensen (2003), denote the maximum likelihood estimates of the mean and coefficient of variation by (μ^*, τ^*) . The likelihood ratio test statistic (G^2) of the hypothesis $H_0 : \mu = \mu_0$ (where μ_0 is a fixed value) equals the drop in $2\log(L)$ between the “full” model and the reduced model with μ fixed at μ_0 , i.e.,

$$G^2 = 2\{\log[L(\mu^*, \tau^*)] - \log[L(\mu_0, \tau_0^*)]\} \quad (12)$$

where τ_0^* is the maximum likelihood estimate of τ for the reduced model (i.e., when $\mu = \mu_0$). Under the null hypothesis, the test statistic G^2 follows a [chi-squared distribution](#) with 1 degree of freedom.

Alternatively, we may express the test statistic in terms of the profile likelihood function L_1 for the mean μ , which is obtained from the usual likelihood function by maximizing over the parameter τ , i.e.,

$$L_1(\mu) = \max_{\tau} L(\mu, \tau) \quad (13)$$

Then we have

$$G^2 = 2\{\log[L_1(\mu^*)] - \log[L_1(\mu_0)]\} \quad (14)$$

A two-sided $(1 - \alpha)100\%$ confidence interval for the mean μ consists of all values of μ_0 for which the test is not significant at level *alpha*:

$$\mu_0 : G^2 \leq \chi_{1, 1-\alpha}^2 \quad (15)$$

where $\chi_{\nu, p}^2$ denotes the p 'th quantile of the [chi-squared distribution](#) with ν degrees of freedom.

Normal Approximation (ci.method="normal.approx")

This method constructs approximate $(1 - \alpha)100\%$ confidence intervals for μ based on the assumption that the estimator of μ is approximately normally distributed. That is, a two-sided $(1 - \alpha)100\%$ confidence interval for μ is constructed as:

$$[\hat{\mu} - t_{1-\alpha/2, m-1} \hat{\sigma}_{\hat{\mu}}, \hat{\mu} + t_{1-\alpha/2, m-1} \hat{\sigma}_{\hat{\mu}}] \quad (16)$$

where $\hat{\mu}$ denotes the estimate of μ , $\hat{\sigma}_{\hat{\mu}}$ denotes the estimated asymptotic standard deviation of the estimator of μ , m denotes the assumed sample size for the confidence interval, and $t_{p, \nu}$ denotes the p 'th quantile of [Student's t-distribution](#) with ν degrees of freedom. One-sided confidence intervals are computed in a similar fashion.

The argument ci.sample.size determines the value of m and by default is equal to the number of uncensored observations. This is simply an ad-hoc method of constructing confidence intervals and is not based on any published theoretical results.

When pivot.statistic="z", the p 'th quantile from the [standard normal distribution](#) is used in place of the p 'th quantile from Student's t-distribution.

The standard deviation of the mle of μ is estimated based on the inverse of the Fisher Information matrix.

Bootstrap and Bias-Corrected Bootstrap Approximation (ci.method="bootstrap")

The bootstrap is a nonparametric method of estimating the distribution (and associated distribution parameters and quantiles) of a sample statistic, regardless of the distribution of the population from which the sample was drawn. The bootstrap was introduced by Efron (1979) and a general reference is Efron and Tibshirani (1993).

In the context of deriving an approximate $(1 - \alpha)100\%$ confidence interval for the population mean μ , the bootstrap can be broken down into the following steps:

1. Create a bootstrap sample by taking a random sample of size N from the observations in \underline{x} , where sampling is done with replacement. Note that because sampling is done with replacement, the same element of \underline{x} can appear more than once in the bootstrap sample. Thus, the bootstrap sample will usually not look exactly like the original sample (e.g., the number of censored observations in the bootstrap sample will often differ from the number of censored observations in the original sample).
2. Estimate μ based on the bootstrap sample created in Step 1, using the same method that was used to estimate μ using the original observations in \underline{x} . Because the bootstrap sample usually does not match the original sample, the estimate of μ based on the bootstrap sample will usually differ from the original estimate based on \underline{x} .
3. Repeat Steps 1 and 2 B times, where B is some large number. For the function `egammaAltCensored`, the number of bootstraps B is determined by the argument `n.bootstraps` (see the section ARGUMENTS above). The default value of `n.bootstraps` is 1000.
4. Use the B estimated values of μ to compute the empirical cumulative distribution function of this estimator of μ (see [ecdfPlot](#)), and then create a confidence interval for μ based on this estimated cdf.

The percentile interval (Efron and Tibshirani, 1993, p.170) is computed as:

$$[\hat{G}^{-1}(\frac{\alpha}{2}), \hat{G}^{-1}(\frac{1-\alpha}{2})] \quad (17)$$

where $\hat{G}(t)$ denotes the empirical cdf evaluated at t and thus $\hat{G}^{-1}(p)$ denotes the p 'th empirical quantile, that is, the p 'th quantile associated with the empirical cdf. The function `egammaAltCensored` calls the R function [quantile](#) to compute the empirical quantiles used in equation (17).

The percentile method bootstrap confidence interval is only first-order accurate (Efron and Tibshirani, 1993, pp.187-188), meaning that the probability that the confidence interval will contain the true value of μ can be off by k/\sqrt{N} , where k is some constant. Efron and Tibshirani (1993, pp.184-188) proposed a bias-corrected and accelerated interval that is second-order accurate, meaning that the probability that the confidence interval will contain the true value of μ may be off by k/N instead of k/\sqrt{N} . The bias-corrected and accelerated confidence interval is computed as:

$$[\hat{G}^{-1}(\alpha_1), \hat{G}^{-1}(\alpha_2)] \quad (18)$$

where

$$\alpha_1 = \Phi\left[\hat{z}_0 + \frac{\hat{z}_0 + z_{\alpha/2}}{1 - \hat{a}(z_0 + z_{\alpha/2})}\right] \quad (19)$$

$$\alpha_2 = \Phi\left[\hat{z}_0 + \frac{\hat{z}_0 + z_{1-\alpha/2}}{1 - \hat{a}(z_0 + z_{1-\alpha/2})}\right] \quad (20)$$

$$\hat{z}_0 = \Phi^{-1}[\hat{G}(\hat{\mu})] \quad (21)$$

$$\hat{a} = \frac{\sum_{i=1}^N (\hat{\mu}_{(\cdot)} - \hat{\mu}_{(i)})^3}{6[\sum_{i=1}^N (\hat{\mu}_{(\cdot)} - \hat{\mu}_{(i)})^2]^{3/2}} \quad (22)$$

where the quantity $\hat{\mu}_{(i)}$ denotes the estimate of μ using all the values in \underline{x} except the i 'th one, and

$$\hat{\mu}_{(\cdot)} = \frac{1}{N} \sum_{i=1}^N \mu_{(i)} \quad (23)$$

The constant \hat{z}_0 incorporates the bias correction, and the constant \hat{a} is the acceleration constant. The term “acceleration” refers to the rate of change of the standard error of the estimate of μ with respect to the true value of μ (Efron and Tibshirani, 1993, p.186). For a normal (Gaussian) distribution, the standard error of the estimate of μ does not depend on the value of μ , hence the acceleration constant is not really necessary. The argument `use.acc.con` (see the section ARGUMENTS above) determines whether the acceleration constant \hat{a} is computed as in equation (22) (`use.acc.con=TRUE`) or is set to 0 (`use.acc.con=FALSE`). The default value is `use.acc.con=FALSE`.

When `ci.method="bootstrap"`, the function `egammaAltCensored` computes both the percentile method and bias-corrected bootstrap confidence intervals.

Value

a list of class “`estimateCensored`” containing the estimated parameters and other information. See [estimateCensored.object](#) for details.

Note

A sample of data contains censored observations if some of the observations are reported only as being below or above some censoring level. In environmental data analysis, Type I left-censored data sets are common, with values being reported as “less than the detection limit” (e.g., Helsel, 2012). Data sets with only one censoring level are called *singly censored*; data sets with multiple censoring levels are called *multiply* or *progressively censored*.

Statistical methods for dealing with censored data sets have a long history in the field of survival analysis and life testing. More recently, researchers in the environmental field have proposed alternative methods of computing estimates and confidence intervals in addition to the classical ones such as maximum likelihood estimation. Helsel (2012, Chapter 6) gives an excellent review of past studies of the properties of various estimators for parameters of a normal or lognormal distribution based on censored environmental data.

In practice, it is better to use a confidence interval for the mean or a joint confidence region for the mean and standard deviation (or coefficient of variation), rather than rely on a single point-estimate of the mean. Few studies have been done to evaluate the performance of methods for constructing confidence intervals for the mean or joint confidence regions for the mean and coefficient of variation of a gamma distribution when data are subjected to single or multiple censoring.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Cohen, A.C. (1963). Progressively Censored Samples in Life Testing. *Technometrics* **5**, 327–339
- Cohen, A.C. (1991). *Truncated and Censored Samples*. Marcel Dekker, New York, New York, 312pp.
- Cox, D.R. (1970). *Analysis of Binary Data*. Chapman & Hall, London. 142pp.
- Efron, B. (1979). Bootstrap Methods: Another Look at the Jackknife. *The Annals of Statistics* **7**, 1–26.
- Efron, B., and R.J. Tibshirani. (1993). *An Introduction to the Bootstrap*. Chapman and Hall, New York, 436pp.
- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions, Fourth Edition*. John Wiley and Sons, Hoboken, NJ.
- Helsel, D.R. (2012). *Statistics for Censored Environmental Data Using Minitab and R, Second Edition*. John Wiley & Sons, Hoboken, New Jersey.
- Johnson, N.L., S. Kotz, and N. Balakrishnan. (1994). *Continuous Univariate Distributions, Volume 1*. Second Edition. John Wiley and Sons, New York, Chapter 17.
- Millard, S.P., P. Dixon, and N.K. Neerchal. (2014; in preparation). *Environmental Statistics with R*. CRC Press, Boca Raton, Florida.
- Nelson, W. (1982). *Applied Life Data Analysis*. John Wiley and Sons, New York, 634pp.
- Royston, P. (2007). Profile Likelihood for Estimation and Confidence Intervals. *The Stata Journal* **7**(3), pp. 376–387.
- Stryhn, H., and J. Christensen. (2003). *Confidence Intervals by the Profile Likelihood Method, with Applications in Veterinary Epidemiology*. Contributed paper at ISVEE X (November 2003, Chile). <http://people.upei.ca/hstryhn/stryhn208.pdf>.
- Venzon, D.J., and S.H. Moolgavkar. (1988). A Method for Computing Profile-Likelihood-Based Confidence Intervals. *Journal of the Royal Statistical Society, Series C (Applied Statistics)* **37**(1), pp. 87–94.

See Also

[egammaCensored](#), [GammaDist](#), [egamma](#), [estimateCensored.object](#).

Examples

```
# Chapter 15 of USEPA (2009) gives several examples of estimating the mean
# and standard deviation of a lognormal distribution on the log-scale using
# manganese concentrations (ppb) in groundwater at five background wells.
# In EnvStats these data are stored in the data frame
# EPA.09.Ex.15.1.manganese.df.
```

```
# Here we will estimate the mean and coefficient of variation
# ON THE ORIGINAL SCALE using the MLE and
# assuming a gamma distribution.
```

```
# First look at the data:
#-----
```

```
EPA.09.Ex.15.1.manganese.df
```

```
# Sample Well Manganese.Orig.ppb Manganese.ppb Censored
#1      1 Well.1          <5          5.0      TRUE
#2      2 Well.1          12.1         12.1     FALSE
#3      3 Well.1          16.9         16.9     FALSE
#...
#23     3 Well.5           3.3          3.3     FALSE
#24     4 Well.5           8.4          8.4     FALSE
#25     5 Well.5          <2           2.0      TRUE
```

```
longToWide(EPA.09.Ex.15.1.manganese.df,
  "Manganese.Orig.ppb", "Sample", "Well",
  paste.row.name = TRUE)
```

```
#      Well.1 Well.2 Well.3 Well.4 Well.5
#Sample.1  <5    <5    <5    6.3   17.9
#Sample.2  12.1   7.7   5.3   11.9  22.7
#Sample.3  16.9  53.6  12.6   10    3.3
#Sample.4  21.6   9.5  106.3  <2    8.4
#Sample.5   <2  45.9  34.5  77.2  <2
```

```
# Now estimate the mean and coefficient of variation
# using the MLE, and compute a confidence interval
# for the mean using the profile-likelihood method.
#-----
```

```
with(EPA.09.Ex.15.1.manganese.df,
  egammaAltCensored(Manganese.ppb, Censored, ci = TRUE))
```

```
#Results of Distribution Parameter Estimation
#Based on Type I Censored Data
```

```
#-----
```

```
#
#Assumed Distribution:          Gamma
#
#Censoring Side:                left
#
#Censoring Level(s):           2 5
#
#Estimated Parameter(s):       mean = 19.664797
#                               cv   =  1.252936
#
#Estimation Method:            MLE
#
#Data:                          Manganese.ppb
#
#Censoring Variable:           Censored
#
```

```
#Sample Size:                25
#
#Percent Censored:           24%
#
#Confidence Interval for:    mean
#
#Confidence Interval Method: Profile Likelihood
#
#Confidence Interval Type:   two-sided
#
#Confidence Level:           95%
#
#Confidence Interval:        LCL = 12.25151
#                             UCL = 34.35332

#-----

# Compare the confidence interval for the mean
# based on assuming a lognormal distribution versus
# assuming a gamma distribution.

with(EPA.09.Ex.15.1.manganese.df,
     elnormAltCensored(Manganese.ppb, Censored,
                       ci = TRUE))$interval$limits
#      LCL      UCL
#12.37629 69.87694

with(EPA.09.Ex.15.1.manganese.df,
     egammaAltCensored(Manganese.ppb, Censored,
                       ci = TRUE))$interval$limits
#      LCL      UCL
#12.25151 34.35332
```

egammaCensored	<i>Estimate Shape and Scale Parameters for a Gamma Distribution Based on Type I Censored Data</i>
----------------	---

Description

Estimate the shape and scale parameters of a [gamma distribution](#) given a sample of data that has been subjected to Type I censoring, and optionally construct a confidence interval for the mean.

Usage

```
egammaCensored(x, censored, method = "mle", censoring.side = "left",
  ci = FALSE, ci.method = "profile.likelihood", ci.type = "two-sided",
  conf.level = 0.95, n.bootstraps = 1000, use.acc.con = FALSE,
  pivot.statistic = "z", ci.sample.size = sum(!censored))
```

Arguments

x numeric vector of observations. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.

censored	numeric or logical vector indicating which values of x are censored. This must be the same length as x . If the mode of censored is "logical", TRUE values correspond to elements of x that are censored, and FALSE values correspond to elements of x that are not censored. If the mode of censored is "numeric", it must contain only 1's and 0's; 1 corresponds to TRUE and 0 corresponds to FALSE. Missing (NA) values are allowed but will be removed.
method	character string specifying the method of estimation. Currently, the only available method is maximum likelihood (method="mle").
censoring.side	character string indicating on which side the censoring occurs. The possible values are "left" (the default) and "right".
ci	logical scalar indicating whether to compute a confidence interval for the mean or variance. The default value is ci=FALSE.
ci.method	character string indicating what method to use to construct the confidence interval for the mean. The possible values are "profile.likelihood" (profile likelihood; the default), "normal.approx" (normal approximation), and "bootstrap" (based on bootstrapping). See the DETAILS section for more information. This argument is ignored if ci=FALSE.
ci.type	character string indicating what kind of confidence interval to compute. The possible values are "two-sided" (the default), "lower", and "upper". This argument is ignored if ci=FALSE.
conf.level	a scalar between 0 and 1 indicating the confidence level of the confidence interval. The default value is conf.level=0.95. This argument is ignored if ci=FALSE.
n.bootstraps	numeric scalar indicating how many bootstraps to use to construct the confidence interval for the mean when ci.type="bootstrap". This argument is ignored if ci=FALSE and/or ci.method does not equal "bootstrap".
use.acc.con	logical scalar indicating whether to use the acceleration constant when computing the bias-corrected bootstrap confidence interval (see the DETAILS section). The default value is FALSE. This argument is ignored if ci=FALSE and/or ci.method does not equal "bootstrap".
pivot.statistic	character string indicating which pivot statistic to use in the construction of the confidence interval for the mean when ci.method="normal.approx" or ci.method="normal.approx.w.cov" (see the DETAILS section). The possible values are pivot.statistic="z" (the default) and pivot.statistic="t". When pivot.statistic="t" you may supply the argument ci.sample.size (see below). The argument pivot.statistic is ignored if ci=FALSE.
ci.sample.size	numeric scalar indicating what sample size to assume to construct the confidence interval for the mean if pivot.statistic="t" and ci.method="normal.approx". The default value is the number of uncensored observations.

Details

If x or censored contain any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

Let \underline{x} denote a vector of N observations from a [gamma distribution](#) with parameters shape= α and scale= β . The relationship between these parameters and the mean μ and coefficient of variation τ of this distribution is given by:

$$\alpha = \tau^{-2} \quad (1)$$

$$\beta = \mu/\alpha \quad (2)$$

$$\mu = \alpha \beta \quad (3)$$

$$\tau = \alpha^{-1/2} \quad (4)$$

Assume n ($0 < n < N$) of these observations are known and c ($c = N - n$) of these observations are all censored below (left-censored) or all censored above (right-censored) at k fixed censoring levels

$$T_1, T_2, \dots, T_k; \quad k \geq 1 \quad (5)$$

For the case when $k \geq 2$, the data are said to be Type I **multiply censored**. For the case when $k = 1$, set $T = T_1$. If the data are left-censored and all n known observations are greater than or equal to T , or if the data are right-censored and all n known observations are less than or equal to T , then the data are said to be Type I **singly censored** (Nelson, 1982, p.7), otherwise they are considered to be Type I multiply censored.

Let c_j denote the number of observations censored below or above censoring level T_j for $j = 1, 2, \dots, k$, so that

$$\sum_{j=1}^k c_j = c \quad (6)$$

Let $x_{(1)}, x_{(2)}, \dots, x_{(N)}$ denote the “ordered” observations, where now “observation” means either the actual observation (for uncensored observations) or the censoring level (for censored observations). For right-censored data, if a censored observation has the same value as an uncensored one, the uncensored observation should be placed first. For left-censored data, if a censored observation has the same value as an uncensored one, the censored observation should be placed first.

Note that in this case the quantity $x_{(i)}$ does not necessarily represent the i ’th “largest” observation from the (unknown) complete sample.

Finally, let Ω (omega) denote the set of n subscripts in the “ordered” sample that correspond to uncensored observations.

Estimation

Maximum Likelihood Estimation (method="mle")

For Type I left censored data, the likelihood function is given by:

$$L(\alpha, \beta | \underline{x}) = \binom{N}{c_1 c_2 \dots c_k n} \prod_{j=1}^k [F(T_j)]^{c_j} \prod_{i \in \Omega} f[x_{(i)}] \quad (7)$$

where f and F denote the probability density function (pdf) and cumulative distribution function (cdf) of the population (Cohen, 1963; Cohen, 1991, pp.6, 50). That is,

$$f(t) = \frac{t^{\alpha-1} e^{-t/\beta}}{\beta^\alpha \Gamma(\alpha)} \quad (8)$$

(Johnson et al., 1994, p.343). For left singly censored data, equation (7) simplifies to:

$$L(\alpha, \beta | \underline{x}) = \binom{N}{c} [F(T)]^c \prod_{i=c+1}^n f[x_{(i)}] \quad (9)$$

Similarly, for Type I right censored data, the likelihood function is given by:

$$L(\alpha, \beta | \underline{x}) = \binom{N}{c_1 c_2 \dots c_k n} \prod_{j=1}^k [1 - F(T_j)]^{c_j} \prod_{i \in \Omega} f[x_{(i)}] \quad (10)$$

and for right singly censored data this simplifies to:

$$L(\alpha, \beta | \underline{x}) = \binom{N}{c} [1 - F(T)]^c \prod_{i=1}^n f[x_{(i)}] \quad (11)$$

The maximum likelihood estimators are computed by minimizing the negative log-likelihood function.

Confidence Intervals

This section explains how confidence intervals for the mean μ are computed. In this section, do not confuse the parameter α used to define the confidence level of the confidence interval with the parameter α that was used earlier to denote the shape parameter of the gamma distribution.

Likelihood Profile (ci.method="profile.likelihood")

This method was proposed by Cox (1970, p.88), and Venzon and Moolgavkar (1988) introduced an efficient method of computation. This method is also discussed by Stryhn and Christensen (2003) and Royston (2007). The idea behind this method is to invert the likelihood-ratio test to obtain a confidence interval for the mean μ while treating the coefficient of variation τ as a nuisance parameter. Equation (7) above shows the form of the likelihood function $L(\mu, \tau | \underline{x})$ for multiply left-censored data, where μ and τ are defined by Equations (3) and (4), and Equation (10) shows the function for multiply right-censored data.

Following Stryhn and Christensen (2003), denote the maximum likelihood estimates of the mean and coefficient of variation by (μ^*, τ^*) . The likelihood ratio test statistic (G^2) of the hypothesis $H_0 : \mu = \mu_0$ (where μ_0 is a fixed value) equals the drop in $2\log(L)$ between the “full” model and the reduced model with μ fixed at μ_0 , i.e.,

$$G^2 = 2\{\log[L(\mu^*, \tau^*)] - \log[L(\mu_0, \tau_0^*)]\} \quad (12)$$

where τ_0^* is the maximum likelihood estimate of τ for the reduced model (i.e., when $\mu = \mu_0$). Under the null hypothesis, the test statistic G^2 follows a [chi-squared distribution](#) with 1 degree of freedom.

Alternatively, we may express the test statistic in terms of the profile likelihood function L_1 for the mean μ , which is obtained from the usual likelihood function by maximizing over the parameter τ , i.e.,

$$L_1(\mu) = \max_{\tau} L(\mu, \tau) \quad (13)$$

Then we have

$$G^2 = 2\{\log[L_1(\mu^*)] - \log[L_1(\mu_0)]\} \quad (14)$$

A two-sided $(1 - \alpha)100\%$ confidence interval for the mean μ consists of all values of μ_0 for which the test is not significant at level *alpha*:

$$\mu_0 : G^2 \leq \chi_{1,1-\alpha}^2 \quad (15)$$

where $\chi_{\nu,p}^2$ denotes the p 'th quantile of the [chi-squared distribution](#) with ν degrees of freedom.

Normal Approximation (ci.method="normal.approx")

This method constructs approximate $(1 - \alpha)100\%$ confidence intervals for μ based on the assumption that the estimator of μ is approximately normally distributed. That is, a two-sided $(1 - \alpha)100\%$ confidence interval for μ is constructed as:

$$[\hat{\mu} - t_{1-\alpha/2, m-1} \hat{\sigma}_{\hat{\mu}}, \hat{\mu} + t_{1-\alpha/2, m-1} \hat{\sigma}_{\hat{\mu}}] \quad (16)$$

where $\hat{\mu}$ denotes the estimate of μ , $\hat{\sigma}_{\hat{\mu}}$ denotes the estimated asymptotic standard deviation of the estimator of μ , m denotes the assumed sample size for the confidence interval, and $t_{p,\nu}$ denotes the p 'th quantile of [Student's t-distribution](#) with ν degrees of freedom. One-sided confidence intervals are computed in a similar fashion.

The argument `ci.sample.size` determines the value of m and by default is equal to the number of uncensored observations. This is simply an ad-hoc method of constructing confidence intervals and is not based on any published theoretical results.

When `pivot.statistic="z"`, the p 'th quantile from the [standard normal distribution](#) is used in place of the p 'th quantile from Student's t-distribution.

The standard deviation of the mle of μ is estimated based on the inverse of the Fisher Information matrix.

Bootstrap and Bias-Corrected Bootstrap Approximation (ci.method="bootstrap")

The bootstrap is a nonparametric method of estimating the distribution (and associated distribution parameters and quantiles) of a sample statistic, regardless of the distribution of the population from which the sample was drawn. The bootstrap was introduced by Efron (1979) and a general reference is Efron and Tibshirani (1993).

In the context of deriving an approximate $(1 - \alpha)100\%$ confidence interval for the population mean μ , the bootstrap can be broken down into the following steps:

1. Create a bootstrap sample by taking a random sample of size N from the observations in \underline{x} , where sampling is done with replacement. Note that because sampling is done with replacement, the same element of \underline{x} can appear more than once in the bootstrap sample. Thus, the bootstrap sample will usually not look exactly like the original sample (e.g., the number of censored observations in the bootstrap sample will often differ from the number of censored observations in the original sample).
2. Estimate μ based on the bootstrap sample created in Step 1, using the same method that was used to estimate μ using the original observations in \underline{x} . Because the bootstrap sample usually does not match the original sample, the estimate of μ based on the bootstrap sample will usually differ from the original estimate based on \underline{x} .
3. Repeat Steps 1 and 2 B times, where B is some large number. For the function `egammaCensored`, the number of bootstraps B is determined by the argument `n.bootstraps` (see the section `ARGUMENTS` above). The default value of `n.bootstraps` is 1000.
4. Use the B estimated values of μ to compute the empirical cumulative distribution function of this estimator of μ (see [ecdfPlot](#)), and then create a confidence interval for μ based on this estimated cdf.

The percentile interval (Efron and Tibshirani, 1993, p.170) is computed as:

$$[\hat{G}^{-1}(\frac{\alpha}{2}), \hat{G}^{-1}(\frac{1-\alpha}{2})] \quad (17)$$

where $\hat{G}(t)$ denotes the empirical cdf evaluated at t and thus $\hat{G}^{-1}(p)$ denotes the p 'th empirical quantile, that is, the p 'th quantile associated with the empirical cdf. The function `egammaCensored` calls the R function [quantile](#) to compute the empirical quantiles used in equation (17).

The percentile method bootstrap confidence interval is only first-order accurate (Efron and Tibshirani, 1993, pp.187-188), meaning that the probability that the confidence interval will contain the true value of μ can be off by k/\sqrt{N} , where k is some constant. Efron and Tibshirani (1993, pp.184-188) proposed a bias-corrected and accelerated interval that is second-order accurate, meaning that the probability that the confidence interval will contain the true value of μ may be off by k/N instead of k/\sqrt{N} . The bias-corrected and accelerated confidence interval is computed as:

$$[\hat{G}^{-1}(\alpha_1), \hat{G}^{-1}(\alpha_2)] \quad (18)$$

where

$$\alpha_1 = \Phi\left[\hat{z}_0 + \frac{\hat{z}_0 + z_{\alpha/2}}{1 - \hat{a}(z_0 + z_{\alpha/2})}\right] \quad (19)$$

$$\alpha_2 = \Phi\left[\hat{z}_0 + \frac{\hat{z}_0 + z_{1-\alpha/2}}{1 - \hat{a}(z_0 + z_{1-\alpha/2})}\right] \quad (20)$$

$$\hat{z}_0 = \Phi^{-1}[\hat{G}(\hat{\mu})] \quad (21)$$

$$\hat{a} = \frac{\sum_{i=1}^N (\hat{\mu}_{(\cdot)} - \hat{\mu}_{(i)})^3}{6[\sum_{i=1}^N (\hat{\mu}_{(\cdot)} - \hat{\mu}_{(i)})^2]^{3/2}} \quad (22)$$

where the quantity $\hat{\mu}_{(i)}$ denotes the estimate of μ using all the values in \underline{x} except the i 'th one, and

$$\hat{\mu}_{(\cdot)} = \frac{1}{N} \sum_{i=1}^N \hat{\mu}_{(i)} \quad (23)$$

The constant \hat{z}_0 incorporates the bias correction, and the constant \hat{a} is the acceleration constant. The term “acceleration” refers to the rate of change of the standard error of the estimate of μ with respect to the true value of μ (Efron and Tibshirani, 1993, p.186). For a normal (Gaussian) distribution, the standard error of the estimate of μ does not depend on the value of μ , hence the acceleration constant is not really necessary. The argument `use.acc.con` (see the section ARGUMENTS above) determines whether the acceleration constant \hat{a} is computed as in equation (22) (`use.acc.con=TRUE`) or is set to 0 (`use.acc.con=FALSE`). The default value is `use.acc.con=FALSE`.

When `ci.method="bootstrap"`, the function `egammaCensored` computes both the percentile method and bias-corrected bootstrap confidence intervals.

Value

a list of class “`estimateCensored`” containing the estimated parameters and other information. See [estimateCensored.object](#) for details.

Note

A sample of data contains censored observations if some of the observations are reported only as being below or above some censoring level. In environmental data analysis, Type I left-censored data sets are common, with values being reported as “less than the detection limit” (e.g., Helsel, 2012). Data sets with only one censoring level are called *singly censored*; data sets with multiple censoring levels are called *multiply* or *progressively censored*.

Statistical methods for dealing with censored data sets have a long history in the field of survival analysis and life testing. More recently, researchers in the environmental field have proposed alternative methods of computing estimates and confidence intervals in addition to the classical ones such as maximum likelihood estimation. Helsel (2012, Chapter 6) gives an excellent review of past studies of the properties of various estimators for parameters of a normal or lognormal distribution based on censored environmental data.

In practice, it is better to use a confidence interval for the mean or a joint confidence region for the mean and standard deviation (or coefficient of variation), rather than rely on a single point-estimate of the mean. Few studies have been done to evaluate the performance of methods for constructing confidence intervals for the mean or joint confidence regions for the mean and coefficient of variation of a gamma distribution when data are subjected to single or multiple censoring.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Cohen, A.C. (1963). Progressively Censored Samples in Life Testing. *Technometrics* **5**, 327–339
- Cohen, A.C. (1991). *Truncated and Censored Samples*. Marcel Dekker, New York, New York, 312pp.
- Cox, D.R. (1970). *Analysis of Binary Data*. Chapman & Hall, London. 142pp.
- Efron, B. (1979). Bootstrap Methods: Another Look at the Jackknife. *The Annals of Statistics* **7**, 1–26.
- Efron, B., and R.J. Tibshirani. (1993). *An Introduction to the Bootstrap*. Chapman and Hall, New York, 436pp.
- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions, Fourth Edition*. John Wiley and Sons, Hoboken, NJ.
- Helsel, D.R. (2012). *Statistics for Censored Environmental Data Using Minitab and R, Second Edition*. John Wiley & Sons, Hoboken, New Jersey.
- Johnson, N.L., S. Kotz, and N. Balakrishnan. (1994). *Continuous Univariate Distributions, Volume 1*. Second Edition. John Wiley and Sons, New York, Chapter 17.
- Millard, S.P., P. Dixon, and N.K. Neerchal. (2014; in preparation). *Environmental Statistics with R*. CRC Press, Boca Raton, Florida.
- Nelson, W. (1982). *Applied Life Data Analysis*. John Wiley and Sons, New York, 634pp.
- Royston, P. (2007). Profile Likelihood for Estimation and Confidence Intervals. *The Stata Journal* **7**(3), pp. 376–387.
- Stryhn, H., and J. Christensen. (2003). *Confidence Intervals by the Profile Likelihood Method, with Applications in Veterinary Epidemiology*. Contributed paper at ISVEE X (November 2003, Chile). <http://people.upei.ca/hstryhn/stryhn208.pdf>.
- Venzon, D.J., and S.H. Moolgavkar. (1988). A Method for Computing Profile-Likelihood-Based Confidence Intervals. *Journal of the Royal Statistical Society, Series C (Applied Statistics)* **37**(1), pp. 87–94.

See Also

[egammaAltCensored](#), [GammaDist](#), [egamma](#), [estimateCensored.object](#).

Examples

```
# Chapter 15 of USEPA (2009) gives several examples of estimating the mean
# and standard deviation of a lognormal distribution on the log-scale using
# manganese concentrations (ppb) in groundwater at five background wells.
# In EnvStats these data are stored in the data frame
# EPA.09.Ex.15.1.manganese.df.

# Here we will estimate the shape and scale parameters using
# the data ON THE ORIGINAL SCALE, using the MLE and
# assuming a gamma distribution.

# First look at the data:
#-----

EPA.09.Ex.15.1.manganese.df

#   Sample   Well Manganese.Orig.ppb Manganese.ppb Censored
#1         1 Well.1                <5             5.0      TRUE
```

```
#2      2 Well.1      12.1      12.1 FALSE
#3      3 Well.1      16.9      16.9 FALSE
#...
#23     3 Well.5       3.3       3.3 FALSE
#24     4 Well.5       8.4       8.4 FALSE
#25     5 Well.5       <2       2.0  TRUE
```

```
longToWide(EPA.09.Ex.15.1.manganese.df,
  "Manganese.Orig.ppb", "Sample", "Well",
  paste.row.name = TRUE)
```

```
#      Well.1 Well.2 Well.3 Well.4 Well.5
#Sample.1    <5    <5    <5    6.3   17.9
#Sample.2   12.1    7.7    5.3   11.9   22.7
#Sample.3   16.9   53.6   12.6    10    3.3
#Sample.4   21.6    9.5  106.3    <2    8.4
#Sample.5    <2   45.9   34.5   77.2    <2
```

```
# Now estimate the shape and scale parameters
# using the MLE, and compute a confidence interval
# for the mean using the profile-likelihood method.
#-----
```

```
with(EPA.09.Ex.15.1.manganese.df,
  egammaCensored(Manganese.ppb, Censored, ci = TRUE))
```

```
#Results of Distribution Parameter Estimation
#Based on Type I Censored Data
#-----
#
#Assumed Distribution:      Gamma
#
#Censoring Side:           left
#
#Censoring Level(s):       2 5
#
#Estimated Parameter(s):   shape =  0.6370043
#                           scale = 30.8707533
#
#Estimation Method:        MLE
#
#Data:                     Manganese.ppb
#
#Censoring Variable:       Censored
#
#Sample Size:              25
#
#Percent Censored:         24%
#
#Confidence Interval for:   mean
#
#Confidence Interval Method: Profile Likelihood
#
#Confidence Interval Type:  two-sided
#
#Confidence Level:         95%
```

```
#
#Confidence Interval:          LCL = 12.25151
#                              UCL = 34.35332

#-----

# Compare the confidence interval for the mean
# based on assuming a lognormal distribution versus
# assuming a gamma distribution.

with(EPA.09.Ex.15.1.manganese.df,
     elnormAltCensored(Manganese.ppb, Censored,
                       ci = TRUE))$interval$limits
#      LCL      UCL
#12.37629 69.87694

with(EPA.09.Ex.15.1.manganese.df,
     egammaCensored(Manganese.ppb, Censored,
                    ci = TRUE))$interval$limits
#      LCL      UCL
#12.25151 34.35332
```

egeom

*Estimate Probability Parameter of a Geometric Distribution***Description**

Estimate the probability parameter of a [geometric distribution](#).

Usage

```
egeom(x, method = "mle/mme")
```

Arguments

x	vector of non-negative integers indicating the number of trials that took place <i>before</i> the first “success” occurred. (The total number of trials that took place is $x+1$). Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed. If $\text{length}(x)=n$ and n is greater than 1, it is assumed that x represents observations from n separate geometric experiments that all had the same probability of success (prob).
method	character string specifying the method of estimation. Possible values are “mle/mme” (maximum likelihood and method of moments; the default) and “mvue” (minimum variance unbiased). You cannot use <code>method=“mvue”</code> if $\text{length}(x)=1$. See the DETAILS section for more information on these estimation methods.

Details

If x contains any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

Let $\underline{x} = (x_1, x_2, \dots, x_n)$ be a vector of n independent observations from a [geometric distribution](#) with parameter $\text{prob}=p$.

It can be shown (e.g., Forbes et al., 2011) that if X is defined as:

$$X = \sum_{i=1}^n x_i$$

then X is an observation from a [negative binomial distribution](#) with parameters `prob=p` and `size=n`.

Estimation

The maximum likelihood and method of moments estimator (mle/mme) of p is given by:

$$\hat{p}_{mle} = \frac{n}{X + n}$$

and the minimum variance unbiased estimator (mvue) of p is given by:

$$\hat{p}_{mvue} = \frac{n - 1}{X + n - 1}$$

(Forbes et al., 2011). Note that the mvue of p is not defined for $n = 1$.

Value

a list of class "estimate" containing the estimated parameters and other information. See [estimate.object](#) for details.

Note

The [geometric distribution](#) with parameter `prob=p` is a special case of the [negative binomial distribution](#) with parameters `size=1` and `prob=p`.

The negative binomial distribution has its roots in a gambling game where participants would bet on the number of tosses of a coin necessary to achieve a fixed number of heads. The negative binomial distribution has been applied in a wide variety of fields, including accident statistics, birth-and-death processes, and modeling spatial distributions of biological organisms.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.
- Johnson, N. L., S. Kotz, and A. Kemp. (1992). *Univariate Discrete Distributions*. Second Edition. John Wiley and Sons, New York, Chapter 5.

See Also

[Geometric](#), [enbinom](#), [NegBinomial](#).

Examples

```
# Generate an observation from a geometric distribution with parameter
# prob=0.2, then estimate the parameter prob.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
```

```

dat <- rgeom(1, prob = 0.2)
dat
#[1] 4

egeom(dat)
#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Geometric
#
#Estimated Parameter(s):       prob = 0.2
#
#Estimation Method:            mle/mme
#
#Data:                          dat
#
#Sample Size:                   1
#-----

# Generate 3 observations from a geometric distribution with parameter
# prob=0.2, then estimate the parameter prob with the mvue.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(200)
dat <- rgeom(3, prob = 0.2)
dat
#[1] 0 1 2

egeom(dat, method = "mvue")
#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Geometric
#
#Estimated Parameter(s):       prob = 0.4
#
#Estimation Method:            mvue
#
#Data:                          dat
#
#Sample Size:                   3
#-----

# Clean up
#-----
rm(dat)

```

Description

Estimate the location, scale and shape parameters of a [generalized extreme value distribution](#), and optionally construct a confidence interval for one of the parameters.

Usage

```
egevd(x, method = "mle", pwme.method = "unbiased", tsoe.method = "med",
      plot.pos.cons = c(a = 0.35, b = 0), ci = FALSE, ci.parameter = "location",
      ci.type = "two-sided", ci.method = "normal.approx", information = "observed",
      conf.level = 0.95)
```

Arguments

<code>x</code>	numeric vector of observations.
<code>method</code>	character string specifying the method of estimation. Possible values are "mle" (maximum likelihood; the default), "pwme" (probability-weighted moments), and "tsoe" (two-stage order-statistics estimator of Castillo and Hadi (1994)). See the DETAILS section for more information on these estimation methods.
<code>pwme.method</code>	character string specifying what method to use to compute the probability-weighted moments when <code>method="pwme"</code> . The possible values are "unbiased" (method based on the U-statistic; the default), or "plotting.position" (method based on the plotting position formula). See the DETAILS section in this help file and the help file for pwMoment for more information. This argument is ignored if <code>method</code> is not equal to "pwme".
<code>tsoe.method</code>	character string specifying the robust function to apply in the second stage of the two-stage order-statistics estimator when <code>method="tsoe"</code> . Possible values are "med" (median; the default), and "lms" (least median of squares). See the DETAILS section for more information on these estimation methods. This argument is ignored if <code>method</code> is not equal to "tsoe".
<code>plot.pos.cons</code>	numeric vector of length 2 specifying the constants used in the formula for the plotting positions when <code>method="pwme"</code> and <code>pwme.method="plotting.position"</code> . The default value is <code>plot.pos.cons=c(a=0.35, b=0)</code> . If this vector has a names attribute with the value <code>c("a", "b")</code> or <code>c("b", "a")</code> , then the elements will be matched by name in the formula for computing the plotting positions. Otherwise, the first element is mapped to the name "a" and the second element to the name "b". See the DETAILS section in this help file and the help file for pwMoment for more information. This argument is used only if <code>method="tsoe"</code> , or if both <code>method="pwme"</code> and <code>pwme.method="plotting.position"</code> .
<code>ci</code>	logical scalar indicating whether to compute a confidence interval for the location, scale, or shape parameter. The default value is FALSE.
<code>ci.parameter</code>	character string indicating the parameter for which the confidence interval is desired. The possible values are "location" (the default), "scale", or "shape". This argument is ignored if <code>ci=FALSE</code> .
<code>ci.type</code>	character string indicating what kind of confidence interval to compute. The possible values are "two-sided" (the default), "lower", and "upper". This argument is ignored if <code>ci=FALSE</code> .
<code>ci.method</code>	character string indicating what method to use to construct the confidence interval for the location or scale parameter. Currently, the only possible value is "normal.approx" (the default). See the DETAILS section for more information. This argument is ignored if <code>ci=FALSE</code> .
<code>information</code>	character string indicating which kind of Fisher information to use when computing the variance-covariance matrix of the maximum likelihood estimators. The possible values are "observed" (the default) and "expected". See the DETAILS section for more information. This argument is used only when <code>method="mle"</code> and <code>ci=TRUE</code> .

`conf.level` a scalar between 0 and 1 indicating the confidence level of the confidence interval. The default value is `conf.level=0.95`. This argument is ignored if `ci=FALSE`.

Details

If `x` contains any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

Let $\underline{x} = (x_1, x_2, \dots, x_n)$ be a vector of n observations from a [generalized extreme value distribution](#) with parameters `location`= η , `scale`= θ , and `shape`= κ .

Estimation

Maximum Likelihood Estimation (`method="mle"`)

The log likelihood function is given by:

$$L(\eta, \theta, \kappa) = -n \log(\theta) - (1 - \kappa) \sum_{i=1}^n y_i - \sum_{i=1}^n e^{y_i}$$

where

$$y_i = -\frac{1}{\kappa} \log\left[\frac{1 - \kappa(x_i - \eta)}{\theta}\right]$$

(see, for example, Jenkinson, 1969; Prescott and Walden, 1980; Prescott and Walden, 1983; Hosking, 1985; MacLeod, 1989). The maximum likelihood estimators (MLE's) of η , θ , and κ are those values that maximize the likelihood function, subject to the following constraints:

$$\theta > 0$$

$$\kappa \leq 1$$

$$x_i < \eta + \frac{\theta}{\kappa} \text{ if } \kappa > 0$$

$$x_i > \eta + \frac{\theta}{\kappa} \text{ if } \kappa < 0$$

Although in theory the value of κ may lie anywhere in the interval $(-\infty, \infty)$ (see [GEVD](#)), the constraint $\kappa \leq 1$ is imposed because when $\kappa > 1$ the likelihood can be made infinite and thus the MLE does not exist (Castillo and Hadi, 1994). Hence, **this method of estimation is not valid when the true value of κ is larger than 1**. Hosking (1985) and Hosking et al. (1985) note that in practice the value of κ tends to lie in the interval $-1/2 < \kappa < 1/2$.

The value of $-L$ is minimized using the R function `nlminb`. Prescott and Walden (1983) give formulas for the gradient and Hessian. Only the gradient is supplied in the call to `nlminb`. The values of the PWME (see below) are used as the starting values. If the starting value of κ is less than 0.001 in absolute value, it is reset to `sign(k) * 0.001`, as suggested by Hosking (1985).

Probability-Weighted Moments Estimation (`method="pwme"`)

The idea of probability-weighted moments was introduced by Greenwood et al. (1979). Landwehr et al. (1979) derived probability-weighted moment estimators (PWME's) for the parameters of the [Type I \(Gumbel\) extreme value distribution](#). Hosking et al. (1985) extended these results to the generalized extreme value distribution. See the [abstract for Hosking et al. \(1985\)](#) for details on how these estimators are computed.

Two-Stage Order Statistics Estimation (`method="tsoe"`)

The two-stage order statistics estimator (TSOE) was introduced by Castillo and Hadi (1994) as an alternative to the MLE and PWME. Unlike the MLE and PWME, the TSOE of κ exists for all combinations of sample values and possible values of κ . See the [abstract for Castillo and Hadi \(1994\)](#) for details on how these estimators are computed. In the second stage, Castillo and Hadi

(1984) suggest using either the median or the least median of squares as the robust function. The function `egevd` allows three options for the robust function: median (`tsoe.method="med"`; see the R help file for [median](#)), least median of squares (`tsoe.method="lms"`; see the help file for [lmsreg](#) in the package **MASS**), and least trimmed squares (`tsoe.method="lts"`; see the help file for [ltsreg](#) in the package **MASS**).

Confidence Intervals

When `ci=TRUE`, an approximate $(1 - \alpha)100\%$ confidence intervals for η can be constructed assuming the distribution of the estimator of η is approximately normally distributed. A two-sided confidence interval is constructed as:

$$[\hat{\eta} - t(n - 1, 1 - \alpha/2)\hat{\sigma}_{\hat{\eta}}, \hat{\eta} + t(n - 1, 1 - \alpha/2)\hat{\sigma}_{\hat{\eta}}]$$

where $t(\nu, p)$ is the p 'th quantile of Student's t -distribution with ν degrees of freedom, and the quantity

$$\hat{\sigma}_{\hat{\eta}}$$

denotes the estimated asymptotic standard deviation of the estimator of η .

Similarly, a two-sided confidence interval for θ is constructed as:

$$[\hat{\theta} - t(n - 1, 1 - \alpha/2)\hat{\sigma}_{\hat{\theta}}, \hat{\theta} + t(n - 1, 1 - \alpha/2)\hat{\sigma}_{\hat{\theta}}]$$

and a two-sided confidence interval for κ is constructed as:

$$[\hat{\kappa} - t(n - 1, 1 - \alpha/2)\hat{\sigma}_{\hat{\kappa}}, \hat{\kappa} + t(n - 1, 1 - \alpha/2)\hat{\sigma}_{\hat{\kappa}}]$$

One-sided confidence intervals for η , θ , and κ are computed in a similar fashion.

Maximum Likelihood Estimator (`method="mle"`)

Prescott and Walden (1980) derive the elements of the Fisher information matrix (the expected information). The inverse of this matrix, evaluated at the values of the MLE, is the estimated asymptotic variance-covariance matrix of the MLE. This method is used to estimate the standard deviations of the estimated distribution parameters when `information="expected"`. The necessary regularity conditions hold for $\kappa < 1/2$. Thus, **this method of constructing confidence intervals is not valid when the true value of κ is greater than or equal to 1/2.**

Prescott and Walden (1983) derive expressions for the observed information matrix (i.e., the Hessian). This matrix is used to compute the estimated asymptotic variance-covariance matrix of the MLE when `information="observed"`.

In computer simulations, Prescott and Walden (1983) found that the variance-covariance matrix based on the observed information gave slightly more accurate estimates of the variance of MLE of κ compared to the estimated variance based on the expected information.

Probability-Weighted Moments Estimator (`method="pwme"`)

Hosking et al. (1985) show that these estimators are asymptotically multivariate normal and derive the asymptotic variance-covariance matrix. See the [abstract for Hosking et al. \(1985\)](#) for details on how this matrix is computed.

Two-Stage Order Statistics Estimator (`method="tsoe"`)

Currently there is no built-in method in **EnvStats** for computing confidence intervals when `method="tsoe"`. [Castillo and Hadi \(1994\)](#) suggest using the bootstrap or jackknife method.

Value

a list of class "estimate" containing the estimated parameters and other information. See [estimate.object](#) for details.

Note

Two-parameter [extreme value distributions](#) (EVD) have been applied extensively since the 1930's to several fields of study, including the distributions of hydrological and meteorological variables, human lifetimes, and strength of materials. The three-parameter [generalized extreme value distribution](#) (GEVD) was introduced by Jenkinson (1955) to model annual maximum and minimum values of meteorological events. Since then, it has been used extensively in the hydrological and meteorological fields.

The three families of EVDs are all special kinds of GEVDs. When the shape parameter $\kappa = 0$, the GEVD reduces to the Type I extreme value (Gumbel) distribution. (The function [zTestGevdShape](#) allows you to test the null hypothesis $H_0 : \kappa = 0$.) When $\kappa > 0$, the GEVD is the same as the Type II extreme value distribution, and when $\kappa < 0$ it is the same as the Type III extreme value distribution.

Hosking et al. (1985) compare the asymptotic and small-sample statistical properties of the PWME with the MLE and Jenkinson's (1969) method of sextiles. Castillo and Hadi (1994) compare the small-sample statistical properties of the MLE, PWME, and TSOE. Hosking and Wallis (1995) compare the small-sample properties of unbiased L -moment estimators vs. plotting-position L -moment estimators. (PWMEs can be written as linear combinations of L -moments and thus have equivalent statistical properties.) Hosking and Wallis (1995) conclude that unbiased estimators should be used for almost all applications.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Castillo, E., and A. Hadi. (1994). Parameter and Quantile Estimation for the Generalized Extreme-Value Distribution. *Environmetrics* **5**, 417–432.
- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). Statistical Distributions. Fourth Edition. John Wiley and Sons, Hoboken, NJ.
- Greenwood, J.A., J.M. Landwehr, N.C. Matalas, and J.R. Wallis. (1979). Probability Weighted Moments: Definition and Relation to Parameters of Several Distributions Expressible in Inverse Form. *Water Resources Research* **15**(5), 1049–1054.
- Hosking, J.R.M. (1984). Testing Whether the Shape Parameter is Zero in the Generalized Extreme-Value Distribution. *Biometrika* **71**(2), 367–374.
- Hosking, J.R.M. (1985). Algorithm AS 215: Maximum-Likelihood Estimation of the Parameters of the Generalized Extreme-Value Distribution. *Applied Statistics* **34**(3), 301–310.
- Hosking, J.R.M., J.R. Wallis, and E.F. Wood. (1985). Estimation of the Generalized Extreme-Value Distribution by the Method of Probability-Weighted Moments. *Technometrics* **27**(3), 251–261.
- Jenkinson, A.F. (1969). Statistics of Extremes. *Technical Note 98*, World Meteorological Office, Geneva.
- Johnson, N. L., S. Kotz, and N. Balakrishnan. (1995). *Continuous Univariate Distributions, Volume 2*. Second Edition. John Wiley and Sons, New York.
- Landwehr, J.M., N.C. Matalas, and J.R. Wallis. (1979). Probability Weighted Moments Compared With Some Traditional Techniques in Estimating Gumbel Parameters and Quantiles. *Water Resources Research* **15**(5), 1055–1064.
- Macleod, A.J. (1989). Remark AS R76: A Remark on Algorithm AS 215: Maximum Likelihood Estimation of the Parameters of the Generalized Extreme-Value Distribution. *Applied Statistics* **38**(1), 198–199.

Prescott, P., and A.T. Walden. (1980). Maximum Likelihood Estimation of the Parameters of the Generalized Extreme-Value Distribution. *Biometrika* **67**(3), 723–724.

Prescott, P., and A.T. Walden. (1983). Maximum Likelihood Estimation of the Three-Parameter Generalized Extreme-Value Distribution from Censored Samples. *Journal of Statistical Computing and Simulation* **16**, 241–250.

See Also

[Generalized Extreme Value Distribution](#), [zTestGevdShape](#), [Extreme Value Distribution](#), [eevd](#).

Examples

```
# Generate 20 observations from a generalized extreme value distribution
# with parameters location=2, scale=1, and shape=0.2, then compute the
# MLE and construct a 90% confidence interval for the location parameter.
# (Note: the call to set.seed simply allows you to reproduce this example.)
```

```
set.seed(498)
dat <- rgevd(20, location = 2, scale = 1, shape = 0.2)
egevd(dat, ci = TRUE, conf.level = 0.9)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Generalized Extreme Value
#
#Estimated Parameter(s):      location = 1.6144631
#                               scale   = 0.9867007
#                               shape    = 0.2632493
#
#Estimation Method:           mle
#
#Data:                         dat
#
#Sample Size:                  20
#
#Confidence Interval for:      location
#
#Confidence Interval Method:    Normal Approximation
#                               (t Distribution) based on
#                               observed information
#
#Confidence Interval Type:      two-sided
#
#Confidence Level:              90%
#
#Confidence Interval:          LCL = 1.225249
#                               UCL = 2.003677
```

```
#-----
```

```
# Compare the values of the different types of estimators:
```

```
egevd(dat, method = "mle")$parameters
# location      scale      shape
#1.6144631 0.9867007 0.2632493
```

```
egevd(dat, method = "pwme")$parameters
# location      scale      shape
#1.5785779 1.0187880 0.2257948

egevd(dat, method = "pwme", pwme.method = "plotting.position")$parameters
# location      scale      shape
#1.5509183 0.9804992 0.1657040

egevd(dat, method = "tsoe")$parameters
# location      scale      shape
#1.5372694 1.0876041 0.2927272

egevd(dat, method = "tsoe", tsoe.method = "lms")$parameters
#location      scale      shape
#1.519469 1.081149 0.284863

egevd(dat, method = "tsoe", tsoe.method = "lts")$parameters
# location      scale      shape
#1.4840198 1.0679549 0.2691914

#-----

# Clean up
#-----
rm(dat)
```

ehyper	<i>Estimate Parameter of a Hypergeometric Distribution</i>
--------	--

Description

Estimate m , the number of white balls in the urn, or $m + n$, the total number of balls in the urn, for a [hypergeometric distribution](#).

Usage

```
ehyper(x, m = NULL, total = NULL, k, method = "mle")
```

Arguments

x	non-negative integer indicating the number of white balls out of a sample of size k drawn without replacement from the urn. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed.
m	non-negative integer indicating the number of white balls in the urn. You must supply m or total, but not both. Missing values (NAs) are not allowed.
total	positive integer indicating the total number of balls in the urn (i.e., m+n). You must supply m or total, but not both. Missing values (NAs) are not allowed.
k	positive integer indicating the number of balls drawn without replacement from the urn. Missing values (NAs) are not allowed.

method character string specifying the method of estimation. Possible values are "mle" (maximum likelihood; the default) and "mvue" (minimum variance unbiased). The mvue method is only available when you are estimating m (i.e., when you supply the argument `total`). See the DETAILS section for more information on these estimation methods.

Details

Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed.

Let x be an observation from a [hypergeometric distribution](#) with parameters $m=M$, $n=N$, and $k=K$. In R nomenclature, x represents the number of white balls drawn out of a sample of K balls drawn *without* replacement from an urn containing M white balls and N black balls. The total number of balls in the urn is thus $M + N$. Denote the total number of balls by $T = M + N$.

Estimation

Estimating M , Given T and K are known

When T and K are known, the maximum likelihood estimator (mle) of M is given by (Forbes et al., 2011):

$$\hat{M}_{mle} = \text{floor}[(T + 1)x/K] \quad (1)$$

where $\text{floor}()$ represents the [floor](#) function. That is, $\text{floor}(y)$ is the largest integer less than or equal to y .

If the quantity $\text{floor}[(T + 1)x/K]$ is an integer, then the mle of M is also given by (Johnson et al., 1992, p.263):

$$\hat{M}_{mle} = [(T + 1)x/K] - 1 \quad (2)$$

which is what the function `ehyper` uses for this case.

The minimum variance unbiased estimator (mvue) of M is given by (Forbes et al., 2011):

$$\hat{M}_{mvue} = (Tx/K) \quad (3)$$

Estimating T , given M and K are known

When M and K are known, the maximum likelihood estimator (mle) of T is given by (Forbes et al., 2011):

$$\hat{T}_{mle} = \text{floor}(KM/x) \quad (4)$$

Value

a list of class "estimate" containing the estimated parameters and other information. See [estimate.object](#) for details.

Note

The [hypergeometric distribution](#) can be described by an urn model with M white balls and N black balls. If K balls are drawn *with* replacement, then the number of white balls in the sample of size K follows a [binomial distribution](#) with parameters `size=K` and `prob=M/(M + N)`. If K balls are drawn *without* replacement, then the number of white balls in the sample of size K follows a [hypergeometric distribution](#) with parameters $m=M$, $n=N$, and $k=K$.

The name "hypergeometric" comes from the fact that the probabilities associated with this distribution can be written as successive terms in the expansion of a function of a Gaussian hypergeometric series.

The hypergeometric distribution is applied in a variety of fields, including quality control and estimation of animal population size. It is also the distribution used to compute probabilities for [Fishers's exact test](#) for a 2x2 contingency table.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.

Johnson, N. L., S. Kotz, and A. Kemp. (1992). *Univariate Discrete Distributions*. Second Edition. John Wiley and Sons, New York, Chapter 6.

See Also

[Hypergeometric](#).

Examples

```
# Generate an observation from a hypergeometric distribution with
# parameters m=10, n=30, and k=5, then estimate the parameter m.
# Note: the call to set.seed simply allows you to reproduce this example.
# Also, the only parameter actually estimated is m; once m is estimated,
# n is computed by subtracting the estimated value of m (8 in this example)
# from the given value of m+n (40 in this example). The parameters
# n and k are shown in the output in order to provide information on
# all of the parameters associated with the hypergeometric distribution.

set.seed(250)
dat <- rhyper(nn = 1, m = 10, n = 30, k = 5)
dat
#[1] 1

ehyper(dat, total = 40, k = 5)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Hypergeometric
#
#Estimated Parameter(s):      m = 8
#                               n = 32
#                               k = 5
#
#Estimation Method:           mle for 'm'
#
#Data:                         dat
#
#Sample Size:                  1
#-----

# Use the same data as in the previous example, but estimate m+n instead.
# Note: The only parameter estimated is m+n. Once this is estimated,
```

```

# n is computed by subtracting the given value of m (10 in this case)
# from the estimated value of m+n (50 in this example).

ehyper(dat, m = 10, k = 5)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Hypergeometric
#
#Estimated Parameter(s):      m = 10
#                             n = 40
#                             k = 5
#
#Estimation Method:           mle for 'm+n'
#
#Data:                        dat
#
#Sample Size:                  1

#-----

# Clean up
#-----
rm(dat)

```

elnorm

Estimate Parameters of a Lognormal Distribution (Log-Scale)

Description

Estimate the mean and standard deviation parameters of the logarithm of a [lognormal distribution](#), and optionally construct a confidence interval for the mean.

Usage

```
elnorm(x, method = "mvue", ci = FALSE, ci.type = "two-sided",
       ci.method = "exact", conf.level = 0.95)
```

Arguments

x	numeric vector of observations.
method	character string specifying the method of estimation. Possible values are "mvue" (minimum variance unbiased; the default), and "mle/mme" (maximum likelihood/method of moments). See the DETAILS section for more information on these estimation methods.
ci	logical scalar indicating whether to compute a confidence interval for the mean. The default value is FALSE.
ci.type	character string indicating what kind of confidence interval to compute. The possible values are "two-sided" (the default), "lower", and "upper". This argument is ignored if ci=FALSE.

<code>ci.method</code>	character string indicating what method to use to construct the confidence interval for the mean or variance. The only possible value is "exact" (the default). See the DETAILS section for more information. This argument is ignored if <code>ci=FALSE</code> .
<code>conf.level</code>	a scalar between 0 and 1 indicating the confidence level of the confidence interval. The default value is <code>conf.level=0.95</code> . This argument is ignored if <code>ci=FALSE</code> .

Details

If x contains any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

Let X denote a random variable with a [lognormal distribution](#) with parameters $\text{meanlog}=\mu$ and $\text{sdlog}=\sigma$. Then $Y = \log(X)$ has a [normal \(Gaussian\) distribution](#) with parameters $\text{mean}=\mu$ and $\text{sd}=\sigma$. Thus, the function `elnorm` simply calls the function `enorm` using the log-transformed values of x .

Value

a list of class "estimate" containing the estimated parameters and other information. See [estimate.object](#) for details.

Note

The normal and lognormal distribution are probably the two most frequently used distributions to model environmental data. In order to make any kind of probability statement about a normally-distributed population (of chemical concentrations for example), you have to first estimate the mean and standard deviation (the population parameters) of the distribution. Once you estimate these parameters, it is often useful to characterize the uncertainty in the estimate of the mean or variance. This is done with confidence intervals.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Aitchison, J., and J.A.C. Brown (1957). *The Lognormal Distribution (with special references to its uses in economics)*. Cambridge University Press, London, Chapter 5.
- Crow, E.L., and K. Shimizu. (1988). *Lognormal Distributions: Theory and Applications*. Marcel Dekker, New York, Chapter 2.
- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.
- Gilbert, R.O. (1987). *Statistical Methods for Environmental Pollution Monitoring*. Van Nostrand Reinhold, New York, NY.
- Johnson, N. L., S. Kotz, and N. Balakrishnan. (1994). *Continuous Univariate Distributions, Volume 1*. Second Edition. John Wiley and Sons, New York.
- Limpert, E., W.A. Stahel, and M. Abbt. (2001). Log-Normal Distributions Across the Sciences: Keys and Clues. *BioScience* **51**, 341–352.
- Millard, S.P., and N.K. Neerchal. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, FL.

Ott, W.R. (1995). *Environmental Statistics and Data Analysis*. Lewis Publishers, Boca Raton, FL.

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.

See Also

[Lognormal](#), [LognormalAlt](#), [Normal](#).

Examples

```
# Using the Reference area TcCB data in the data frame EPA.94b.tccb.df,
# estimate the mean and standard deviation of the log-transformed distribution,
# and construct a 95% confidence interval for the mean.

with(EPA.94b.tccb.df, elnorm(TcCB[Area == "Reference"], ci = TRUE))

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:      Lognormal
#
#Estimated Parameter(s):   meanlog = -0.6195712
#                           sdlog   =  0.4679530
#
#Estimation Method:        mvue
#
#Data:                     TcCB[Area == "Reference"]
#
#Sample Size:              47
#
#Confidence Interval for:   mean
#
#Confidence Interval Method: Exact
#
#Confidence Interval Type:  two-sided
#
#Confidence Level:         95%
#
#Confidence Interval:      LCL = -0.7569673
#                           UCL = -0.4821751
```

elnorm3	<i>Estimate Parameters of a Three-Parameter Lognormal Distribution (Log-Scale)</i>
---------	--

Description

Estimate the mean, standard deviation, and threshold parameters for a [three-parameter lognormal distribution](#), and optionally construct a confidence interval for the threshold or the median of the distribution.

Usage

```
elnorm3(x, method = "lmle", ci = FALSE, ci.parameter = "threshold",
        ci.method = "avar", ci.type = "two-sided", conf.level = 0.95,
        threshold.lb.sd = 100)
```

Arguments

<code>x</code>	numeric vector of observations.
<code>method</code>	character string specifying the method of estimation. Possible values are "lmle" (local maximum likelihood; the default), "mme" (method of moments), "mmue" (method of moments using an unbiased estimate of variance), "mmme" (modified method of moments due to Cohen and Whitten (1980)), "zero.skew" (zero-skewness estimator due to Griffiths (1980)), and "royston.skew" (estimator based on Royston's (1992b) index of skewness). See the DETAILS section for more information on these estimation methods.
<code>ci</code>	logical scalar indicating whether to compute a confidence interval for either the threshold or median of the distribution. The default value is FALSE.
<code>ci.parameter</code>	character string indicating the parameter for which the confidence interval is desired. The possible values are "threshold" (the default) and "median". This argument is ignored if <code>ci=FALSE</code> .
<code>ci.method</code>	character string indicating the method to use to construct the confidence interval. The possible values are "avar" (asymptotic variance; the default), "likelihood.profile", and "skewness" (method suggested by Royston (1992b) for <code>method="zero.skew"</code>). This argument is ignored if <code>ci=FALSE</code> .
<code>ci.type</code>	character string indicating what kind of confidence interval to compute. The possible values are "two-sided" (the default), "lower", and "upper". This argument is ignored if <code>ci=FALSE</code> .
<code>conf.level</code>	a scalar between 0 and 1 indicating the confidence level of the confidence interval. The default value is <code>conf.level=0.95</code> . This argument is ignored if <code>ci=FALSE</code> .
<code>threshold.lb.sd</code>	a positive numeric scalar specifying the range over which to look for the local maximum likelihood (<code>method="lmle"</code>) or zero-skewness (<code>method="zero.skewness"</code>) estimator of threshold. The range is set to $[\text{mean}(x) - \text{threshold.lb.sd} * \text{sd}(x), \text{min}(x)]$. If you receive a warning message that <code>elnorm3</code> is unable to find an acceptable estimate of threshold in this range, it may be because of convergence problems specific to the data in <code>x</code> . When this occurs, try changing the value of <code>threshold.lb.sd</code> . This same range is used in constructing confidence intervals for the threshold parameter. The default value is <code>threshold.lb.sd=100</code> . This argument is relevant only if <code>method="lmle"</code> , <code>method="zero.skew"</code> , <code>ci.method="likelihood.profile"</code> and/or <code>ci.method="skewness"</code> .

Details

If `x` contains any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

Let X denote a random variable from a [three-parameter lognormal distribution](#) with parameters $\text{meanlog}=\mu$, $\text{sdlog}=\sigma$, and $\text{threshold}=\gamma$. Let \underline{x} denote a vector of n observations from this distribution. Furthermore, let $x_{(i)}$ denote the i 'th order statistic in the sample, so that $x_{(1)}$ denotes the

smallest value and $x_{(n)}$ denote the largest value in \underline{x} . Finally, denote the sample mean and variance by:

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i \quad (1)$$

$$s^2 = \frac{1}{n-1} \sum_{i=1}^n (x_i - \bar{x})^2 \quad (2)$$

Note that the sample variance is the unbiased version. Denote the method of moments estimator of variance by:

$$s_m^2 = \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2 \quad (3)$$

Estimation

Local Maximum Likelihood Estimation (method="lmle")

Hill (1963) showed that the likelihood function approaches infinity as γ approaches $x_{(1)}$, so that the global maximum likelihood estimators of (μ, σ, γ) are $(-\infty, \infty, x_{(1)})$, which are inadmissible, since γ must be smaller than $x_{(1)}$. Cohen (1951) suggested using local maximum likelihood estimators (lmle's), derived by equating partial derivatives of the log-likelihood function to zero. These estimators were studied by Harter and Moore (1966), Calitz (1973), Cohen and Whitten (1980), and Griffiths (1980), and appear to possess most of the desirable properties ordinarily associated with maximum likelihood estimators.

Cohen (1951) showed that the lmle of γ is given by the solution to the following equation:

$$\left[\sum_{i=1}^n w_i \right] \left\{ \sum_{i=1}^n y_i - \sum_{i=1}^n y_i^2 + \frac{1}{n} \left[\sum_{i=1}^n y_i \right]^2 \right\} - n \sum_{i=1}^n \frac{y_i}{w_i} = 0 \quad (4)$$

where

$$w_i = x_i - \hat{\gamma} \quad (5)$$

$$y_i = \log(x_i - \hat{\gamma}) = \log(w_i) \quad (6)$$

and that the lmle's of μ and σ then follow as:

$$\hat{\mu} = \frac{1}{n} \sum_{i=1}^n y_i = \bar{y} \quad (7)$$

$$\hat{\sigma}^2 = \frac{1}{n} \sum_{i=1}^n (y_i - \bar{y})^2 \quad (8)$$

Unfortunately, while equation (4) simplifies the task of computing the lmle's, for certain data sets there still may be convergence problems (Calitz, 1973), and occasionally multiple roots of equation (4) may exist. When multiple roots to equation (4) exist, Cohen and Whitten (1980) recommend using the one that results in closest agreement between the mle of μ (equation (7)) and the sample mean (equation (1)).

On the other hand, Griffiths (1980) showed that for a given value of the threshold parameter γ , the maximized value of the log-likelihood (the "profile likelihood" for γ) is given by:

$$\log[L(\gamma)] = \frac{-n}{2} [1 + \log(2\pi) + 2\hat{\mu} + \log(\hat{\sigma}^2)] \quad (9)$$

where the estimates of μ and σ are defined in equations (7) and (8), so the lmle of γ reduces to an iterative search over the values of γ . Griffiths (1980) noted that the distribution of the lmle of

γ is far from normal and that $\log[L(\gamma)]$ is not quadratic near the lmle of γ . He suggested a better parameterization based on

$$\eta = -\log(x_{(1)} - \gamma) \quad (10)$$

Thus, once the lmle of η is found using equations (9) and (10), the lmle of γ is given by:

$$\hat{\gamma} = x_{(1)} - \exp(-\hat{\eta}) \quad (11)$$

When method="lmle", the function `elnorm3` uses the function `nlminb` to search for the minimum of $-2\log[L(\eta)]$, using the modified method of moments estimator (method="mme"; see below) as the starting value for γ . Equation (11) is then used to solve for the lmle of γ , and equation (4) is used to "fine tune" the estimated value of γ . The lmle's of μ and σ are then computed using equations (6)-(8).

Method of Moments Estimation (method="mme")

Denote the r 'th sample central moment by:

$$m_r = \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^r \quad (12)$$

and note that

$$s_m^2 = m_2 \quad (13)$$

Equating the sample first moment (the sample mean) with its population value (the population mean), and equating the second and third sample central moments with their population values yields (Johnson et al., 1994, p.228):

$$\bar{x} = \gamma + \beta\sqrt{\omega} \quad (14)$$

$$m_2 = s_m^2 = \beta^2\omega(\omega - 1) \quad (15)$$

$$m_3 = \beta^3\omega^{3/2}(\omega - 1)^2(\omega + 2) \quad (16)$$

where

$$\beta = \exp(\mu) \quad (17)$$

$$\omega = \exp(\sigma^2) \quad (18)$$

Combining equations (15) and (16) yields:

$$b_1 = \frac{m_3}{m_2^{3/2}} = (\omega + 2)\sqrt{\omega - 1} \quad (19)$$

The quantity on the left-hand side of equation (19) is the usual estimator of skewness. Solving equation (19) for ω yields:

$$\hat{\omega} = (d + h)^{1/3} + (d - h)^{1/3} - 1 \quad (20)$$

where

$$d = 1 + \frac{b_1}{2} \quad (21)$$

$$h = \sqrt{d^2 - 1} \quad (22)$$

Using equation (18), the method of moments estimator of σ is then computed as:

$$\hat{\sigma}^2 = \log(\hat{\omega}) \quad (23)$$

Combining equations (15) and (17), the method of moments estimator of μ is computed as:

$$\hat{\mu} = \frac{1}{2} \log \left[\frac{s_m^2}{\hat{\omega}(\hat{\omega} - 1)} \right] \quad (24)$$

Finally, using equations (14), (17), and (18), the method of moments estimator of γ is computed as:

$$\bar{x} - \exp\left(\hat{\mu} + \frac{\hat{\sigma}^2}{2}\right) \quad (25)$$

There are two major problems with using method of moments estimators for the three-parameter lognormal distribution. First, they are subject to very large sampling error due to the use of second and third sample moments (Cohen, 1988, p.121; Johnson et al., 1994, p.228). Second, Heyde (1963) showed that the lognormal distribution is not uniquely determined by its moments.

Method of Moments Estimators Using an Unbiased Estimate of Variance (method="mmue")

This method of estimation is exactly the same as the method of moments (method="mme"), except that the unbiased estimator of variance (equation (3)) is used in place of the method of moments one (equation (4)). This modification is given in Cohen (1988, pp.119-120).

Modified Method of Moments Estimation (method="mmme")

This method of estimation is described by Cohen (1988, pp.125-132). It was introduced by Cohen and Whitten (1980; their MME-II with $r=1$) and was further investigated by Cohen et al. (1985). It is motivated by the fact that the first order statistic in the sample, $x_{(1)}$, contains more information about the threshold parameter γ than any other observation and often more information than all of the other observations combined (Cohen, 1988, p.125).

The first two sets of equations are the same as for the modified method of moments estimators (method="mmme"), i.e., equations (14) and (15) with the unbiased estimator of variance (equation (3)) used in place of the method of moments one (equation (4)). The third equation replaces equation (16) by equating a function of the first order statistic with its expected value:

$$\log(x_{(1)} - \gamma) = \mu + \sigma E[Z_{(1,n)}] \quad (26)$$

where $E[Z_{(i,n)}]$ denotes the expected value of the i 'th order statistic in a random sample of n observations from a standard normal distribution. (See the help file for [evNormOrdStats](#) for information on how $E[Z_{(i,n)}]$ is computed.) Using equations (17) and (18), equation (26) can be rewritten as:

$$x_{(1)} = \gamma + \beta \exp\{\sqrt{\log(\omega)} E[Z_{(i,n)}]\} \quad (27)$$

Combining equations (14), (15), (17), (18), and (27) yields the following equation for the estimate of ω :

$$\frac{s^2}{[\bar{x} - x_{(1)}]^2} = \frac{\hat{\omega}(\hat{\omega} - 1)}{[\sqrt{\hat{\omega}} - \exp\{\sqrt{\log(\omega)} E[Z_{(i,n)}]\}]^2} \quad (28)$$

After equation (28) is solved for $\hat{\omega}$, the estimate of σ is again computed using equation (23), and the estimate of μ is computed using equation (24), where the unbiased estimate of variance is used in place of the biased one (just as for method="mmue").

Zero-Skewness Estimation (method="zero.skew")

This method of estimation was introduced by Griffiths (1980), and elaborated upon by Royston (1992b). The idea is that if the threshold parameter γ were known, then the distribution of:

$$Y = \log(X - \gamma) \quad (29)$$

is normal, so the skew of Y is 0. Thus, the threshold parameter γ is estimated as that value that forces the sample skew (defined in equation (19)) of the observations defined in equation (6) to be 0. That is, the zero-skewness estimator of γ is the value that satisfies the following equation:

$$0 = \frac{\frac{1}{n} \sum_{i=1}^n (y_i - \bar{y})^3}{[\frac{1}{n} \sum_{i=1}^n (y_i - \bar{y})^2]^{3/2}} \quad (30)$$

where

$$y_i = \log(x_i - \hat{\gamma}) \quad (31)$$

Note that since the denominator in equation (30) is always positive (assuming there are at least two unique values in \underline{x}), only the numerator needs to be used to determine the value of $\hat{\gamma}$.

Once the value of $\hat{\gamma}$ has been determined, μ and σ are estimated using equations (7) and (8), except the unbiased estimator of variance is used in equation (8).

Royston (1992b) developed a modification of the Shapiro-Wilk goodness-of-fit test for normality based on transforming the data using equation (6) and the zero-skewness estimator of γ (see [gofTest](#)).

Estimators Based on Royston's Index of Skewness (method="royston.skew")

This method of estimation is discussed by Royston (1992b), and is similar to the zero-skewness method discussed above, except a different measure of skewness is used. Royston's (1992b) index of skewness is given by:

$$q = \frac{y_{(n)} - \tilde{y}}{\tilde{y} - y_{(1)}} \quad (32)$$

where $y_{(i)}$ denotes the i 'th order statistic of y and y is defined in equation (31) above, and \tilde{y} denotes the median of y . Royston (1992b) shows that the value of γ that yields a value of $q = 0$ is given by:

$$\hat{\gamma} = \frac{y_{(1)}y_{(n)} - \tilde{y}^2}{y_{(1)} + y_{(n)} - 2\tilde{y}} \quad (33)$$

Again, as for the zero-skewness method, once the value of $\hat{\gamma}$ has been determined, μ and σ are estimated using equations (7) and (8), except the unbiased estimator of variance is used in equation (8).

Royston (1992b) developed this estimator as a quick way to estimate γ .

Confidence Intervals

This section explains three different methods for constructing confidence intervals for the threshold parameter γ , or the median of the three-parameter lognormal distribution, which is given by:

$$\text{Med}[X] = \gamma + \exp(\mu) = \gamma + \beta \quad (34)$$

Normal Approximation Based on Asymptotic Variances and Covariances (ci.method="avar")

Formulas for asymptotic variances and covariances for the three-parameter lognormal distribution, based on the information matrix, are given in Cohen (1951), Cohen and Whitten (1980), Cohen et al., (1985), and Cohen (1988). The relevant quantities for γ and the median are:

$$\text{Var}(\hat{\gamma}) = \sigma_{\hat{\gamma}}^2 = \frac{\sigma^2}{n} \left(\frac{\beta^2}{\omega} \right) H \quad (35)$$

$$\text{Var}(\hat{\beta}) = \sigma_{\hat{\beta}}^2 = \frac{\sigma^2}{n} \beta^2 (1 + H) \quad (36)$$

$$Cov(\hat{\gamma}, \hat{\beta}) = \sigma_{\hat{\gamma}, \hat{\beta}} = \frac{-\sigma^3}{n} \left(\frac{\beta^2}{\sqrt{\omega}} \right) H \quad (37)$$

where

$$H = [\omega(1 + \sigma^2) - 2\sigma^2 - 1]^{-1} \quad (38)$$

A two-sided $(1 - \alpha)100\%$ confidence interval for γ is computed as:

$$\hat{\gamma} - t_{n-2, 1-\alpha/2} \hat{\sigma}_{\hat{\gamma}}, \hat{\gamma} + t_{n-2, 1-\alpha/2} \hat{\sigma}_{\hat{\gamma}} \quad (39)$$

where $t_{\nu, p}$ denotes the p 'th quantile of [Student's t-distribution](#) with n degrees of freedom, and the quantity $\hat{\sigma}_{\hat{\gamma}}$ is computed using equations (35) and (38) and substituting estimated values of β , ω , and σ . One-sided confidence intervals are computed in a similar manner.

A two-sided $(1 - \alpha)100\%$ confidence interval for the median (see equation (34) above) is computed as:

$$\hat{\gamma} + \hat{\beta} - t_{n-2, 1-\alpha/2} \hat{\sigma}_{\hat{\gamma} + \hat{\beta}}, \hat{\gamma} + \hat{\beta} + t_{n-2, 1-\alpha/2} \hat{\sigma}_{\hat{\gamma} + \hat{\beta}} \quad (40)$$

where

$$\hat{\sigma}_{\hat{\gamma} + \hat{\beta}}^2 = \hat{\sigma}_{\hat{\gamma}}^2 + \hat{\sigma}_{\hat{\beta}}^2 + \hat{\sigma}_{\hat{\gamma}, \hat{\beta}} \quad (41)$$

is computed using equations (35)-(38) and substituting estimated values of β , ω , and σ . One-sided confidence intervals are computed in a similar manner.

This method of constructing confidence intervals is analogous to using the Wald test (e.g., Silvey, 1975, pp.115-118) to test hypotheses on the parameters.

Because of the regularity problems associated with the global maximum likelihood estimators, it is questionable whether the asymptotic variances and covariances shown above apply to local maximum likelihood estimators. Simulation studies, however, have shown that these estimates of variance and covariance perform reasonably well (Harter and Moore, 1966; Cohen and Whitten, 1980).

Note that this method of constructing confidence intervals can be used with estimators other than the lmle's. Cohen and Whitten (1980) and Cohen et al. (1985) found that the asymptotic variances and covariances are reasonably close to corresponding simulated variances and covariances for the modified method of moments estimators (method="mmme").

Likelihood Profile (ci.method="likelihood.profile")

Griffiths (1980) suggested constructing confidence intervals for the threshold parameter γ based on the profile likelihood function given in equations (9) and (10). Royston (1992b) further elaborated upon this procedure. A two-sided $(1 - \alpha)100\%$ confidence interval for η is constructed as:

$$[\eta_{LCL}, \eta_{UCL}] \quad (42)$$

by finding the two values of η (one larger than the lmle of η and one smaller than the lmle of η) that satisfy:

$$\log[L(\eta)] = \log[L(\hat{\eta}_{lmle})] - \frac{1}{2} \chi_{1, \alpha/2}^2 \quad (43)$$

where $\chi_{\nu, p}^2$ denotes the p 'th quantile of the [chi-square distribution](#) with ν degrees of freedom. Once these values are found, the two-sided confidence for γ is computed as:

$$[\gamma_{LCL}, \gamma_{UCL}] \quad (44)$$

where

$$\gamma_{LCL} = x_{(1)} - \exp(-\eta_{LCL}) \quad (45)$$

$$\gamma_{UCL} = x_{(1)} - \exp(-\eta_{UCL}) \quad (46)$$

One-sided intervals are constructed in a similar manner.

This method of constructing confidence intervals is analogous to using the likelihood-ratio test (e.g., Silvey, 1975, pp.108-115) to test hypotheses on the parameters.

To construct a two-sided $(1 - \alpha)100\%$ confidence interval for the median (see equation (34)), Royston (1992b) suggested the following procedure:

1. Construct a confidence interval for γ using the likelihood profile procedure.
2. Construct a confidence interval for β as:

$$[\beta_{LCL}, \beta_{UCL}] = [\exp(\hat{\mu} - t_{n-2, 1-\alpha/2} \frac{\hat{\sigma}}{n}), \exp(\hat{\mu} + t_{n-2, 1-\alpha/2} \frac{\hat{\sigma}}{n})] \quad (47)$$

3. Construct the confidence interval for the median as:

$$[\gamma_{LCL} + \beta_{LCL}, \gamma_{UCL} + \beta_{UCL}] \quad (48)$$

Royston (1992b) actually suggested using the quantile from the standard normal distribution instead of Student's t-distribution in step 2 above. The function `elnorm3`, however, uses the Student's t quantile.

Note that this method of constructing confidence intervals can be used with estimators other than the `lmle`'s.

Royston's Confidence Interval Based on Significant Skewness (`ci.method="skewness"`)

Royston (1992b) suggested constructing confidence intervals for the threshold parameter γ based on the idea behind the zero-skewness estimator (`method="zero.skew"`). A two-sided $(1 - \alpha)100\%$ confidence interval for γ is constructed by finding the two values of γ that yield a p-value of $\alpha/2$ for the test of zero-skewness on the observations y defined in equation (6) (see `gofTest`). One-sided confidence intervals are constructed in a similar manner.

To construct $(1 - \alpha)100\%$ confidence intervals for the median (see equation (34)), the exact same procedure is used as for `ci.method="likelihood.profile"`, except that the confidence interval for γ is based on the zero-skewness method just described instead of the likelihood profile method.

Value

a list of class "estimate" containing the estimated parameters and other information. See `estimate.object` for details.

Note

The problem of estimating the parameters of a three-parameter lognormal distribution has been extensively discussed by Aitchison and Brown (1957, Chapter 6), Calitz (1973), Cohen (1951), Cohen (1988), Cohen and Whitten (1980), Cohen et al. (1985), Griffiths (1980), Harter and Moore (1966), Hill (1963), and Royston (1992b). Stedinger (1980) and Hoshi et al. (1984) discuss fitting the three-parameter lognormal distribution to hydrologic data.

The global maximum likelihood estimates are inadmissible. In the past, several researchers have found that the local maximum likelihood estimates (`lmle`'s) occasionally fail because of convergence problems, but they were not using the likelihood profile and reparameterization of Griffiths (1980). Cohen (1988) recommends the modified methods of moments estimators over `lmle`'s because they are easy to compute, they are unbiased with respect to μ and σ^2 (the mean and standard deviation on the log-scale), their variances are minimal or near minimal, and they do not suffer from regularity problems.

Because the distribution of the `lmle` of the threshold parameter γ is far from normal for moderate sample sizes (Griffiths, 1980), it is questionable whether confidence intervals for γ or the

median based on asymptotic variances and covariances will perform well. Cohen and Whitten (1980) and Cohen et al. (1985), however, found that the asymptotic variances and covariances are reasonably close to corresponding simulated variances and covariances for the modified method of moments estimators (method="mmme"). In a simulation study (5000 monte carlo trials), Royston (1992b) found that the coverage of confidence intervals for γ based on the likelihood profile (ci.method="likelihood.profile") was very close the nominal level (94.1% for a nominal level of 95%), although not symmetric. Royston (1992b) also found that the coverage of confidence intervals for γ based on the skewness method (ci.method="skewness") was also very close (95.4%) and symmetric.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Aitchison, J., and J.A.C. Brown (1957). *The Lognormal Distribution (with special references to its uses in economics)*. Cambridge University Press, London, Chapter 5.
- Calitz, F. (1973). Maximum Likelihood Estimation of the Parameters of the Three-Parameter Lognormal Distribution—a Reconsideration. *Australian Journal of Statistics* **15**(3), 185–190.
- Cohen, A.C. (1951). Estimating Parameters of Logarithmic-Normal Distributions by Maximum Likelihood. *Journal of the American Statistical Association* **46**, 206–212.
- Cohen, A.C. (1988). Three-Parameter Estimation. In Crow, E.L., and K. Shimizu, eds. *Lognormal Distributions: Theory and Applications*. Marcel Dekker, New York, Chapter 4.
- Cohen, A.C., and B.J. Whitten. (1980). Estimation in the Three-Parameter Lognormal Distribution. *Journal of the American Statistical Association* **75**, 399–404.
- Cohen, A.C., B.J. Whitten, and Y. Ding. (1985). Modified Moment Estimation for the Three-Parameter Lognormal Distribution. *Journal of Quality Technology* **17**, 92–99.
- Crow, E.L., and K. Shimizu. (1988). *Lognormal Distributions: Theory and Applications*. Marcel Dekker, New York, Chapter 2.
- Griffiths, D.A. (1980). Interval Estimation for the Three-Parameter Lognormal Distribution via the Likelihood Function. *Applied Statistics* **29**, 58–68.
- Harter, H.L., and A.H. Moore. (1966). Local-Maximum-Likelihood Estimation of the Parameters of Three-Parameter Lognormal Populations from Complete and Censored Samples. *Journal of the American Statistical Association* **61**, 842–851.
- Heyde, C.C. (1963). On a Property of the Lognormal Distribution. *Journal of the Royal Statistical Society, Series B* **25**, 392–393.
- Hill, .B.M. (1963). The Three-Parameter Lognormal Distribution and Bayesian Analysis of a Point-Source Epidemic. *Journal of the American Statistical Association* **58**, 72–84.
- Hoshi, K., J.R. Stedinger, and J. Burges. (1984). Estimation of Log-Normal Quantiles: Monte Carlo Results and First-Order Approximations. *Journal of Hydrology* **71**, 1–30.
- Johnson, N. L., S. Kotz, and N. Balakrishnan. (1994). *Continuous Univariate Distributions, Volume 1*. Second Edition. John Wiley and Sons, New York.
- Royston, J.P. (1992b). Estimation, Reference Ranges and Goodness of Fit for the Three-Parameter Log-Normal Distribution. *Statistics in Medicine* **11**, 897–912.
- Stedinger, J.R. (1980). Fitting Lognormal Distributions to Hydrologic Data. *Water Resources Research* **16**(3), 481–490.

See Also

[Lognormal3](#), [Lognormal](#), [LognormalAlt](#), [Normal](#).

Examples

```
# Generate 20 observations from a 3-parameter lognormal distribution
# with parameters meanlog=1.5, sdlog=1, and threshold=10, then use
# Cohen and Whitten's (1980) modified moments estimators to estimate
# the parameters, and construct a confidence interval for the
# threshold based on the estimated asymptotic variance.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
dat <- rlnorm3(20, meanlog = 1.5, sdlog = 1, threshold = 10)
elnorm3(dat, method = "mmme", ci = TRUE)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          3-Parameter Lognormal
#
#Estimated Parameter(s):      meanlog   = 1.5206664
#                               sdlog     = 0.5330974
#                               threshold  = 9.6620403
#
#Estimation Method:           mmme
#
#Data:                         dat
#
#Sample Size:                  20
#
#Confidence Interval for:      threshold
#
#Confidence Interval Method:    Normal Approximation
#                               Based on Asymptotic Variance
#
#Confidence Interval Type:     two-sided
#
#Confidence Level:             95%
#
#Confidence Interval:          LCL =  6.985258
#                               UCL = 12.338823
#
#-----

# Repeat the above example using the other methods of estimation
# and compare.

round(elnorm3(dat, "lmle")$parameters, 1)
#meanlog      sdlog threshold
#   1.3        0.7      10.5

round(elnorm3(dat, "mme")$parameters, 1)
#meanlog      sdlog threshold
#   2.1        0.3       6.0
```

```

round(elnorm3(dat, "mmue")$parameters, 1)
#meanlog      sdlog threshold
#   2.2        0.3        5.8

round(elnorm3(dat, "mmme")$parameters, 1)
#meanlog      sdlog threshold
#   1.5        0.5        9.7

round(elnorm3(dat, "zero.skew")$parameters, 1)
#meanlog      sdlog threshold
#   1.3        0.6       10.3

round(elnorm3(dat, "royston")$parameters, 1)
#meanlog      sdlog threshold
#   1.4        0.6       10.1

#-----

# Compare methods for computing a two-sided 95% confidence interval
# for the threshold:
# modified method of moments estimator using asymptotic variance,
# lmle using asymptotic variance,
# lmle using likelihood profile, and
# zero-skewness estimator using the skewness method.

elnorm3(dat, method = "mmme", ci = TRUE,
  ci.method = "avar")$interval$limits
#      LCL      UCL
# 6.985258 12.338823

elnorm3(dat, method = "lmle", ci = TRUE,
  ci.method = "avar")$interval$limits
#      LCL      UCL
# 9.017223 11.980107

elnorm3(dat, method = "lmle", ci = TRUE,
  ci.method="likelihood.profile")$interval$limits
#      LCL      UCL
# 3.699989 11.266029

elnorm3(dat, method = "zero.skew", ci = TRUE,
  ci.method = "skewness")$interval$limits
#      LCL      UCL
#-25.18851  11.18652

#-----

# Now construct a confidence interval for the median of the distribution
# based on using the modified method of moments estimator for threshold
# and the asymptotic variances and covariances. Note that the true median
# is given by threshold + exp(meanlog) = 10 + exp(1.5) = 14.48169.

elnorm3(dat, method = "mmme", ci = TRUE, ci.parameter = "median")

#Results of Distribution Parameter Estimation
#-----

```

```

#
#Assumed Distribution:      3-Parameter Lognormal
#
#Estimated Parameter(s):   meanlog   = 1.5206664
#                           sdlog     = 0.5330974
#                           threshold = 9.6620403
#
#Estimation Method:        mmme
#
#Data:                     dat
#
#Sample Size:              20
#
#Confidence Interval for:   median
#
#Confidence Interval Method: Normal Approximation
#                           Based on Asymptotic Variance
#
#Confidence Interval Type:  two-sided
#
#Confidence Level:         95%
#
#Confidence Interval:      LCL = 11.20541
#                           UCL = 17.26922
#
#-----

# Compare methods for computing a two-sided 95% confidence interval
# for the median:
# modified method of moments estimator using asymptotic variance,
# lmle using asymptotic variance,
# lmle using likelihood profile, and
# zero-skewness estimator using the skewness method.

elnorm3(dat, method = "mmme", ci = TRUE, ci.parameter = "median",
  ci.method = "avar")$interval$limits
#      LCL      UCL
#11.20541 17.26922

elnorm3(dat, method = "lmle", ci = TRUE, ci.parameter = "median",
  ci.method = "avar")$interval$limits
#      LCL      UCL
#12.28326 15.87233

elnorm3(dat, method = "lmle", ci = TRUE, ci.parameter = "median",
  ci.method = "likelihood.profile")$interval$limits
#      LCL      UCL
# 6.314583 16.165525

elnorm3(dat, method = "zero.skew", ci = TRUE, ci.parameter = "median",
  ci.method = "skewness")$interval$limits
#      LCL      UCL
#-22.38322 16.33569
#-----

# Clean up

```

```
#-----
rm(dat)
```

elnormAlt

Estimate Parameters of a Lognormal Distribution (Original Scale)

Description

Estimate the mean and coefficient of variation of a [lognormal distribution](#), and optionally construct a confidence interval for the mean.

Usage

```
elnormAlt(x, method = "mvue", ci = FALSE, ci.type = "two-sided",
  ci.method = "land", conf.level = 0.95, parkin.list = NULL)
```

Arguments

x	numeric vector of positive observations.
method	character string specifying the method of estimation. Possible values are "mvue" (minimum variance unbiased; the default), "qmlle" (quasi maximum likelihood), "mle" (maximum likelihood), "mme" (method of moments), and "mmue" (method of moments based on the unbiased estimate of variance). See the DETAILS section for more information on these estimation methods.
ci	logical scalar indicating whether to compute a confidence interval for the mean. The default value is FALSE.
ci.type	character string indicating what kind of confidence interval to compute. The possible values are "two-sided" (the default), "lower", and "upper". This argument is ignored if ci=FALSE.
ci.method	character string indicating what method to use to construct the confidence interval for the mean. The possible values are "land" (Land's method; the default), "zou" (Zou et al.'s method), "parkin" (Parkin et al.'s method), "cox" (Cox's approximation), and "normal.approx" (normal approximation). See the DETAILS section for more information. This argument is ignored if ci=FALSE.
conf.level	a scalar between 0 and 1 indicating the confidence level of the confidence interval. The default value is conf.level=0.95. This argument is ignored if ci=FALSE.
parkin.list	a list containing arguments for the function eqnpar . The components of this list are lcl.rank (set to NULL by default), ucl.rank (set to NULL by default), ci.method (set to "exact" if the sample size is ≤ 20 , otherwise set to "normal.approx"), and approx.conf.level (set to the value of conf.level). This argument is ignored unless ci=TRUE and ci.method="parkin".

Details

If \mathbf{x} contains any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

Let \underline{x} be a vector of n observations from a [lognormal distribution](#) with parameters $\text{mean}=\theta$ and $\text{cv}=\tau$. Let η denote the standard deviation of this distribution, so that $\eta = \theta\tau$. Set $\underline{y} = \log(\underline{x})$. Then \underline{y} is a vector of observations from a normal distribution with parameters $\text{mean}=\mu$ and $\text{sd}=\sigma$. See the help file for [LognormalAlt](#) for the relationship between θ, τ, η, μ , and σ .

Estimation

This section explains how each of the estimators of $\text{mean}=\theta$ and $\text{cv}=\tau$ are computed. The approach is to first compute estimates of θ and η^2 (the mean and variance of the lognormal distribution), say $\hat{\theta}$ and $\hat{\eta}^2$, then compute the estimate of the cv τ by $\hat{\tau} = \hat{\eta}/\hat{\theta}$.

Minimum Variance Unbiased Estimation (method="mvue")

The minimum variance unbiased estimators (mvue's) of θ and η^2 were derived by Finney (1941) and are discussed in Gilbert (1987, pp. 164-167) and Cohn et al. (1989). These estimators are computed as:

$$\hat{\theta}_{mvue} = e^{\bar{y}} g_{n-1}\left(\frac{s^2}{2}\right) \quad (1)$$

$$\hat{\eta}_{mvue}^2 = e^{2\bar{y}} \{g_{n-1}(2s^2) - g_{n-1}\left[\frac{(n-2)s^2}{n-1}\right]\} \quad (2)$$

where

$$\bar{y} = \frac{1}{n} \sum_{i=1}^n y_i \quad (3)$$

$$s^2 = \frac{1}{n-1} \sum_{i=1}^n (y_i - \bar{y})^2 \quad (4)$$

$$g_m(z) = \sum_{i=0}^{\infty} \frac{m^i(m+2i)}{m(m+2) \cdots (m+2i)} \left(\frac{m}{m+1}\right)^i \left(\frac{z}{i!}\right)^i \quad (5)$$

The expected value and variance of the mvue of θ are (Bradru and Mundlak, 1970; Cohn et al., 1989):

$$E[\hat{\theta}_{mvue}] = \theta \quad (6)$$

$$\text{Var}[\hat{\theta}_{mvue}] = e^{2\mu} \{e^{[(2+n-1)\sigma^2]/n} g_{n-1}\left(\frac{\sigma^4}{4n}\right) - e^{\sigma^2}\} \quad (7)$$

Maximum Likelihood Estimation (method="mle")

The maximum likelihood estimators (mle's) of θ and η^2 are given by:

$$\hat{\theta}_{mle} = \exp\left(\bar{y} + \frac{\hat{\sigma}_{mle}^2}{2}\right) \quad (8)$$

$$\hat{\eta}_{mle}^2 = \hat{\theta}_{mle}^2 \hat{\tau}_{mle}^2 \quad (9)$$

where

$$\hat{\tau}_{mle}^2 = \exp(\hat{\sigma}_{mle}^2) - 1 \quad (10)$$

$$\hat{\sigma}_{mle}^2 = \frac{n-1}{n} s^2 \quad (11)$$

The expected value and variance of the mle of θ are (after Cohn et al., 1989):

$$E[\hat{\theta}_{mle}] = \theta \exp\left[\frac{-(n-1)\sigma^2}{2n}\right] \left(1 - \frac{\sigma^2}{n}\right)^{-(n-1)/2} \quad (12)$$

$$Var[\hat{\theta}_{mle}] = \exp(2\mu + \frac{\sigma^2}{n}) \{ \exp(\frac{\sigma^2}{n}) [1 - \frac{2\sigma^2}{n}]^{-(n-1)/2} - [1 - \frac{\sigma^2}{n}]^{-(n-1)} \} \quad (13)$$

As can be seen from equation (12), the expected value of the mle of θ does not exist when $\sigma^2 > n$. In general, the p 'th moment of the mle of θ does not exist when $\sigma^2 > n/p$.

Quasi Maximum Likelihood Estimation (method="qml")

The quasi maximum likelihood estimators (qml's; Cohn et al., 1989; Gilbert, 1987, p.167) of θ and η^2 are the same as the mle's, except the mle of σ^2 in equations (8) and (10) is replaced with the more commonly used mvue of σ^2 shown in equation (4):

$$\hat{\theta}_{qml} = \exp(\bar{y} + \frac{s^2}{2}) \quad (14)$$

$$\hat{\eta}_{qml}^2 = \hat{\theta}_{qml}^2 \hat{\tau}_{qml}^2 \quad (15)$$

$$\hat{\tau}_{qml}^2 = \exp(s^2) - 1 \quad (16)$$

The expected value and variance of the qml of θ are (Cohn et al., 1989):

$$E[\hat{\theta}_{qml}] = \theta \exp[\frac{-(n-1)\sigma^2}{2n}] (1 - \frac{\sigma^2}{n-1})^{-(n-1)/2} \quad (17)$$

$$Var[\hat{\theta}_{qml}] = \exp(2\mu + \frac{\sigma^2}{n}) \{ \exp(\frac{\sigma^2}{n}) [1 - \frac{2\sigma^2}{n-1}]^{-(n-1)/2} - [1 - \frac{\sigma^2}{n-1}]^{-(n-1)} \} \quad (18)$$

As can be seen from equation (17), the expected value of the qml of θ does not exist when $\sigma^2 > (n-1)$. In general, the p 'th moment of the mle of θ does not exist when $\sigma^2 > (n-1)/p$.

Note that Gilbert (1987, p. 167) incorrectly presents equation (12) rather than equation (17) as the expected value of the qml of θ . For large values of n relative to σ^2 , however, equations (12) and (17) are virtually identical.

Method of Moments Estimation (method="mme")

The method of moments estimators (mme's) of θ and η^2 are found by equating the sample mean and variance with their population values:

$$\hat{\theta}_{mme} = \bar{x} = \frac{1}{n} \sum_{i=1}^n x_i \quad (19)$$

$$\hat{\eta}_{mme}^2 = \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2 \quad (20)$$

Note that the estimator of variance in equation (20) is biased.

The expected value and variance of the mme of θ are:

$$E[\hat{\theta}_{mme}] = \theta \quad (21)$$

$$Var[\hat{\theta}_{mme}] = \frac{\eta^2}{n} = \frac{1}{n} \exp(2\mu + \sigma^2) [\exp(\sigma^2) - 1] \quad (22)$$

Method of Moments Estimation Based on the Unbiased Estimate of Variance (method="mmue")

These estimators are exactly the same as the method of moments estimators described above, except that the usual unbiased estimate of variance is used:

$$\hat{\theta}_{mmue} = \bar{x} = \frac{1}{n} \sum_{i=1}^n x_i \quad (23)$$

$$\hat{\eta}_{mmue} = \frac{1}{n-1} \sum_{i=1}^n (x_i - \bar{x})^2 \quad (24)$$

Since the mmue of θ is equivalent to the mme of θ , so are its mean and variance.

Confidence Intervals

This section explains the different methods for constructing confidence intervals for θ , the mean of the lognormal distribution.

Land's Method (ci.method="land")

Land (1971, 1975) derived a method for computing one-sided (lower or upper) uniformly most accurate unbiased confidence intervals for θ . A two-sided confidence interval can be constructed by combining an optimal lower confidence limit with an optimal upper confidence limit. This procedure for two-sided confidence intervals is only asymptotically optimal, but for most purposes should be acceptable (Land, 1975, p.387).

As shown in equation (3) in the help file for [LognormalAlt](#), the mean θ of a lognormal random variable is related to the mean μ and standard deviation σ of the log-transformed random variable by the following relationship:

$$\theta = e^\beta \quad (25)$$

where

$$\beta = \mu + \frac{\sigma^2}{2} \quad (26)$$

Land (1971) developed confidence bounds for the quantity β . The mvue of β is given by:

$$\hat{\beta}_{mvue} = \bar{y} + \frac{s^2}{2} \quad (27)$$

Note that $\hat{\theta}_{qmle} = \exp(\hat{\beta}_{mvue})$. The $(1 - \alpha)100\%$ two-sided confidence interval for β is given by:

$$[\hat{\beta}_{mvue} + s \frac{C_{\alpha/2}}{\sqrt{n-1}}, \hat{\beta}_{mvue} + s \frac{C_{1-\alpha/2}}{\sqrt{n-1}}] \quad (28)$$

the $(1 - \alpha)100\%$ one-sided upper confidence interval for β is given by:

$$[-\infty, \hat{\beta}_{mvue} + s \frac{C_{1-\alpha}}{\sqrt{n-1}}] \quad (29)$$

and the $(1 - \alpha)100\%$ one-sided lower confidence interval for β is given by:

$$[\hat{\beta}_{mvue} + s \frac{C_\alpha}{\sqrt{n-1}}, \infty] \quad (30)$$

where s is the estimate of σ (see equation (4) above), and the factor C is given in tables in Land (1975).

Thus, by equations (25)-(30), the two-sided $(1 - \alpha)100\%$ confidence interval for θ is given by:

$$\{\hat{\theta}_{qmle} \exp[s \frac{C_{\alpha/2}}{\sqrt{n-1}}], \hat{\theta}_{qmle} \exp[s \frac{C_{1-\alpha/2}}{\sqrt{n-1}}]\} \quad (31)$$

the $(1 - \alpha)100\%$ one-sided upper confidence interval for θ is given by:

$$\{0, \hat{\theta}_{qmle} \exp[s \frac{C_{1-\alpha}}{\sqrt{n-1}}]\} \quad (32)$$

and the $(1 - \alpha)100\%$ one-sided lower confidence interval for θ is given by:

$$\{\hat{\theta}_{qmle} \exp[s \frac{C_\alpha}{\sqrt{n-1}}], \infty\} \quad (33)$$

Note that Gilbert (1987, pp. 169-171, 264-265) denotes the quantity C above as H and reproduces a subset of Land's (1975) tables. Some guidance documents (e.g., USEPA, 1992d) refer to this quantity as the H -statistic.

Zou et al.'s Method (ci.method="zou")

Zou et al. (2009) proposed the following approximation for the two-sided $(1 - \alpha)100\%$ confidence intervals for θ . The lower limit LL is given by:

$$LL = \hat{\theta}_{qmle} \exp\left\{-\left[\frac{z_{1-\alpha/2}^2 s^2}{n} + \left(\frac{s^2}{2} - \frac{(n-1)s^2}{2\chi_{1-\alpha/2, n-1}^2}\right)^2\right]^{1/2}\right\} \quad (34)$$

and the upper limit UL is given by:

$$UL = \hat{\theta}_{qmle} \exp\left\{\left[\frac{z_{1-\alpha/2}^2 s^2}{n} + \left(\frac{(n-1)s^2}{2\chi_{\alpha/2, n-1}^2} - \frac{s^2}{2}\right)^2\right]^{1/2}\right\} \quad (35)$$

where z_p denotes the p 'th quantile of the standard normal distribution, and $\chi_{p, \nu}$ denotes the p 'th quantile of the chi-square distribution with ν degrees of freedom. The $(1 - \alpha)100\%$ one-sided lower confidence limit and one-sided upper confidence limit are given by equations (34) and (35), respectively, with $\alpha/2$ replaced by α .

Parkin et al.'s Method (ci.method="parkin")

This method was developed by Parkin et al. (1990). It can be shown that the mean of a lognormal distribution corresponds to the p 'th quantile, where

$$p = \Phi\left(\frac{\sigma}{2}\right) \quad (36)$$

and Φ denotes the cumulative distribution function of the standard normal distribution. Parkin et al. (1990) suggested estimating p by replacing σ in equation (36) with the estimate s as computed in equation (4). Once an estimate of p is obtained, a nonparametric confidence interval can be constructed for p , assuming p is equal to its estimated value (see eqnpar).

Cox's Method (ci.method="cox")

This method was suggested by Professor D.R. Cox and is illustrated in Land (1972). El-Shaarawi (1989) adapts this method to the case of censored water quality data. Cox's idea is to construct an approximate $(1 - \alpha)100\%$ confidence interval for the quantity β defined in equation (26) above assuming the estimate of β is approximately normally distributed, and then exponentiate the confidence limits. That is, a two-sided $(1 - \alpha)100\%$ confidence interval for θ is constructed as:

$$[\exp(\hat{\beta} - t_{1-\alpha/2, n-1} \hat{\sigma}_{\hat{\beta}}), \exp(\hat{\beta} + t_{1-\alpha/2, n-1} \hat{\sigma}_{\hat{\beta}})] \quad (37)$$

where $t(p, \nu)$ denotes the p 'th quantile of Student's t-distribution with ν degrees of freedom. Note that this method, unlike the normal approximation method discussed below, guarantees a positive value for the lower confidence limit. One-sided confidence intervals are computed in a similar fashion.

Define an estimator of β by:

$$\hat{\beta} = \hat{\mu} + \frac{\hat{\sigma}^2}{2} \quad (38)$$

Then the variance of this estimator is given by:

$$Var(\hat{\beta}) = Var(\hat{\mu}) + Cov(\hat{\mu}, \hat{\sigma}^2) + \frac{1}{4}Var(\hat{\sigma}^2) \quad (39)$$

The function `elnormAlt` follows Land (1972) and uses the minimum variance unbiased estimator for β shown in equation (27) above, so the variance and estimated variance of this estimator are:

$$\text{Var}(\hat{\beta}_{mvue}) = \frac{\sigma^2}{n} + \frac{\sigma^4}{2(n-1)} \quad (40)$$

$$\hat{\sigma}_{\hat{\beta}}^2 = \frac{s^2}{n} + \frac{s^4}{2(n+1)} \quad (41)$$

Note that El-Shaarawi (1989, equation 5) simply replaces the value of s^2 in equation (41) with some estimator of σ^2 (the mle or mvue of σ^2), rather than using the mvue of the variance of β as shown in equation (41).

Normal Approximation (`ci.method="normal.approx"`) This method constructs approximate $(1 - \alpha)100\%$ confidence intervals for θ based on the assumption that the estimator of θ is approximately normally distributed. That is, a two-sided $(1 - \alpha)100\%$ confidence interval for θ is constructed as:

$$[\hat{\theta} - t_{1-\alpha/2, n-1} \hat{\sigma}_{\hat{\theta}}, \hat{\theta} + t_{1-\alpha/2, n-1} \hat{\sigma}_{\hat{\theta}}] \quad (42)$$

One-sided confidence intervals are computed in a similar fashion.

When `method="mvue"` is used to estimate θ , an unbiased estimate of the variance of the estimator of θ is used in equation (42) (Bradu and Mundlak, 1970, equation 4.3; Gilbert, 1987, equation 13.5):

$$\hat{\sigma}_{\hat{\theta}}^2 = e^{2\bar{y}} \{ [g_{n-1}(\frac{s^2}{2})]^2 - g_{n-1}[\frac{s^2(n-2)}{n-1}] \} \quad (43)$$

When `method="mle"` is used to estimate θ , the estimate of the variance of the estimator of θ is computed by replacing μ and σ^2 in equation (13) with their mle's:

$$\hat{\sigma}_{\hat{\theta}}^2 = \exp(2\bar{y} + \frac{\hat{\sigma}_{mle}^2}{n}) \{ \exp(\frac{\hat{\sigma}_{mle}^2}{n}) [1 - \frac{2\hat{\sigma}_{mle}^2}{n}]^{-(n-1)/2} - [1 - \frac{\hat{\sigma}_{mle}^2}{n}]^{-(n-1)} \} \quad (44)$$

When `method="qml"` is used to estimate θ , the estimate of the variance of the estimator of θ is computed by replacing μ and σ^2 in equation (18) with their mvue's:

$$\hat{\sigma}_{\hat{\theta}}^2 = \exp(2\bar{y} + \frac{s^2}{n}) \{ \exp(\frac{s^2}{n}) [1 - \frac{2s^2}{n-1}]^{-(n-1)/2} - [1 - \frac{s^2}{n-1}]^{-(n-1)} \} \quad (45)$$

Note that equation (45) is exactly the same as Gilbert's (1987, p. 167) equation 13.8a, except that Gilbert (1987) erroneously uses n where he should use $n - 1$ instead. For large values of n relative to s^2 , however, this makes little difference.

When `method="mme"`, the estimate of the variance of the estimator of θ is computed by replacing η^2 in equation (22) with the mme of η^2 defined in equation (20):

$$\hat{\sigma}_{\hat{\theta}}^2 = \frac{\hat{\eta}_{mme}^2}{n} = \frac{1}{n^2} \sum_{i=1}^n (x_i - \bar{x})^2 \quad (46)$$

When `method="mmue"`, the estimate of the variance of the estimator of θ is computed by replacing η^2 in equation (22) with the mmue of η^2 defined in equation (24):

$$\hat{\sigma}_{\hat{\theta}}^2 = \frac{\hat{\eta}_{mmue}^2}{n} = \frac{1}{n(n-1)} \sum_{i=1}^n (x_i - \bar{x})^2 \quad (47)$$

Value

a list of class "estimate" containing the estimated parameters and other information. See `estimate.object` for details.

Note

The normal and lognormal distribution are probably the two most frequently used distributions to model environmental data. In order to make any kind of probability statement about a normally-distributed population (of chemical concentrations for example), you have to first estimate the mean and standard deviation (the population parameters) of the distribution. Once you estimate these parameters, it is often useful to characterize the uncertainty in the estimate of the mean or variance. This is done with confidence intervals.

Some EPA guidance documents (e.g., Singh et al., 2002; Singh et al., 2010a,b) strongly recommend against using a lognormal model for environmental data and recommend trying a gamma distribution instead.

USEPA (1992d) directs persons involved in risk assessment for Superfund sites to use Land's (1971, 1975) method (`ci.method="land"`) for computing the upper 95% confidence interval for the mean, assuming the data follow a lognormal distribution (the guidance document cites Gilbert (1987) as a source of descriptions and tables for this method). The last example in the EXAMPLES section below reproduces an example from this guidance document.

In the past, some authors suggested using the geometric mean, also called the "rating curve" estimator (Cohn et al., 1989), as the estimator of the mean, θ . This estimator is computed as:

$$\hat{\theta}_{rc} = e^{\bar{g}} \quad (48)$$

Cohn et al. (1989) cite several authors who have pointed out this estimator is biased and is not even a consistent estimator of the mean. In fact, it is the maximum likelihood estimator of the median of the distribution (see [eqlnorm.](#))

Finney (1941) computed the efficiency of the method of moments estimators of the mean (θ) and variance (η^2) of the lognormal distribution (equations (19)-(20)) relative to the mvue's (equations (1)-(2)) as a function of σ^2 (the variance of the log-transformed observations), and found that while the mme of θ is reasonably efficient compared to the mvue of θ , the mme of η^2 performs quite poorly relative to the mvue of η^2 .

Cohn et al. (1989) and Parkin et al. (1988) have shown that the qmle and the mle of the mean can be severely biased for typical environmental data, and suggest always using the mvue.

Parkin et al. (1990) studied the performance of various methods for constructing a confidence interval for the mean via Monte Carlo simulation. They compared approximate methods to Land's optimal method (`ci.method="land"`). They used four parent lognormal distributions to generate observations; all had mean 10, but differed in coefficient of variation: 50, 100, 200, and 500%. They also generated sample sizes from 6 to 100 in increments of 2. For each combination of parent distribution and sample size, they generated 25,000 Monte Carlo trials. Parkin et al. found that for small sample sizes ($n < 20$), none of the approximate methods ("parkin", "cox", "normal.approx") worked very well. For $n > 20$, their method ("parkin") provided reasonably accurate coverage. Cox's method ("cox") worked well for $n > 60$, and performed slightly better than Parkin et al.'s method ("parkin") for highly skewed populations.

Zou et al. (2009) used Monte Carlo simulation to compare the performance of their method with the CGI method of Krishnamoorthy and Mathew (2003) and the modified Cox method of Armstrong (1992) and El-Shaarawi and Lin (2007). Performance was assessed based on 1) percentage of times the interval contained the parameter value (coverage%), 2) balance between left and right tail errors, and 3) confidence interval width. All three methods showed acceptable coverage percentages. The modified Cox method showed unbalanced tail errors, and Zou et al.'s method showed consistently narrower average width.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Aitchison, J., and J.A.C. Brown (1957). *The Lognormal Distribution (with special references to its uses in economics)*. Cambridge University Press, London, Chapter 5.
- Armstrong, B.G. (1992). Confidence Intervals for Arithmetic Means of Lognormally Distributed Exposures. *American Industrial Hygiene Association Journal* **53**, 481–485.
- Bradru, D., and Y. Mundlak. (1970). Estimation in Lognormal Linear Models. *Journal of the American Statistical Association* **65**, 198–211.
- Cohn, T.A., L.L. DeLong, E.J. Gilroy, R.M. Hirsch, and D.K. Wells. (1989). Estimating Constituent Loads. *Water Resources Research* **25**(5), 937–942.
- Crow, E.L., and K. Shimizu. (1988). *Lognormal Distributions: Theory and Applications*. Marcel Dekker, New York, Chapter 2.
- El-Shaarawi, A.H., and J. Lin. (2007). Interval Estimation for Log-Normal Mean with Applications to Water Quality. *Environmetrics* **18**, 1–10.
- El-Shaarawi, A.H., and R. Viveros. (1997). Inference About the Mean in Log-Regression with Environmental Applications. *Environmetrics* **8**, 569–582.
- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.
- Finney, D.J. (1941). On the Distribution of a Variate Whose Logarithm is Normally Distributed. *Supplement to the Journal of the Royal Statistical Society* **7**, 155–161.
- Gilbert, R.O. (1987). *Statistical Methods for Environmental Pollution Monitoring*. Van Nostrand Reinhold, New York, NY.
- Johnson, N. L., S. Kotz, and N. Balakrishnan. (1994). *Continuous Univariate Distributions, Volume 1*. Second Edition. John Wiley and Sons, New York.
- Krishnamoorthy, K., and T.P. Mathew. (2003). Inferences on the Means of Lognormal Distributions Using Generalized p-Values and Generalized Confidence Intervals. *Journal of Statistical Planning and Inference* **115**, 103–121.
- Land, C.E. (1971). Confidence Intervals for Linear Functions of the Normal Mean and Variance. *The Annals of Mathematical Statistics* **42**(4), 1187–1205.
- Land, C.E. (1972). An Evaluation of Approximate Confidence Interval Estimation Methods for Lognormal Means. *Technometrics* **14**(1), 145–158.
- Land, C.E. (1973). Standard Confidence Limits for Linear Functions of the Normal Mean and Variance. *Journal of the American Statistical Association* **68**(344), 960–963.
- Land, C.E. (1975). Tables of Confidence Limits for Linear Functions of the Normal Mean and Variance, in *Selected Tables in Mathematical Statistics, Vol. III*. American Mathematical Society, Providence, RI, pp. 385–419.
- Likes, J. (1980). Variance of the MVUE for Lognormal Variance. *Technometrics* **22**(2), 253–258.
- Limpert, E., W.A. Stahel, and M. Abbt. (2001). Log-Normal Distributions Across the Sciences: Keys and Clues. *BioScience* **51**, 341–352.
- Millard, S.P., and N.K. Neerchal. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, FL.
- Ott, W.R. (1995). *Environmental Statistics and Data Analysis*. Lewis Publishers, Boca Raton, FL.
- Parkin, T.B., J.J. Meisinger, S.T. Chester, J.L. Starr, and J.A. Robinson. (1988). Evaluation of Statistical Estimation Methods for Lognormally Distributed Variables. *Journal of the Soil Science Society of America* **52**, 323–329.

Parkin, T.B., S.T. Chester, and J.A. Robinson. (1990). Calculating Confidence Intervals for the Mean of a Lognormally Distributed Variable. *Journal of the Soil Science Society of America* **54**, 321–326.

Singh, A., A.K. Singh, and R.J. Iaci. (2002). *Estimation of the Exposure Point Concentration Term Using a Gamma Distribution*. EPA/600/R-02/084. October 2002. Technology Support Center for Monitoring and Site Characterization, Office of Research and Development, Office of Solid Waste and Emergency Response, U.S. Environmental Protection Agency, Washington, D.C.

Singh, A., R. Maichle, and N. Armbya. (2010a). *ProUCL Version 4.1.00 User Guide (Draft)*. EPA/600/R-07/041, May 2010. Office of Research and Development, U.S. Environmental Protection Agency, Washington, D.C.

Singh, A., N. Armbya, and A. Singh. (2010b). *ProUCL Version 4.1.00 Technical Guide (Draft)*. EPA/600/R-07/041, May 2010. Office of Research and Development, U.S. Environmental Protection Agency, Washington, D.C.

USEPA. (1992d). *Supplemental Guidance to RAGS: Calculating the Concentration Term*. Publication 9285.7-081, May 1992. Intermittenet Bulletin, Volume 1, Number 1. Office of Emergency and Remedial Response, Hazardous Site Evaluation Division, OS-230. Office of Solid Waste and Emergency Response, U.S. Environmental Protection Agency, Washington, D.C.

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.

Zou, G.Y., C.Y. Huo, and J. Taleban. (2009). Simple Confidence Intervals for Lognormal Means and their Differences with Environmental Applications. *Environmetrics* **20**, 172–180.

See Also

[LognormalAlt](#), [Lognormal](#), [Normal](#).

Examples

```
# Using the Reference area TcCB data in the data frame EPA.94b.tccb.df,
# estimate the mean and coefficient of variation,
# and construct a 95% confidence interval for the mean.

with(EPA.94b.tccb.df, elnormAlt(TcCB[Area == "Reference"], ci = TRUE))

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Lognormal
#
#Estimated Parameter(s):      mean = 0.5989072
#                               cv   = 0.4899539
#
#Estimation Method:           mvue
#
#Data:                         TcCB[Area == "Reference"]
#
#Sample Size:                  47
#
#Confidence Interval for:      mean
#
#Confidence Interval Method:   Land
```

```

#
#Confidence Interval Type:      two-sided
#
#Confidence Level:             95%
#
#Confidence Interval:          LCL = 0.5243787
#                               UCL = 0.7016992

#-----

# Compare the different methods of estimating the distribution parameters using the
# Reference area TcCB data.

with(EPA.94b.tccb.df, elnormAlt(TcCB[Area == "Reference"], method = "mvue"))$parameters
#      mean      cv
#0.5989072 0.4899539

with(EPA.94b.tccb.df, elnormAlt(TcCB[Area == "Reference"], method = "qmle"))$parameters
#      mean      cv
#0.6004468 0.4947791

with(EPA.94b.tccb.df, elnormAlt(TcCB[Area == "Reference"], method = "mle"))$parameters
#      mean      cv
#0.5990497 0.4888968

with(EPA.94b.tccb.df, elnormAlt(TcCB[Area == "Reference"], method = "mme"))$parameters
#      mean      cv
#0.5985106 0.4688423

with(EPA.94b.tccb.df, elnormAlt(TcCB[Area == "Reference"], method = "mmue"))$parameters
#      mean      cv
#0.5985106 0.4739110

#-----

# Compare the different methods of constructing the confidence interval for
# the mean using the Reference area TcCB data.

with(EPA.94b.tccb.df, elnormAlt(TcCB[Area == "Reference"],
  method = "mvue", ci = TRUE, ci.method = "land"))$interval$limits
#      LCL      UCL
#0.5243787 0.7016992

with(EPA.94b.tccb.df, elnormAlt(TcCB[Area == "Reference"],
  method = "mvue", ci = TRUE, ci.method = "zou"))$interval$limits
#      LCL      UCL
#0.5230444 0.6962071

with(EPA.94b.tccb.df, elnormAlt(TcCB[Area == "Reference"],
  method = "mvue", ci = TRUE, ci.method = "parkin"))$interval$limits
# LCL  UCL
#0.50 0.74

with(EPA.94b.tccb.df, elnormAlt(TcCB[Area == "Reference"],
  method = "mvue", ci = TRUE, ci.method = "cox"))$interval$limits
#      LCL      UCL
#0.5196213 0.6938444

```

```

with(EPA.94b.tccb.df, elnormAlt(TcCB[Area == "Reference"],
  method = "mvue", ci = TRUE, ci.method = "normal.approx"))$interval$limits
#      LCL      UCL
#0.5130160 0.6847984

#-----

# Reproduce the example in Highlights 7 and 8 of USEPA (1992d). This example shows
# how to compute the upper 95% confidence limit of the mean of a lognormal distribution
# and compares it to the result of computing the upper 95% confidence limit assuming a
# normal distribution. The data for this example are chromium concentrations (mg/kg) in
# soil samples collected randomly over a Superfund site, and are stored in the data frame
# EPA.92d.chromium.vec.

# First look at the data

EPA.92d.chromium.vec
# [1]  10  13  20  36  41  59  67 110 110 136 140 160 200 230 1300

stripChart(EPA.92d.chromium.vec, ylab = "Chromium (mg/kg)")

# Note there is one very large "outlier" (1300).
# Perform a goodness-of-fit test to determine whether a lognormal distribution
# is appropriate:

gof.list <- gofTest(EPA.92d.chromium.vec, dist = 'lnormAlt')
gof.list

#Results of Goodness-of-Fit Test
#-----
#
#Test Method:                Shapiro-Wilk GOF
#
#Hypothesized Distribution:   Lognormal
#
#Estimated Parameter(s):     mean = 159.855185
#                             cv   =  1.493994
#
#Estimation Method:          mvue
#
#Data:                        EPA.92d.chromium.vec
#
#Sample Size:                 15
#
#Test Statistic:              W = 0.9607179
#
#Test Statistic Parameter:    n = 15
#
#P-value:                     0.7048747
#
#Alternative Hypothesis:      True cdf does not equal the
#                             Lognormal Distribution.

plot(gof.list, digits = 2)

# The lognormal distribution seems to provide an adequate fit, although the largest

```

```
# observation (1300) is somewhat suspect, and given the small sample size there is
# not much power to detect any kind of mild deviation from a lognormal distribution.
```

```
# Now compute the one-sided 95% upper confidence limit for the mean.
# Note that the value of 502 mg/kg shown in Hightlight 7 of USEPA (1992d) is a bit
# larger than the exact value of 496.6 mg/kg shown below.
# This is simply due to rounding error.
```

```
elnormAlt(EPA.92d.chromium.vec, ci = TRUE, ci.type = "upper")
```

```
#Results of Distribution Parameter Estimation
```

```
#-----
```

```
#
#Assumed Distribution:      Lognormal
#
#Estimated Parameter(s):   mean = 159.855185
#                           cv   =  1.493994
#
#Estimation Method:       mvue
#
#Data:                     EPA.92d.chromium.vec
#
#Sample Size:              15
#
#Confidence Interval for:  mean
#
#Confidence Interval Method: Land
#
#Confidence Interval Type: upper
#
#Confidence Level:         95%
#
#Confidence Interval:      LCL =  0
#                           UCL = 496.6282
```

```
# Now compare this result with the upper 95% confidence limit based on assuming
# a normal distribution. Again note that the value of 325 mg/kg shown in
# Hightlight 8 is slightly larger than the exact value of 320.3 mg/kg shown below.
# This is simply due to rounding error.
```

```
enorm(EPA.92d.chromium.vec, ci = TRUE, ci.type = "upper")
```

```
#Results of Distribution Parameter Estimation
```

```
#-----
```

```
#
#Assumed Distribution:      Normal
#
#Estimated Parameter(s):   mean = 175.4667
#                           sd   = 318.5440
#
#Estimation Method:       mvue
#
#Data:                     EPA.92d.chromium.vec
#
#Sample Size:              15
#
#Confidence Interval for:  mean
```



```
#
#Confidence Interval Method:      Exact
#
#Confidence Interval Type:       upper
#
#Confidence Level:               95%
#
#Confidence Interval:            LCL =      -Inf
#                                UCL = 320.3304

#-----

# Clean up
#-----

rm(gof.list)
```

elnormAltCensored	<i>Estimate Parameters for a Lognormal Distribution (Original Scale) Based on Type I Censored Data</i>
-------------------	--

Description

Estimate the mean and coefficient of variation of a [lognormal distribution](#) given a sample of data that has been subjected to Type I censoring, and optionally construct a confidence interval for the mean.

Usage

```
elnormAltCensored(x, censored, method = "mle", censoring.side = "left",
  ci = FALSE, ci.method = "profile.likelihood", ci.type = "two-sided",
  conf.level = 0.95, n.bootstraps = 1000, use.acc.con = FALSE,
  pivot.statistic = "z", ...)
```

Arguments

x	numeric vector of observations. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
censored	numeric or logical vector indicating which values of x are censored. This must be the same length as x. If the mode of censored is "logical", TRUE values correspond to elements of x that are censored, and FALSE values correspond to elements of x that are not censored. If the mode of censored is "numeric", it must contain only 1's and 0's; 1 corresponds to TRUE and 0 corresponds to FALSE. Missing (NA) values are allowed but will be removed.
method	character string specifying the method of estimation. For singly censored data, the possible values are: "mle" (maximum likelihood; the default), "qmvue" (quasi minimum variance unbiased estimation) "bcmle" (bias-corrected maximum likelihood), "impute.w.qq.reg" (moment estimation based on imputation using the qq.reg method), "impute.w.qq.reg.w.cen.level" (moment estimation based on imputation using the qq.reg.w.cen.level method), "impute.w.mle" (moment estimation based on imputation using the mle), and

"half.cen.level" (moment estimation based on setting the censored observations to half the censoring level).

For multiply censored data, the possible values are: "mle" (maximum likelihood; the default), "qmvue" (quasi minimum variance unbiased estimation), "bcmle" (bias-corrected maximum likelihood), "impute.w.qq.reg" (moment estimation based on imputation using quantile-quantile regression), and "half.cen.level" (moment estimation based on setting the censored observations to half the censoring level).

See the DETAILS section for more information.

censoring.side	character string indicating on which side the censoring occurs. The possible values are "left" (the default) and "right".
ci	logical scalar indicating whether to compute a confidence interval for the mean or variance. The default value is ci=FALSE.
ci.method	<p>character string indicating what method to use to construct the confidence interval for the mean. The possible values are "profile.likelihood" (profile likelihood; the default), "cox" (Cox's approximation), "delta" (normal approximation based on the delta method), "normal.approx" (normal approximation), and "bootstrap" (based on bootstrapping).</p> <p>The confidence interval methods "delta" and "cox" are valid only when method is one of "mle", "bcmle", or "qmvue". The confidence interval method "normal.approx" is valid only when method is one of "impute.w.qq.reg", "impute.w.qq.reg.w.cen.level", "impute.w.mle", or "half.cen.level".</p> <p>See the DETAILS section for more information. This argument is ignored if ci=FALSE.</p>
ci.type	character string indicating what kind of confidence interval to compute. The possible values are "two-sided" (the default), "lower", and "upper". This argument is ignored if ci=FALSE.
conf.level	a scalar between 0 and 1 indicating the confidence level of the confidence interval. The default value is conf.level=0.95. This argument is ignored if ci=FALSE.
n.bootstraps	numeric scalar indicating how many bootstraps to use to construct the confidence interval for the mean when ci.type="bootstrap". This argument is ignored if ci=FALSE and/or ci.method does not equal "bootstrap".
use.acc.con	logical scalar indicating whether to use the acceleration constant when computing the bias-corrected bootstrap confidence interval (see the DETAILS section). The default value is FALSE. This argument is ignored if ci=FALSE and/or ci.method does not equal "bootstrap".
pivot.statistic	character string indicating which pivot statistic to use in the construction of the confidence interval for the mean when ci.method is equal to "delta", "cox", or "normal.approx" (see the DETAILS section). The possible values are pivot.statistic="z" (the default) and pivot.statistic="t". When pivot.statistic="t" you may supply the argument ci.sample size (see below). The argument pivot.statistic is ignored if ci=FALSE.
...	<p>additional arguments to pass to other functions.</p> <ul style="list-style-type: none"> • prob.method. Character string indicating what method to use to compute the plotting positions (empirical probabilities) when method is one of "impute.w.qq.reg", "impute.w.qq.reg.w.cen.level", or "impute.w.mle". Possible values are "kaplan-meier" (product-limit method of Kaplan and

Meier (1958)), "nelson" (hazard plotting method of Nelson (1972)), "michael-schucany" (generalization of the product-limit method due to Michael and Schucany (1986)), and "hirsch-stedinger" (generalization of the product-limit method due to Hirsch and Stedinger (1987)). The default value is `prob.method="michael-schucany"`. The "nelson" method is only available for `censoring.side="right"`. See the DETAILS section and the help file for [ppointsCensored](#) for more information.

- `plot.pos.con`. Numeric scalar between 0 and 1 containing the value of the plotting position constant to use when method is one of "impute.w.qq.reg", "impute.w.qq.reg.w.cen.level", or "impute.w.mle". The default value is `plot.pos.con=0.375`. See the DETAILS section and the help file for [ppointsCensored](#) for more information.
- `ci.sample.size`. Numeric scalar indicating what sample size to assume to construct the confidence interval for the mean if `pivot.statistic="t"` and `ci.method` is equal to "delta", "cox", or "normal.approx". When method equals "mle", "bcmle", or "qmvue", the default value is the expected number of uncensored observations, otherwise it is the observed number of uncensored observations.
- `lb.impute`. Numeric scalar indicating the lower bound for imputed observations when method is one of "impute.w.qq.reg", "impute.w.qq.reg.w.cen.level", or "impute.w.mle". Imputed values smaller than this value will be set to this value. The default is `lb.impute=-Inf`.
- `ub.impute`. Numeric scalar indicating the upper bound for imputed observations when method is one of "impute.w.qq.reg", "impute.w.qq.reg.w.cen.level", or "impute.w.mle". Imputed values larger than this value will be set to this value. The default is `ub.impute=Inf`.

Details

If `x` or `censored` contain any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

Let \underline{x} be a vector of n observations from a [lognormal distribution](#) with parameters $\text{mean}=\theta$ and $\text{cv}=\tau$. Let η denote the standard deviation of this distribution, so that $\eta = \theta\tau$. Set $\underline{y} = \log(\underline{x})$. Then \underline{y} is a vector of observations from a normal distribution with parameters $\text{mean}=\mu$ and $\text{sd}=\sigma$. See the help file for [LognormalAlt](#) for the relationship between θ, τ, η, μ , and σ .

Let \underline{x} denote a vector of N observations from a [lognormal distribution](#) with parameters $\text{mean}=\theta$ and $\text{cv}=\tau$. Let η denote the standard deviation of this distribution, so that $\eta = \theta\tau$. Set $\underline{y} = \log(\underline{x})$. Then \underline{y} is a vector of observations from a normal distribution with parameters $\text{mean}=\mu$ and $\text{sd}=\sigma$. See the help file for [LognormalAlt](#) for the relationship between θ, τ, η, μ , and σ .

Assume n ($0 < n < N$) of the N observations are known and c ($c = N - n$) of the observations are all censored below (left-censored) or all censored above (right-censored) at k fixed censoring levels

$$T_1, T_2, \dots, T_k; \quad k \geq 1 \quad (1)$$

For the case when $k \geq 2$, the data are said to be Type I **multiply censored**. For the case when $k = 1$, set $T = T_1$. If the data are left-censored and all n known observations are greater than or equal to T , or if the data are right-censored and all n known observations are less than or equal to T , then the data are said to be Type I **singly censored** (Nelson, 1982, p.7), otherwise they are considered to be Type I multiply censored.

Let c_j denote the number of observations censored below or above censoring level T_j for $j =$

1, 2, ..., k, so that

$$\sum_{j=1}^k c_j = c \quad (2)$$

Let $x_{(1)}, x_{(2)}, \dots, x_{(N)}$ denote the “ordered” observations, where now “observation” means either the actual observation (for uncensored observations) or the censoring level (for censored observations). For right-censored data, if a censored observation has the same value as an uncensored one, the uncensored observation should be placed first. For left-censored data, if a censored observation has the same value as an uncensored one, the censored observation should be placed first.

Note that in this case the quantity $x_{(i)}$ does not necessarily represent the i ’th “largest” observation from the (unknown) complete sample.

Finally, let Ω (omega) denote the set of n subscripts in the “ordered” sample that correspond to uncensored observations.

ESTIMATION

This section explains how each of the estimators of $\text{mean}=\theta$ and $\text{cv}=\tau$ are computed. The approach is to first compute estimates of θ and η^2 (the mean and variance of the lognormal distribution), say $\hat{\theta}$ and $\hat{\eta}^2$, then compute the estimate of the cv τ by $\hat{\tau} = \hat{\eta}/\hat{\theta}$.

Maximum Likelihood Estimation (method="mle")

The maximum likelihood estimators of θ , τ , and η are computed as:

$$\hat{\theta}_{mle} = \exp(\hat{\mu}_{mle} + \frac{\hat{\sigma}_{mle}^2}{2}) \quad (3)$$

$$\hat{\tau}_{mle} = [\exp(\hat{\sigma}_{mle}^2) - 1]^{1/2} \quad (4)$$

$$\hat{\eta}_{mle} = \hat{\theta}_{mle} \hat{\tau}_{mle} \quad (5)$$

where $\hat{\mu}_{mle}$ and $\hat{\sigma}_{mle}$ denote the maximum likelihood estimators of μ and σ . See the help for [enormCensored](#) for information on how $\hat{\mu}_{mle}$ and $\hat{\sigma}_{mle}$ are computed.

Quasi Minimum Variance Unbiased Estimation Based on the MLE's (method="qmvue")

The maximum likelihood estimators of θ and η^2 are biased. Even for complete (uncensored) samples these estimators are biased (see equation (12) in the help file for [elnormAlt](#)). The bias tends to 0 as the sample size increases, but it can be considerable for small sample sizes. (Cohn et al., 1989, demonstrate the bias for complete data sets.) For the case of complete samples, the minimum variance unbiased estimators (mvue's) of θ and η^2 were derived by Finney (1941) and are discussed in Gilbert (1987, pp.164-167) and Cohn et al. (1989). These estimators are computed as:

$$\hat{\theta}_{mvue} = e^{\bar{y}} g_{n-1}(\frac{s^2}{2}) \quad (6)$$

$$\hat{\eta}_{mvue}^2 = e^{2\bar{y}} \{g_{n-1}(2s^2) - g_{n-1}[\frac{(n-2)s^2}{n-1}]\} \quad (7)$$

where

$$\bar{y} = \frac{1}{n} \sum_{i=1}^n y_i \quad (8)$$

$$s^2 = \frac{1}{n-1} \sum_{i=1}^n (y_i - \bar{y})^2 \quad (9)$$

$$g_m(z) = \sum_{i=0}^{\infty} \frac{m^i (m+2i)}{m(m+2) \cdots (m+2i)} \left(\frac{m}{m+1}\right)^i \left(\frac{z}{i!}\right) \quad (10)$$

(see the help file for [elnormAlt](#)).

For Type I censored samples, the quasi minimum variance unbiased estimators (qmvue's) of θ and η^2 are computed using equations (6) and (7) and estimating μ and σ with their mle's (see [elnormCensored](#)).

For singly censored data, this is apparently the LM method of Gilliom and Helsel (1986, p.137) (it is not clear from their description on page 137 whether their LM method is the straight method="mle" described above or method="qmvue" described here). This method was also used by Newman et al. (1989, p.915, equations 10-11).

For multiply censored data, this is apparently the MM method of Helsel and Cohn (1988, p.1998). (It is not clear from their description on page 1998 and the description in Gilliom and Helsel, 1986, page 137 whether Helsel and Cohn's (1988) MM method is the straight method="mle" described above or method="qmvue" described here.)

Bias-Corrected Maximum Likelihood Estimation (method="bcmle")

This method was derived by El-Shaarawi (1989) and can be applied to complete or censored data sets. For complete data, the exact relative bias of the mle of the mean θ is given as:

$$B_{mle} = \frac{E[\hat{\theta}_{mle}]}{\theta} = \exp\left[\frac{-(n-1)\sigma^2}{2n}\right] \left(1 - \frac{\sigma^2}{n}\right)^{-(n-1)/2} \quad (11)$$

(see equation (12) in the help file for [elnormAlt](#)).

For the case of complete or censored data, El-Shaarawi (1989) proposed the following "bias-corrected" maximum likelihood estimator:

$$\hat{\theta}_{bcmle} = \frac{\hat{\theta}_{mle}}{\hat{B}_{mle}} \quad (12)$$

where

$$\hat{B}_{mle} = \exp\left[\frac{1}{2}(\hat{V}_{11} + 2\hat{\sigma}_{mle}\hat{V}_{12} + \hat{\sigma}_{mle}^2\hat{V}_{22})\right] \quad (13)$$

and V denotes the asymptotic variance-covariance of the mle's of μ and σ , which is based on the observed information matrix, formulas for which are given in Cohen (1991). El-Shaarawi (1989) does not propose a bias-corrected estimator of the variance η^2 , so the mle of η is computed when method="bcmle".

Imputation Using Quantile-Quantile Regression (method = "impute.w.qq.reg")

This method involves using quantile-quantile regression on the log-transformed observations to fit a regression line (and thus initially estimate the mean μ and standard deviation σ in log-space), imputing the log-transformed values of the c censored observations by predicting them from the regression equation, transforming the log-scale imputed values back to the original scale, and then computing the method of moments estimates of the mean and standard deviation based on the observed and imputed values.

The steps are:

1. Estimate μ and σ by computing the least-squares estimates in the following model:

$$y_{(i)} = \mu + \sigma\Phi^{-1}(p_i) + \epsilon_i, \quad i \in \Omega \quad (14)$$

where p_i denotes the plotting position associated with the i 'th largest value, a is a constant such that $0 \leq a \leq 1$ (the default value is 0.375), Φ denotes the cumulative distribution function (cdf) of the standard normal distribution and Ω denotes the set of n subscripts associated with the uncensored observations in the ordered sample. The plotting positions are computed by calling the function [ppointsCensored](#).

2. Compute the log-scale imputed values as:

$$\hat{y}_{(i)} = \hat{\mu}_{qqreg} + \hat{\sigma}_{qqreg} \Phi^{-1}(p_i), \quad i \notin \Omega \quad (15)$$

3. Retransform the log-scale imputed values:

$$\hat{x}_{(i)} = \exp[\hat{y}_{(i)}], \quad i \notin \Omega \quad (16)$$

4. Compute the usual method of moments estimates of the mean and variance.

$$\hat{\theta} = \frac{1}{N} \left[\sum_{i \notin \Omega} \hat{x}_{(i)} + \sum_{i \in \Omega} x_{(i)} \right] \quad (17)$$

$$\hat{\eta}^2 = \frac{1}{N-1} \left[\sum_{i \notin \Omega} (\hat{x}_{(i)} - \hat{\theta})^2 + \sum_{i \in \Omega} (x_{(i)} - \hat{\theta})^2 \right] \quad (18)$$

Note that the estimate of variance is actually the usual unbiased one (not the method of moments one) in the case of complete data.

For singly censored data, this method is discussed by Hashimoto and Trussell (1983), Gilliom and Helsel (1986), and El-Shaarawi (1989), and is referred to as the LR (Log-Regression) or Log-Probability Method.

For multiply censored data, this is the MR method of Helsel and Cohn (1988, p.1998). They used it with the probability method of Hirsch and Stedinger (1987) and Weibull plotting positions (i.e., `prob.method="hirsch-stedinger"` and `plot.pos.con=0`).

The argument `plot.pos.con` (see the entry for ... in the ARGUMENTS section above) determines the value of the plotting positions computed in equations (14) and (15) when `method` equals "hirsch-stedinger" or "michael-schucany". The default value is `plot.pos.con=0.375`. See the help file for [ppointsCensored](#) for more information.

The arguments `lb.impute` and `ub.impute` (see the entry for ... in the ARGUMENTS section above) determine the lower and upper bounds for the imputed values. Imputed values smaller than `lb.impute` are set to this value. Imputed values larger than `ub.impute` are set to this value. The default values are `lb.impute=0` and `ub.impute=Inf`.

Imputation Using Quantile-Quantile Regression Including the Censoring Level (`method="impute.w.qq.reg.w.cen.1"`)

This method is only available for singly censored data. This method was proposed by El-Shaarawi (1989), which he denoted as the Modified LR Method. It is exactly the same method as imputation using quantile-quantile regression (`method="impute.w.qq.reg"`), except that the quantile-quantile regression includes the censoring level. For left singly censored data, the modification involves adding the point $[\Phi^{-1}(p_c), T]$ to the plot before fitting the least-squares line. For right singly censored data, the point $[\Phi^{-1}(p_{n+1}), T]$ is added to the plot before fitting the least-squares line.

Imputation Using Maximum Likelihood (`method="impute.w.mle"`)

This method is only available for singly censored data. This is exactly the same method as imputation with quantile-quantile regression (`method="impute.w.qq.reg"`), except that the maximum likelihood method (`method="mle"`) is used to compute the initial estimates of the mean and standard deviation. In the context of lognormal data, this method is discussed by El-Shaarawi (1989),

which he denotes as the Modified Maximum Likelihood Method.

Setting Censored Observations to Half the Censoring Level (method="half.cen.level")

This method is applicable only to left censored data that is bounded below by 0. This method involves simply replacing all the censored observations with half their detection limit, and then computing the usual moment estimators of the mean and variance. That is, all censored observations are imputed to be half the detection limit, and then Equations (17) and (18) are used to estimate the mean and variance.

This method is included only to allow comparison of this method to other methods. **Setting left-censored observations to half the censoring level is not recommended.** In particular, El-Shaarawi and Esterby (1992) show that these estimators are biased and inconsistent (i.e., the bias remains even as the sample size increases).

CONFIDENCE INTERVALS

This section explains how confidence intervals for the mean θ are computed.

Likelihood Profile (ci.method="profile.likelihood")

This method was proposed by Cox (1970, p.88), and Venzon and Moolgavkar (1988) introduced an efficient method of computation. This method is also discussed by Stryhn and Christensen (2003) and Royston (2007). The idea behind this method is to invert the likelihood-ratio test to obtain a confidence interval for the mean θ while treating the coefficient of variation τ as a nuisance parameter.

For Type I left censored data, the likelihood function is given by:

$$L(\theta, \tau | \underline{x}) = \binom{N}{c_1 c_2 \dots c_k n} \prod_{j=1}^k [F(T_j)]^{c_j} \prod_{i \in \Omega} f[x_{(i)}] \quad (19)$$

where f and F denote the probability density function (pdf) and cumulative distribution function (cdf) of the population. That is,

$$f(t) = \phi\left(\frac{t - \mu}{\sigma}\right) \quad (20)$$

$$F(t) = \Phi\left(\frac{t - \mu}{\sigma}\right) \quad (21)$$

where

$$\mu = \log\left(\frac{\theta}{\sqrt{\tau^2 + 1}}\right) \quad (22)$$

$$\sigma = [\log(\tau^2 + 1)]^{1/2} \quad (23)$$

and ϕ and Φ denote the pdf and cdf of the standard normal distribution, respectively (Cohen, 1963; 1991, pp.6, 50). For left singly censored data, equation (3) simplifies to:

$$L(\mu, \sigma | \underline{x}) = \binom{N}{c} [F(T)]^c \prod_{i=c+1}^n f[x_{(i)}] \quad (24)$$

Similarly, for Type I right censored data, the likelihood function is given by:

$$L(\mu, \sigma | \underline{x}) = \binom{N}{c_1 c_2 \dots c_k n} \prod_{j=1}^k [1 - F(T_j)]^{c_j} \prod_{i \in \Omega} f[x_{(i)}] \quad (25)$$

and for right singly censored data this simplifies to:

$$L(\mu, \sigma | \underline{x}) = \binom{N}{c} [1 - F(T)]^c \prod_{i=1}^n f[x_{(i)}] \quad (26)$$

Following Stryhn and Christensen (2003), denote the maximum likelihood estimates of the mean and coefficient of variation by (θ^*, τ^*) . The likelihood ratio test statistic (G^2) of the hypothesis $H_0 : \theta = \theta_0$ (where θ_0 is a fixed value) equals the drop in $2\log(L)$ between the “full” model and the reduced model with θ fixed at θ_0 , i.e.,

$$G^2 = 2\{\log[L(\theta^*, \tau^*)] - \log[L(\theta_0, \tau_0^*)]\} \quad (27)$$

where τ_0^* is the maximum likelihood estimate of τ for the reduced model (i.e., when $\theta = \theta_0$). Under the null hypothesis, the test statistic G^2 follows a [chi-squared distribution](#) with 1 degree of freedom.

Alternatively, we may express the test statistic in terms of the profile likelihood function L_1 for the mean θ , which is obtained from the usual likelihood function by maximizing over the parameter τ , i.e.,

$$L_1(\theta) = \max_{\tau} L(\theta, \tau) \quad (28)$$

Then we have

$$G^2 = 2\{\log[L_1(\theta^*)] - \log[L_1(\theta_0)]\} \quad (29)$$

A two-sided $(1 - \alpha)100\%$ confidence interval for the mean θ consists of all values of θ_0 for which the test is not significant at level *alpha*:

$$\theta_0 : G^2 \leq \chi_{1,1-\alpha}^2 \quad (30)$$

where $\chi_{\nu,p}^2$ denotes the p 'th quantile of the [chi-squared distribution](#) with ν degrees of freedom.

Direct Normal Approximations (ci.method="delta" or ci.method="normal.approx")

An approximate $(1 - \alpha)100\%$ confidence interval for θ can be constructed assuming the distribution of the estimator of θ is approximately normally distributed. That is, a two-sided $(1 - \alpha)100\%$ confidence interval for θ is constructed as:

$$[\hat{\theta} - t_{1-\alpha/2, m-1} \hat{\sigma}_{\hat{\theta}}, \hat{\theta} + t_{1-\alpha/2, m-1} \hat{\sigma}_{\hat{\theta}}] \quad (31)$$

where $\hat{\theta}$ denotes the estimate of θ , $\hat{\sigma}_{\hat{\theta}}$ denotes the estimated asymptotic standard deviation of the estimator of θ , m denotes the assumed sample size for the confidence interval, and $t_{p,\nu}$ denotes the p 'th quantile of [Student's t-distribution](#) with ν degrees of freedom. One-sided confidence intervals are computed in a similar fashion.

The argument ci.sample.size determines the value of m (see the entry for ... in the ARGUMENTS section above). When method equals "mle", "qmvue", or "bcmle" and the data are singly censored, the default value is the expected number of uncensored observations, otherwise it is n , the observed number of uncensored observations. This is simply an ad-hoc method of constructing confidence intervals and is not based on any published theoretical results.

When pivot.statistic="z", the p 'th quantile from the [standard normal distribution](#) is used in place of the p 'th quantile from Student's t-distribution.

Direct Normal Approximation Based on the Delta Method (ci.method="delta")

This method is usually applied with the maximum likelihood estimators (method="mle"). It should also work approximately for the quasi minimum variance unbiased estimators (method="qmvue") and the bias-corrected maximum likelihood estimators (method="bcmle").

When `method="mle"`, the variance of the mle of θ can be estimated based on the variance-covariance matrix of the mle's of μ and σ (denoted V), and the delta method:

$$\hat{\sigma}_{\theta}^2 = \left(\frac{\partial \theta}{\partial \underline{\lambda}} \right)'_{\underline{\lambda}} \hat{V} \left(\frac{\partial \theta}{\partial \underline{\lambda}} \right)_{\underline{\lambda}} \quad (32)$$

where

$$\underline{\lambda}' = (\mu, \sigma) \quad (33)$$

$$\frac{\partial \theta}{\partial \mu} = \exp\left(\mu + \frac{\sigma^2}{2}\right) \quad (34)$$

$$\frac{\partial \theta}{\partial \sigma} = \sigma \exp\left(\mu + \frac{\sigma^2}{2}\right) \quad (35)$$

(Shumway et al., 1989). The variance-covariance matrix V of the mle's of μ and σ is estimated based on the inverse of the observed Fisher Information matrix, formulas for which are given in Cohen (1991).

Direct Normal Approximation Based on the Moment Estimators (`ci.method="normal.approx"`)

This method is valid only for the moment estimators based on imputed values (i.e., `method="impute.w.qq.reg"` or `method="half.cen.level"`). For these cases, the standard deviation of the estimated mean is assumed to be approximated by

$$\hat{\sigma}_{\theta} = \frac{\hat{\eta}}{\sqrt{m}} \quad (36)$$

where, as already noted, m denotes the assumed sample size. This is simply an ad-hoc method of constructing confidence intervals and is not based on any published theoretical results.

Cox's Method (`ci.method="cox"`)

This method may be applied with the maximum likelihood estimators (`method="mle"`), the quasi minimum variance unbiased estimators (`method="qmvue"`), and the bias-corrected maximum likelihood estimators (`method="bcmle"`).

This method was proposed by El-Shaarawi (1989) and is an extension of the method derived by Cox and presented in Land (1972) for the case of complete data (see the explanation of `ci.method="cox"` in the help file for [elnormAlt](#)). The idea is to construct an approximate $(1 - \alpha)100\%$ confidence interval for the quantity

$$\beta = \exp\left(\mu + \frac{\sigma^2}{2}\right) \quad (37)$$

assuming the estimate of β

$$\hat{\beta} = \exp\left(\hat{\mu} + \frac{\hat{\sigma}^2}{2}\right) \quad (38)$$

is approximately normally distributed, and then exponentiate the confidence limits. That is, a two-sided $(1 - \alpha)100\%$ confidence interval for θ is constructed as:

$$[\exp(\hat{\beta} - h), \exp(\hat{\beta} + h)] \quad (39)$$

where

$$h = t_{1-\alpha/2, m-1} \hat{\sigma}_{\hat{\beta}} \quad (40)$$

and $\hat{\sigma}_{\hat{\beta}}$ denotes the estimated asymptotic standard deviation of the estimator of β , m denotes the assumed sample size for the confidence interval, and $t_{p, \nu}$ denotes the p 'th quantile of [Student's t-distribution](#) with ν degrees of freedom.

El-Shaarawi (1989) shows that the standard deviation of the mle of β can be estimated by:

$$\hat{\sigma}_{\hat{\beta}} = \sqrt{\hat{V}_{11} + 2\hat{\sigma}\hat{V}_{12} + \hat{\sigma}^2\hat{V}_{22}} \quad (41)$$

where V denotes the variance-covariance matrix of the mle's of μ and σ and is estimated based on the inverse of the Fisher Information matrix.

One-sided confidence intervals are computed in a similar fashion.

Bootstrap and Bias-Corrected Bootstrap Approximation (ci.method="bootstrap")

The bootstrap is a nonparametric method of estimating the distribution (and associated distribution parameters and quantiles) of a sample statistic, regardless of the distribution of the population from which the sample was drawn. The bootstrap was introduced by Efron (1979) and a general reference is Efron and Tibshirani (1993).

In the context of deriving an approximate $(1 - \alpha)100\%$ confidence interval for the population mean θ , the bootstrap can be broken down into the following steps:

1. Create a bootstrap sample by taking a random sample of size N from the observations in \underline{x} , where sampling is done with replacement. Note that because sampling is done with replacement, the same element of \underline{x} can appear more than once in the bootstrap sample. Thus, the bootstrap sample will usually not look exactly like the original sample (e.g., the number of censored observations in the bootstrap sample will often differ from the number of censored observations in the original sample).
2. Estimate θ based on the bootstrap sample created in Step 1, using the same method that was used to estimate θ using the original observations in \underline{x} . Because the bootstrap sample usually does not match the original sample, the estimate of θ based on the bootstrap sample will usually differ from the original estimate based on \underline{x} .
3. Repeat Steps 1 and 2 B times, where B is some large number. The number of bootstraps B is determined by the argument n.bootstraps (see the section ARGUMENTS above). The default value of n.bootstraps is 1000.
4. Use the B estimated values of θ to compute the empirical cumulative distribution function of this estimator of θ (see [ecdfPlot](#)), and then create a confidence interval for θ based on this estimated cdf.

The percentile interval (Efron and Tibshirani, 1993, p.170) is computed as:

$$[\hat{G}^{-1}(\frac{\alpha}{2}), \hat{G}^{-1}(\frac{1-\alpha}{2})] \quad (42)$$

where $\hat{G}(t)$ denotes the empirical cdf evaluated at t and thus $\hat{G}^{-1}(p)$ denotes the p 'th empirical quantile, that is, the p 'th quantile associated with the empirical cdf. The function `elnormCensored` calls the R function `quantile` to compute the empirical quantiles used in equation (42).

The percentile method bootstrap confidence interval is only first-order accurate (Efron and Tibshirani, 1993, pp.187-188), meaning that the probability that the confidence interval will contain the true value of θ can be off by k/\sqrt{N} , where k is some constant. Efron and Tibshirani (1993, pp.184-188) proposed a bias-corrected and accelerated interval that is second-order accurate, meaning that the probability that the confidence interval will contain the true value of θ may be off by k/N instead of k/\sqrt{N} . The bias-corrected and accelerated confidence interval is computed as:

$$[\hat{G}^{-1}(\alpha_1), \hat{G}^{-1}(\alpha_2)] \quad (43)$$

where

$$\alpha_1 = \Phi[\hat{z}_0 + \frac{\hat{z}_0 + z_{\alpha/2}}{1 - \hat{a}(z_0 + z_{\alpha/2})}] \quad (44)$$

$$\alpha_2 = \Phi[\hat{z}_0 + \frac{\hat{z}_0 + z_{1-\alpha/2}}{1 - \hat{a}(z_0 + z_{1-\alpha/2})}] \quad (45)$$

$$\hat{z}_0 = \Phi^{-1}[\hat{G}(\hat{\theta})] \quad (46)$$

$$\hat{a} = \frac{\sum_{i=1}^N (\hat{\theta}_{(\cdot)} - \hat{\theta}_{(i)})^3}{6[\sum_{i=1}^N (\hat{\theta}_{(\cdot)} - \hat{\theta}_{(i)})^2]^{3/2}} \quad (47)$$

where the quantity $\hat{\theta}_{(i)}$ denotes the estimate of θ using all the values in \underline{x} except the i 'th one, and

$$\hat{\theta}_{(\cdot)} = \frac{1}{N} \sum_{i=1}^N \hat{\theta}_{(i)} \quad (48)$$

The constant \hat{z}_0 incorporates the bias correction, and the constant \hat{a} is the acceleration constant. The term “acceleration” refers to the rate of change of the standard error of the estimate of θ with respect to the true value of θ (Efron and Tibshirani, 1993, p.186). For a normal (Gaussian) distribution, the standard error of the estimate of θ does not depend on the value of θ , hence the acceleration constant is not really necessary. The argument `use.acc.con` (see the section ARGUMENTS above) determines whether the acceleration constant \hat{a} is computed as in equation (47) (`use.acc.con=TRUE`) or is set to 0 (`use.acc.con=FALSE`). The default value is `use.acc.con=FALSE`.

When `ci.method="bootstrap"`, the function `enormCensored` computes both the percentile method and bias-corrected bootstrap confidence intervals.

This method of constructing confidence intervals for censored data was studied by Shumway et al. (1989).

Value

a list of class “`estimateCensored`” containing the estimated parameters and other information. See [estimateCensored.object](#) for details.

Note

A sample of data contains censored observations if some of the observations are reported only as being below or above some censoring level. In environmental data analysis, Type I left-censored data sets are common, with values being reported as “less than the detection limit” (e.g., Helsel, 2012). Data sets with only one censoring level are called *singly censored*; data sets with multiple censoring levels are called *multiply* or *progressively censored*.

Statistical methods for dealing with censored data sets have a long history in the field of survival analysis and life testing. More recently, researchers in the environmental field have proposed alternative methods of computing estimates and confidence intervals in addition to the classical ones such as maximum likelihood estimation.

Helsel (2012, Chapter 6) gives an excellent review of past studies of the properties of various estimators based on censored environmental data.

In practice, it is better to use a confidence interval for the mean or a joint confidence region for the mean and standard deviation, rather than rely on a single point-estimate of the mean. Since confidence intervals and regions depend on the properties of the estimators for both the mean and standard deviation, the results of studies that simply evaluated the performance of the mean and standard deviation separately cannot be readily extrapolated to predict the performance of various methods of constructing confidence intervals and regions. Furthermore, for several of the methods that have been proposed to estimate the mean based on type I left-censored data, standard errors of the estimates are not available, hence it is not possible to construct confidence intervals (El-Shaarawi and Dolan, 1989).

Few studies have been done to evaluate the performance of methods for constructing confidence intervals for the mean or joint confidence regions for the mean and standard deviation **on the original scale, not the log-scale** when data are subjected to single or multiple censoring.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Bain, L.J., and M. Engelhardt. (1991). *Statistical Analysis of Reliability and Life-Testing Models*. Marcel Dekker, New York, 496pp.
- Cohen, A.C. (1959). Simplified Estimators for the Normal Distribution When Samples are Singly Censored or Truncated. *Technometrics* **1**(3), 217–237.
- Cohen, A.C. (1963). Progressively Censored Samples in Life Testing. *Technometrics* **5**, 327–339.
- Cohen, A.C. (1991). *Truncated and Censored Samples*. Marcel Dekker, New York, New York, 312pp.
- Cox, D.R. (1970). *Analysis of Binary Data*. Chapman & Hall, London. 142pp.
- Efron, B. (1979). Bootstrap Methods: Another Look at the Jackknife. *The Annals of Statistics* **7**, 1–26.
- Efron, B., and R.J. Tibshirani. (1993). *An Introduction to the Bootstrap*. Chapman and Hall, New York, 436pp.
- El-Shaarawi, A.H. (1989). Inferences About the Mean from Censored Water Quality Data. *Water Resources Research* **25**(4) 685–690.
- El-Shaarawi, A.H., and D.M. Dolan. (1989). Maximum Likelihood Estimation of Water Quality Concentrations from Censored Data. *Canadian Journal of Fisheries and Aquatic Sciences* **46**, 1033–1039.
- El-Shaarawi, A.H., and S.R. Esterby. (1992). Replacement of Censored Observations by a Constant: An Evaluation. *Water Research* **26**(6), 835–844.
- El-Shaarawi, A.H., and A. Naderi. (1991). Statistical Inference from Multiply Censored Environmental Data. *Environmental Monitoring and Assessment* **17**, 339–347.
- Gibbons, R.D., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*, Second Edition. John Wiley & Sons, Hoboken.
- Gilliom, R.J., and D.R. Helsel. (1986). Estimation of Distributional Parameters for Censored Trace Level Water Quality Data: 1. Estimation Techniques. *Water Resources Research* **22**, 135–146.
- Gleit, A. (1985). Estimation for Small Normal Data Sets with Detection Limits. *Environmental Science and Technology* **19**, 1201–1206.
- Haas, C.N., and P.A. Scheff. (1990). Estimation of Averages in Truncated Samples. *Environmental Science and Technology* **24**(6), 912–919.
- Hashimoto, L.K., and R.R. Trussell. (1983). Evaluating Water Quality Data Near the Detection Limit. Paper presented at the Advanced Technology Conference, American Water Works Association, Las Vegas, Nevada, June 5-9, 1983.
- Helsel, D.R. (1990). Less than Obvious: Statistical Treatment of Data Below the Detection Limit. *Environmental Science and Technology* **24**(12), 1766–1774.
- Helsel, D.R. (2012). *Statistics for Censored Environmental Data Using Minitab and R, Second Edition*. John Wiley & Sons, Hoboken, New Jersey.
- Helsel, D.R., and T.A. Cohn. (1988). Estimation of Descriptive Statistics for Multiply Censored Water Quality Data. *Water Resources Research* **24**(12), 1997–2004.
- Hirsch, R.M., and J.R. Stedinger. (1987). Plotting Positions for Historical Floods and Their Precision. *Water Resources Research* **23**(4), 715–727.

- Korn, L.R., and D.E. Tyler. (2001). Robust Estimation for Chemical Concentration Data Subject to Detection Limits. In Fernholz, L., S. Morgenthaler, and W. Stahel, eds. *Statistics in Genetics and in the Environmental Sciences*. Birkhauser Verlag, Basel, pp.41–63.
- Krishnamoorthy K., and T. Mathew. (2009). *Statistical Tolerance Regions: Theory, Applications, and Computation*. John Wiley and Sons, Hoboken.
- Michael, J.R., and W.R. Schucany. (1986). Analysis of Data from Censored Samples. In D’Agostino, R.B., and M.A. Stephens, eds. *Goodness-of Fit Techniques*. Marcel Dekker, New York, 560pp, Chapter 11, 461–496.
- Millard, S.P., P. Dixon, and N.K. Neerchal. (2014; in preparation). *Environmental Statistics with R*. CRC Press, Boca Raton, Florida.
- Nelson, W. (1982). *Applied Life Data Analysis*. John Wiley and Sons, New York, 634pp.
- Newman, M.C., P.M. Dixon, B.B. Looney, and J.E. Pinder. (1989). Estimating Mean and Variance for Environmental Samples with Below Detection Limit Observations. *Water Resources Bulletin* **25**(4), 905–916.
- Pettitt, A. N. (1983). Re-Weighted Least Squares Estimation with Censored and Grouped Data: An Application of the EM Algorithm. *Journal of the Royal Statistical Society, Series B* **47**, 253–260.
- Regal, R. (1982). Applying Order Statistic Censored Normal Confidence Intervals to Time Censored Data. Unpublished manuscript, University of Minnesota, Duluth, Department of Mathematical Sciences.
- Royston, P. (2007). Profile Likelihood for Estimation and Confidence Intervals. *The Stata Journal* **7**(3), pp. 376–387.
- Saw, J.G. (1961b). The Bias of the Maximum Likelihood Estimators of Location and Scale Parameters Given a Type II Censored Normal Sample. *Biometrika* **48**, 448–451.
- Schmee, J., D.Gladstein, and W. Nelson. (1985). Confidence Limits for Parameters of a Normal Distribution from Singly Censored Samples, Using Maximum Likelihood. *Technometrics* **27**(2) 119–128.
- Schneider, H. (1986). *Truncated and Censored Samples from Normal Populations*. Marcel Dekker, New York, New York, 273pp.
- Shumway, R.H., A.S. Azari, and P. Johnson. (1989). Estimating Mean Concentrations Under Transformations for Environmental Data With Detection Limits. *Technometrics* **31**(3), 347–356.
- Stryhn, H., and J. Christensen. (2003). *Confidence Intervals by the Profile Likelihood Method, with Applications in Veterinary Epidemiology*. Contributed paper at ISVEE X (November 2003, Chile). <http://people.upei.ca/hstryhn/stryhn208.pdf>.
- Travis, C.C., and M.L. Land. (1990). Estimating the Mean of Data Sets with Nondetectable Values. *Environmental Science and Technology* **24**, 961–962.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. Chapter 15.
- USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.
- Venzon, D.J., and S.H. Moolgavkar. (1988). A Method for Computing Profile-Likelihood-Based Confidence Intervals. *Journal of the Royal Statistical Society, Series C (Applied Statistics)* **37**(1), pp. 87–94.

See Also

[LognormalAlt](#), [elnormAlt](#), [elnormCensored](#), [enormCensored](#), [estimateCensored.object](#).

Examples

```
# Chapter 15 of USEPA (2009) gives several examples of estimating the mean
# and standard deviation of a lognormal distribution on the log-scale using
# manganese concentrations (ppb) in groundwater at five background wells.
# In EnvStats these data are stored in the data frame
# EPA.09.Ex.15.1.manganese.df.
```

```
# Here we will estimate the mean and coefficient of variation
# ON THE ORIGINAL SCALE using the MLE, QMVUE,
# and imputation with Q-Q regression (also called robust ROS).
```

```
# First look at the data:
#-----
```

```
EPA.09.Ex.15.1.manganese.df
```

```
#   Sample  Well Manganese.Orig.ppb Manganese.ppb Censored
#1      1 Well.1             <5           5.0      TRUE
#2      2 Well.1             12.1          12.1     FALSE
#3      3 Well.1             16.9          16.9     FALSE
#...
#23     3 Well.5              3.3           3.3     FALSE
#24     4 Well.5              8.4           8.4     FALSE
#25     5 Well.5              <2           2.0      TRUE
```

```
longToWide(EPA.09.Ex.15.1.manganese.df,
  "Manganese.Orig.ppb", "Sample", "Well",
  paste.row.name = TRUE)
```

```
#           Well.1 Well.2 Well.3 Well.4 Well.5
#Sample.1      <5      <5      <5      6.3     17.9
#Sample.2     12.1      7.7      5.3     11.9     22.7
#Sample.3     16.9     53.6     12.6      10      3.3
#Sample.4     21.6      9.5    106.3      <2      8.4
#Sample.5      <2     45.9     34.5     77.2     <2
```

```
# Now estimate the mean and coefficient of variation
# using the MLE:
#-----
```

```
with(EPA.09.Ex.15.1.manganese.df,
  elnormAltCensored(Manganese.ppb, Censored))
```

```
#Results of Distribution Parameter Estimation
#Based on Type I Censored Data
#-----
#
#Assumed Distribution:          Lognormal
#
#Censoring Side:               left
#
```

```

#Censoring Level(s):          2 5
#
#Estimated Parameter(s):      mean = 23.003987
#                              cv   = 2.300772
#
#Estimation Method:           MLE
#
#Data:                         Manganese.ppb
#
#Censoring Variable:          Censored
#
#Sample Size:                  25
#
#Percent Censored:             24%

# Now compare the MLE with the QMVUE and the
# estimator based on imputation with Q-Q regression
#-----

with(EPA.09.Ex.15.1.manganese.df,
     elnormAltCensored(Manganese.ppb, Censored))$parameters
#      mean      cv
#23.003987  2.300772

with(EPA.09.Ex.15.1.manganese.df,
     elnormAltCensored(Manganese.ppb, Censored,
                       method = "qmvue"))$parameters
#      mean      cv
#21.566945  1.841366

with(EPA.09.Ex.15.1.manganese.df,
     elnormAltCensored(Manganese.ppb, Censored,
                       method = "impute.w.qq.reg"))$parameters
#      mean      cv
#19.886180  1.298868

#-----

# The method used to estimate quantiles for a Q-Q plot is
# determined by the argument prob.method. For the function
# elnormCensoredAlt, for any estimation method that involves
# Q-Q regression, the default value of prob.method is
# "hirsch-stedinger" and the default value for the
# plotting position constant is plot.pos.con=0.375.

# Both Helsel (2012) and USEPA (2009) also use the Hirsch-Stedinger
# probability method but set the plotting position constant to 0.

with(EPA.09.Ex.15.1.manganese.df,
     elnormAltCensored(Manganese.ppb, Censored,
                       method = "impute.w.qq.reg", plot.pos.con = 0))$parameters
#      mean      cv
#19.827673  1.304725

#-----

# Using the same data as above, compute a confidence interval

```

```
# for the mean using the profile-likelihood method.

with(EPA.09.Ex.15.1.manganese.df,
     elnormAltCensored(Manganese.ppb, Censored, ci = TRUE))

#Results of Distribution Parameter Estimation
#Based on Type I Censored Data
#-----
#
#Assumed Distribution:      Lognormal
#
#Censoring Side:           left
#
#Censoring Level(s):       2 5
#
#Estimated Parameter(s):   mean = 23.003987
#                           cv   = 2.300772
#
#Estimation Method:        MLE
#
#Data:                     Manganese.ppb
#
#Censoring Variable:       Censored
#
#Sample Size:              25
#
#Percent Censored:         24%
#
#Confidence Interval for:   mean
#
#Confidence Interval Method: Profile Likelihood
#
#Confidence Interval Type:  two-sided
#
#Confidence Level:         95%
#
#Confidence Interval:      LCL = 12.37629
#                           UCL = 69.87694
```

elnormCensored	<i>Estimate Parameters for a Lognormal Distribution (Log-Scale) Based on Type I Censored Data</i>
----------------	---

Description

Estimate the mean and standard deviation parameters of the logarithm of a [lognormal distribution](#) given a sample of data that has been subjected to Type I censoring, and optionally construct a confidence interval for the mean.

Usage

```
elnormCensored(x, censored, method = "mle", censoring.side = "left",
               ci = FALSE, ci.method = "profile.likelihood", ci.type = "two-sided",
               conf.level = 0.95, n.bootstraps = 1000, use.acc.con = FALSE,
               pivot.statistic = "z", nmc = 1000, seed = NULL, ...)
```


Arguments

<code>x</code>	numeric vector of observations. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
<code>censored</code>	numeric or logical vector indicating which values of <code>x</code> are censored. This must be the same length as <code>x</code> . If the mode of <code>censored</code> is "logical", TRUE values correspond to elements of <code>x</code> that are censored, and FALSE values correspond to elements of <code>x</code> that are not censored. If the mode of <code>censored</code> is "numeric", it must contain only 1's and 0's; 1 corresponds to TRUE and 0 corresponds to FALSE. Missing (NA) values are allowed but will be removed.
<code>method</code>	character string specifying the method of estimation. For singly censored data, the possible values are: "mle" (maximum likelihood; the default), "bcmle" (bias-corrected maximum likelihood), "qq.reg" (quantile-quantile regression), "qq.reg.w.cen.level" (quantile-quantile regression including the censoring level), "impute.w.qq.reg" (moment estimation based on imputation using the qq.reg method), "impute.w.qq.reg.w.cen.level" (moment estimation based on imputation using the qq.reg.w.cen.level method), "impute.w.mle" (moment estimation based on imputation using the mle), "iterative.impute.w.qq.reg" (moment estimation based on iterative imputation using the qq.reg method), "m.est" (robust M-estimation), and "half.cen.level" (moment estimation based on setting the censored observations to half the censoring level). For multiply censored data, the possible values are: "mle" (maximum likelihood; the default), "qq.reg" (quantile-quantile regression), "impute.w.qq.reg" (moment estimation based on imputation using the qq.reg method), and "half.cen.level" (moment estimation based on setting the censored observations to half the censoring level). See the DETAILS section for more information.
<code>censoring.side</code>	character string indicating on which side the censoring occurs. The possible values are "left" (the default) and "right".
<code>ci</code>	logical scalar indicating whether to compute a confidence interval for the mean or variance. The default value is <code>ci=FALSE</code> .
<code>ci.method</code>	character string indicating what method to use to construct the confidence interval for the mean. The possible values are "profile.likelihood" (profile likelihood; the default), "normal.approx" (normal approximation), "normal.approx.w.cov" (normal approximation taking into account the covariance between the estimated mean and standard deviation; only available for singly censored data), "gppq" (generalized pivotal quantity), and "bootstrap" (based on bootstrapping). See the DETAILS section for more information. This argument is ignored if <code>ci=FALSE</code> .
<code>ci.type</code>	character string indicating what kind of confidence interval to compute. The possible values are "two-sided" (the default), "lower", and "upper". This argument is ignored if <code>ci=FALSE</code> .
<code>conf.level</code>	a scalar between 0 and 1 indicating the confidence level of the confidence interval. The default value is <code>conf.level=0.95</code> . This argument is ignored if <code>ci=FALSE</code> .
<code>n.bootstraps</code>	numeric scalar indicating how many bootstraps to use to construct the confidence interval for the mean when <code>ci.type="bootstrap"</code> . This argument is ignored if <code>ci=FALSE</code> and/or <code>ci.method</code> does not equal "bootstrap".
<code>use.acc.con</code>	logical scalar indicating whether to use the acceleration constant when computing the bias-corrected bootstrap confidence interval (see the DETAILS sec-

tion). The default value is FALSE. This argument is ignored if `ci=FALSE` and/or `ci.method` does not equal "bootstrap".

`pivot.statistic`

character string indicating which pivot statistic to use in the construction of the confidence interval for the mean when `ci.method="normal.approx"` or `ci.method="normal.approx.w.cov"` (see the DETAILS section). The possible values are `pivot.statistic="z"` (the default) and `pivot.statistic="t"`. When `pivot.statistic="t"` you may supply the argument `ci.sample.size` (see below). The argument `pivot.statistic` is ignored if `ci=FALSE`.

`nmc`

numeric scalar indicating the number of Monte Carlo simulations to run when `ci.method="gpq"`. The default is `nmc=1000`. This argument is ignored if `ci=FALSE`.

`seed`

integer supplied to the function [set.seed](#) and used when `ci.method="bootstrap"` or `ci.method="gpq"`. The default value is `seed=NULL`, in which case the current value of `.Random.seed` is used. This argument is ignored when `ci=FALSE`.

...

additional arguments to pass to other functions.

- `prob.method`. Character string indicating what method to use to compute the plotting positions (empirical probabilities) when method is one of "qq.reg", "qq.reg.w.cen.level", "impute.w.qq.reg", "impute.w.qq.reg.w.cen.level", "impute.w.mle", or "iterative.impute.w.qq.reg". Possible values are "kaplan-meier" (product-limit method of Kaplan and Meier (1958)), "nelson" (hazard plotting method of Nelson (1972)), "michael-schucany" (generalization of the product-limit method due to Michael and Schucany (1986)), and "hirsch-stedinger" (generalization of the product-limit method due to Hirsch and Stedinger (1987)). The default value is `prob.method="michael-schucany"`. The "nelson" method is only available for `censoring.side="right"`. See the DETAILS section and the help file for [ppointsCensored](#) for more information.
- `plot.pos.con`. Numeric scalar between 0 and 1 containing the value of the plotting position constant to use when method is one of "qq.reg", "qq.reg.w.cen.level", "impute.w.qq.reg", "impute.w.qq.reg.w.cen.level", "impute.w.mle", or "iterative.impute.w.qq.reg". The default value is `plot.pos.con=0.375`. See the DETAILS section and the help file for [ppointsCensored](#) for more information.
- `ci.sample.size`. Numeric scalar indicating what sample size to assume to construct the confidence interval for the mean if `pivot.statistic="t"` and `ci.method="normal.approx"` or `ci.method="normal.approx.w.cov"`. When method equals "mle" or "bcmle", the default value is the expected number of uncensored observations, otherwise it is the observed number of uncensored observations.
- `lb.impute`. Numeric scalar indicating the lower bound for imputed observations when method is one of "impute.w.qq.reg", "impute.w.qq.reg.w.cen.level", "impute.w.mle", or "iterative.impute.w.qq.reg". Imputed values smaller than this value will be set to this value. The default is `lb.impute=-Inf`.
- `ub.impute`. Numeric scalar indicating the upper bound for imputed observations when method is one of "impute.w.qq.reg", "impute.w.qq.reg.w.cen.level", "impute.w.mle", or "iterative.impute.w.qq.reg". Imputed values larger than this value will be set to this value. The default is `ub.impute=Inf`.
- `convergence`. Character string indicating the kind of convergence criterion when method="iterative.impute.w.qq.reg". The possible values are "relative" (the default) and "absolute". See the DETAILS section for more information.

- `tol`. Numeric scalar indicating the convergence tolerance when `method="iterative.impute.w"`. The default value is `tol=1e-6`. If `convergence="relative"`, then the relative difference in the old and new estimates of the mean and the relative difference in the old and new estimates of the standard deviation must be less than `tol` for convergence to be achieved. If `convergence="absolute"`, then the absolute difference in the old and new estimates of the mean and the absolute difference in the old and new estimates of the standard deviation must be less than `tol` for convergence to be achieved.
- `max.iter`. Numeric scalar indicating the maximum number of iterations when `method="iterative.impute.w.qq.reg"`.
- `t.df`. Numeric scalar greater than or equal to 1 that determines the robustness and efficiency properties of the estimator when `method="m.est"`. The default value is `t.df=3`.

Details

If `x` or `censored` contain any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

Let X denote a random variable with a [lognormal distribution](#) with parameters $\text{meanlog}=\mu$ and $\text{sdlog}=\sigma$. Then $Y = \log(X)$ has a [normal \(Gaussian\) distribution](#) with parameters $\text{mean}=\mu$ and $\text{sd}=\sigma$. Thus, the function `elnormCensored` simply calls the function `enormCensored` using the log-transformed values of `x`.

Value

a list of class `"estimateCensored"` containing the estimated parameters and other information. See [estimateCensored.object](#) for details.

Note

A sample of data contains censored observations if some of the observations are reported only as being below or above some censoring level. In environmental data analysis, Type I left-censored data sets are common, with values being reported as "less than the detection limit" (e.g., Helsel, 2012). Data sets with only one censoring level are called *singly censored*; data sets with multiple censoring levels are called *multiply* or *progressively censored*.

Statistical methods for dealing with censored data sets have a long history in the field of survival analysis and life testing. More recently, researchers in the environmental field have proposed alternative methods of computing estimates and confidence intervals in addition to the classical ones such as maximum likelihood estimation.

Helsel (2012, Chapter 6) gives an excellent review of past studies of the properties of various estimators based on censored environmental data.

In practice, it is better to use a confidence interval for the mean or a joint confidence region for the mean and standard deviation, rather than rely on a single point-estimate of the mean. Since confidence intervals and regions depend on the properties of the estimators for both the mean and standard deviation, the results of studies that simply evaluated the performance of the mean and standard deviation separately cannot be readily extrapolated to predict the performance of various methods of constructing confidence intervals and regions. Furthermore, for several of the methods that have been proposed to estimate the mean based on type I left-censored data, standard errors of the estimates are not available, hence it is not possible to construct confidence intervals (El-Shaarawi and Dolan, 1989).

Few studies have been done to evaluate the performance of methods for constructing confidence intervals for the mean or joint confidence regions for the mean and standard deviation when data are subjected to single or multiple censoring.

Schmee et al. (1985) studied Type II censoring for a normal distribution and noted that the bias and variances of the maximum likelihood estimators are of the order $1/N$, and that the bias is negligible for $N = 100$ and as much as 90% censoring. (If the proportion of censored observations is less than 90%, the bias becomes negligible for smaller sample sizes.) For small samples with moderate to high censoring, however, the bias of the mle's causes confidence intervals based on them using a normal approximation (e.g., `method="mle"` and `ci.method="normal.approx"`) to be too short. Schmee et al. (1985) provide tables for exact confidence intervals for sample sizes up to $N = 100$ that were created based on Monte Carlo simulation. Schmee et al. (1985) state that these tables should work well for Type I censored data as well.

Shumway et al. (1989) evaluated the coverage of 90% confidence intervals for the mean based on using a Box-Cox transformation to induce normality, computing the mle's based on the normal distribution, then computing the mean in the original scale. They considered three methods of constructing confidence intervals: the delta method, the bootstrap, and the bias-corrected bootstrap. Shumway et al. (1989) used three parent distributions in their study: Normal(3,1), the square of this distribution, and the exponentiation of this distribution (i.e., a lognormal distribution). Based on sample sizes of 10 and 50 with a censoring level at the 10'th or 20'th percentile, Shumway et al. (1989) found that the delta method performed quite well and was superior to the bootstrap method.

Millard et al. (2014; in preparation) show that the coverage of profile likelihood method is excellent.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Bain, L.J., and M. Engelhardt. (1991). *Statistical Analysis of Reliability and Life-Testing Models*. Marcel Dekker, New York, 496pp.
- Cohen, A.C. (1959). Simplified Estimators for the Normal Distribution When Samples are Singly Censored or Truncated. *Technometrics* **1**(3), 217–237.
- Cohen, A.C. (1963). Progressively Censored Samples in Life Testing. *Technometrics* **5**, 327–339
- Cohen, A.C. (1991). *Truncated and Censored Samples*. Marcel Dekker, New York, New York, 312pp.
- Cox, D.R. (1970). *Analysis of Binary Data*. Chapman & Hall, London. 142pp.
- Efron, B. (1979). Bootstrap Methods: Another Look at the Jackknife. *The Annals of Statistics* **7**, 1–26.
- Efron, B., and R.J. Tibshirani. (1993). *An Introduction to the Bootstrap*. Chapman and Hall, New York, 436pp.
- El-Shaarawi, A.H. (1989). Inferences About the Mean from Censored Water Quality Data. *Water Resources Research* **25**(4) 685–690.
- El-Shaarawi, A.H., and D.M. Dolan. (1989). Maximum Likelihood Estimation of Water Quality Concentrations from Censored Data. *Canadian Journal of Fisheries and Aquatic Sciences* **46**, 1033–1039.
- El-Shaarawi, A.H., and S.R. Esterby. (1992). Replacement of Censored Observations by a Constant: An Evaluation. *Water Research* **26**(6), 835–844.
- El-Shaarawi, A.H., and A. Naderi. (1991). Statistical Inference from Multiply Censored Environmental Data. *Environmental Monitoring and Assessment* **17**, 339–347.

- Gibbons, R.D., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*, Second Edition. John Wiley & Sons, Hoboken.
- Gilliom, R.J., and D.R. Helsel. (1986). Estimation of Distributional Parameters for Censored Trace Level Water Quality Data: 1. Estimation Techniques. *Water Resources Research* **22**, 135–146.
- Gleit, A. (1985). Estimation for Small Normal Data Sets with Detection Limits. *Environmental Science and Technology* **19**, 1201–1206.
- Haas, C.N., and P.A. Scheff. (1990). Estimation of Averages in Truncated Samples. *Environmental Science and Technology* **24**(6), 912–919.
- Hashimoto, L.K., and R.R. Trussell. (1983). Evaluating Water Quality Data Near the Detection Limit. Paper presented at the Advanced Technology Conference, American Water Works Association, Las Vegas, Nevada, June 5-9, 1983.
- Helsel, D.R. (1990). Less than Obvious: Statistical Treatment of Data Below the Detection Limit. *Environmental Science and Technology* **24**(12), 1766–1774.
- Helsel, D.R. (2012). *Statistics for Censored Environmental Data Using Minitab and R, Second Edition*. John Wiley & Sons, Hoboken, New Jersey.
- Helsel, D.R., and T.A. Cohn. (1988). Estimation of Descriptive Statistics for Multiply Censored Water Quality Data. *Water Resources Research* **24**(12), 1997–2004.
- Hirsch, R.M., and J.R. Stedinger. (1987). Plotting Positions for Historical Floods and Their Precision. *Water Resources Research* **23**(4), 715–727.
- Korn, L.R., and D.E. Tyler. (2001). Robust Estimation for Chemical Concentration Data Subject to Detection Limits. In Fernholz, L., S. Morgenthaler, and W. Stahel, eds. *Statistics in Genetics and in the Environmental Sciences*. Birkhauser Verlag, Basel, pp.41–63.
- Krishnamoorthy K., and T. Mathew. (2009). *Statistical Tolerance Regions: Theory, Applications, and Computation*. John Wiley and Sons, Hoboken.
- Michael, J.R., and W.R. Schucany. (1986). Analysis of Data from Censored Samples. In D'Agostino, R.B., and M.A. Stephens, eds. *Goodness-of Fit Techniques*. Marcel Dekker, New York, 560pp, Chapter 11, 461–496.
- Millard, S.P., P. Dixon, and N.K. Neerchal. (2014; in preparation). *Environmental Statistics with R*. CRC Press, Boca Raton, Florida.
- Nelson, W. (1982). *Applied Life Data Analysis*. John Wiley and Sons, New York, 634pp.
- Newman, M.C., P.M. Dixon, B.B. Looney, and J.E. Pinder. (1989). Estimating Mean and Variance for Environmental Samples with Below Detection Limit Observations. *Water Resources Bulletin* **25**(4), 905–916.
- Pettitt, A. N. (1983). Re-Weighted Least Squares Estimation with Censored and Grouped Data: An Application of the EM Algorithm. *Journal of the Royal Statistical Society, Series B* **47**, 253–260.
- Regal, R. (1982). Applying Order Statistic Censored Normal Confidence Intervals to Time Censored Data. Unpublished manuscript, University of Minnesota, Duluth, Department of Mathematical Sciences.
- Royston, P. (2007). Profile Likelihood for Estimation and Confidence Intervals. *The Stata Journal* **7**(3), pp. 376–387.
- Saw, J.G. (1961b). The Bias of the Maximum Likelihood Estimators of Location and Scale Parameters Given a Type II Censored Normal Sample. *Biometrika* **48**, 448–451.
- Schmee, J., D.Gladstein, and W. Nelson. (1985). Confidence Limits for Parameters of a Normal Distribution from Singly Censored Samples, Using Maximum Likelihood. *Technometrics* **27**(2) 119–128.

- Schneider, H. (1986). *Truncated and Censored Samples from Normal Populations*. Marcel Dekker, New York, New York, 273pp.
- Shumway, R.H., A.S. Azari, and P. Johnson. (1989). Estimating Mean Concentrations Under Transformations for Environmental Data With Detection Limits. *Technometrics* **31**(3), 347–356.
- Stryhn, H., and J. Christensen. (2003). *Confidence Intervals by the Profile Likelihood Method, with Applications in Veterinary Epidemiology*. Contributed paper at ISVEE X (November 2003, Chile). <http://people.upei.ca/hstryhn/stryhn208.pdf>.
- Travis, C.C., and M.L. Land. (1990). Estimating the Mean of Data Sets with Nondetectable Values. *Environmental Science and Technology* **24**, 961–962.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. Chapter 15.
- USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.
- Venzon, D.J., and S.H. Moolgavkar. (1988). A Method for Computing Profile-Likelihood-Based Confidence Intervals. *Journal of the Royal Statistical Society, Series C (Applied Statistics)* **37**(1), pp. 87–94.

See Also

[enormCensored](#), [Lognormal](#), [elnorm](#), [estimateCensored.object](#).

Examples

```
# Chapter 15 of USEPA (2009) gives several examples of estimating the mean
# and standard deviation of a lognormal distribution on the log-scale using
# manganese concentrations (ppb) in groundwater at five background wells.
# In EnvStats these data are stored in the data frame
# EPA.09.Ex.15.1.manganese.df.

# Here we will estimate the mean and standard deviation using the MLE,
# Q-Q regression (also called parametric regression on order statistics
# or ROS; e.g., USEPA, 2009 and Helsel, 2012), and imputation with Q-Q
# regression (also called robust ROS).

# First look at the data:
#-----

EPA.09.Ex.15.1.manganese.df

#   Sample  Well Manganese.Orig.ppb Manganese.ppb Censored
#1      1 Well.1          <5          5.0      TRUE
#2      2 Well.1         12.1         12.1     FALSE
#3      3 Well.1         16.9         16.9     FALSE
#...
#23     3 Well.5          3.3          3.3     FALSE
#24     4 Well.5          8.4          8.4     FALSE
#25     5 Well.5          <2          2.0      TRUE

longToWide(EPA.09.Ex.15.1.manganese.df,
           "Manganese.Orig.ppb", "Sample", "Well",
```

```

paste.row.name = TRUE)

#           Well.1 Well.2 Well.3 Well.4 Well.5
#Sample.1    <5    <5    <5    6.3   17.9
#Sample.2   12.1    7.7    5.3   11.9   22.7
#Sample.3   16.9   53.6   12.6    10    3.3
#Sample.4   21.6    9.5  106.3    <2    8.4
#Sample.5    <2   45.9   34.5   77.2    <2

# Now estimate the mean and standard deviation on the log-scale
# using the MLE:
#-----

with(EPA.09.Ex.15.1.manganese.df,
     eInormCensored(Manganese.ppb, Censored))

#Results of Distribution Parameter Estimation
#Based on Type I Censored Data
#-----
#
#Assumed Distribution:          Lognormal
#
#Censoring Side:                left
#
#Censoring Level(s):           2 5
#
#Estimated Parameter(s):       meanlog = 2.215905
#                               sdlog   = 1.356291
#
#Estimation Method:            MLE
#
#Data:                          Manganese.ppb
#
#Censoring Variable:           Censored
#
#Sample Size:                  25
#
#Percent Censored:             24%

# Now compare the MLE with the estimators based on
# Q-Q regression and imputation with Q-Q regression
#-----

with(EPA.09.Ex.15.1.manganese.df,
     eInormCensored(Manganese.ppb, Censored))$parameters
# meanlog    sdlog
#2.215905 1.356291

with(EPA.09.Ex.15.1.manganese.df,
     eInormCensored(Manganese.ppb, Censored,
                     method = "qq.reg"))$parameters
# meanlog    sdlog
#2.293742 1.283635

with(EPA.09.Ex.15.1.manganese.df,
     eInormCensored(Manganese.ppb, Censored,

```

```

    method = "impute.w.qq.reg"))$parameters
# meanlog    sdlog
#2.298656 1.238104

#-----

# The method used to estimate quantiles for a Q-Q plot is
# determined by the argument prob.method. For the functions
# enormCensored and elnormCensored, for any estimation
# method that involves Q-Q regression, the default value of
# prob.method is "hirsch-stedinger" and the default value for the
# plotting position constant is plot.pos.con=0.375.

# Both Helsel (2012) and USEPA (2009) also use the Hirsch-Stedinger
# probability method but set the plotting position constant to 0.

with(EPA.09.Ex.15.1.manganese.df,
     elnormCensored(Manganese.ppb, Censored,
                    method = "impute.w.qq.reg", plot.pos.con = 0))$parameters
# meanlog    sdlog
#2.277175 1.261431

#-----

# Using the same data as above, compute a confidence interval
# for the mean on the log-scale using the profile-likelihood
# method.

with(EPA.09.Ex.15.1.manganese.df,
     elnormCensored(Manganese.ppb, Censored, ci = TRUE))

#Results of Distribution Parameter Estimation
#Based on Type I Censored Data
#-----
#
#Assumed Distribution:      Lognormal
#
#Censoring Side:           left
#
#Censoring Level(s):       2 5
#
#Estimated Parameter(s):   meanlog = 2.215905
#                           sdlog   = 1.356291
#
#Estimation Method:        MLE
#
#Data:                     Manganese.ppb
#
#Censoring Variable:       Censored
#
#Sample Size:              25
#
#Percent Censored:         24%
#
#Confidence Interval for:   meanlog
#
#Confidence Interval Method: Profile Likelihood

```



```
#
#Confidence Interval Type:      two-sided
#
#Confidence Level:             95%
#
#Confidence Interval:          LCL = 1.595062
#                               UCL = 2.771197
```

elogis

Estimate Parameters of a Logistic Distribution

Description

Estimate the location and scale parameters of a [logistic distribution](#), and optionally construct a confidence interval for the location parameter.

Usage

```
elogis(x, method = "mle", ci = FALSE, ci.type = "two-sided",
       ci.method = "normal.approx", conf.level = 0.95)
```

Arguments

x	numeric vector of observations.
method	character string specifying the method of estimation. Possible values are "mle" (maximum likelihood; the default), "mme" (methods of moments), and "mmue" (method of moments based on the unbiased estimator of variance). See the DETAILS section for more information on these estimation methods.
ci	logical scalar indicating whether to compute a confidence interval for the location or scale parameter. The default value is FALSE.
ci.type	character string indicating what kind of confidence interval to compute. The possible values are "two-sided" (the default), "lower", and "upper". This argument is ignored if ci=FALSE.
ci.method	character string indicating what method to use to construct the confidence interval for the location or scale parameter. Currently, the only possible value is "normal.approx" (the default). See the DETAILS section for more information. This argument is ignored if ci=FALSE.
conf.level	a scalar between 0 and 1 indicating the confidence level of the confidence interval. The default value is conf.level=0.95. This argument is ignored if ci=FALSE.

Details

If x contains any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

Let $\underline{x} = (x_1, x_2, \dots, x_n)$ be a vector of n observations from an [logistic distribution](#) with parameters location= η and scale= θ .

Estimation

Maximum Likelihood Estimation (method="mle")

The maximum likelihood estimators (mle's) of η and θ are the solutions of the simultaneous equations (Forbes et al., 2011):

$$\sum_{i=1}^n \frac{1}{1 + e^{z_i}} = \frac{n}{2} \quad (1)$$

$$\sum_{i=1}^n z_i \left[\frac{1 - e^{z_i}}{1 + e^{z_i}} \right] = n \quad (2)$$

where

$$z_i = \frac{x_i - \hat{e}ta_{mle}}{\hat{\theta}_{mle}} \quad (3)$$

Method of Moments Estimation (method="mme")

The method of moments estimators (mme's) of η and θ are given by:

$$\hat{\eta}_{mme} = \bar{x} \quad (4)$$

$$\hat{\theta}_{mme} = \frac{\sqrt{3}}{\pi} s_m \quad (5)$$

where

$$\bar{x} = \sum_{i=1}^n x_i \quad (6)$$

$$s_m^2 = \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2 \quad (7)$$

that is, s_m denotes the square root of the method of moments estimator of variance.

Method of Moments Estimators Based on the Unbiased Estimator of Variance (method="mmue")

These estimators are exactly the same as the method of moments estimators given in equations (4-7) above, except that the method of moments estimator of variance in equation (7) is replaced with the unbiased estimator of variance:

$$s^2 = \frac{1}{n-1} \sum_{i=1}^n (x_i - \bar{x})^2 \quad (8)$$

Confidence Intervals

When ci=TRUE, an approximate $(1 - \alpha)100\%$ confidence intervals for η can be constructed assuming the distribution of the estimator of η is approximately normally distributed. A two-sided confidence interval is constructed as:

$$[\hat{\eta} - t(n-1, 1 - \alpha/2) \hat{\sigma}_{\hat{\eta}}, \hat{\eta} + t(n-1, 1 - \alpha/2) \hat{\sigma}_{\hat{\eta}}]$$

where $t(\nu, p)$ is the p 'th quantile of [Student's t-distribution](#) with ν degrees of freedom, and the quantity

$$\hat{\sigma}_{\hat{\eta}} = \frac{\pi \hat{\theta}}{\sqrt{3n}} \quad (9)$$

denotes the estimated asymptotic standard deviation of the estimator of η .

One-sided confidence intervals for η and θ are computed in a similar fashion.

Value

a list of class "estimate" containing the estimated parameters and other information. See [estimate.object](#) for details.

Note

The [logistic distribution](#) is defined on the real line and is unimodal and symmetric about its location parameter (the mean). It has longer tails than a normal (Gaussian) distribution. It is used to model growth curves and bioassay data.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.

Johnson, N. L., S. Kotz, and N. Balakrishnan. (1995). *Continuous Univariate Distributions, Volume 2*. Second Edition. John Wiley and Sons, New York.

See Also

[Logistic](#).

Examples

```
# Generate 20 observations from a logistic distribution with
# parameters location=0 and scale=1, then estimate the parameters
# and construct a 90% confidence interval for the location parameter.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
dat <- rlogis(20)
elogs(dat, ci = TRUE, conf.level = 0.9)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:      Logistic
#
#Estimated Parameter(s):   location = -0.2181845
#                           scale    =  0.8152793
#
#Estimation Method:        mle
#
#Data:                     dat
#
#Sample Size:              20
#
#Confidence Interval for:   location
#
#Confidence Interval Method: Normal Approximation
#                           (t Distribution)
#
```

```

#Confidence Interval Type:      two-sided
#
#Confidence Level:             90%
#
#Confidence Interval:          LCL = -0.7899382
#                               UCL =  0.3535693

#-----

# Clean up
#-----
rm(dat)

```

Empirical

The Empirical Distribution Based on a Set of Observations

Description

Density, distribution function, quantile function, and random generation for the empirical distribution based on a set of observations

Usage

```

demp(x, obs, discrete = FALSE, density.arg.list = NULL)
pemp(q, obs, discrete = FALSE,
     prob.method = ifelse(discrete, "emp.probs", "plot.pos"),
     plot.pos.con = 0.375)
qemp(p, obs, discrete = FALSE,
     prob.method = ifelse(discrete, "emp.probs", "plot.pos"),
     plot.pos.con = 0.375)
remp(n, obs)

```

Arguments

x	vector of quantiles.
q	vector of quantiles.
p	vector of probabilities between 0 and 1.
n	sample size. If length(n) is larger than 1, then length(n) random values are returned.
obs	numeric vector of observations. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
discrete	logical scalar indicating whether the assumed parent distribution of x is discrete (discrete=TRUE) or continuous (discrete=FALSE). The default value is FALSE.
density.arg.list	list with arguments to the R density function. The default value is NULL. (See the help file for density for more information on the arguments to density.) The argument density.arg.list is ignored if discrete=TRUE.
prob.method	character string indicating what method to use to compute the empirical probabilities. Possible values are "emp.probs" (empirical probabilities, default if discrete=TRUE) and "plot.pos" (plotting positions, default if discrete=FALSE). See the DETAILS section for more explanation.

`plot.pos.con` numeric scalar between 0 and 1 containing the value of the plotting position constant. The default value is `plot.pos.con=0.375`. See the DETAILS section for more information. This argument is ignored if `prob.method="emp.probs"`.

Details

Let x_1, x_2, \dots, x_n denote a random sample of n observations from some unknown probability distribution (i.e., the elements of the argument `obs`), and let $x_{(i)}$ denote the i^{th} order statistic, that is, the i^{th} largest observation, for $i = 1, 2, \dots, n$.

Estimating Density

The function `demp` computes the empirical probability density function. If the observations are assumed to come from a discrete distribution, the probability density (mass) function is estimated by:

$$\hat{f}(x) = \widehat{Pr}(X = x) = \frac{\sum_{i=1}^n I_{[x]}(x_i)}{n}$$

where I is the indicator function:

$$I_{[x]}(y) = \begin{cases} 1 & \text{if } y = x, \\ 0 & \text{if } y \neq x \end{cases}$$

That is, the estimated probability of observing the value x is simply the observed proportion of observations equal to x .

If the observations are assumed to come from a continuous distribution, the function `demp` calls the R function `density` to compute the estimated density based on the values specified in the argument `obs`, and then uses linear interpolation to estimate the density at the values specified in the argument `x`. See the R help file for `density` for more information on how the empirical density is computed in the continuous case.

Estimating Probabilities

The function `pemp` computes the estimated cumulative distribution function (cdf), also called the empirical cdf (ecdf). If the observations are assumed to come from a discrete distribution, the value of the cdf evaluated at the i^{th} order statistic is usually estimated by:

$$\hat{F}[x_{(i)}] = \widehat{Pr}(X \leq x_{(i)}) = \hat{p}_i = \frac{\sum_{j=1}^n I_{(-\infty, x_{(i)}]}(x_j)}{n}$$

where:

$$I_{(-\infty, x]}(y) = \begin{cases} 1 & \text{if } y \leq x, \\ 0 & \text{if } y > x \end{cases}$$

(D'Agostino, 1986a). That is, the estimated value of the cdf at the i^{th} order statistic is simply the observed proportion of observations less than or equal to the i^{th} order statistic. This estimator is sometimes called the “empirical probabilities” estimator and is intuitively appealing. The function `pemp` uses the above equations to compute the empirical cdf when `prob.method="emp.probs"`.

For any general value of x , when the observations are assumed to come from a discrete distribution, the value of the cdf is estimated by:

$$\hat{F}(x) = \begin{cases} 0 & \text{if } x < x_{(1)}, \\ \hat{p}_i & \text{if } x_{(i)} \leq x < x_{(i+1)}, \\ 1 & \text{if } x \geq x_{(n)} \end{cases}$$

The function `pemp` uses the above equation when `discrete=TRUE`.

If the observations are assumed to come from a continuous distribution, the value of the cdf evaluated at the i^{th} order statistic is usually estimated by:

$$\hat{F}[x_{(i)}] = \hat{p}_i = \frac{i - a}{n - 2a + 1}$$

where a denotes the plotting position constant and $0 \leq a \leq 1$ (Cleveland, 1993, p.18; D'Agostino, 1986a, pp.8,25). The estimators defined by the above equation are called *plotting positions* and are used to construct **probability plots**. The function `pemp` uses the above equation when `prob.method="plot.pos"`.

For any general value of x , the value of the cdf is estimated by linear interpolation:

$$\hat{F}(x) = \begin{cases} \hat{p}_1 & \text{if } x < x_{(1)}, \\ (1-r)\hat{p}_i + r\hat{p}_{i+1} & \text{if } x_{(i)} \leq x < x_{(i+1)}, \\ \hat{p}_n & \text{if } x \geq x_{(n)} \end{cases}$$

where

$$r = \frac{x - x_{(i)}}{x_{(i+1)} - x_{(i)}}$$

(Chambers et al., 1983). The function `pemp` uses the above two equations when `discrete=FALSE`.

Estimating Quantiles

The function `qemp` computes the estimated quantiles based on the observed data. If the observations are assumed to come from a discrete distribution, the p^{th} quantile is usually estimated by:

$$\hat{x}_p = \begin{cases} x_{(1)} & \text{if } p \leq \hat{p}_1, \\ x_{(i)} & \text{if } \hat{p}_{i-1} < p \leq \hat{p}_i, \\ x_n & \text{if } p > \hat{p}_n \end{cases}$$

The function `qemp` uses the above equation when `discrete=TRUE`.

If the observations are assumed to come from a continuous distribution, the p^{th} quantile is usually estimated by linear interpolation:

$$\hat{x}_p = \begin{cases} x_{(1)} & \text{if } p \leq \hat{p}_1, \\ (1-r)x_{(i-1)} + rx_{(i)} & \text{if } \hat{p}_{i-1} < p \leq \hat{p}_i, \\ x_n & \text{if } p > \hat{p}_n \end{cases}$$

where

$$r = \frac{p - \hat{p}_{i-1}}{\hat{p}_i - \hat{p}_{i-1}}$$

The function `qemp` uses the above two equations when `discrete=FALSE`.

Generating Random Numbers From the Empirical Distribution

The function `rem` simply calls the R function `sample` to sample the elements of `obs` with replacement.

Value

density (`demp`), probability (`pemp`), quantile (`qemp`), or random sample (`rem`) for the empirical distribution based on the data contained in the vector `obs`.

Note

The function `demp` let's you perform nonparametric density estimation. The function `pemp` computes the value of the empirical cumulative distribution function (`ecdf`) for user-specified quantiles. The `ecdf` is a nonparametric estimate of the true cdf (see [ecdfPlot](#)). The function `qemp` computes nonparametric estimates of quantiles (see the help files for [eqnpar](#) and [quantile](#)). The function `remp` let's you sample a set of observations with replacement, which is often done while bootstrapping or performing some other kind of Monte Carlo simulation.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Chambers, J.M., W.S. Cleveland, B. Kleiner, and P.A. Tukey. (1983). *Graphical Methods for Data Analysis*. Duxbury Press, Boston, MA, pp.11–16.
- Cleveland, W.S. (1993). *Visualizing Data*. Hobart Press, Summit, New Jersey, 360pp.
- D'Agostino, R.B. (1986a). Graphical Analysis. In: D'Agostino, R.B., and M.A. Stephens, eds. *Goodness-of-Fit Techniques*. Marcel Dekker, New York, Chapter 2, pp.7–62.
- Scott, D. W. (1992). *Multivariate Density Estimation: Theory, Practice and Visualization*. John Wiley and Sons, New York.
- Sheather, S. J. and Jones M. C. (1991). A Reliable Data-Based Bandwidth Selection Method for Kernel Density Estimation. *Journal of the Royal Statistical Society B*, 683–690.
- Silverman, B.W. (1986). *Density Estimation for Statistics and Data Analysis*. Chapman and Hall, London.
- Wegman, E.J. (1972). Nonparametric Probability Density Estimation. *Technometrics* **14**, 533-546.

See Also

[density](#), [approx](#), [epdfPlot](#), [ecdfPlot](#), [cdfCompare](#), [qqplot](#), [eqnpar](#), [quantile](#), [sample](#), [simulateVector](#), [simulateMvMatrix](#).

Examples

```
# Create a set of 100 observations from a gamma distribution with
# parameters shape=4 and scale=5.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(3)
obs <- rgamma(100, shape=4, scale=5)

# Now plot the empirical distribution (with a histogram) and the true distribution:

windows()
hist(obs, col = "cyan", xlim = c(0, 65), freq = FALSE,
     ylab = "Relative Frequency")

pdfPlot('gamma', list(shape = 4, scale = 5), add = TRUE)

box()

# Now plot the empirical distribution (based on demp) with the
```

```

# true distribution:

x <- qemp(p = seq(0, 1, len = 100), obs = obs)
y <- demp(x, obs)

windows()
plot(x, y, xlim = c(0, 65), type = "n",
     xlab = "Value of Random Variable",
     ylab = "Relative Frequency")
lines(x, y, lwd = 2, col = "cyan")

pdfPlot('gamma', list(shape = 4, scale = 5), add = TRUE)

# Alternatively, you can create the above plot with the function
# epdfPlot:

windows()
epdfPlot(obs, xlim = c(0, 65), epdf.col = "cyan",
         xlab = "Value of Random Variable",
         main = "Empirical and Theoretical PDFs")

pdfPlot('gamma', list(shape = 4, scale = 5), add = TRUE)

# Clean Up
#-----
rm(obs, x, y)

```

enbinom

Estimate Probability Parameter of a Negative Binomial Distribution

Description

Estimate the probability parameter of a [negative binomial distribution](#).

Usage

```
enbinom(x, size, method = "mle/mme")
```

Arguments

x	vector of non-negative integers indicating the number of trials that took place <i>before</i> size “successes” occurred. (The total number of trials that took place is $x+1$). Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed. If $\text{length}(x)=n$ and n is greater than 1, it is assumed that x represents observations from n separate negative binomial experiments that all had the same probability of success (prob), but possibly different values of size.
---	--

size	vector of positive integers indicating the number of “successes” that must be observed before the trials are stopped. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed. The length of size must be 1 or else the same length as x.
method	character string specifying the method of estimation. Possible values are "mle/mme" (maximum likelihood and method of moments; the default) and "mvue" (minimum variance unbiased). You cannot use method="mvue" if the sum of the elements in size is 1. See the DETAILS section for more information on these estimation methods.

Details

If x contains any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

Let $\underline{x} = (x_1, x_2, \dots, x_n)$ be a vector of n independent observations from [negative binomial distributions](#) with parameters $\text{prob}=p$ and $\text{size}=\underline{k}$, where $\underline{k} = c(k_1, k_2, \dots, k_n)$ is a vector of n (possibly different) values.

It can be shown (e.g., Forbes et al., 2011) that if X is defined as:

$$X = \sum_{i=1}^n x_i$$

then X is an observation from a [negative binomial distribution](#) with parameters $\text{prob}=p$ and $\text{size}=K$, where

$$K = \sum_{i=1}^n k_i$$

Estimation

The maximum likelihood and method of moments estimator (mle/mme) of p is given by:

$$\hat{p}_{mle} = \frac{K}{X + K}$$

and the minimum variance unbiased estimator (mvue) of p is given by:

$$\hat{p}_{mvue} = \frac{K - 1}{X + K - 1}$$

(Forbes et al., 2011). Note that the mvue of p is not defined for $K = 1$.

Value

a list of class "estimate" containing the estimated parameters and other information. See [estimate.object](#) for details.

Note

The [negative binomial distribution](#) has its roots in a gambling game where participants would bet on the number of tosses of a coin necessary to achieve a fixed number of heads. The negative binomial distribution has been applied in a wide variety of fields, including accident statistics, birth-and-death processes, and modeling spatial distributions of biological organisms.

The [geometric distribution](#) with parameter $\text{prob}=p$ is a special case of the negative binomial distribution with parameters $\text{size}=1$ and $\text{prob}=p$.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.

Johnson, N. L., S. Kotz, and A. Kemp. (1992). *Univariate Discrete Distributions*. Second Edition. John Wiley and Sons, New York, Chapter 5.

See Also

[NegBinomial](#), [egeom](#), [Geometric](#).

Examples

```
# Generate an observation from a negative binomial distribution with
# parameters size=2 and prob=0.2, then estimate the parameter prob.
# Note: the call to set.seed simply allows you to reproduce this example.
# Also, the only parameter that is estimated is prob; the parameter
# size is supplied in the call to enbinom. The parameter size is printed in
# order to show all of the parameters associated with the distribution.

set.seed(250)
dat <- rnbinom(1, size = 2, prob = 0.2)
dat
#[1] 5

enbinom(dat, size = 2)
#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Negative Binomial
#
#Estimated Parameter(s):      size = 2.0000000
#                              prob = 0.2857143
#
#Estimation Method:           mle/mme for 'prob'
#
#Data:                        dat, 2
#
#Sample Size:                  1

#-----

# Generate 3 observations from negative binomial distributions with
# parameters size=c(2,3,4) and prob=0.2, then estimate the parameter
# prob using the mvue.
# (Note: the call to set.seed simply allows you to reproduce this example.)

size.vec <- 2:4
set.seed(250)
dat <- rnbinom(3, size = size.vec, prob = 0.2)
dat
#[1] 5 19 12
```

```

enbinom(dat, size = size.vec, method = "mvue")
#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:      Negative Binomial
#
#Estimated Parameter(s):   size = 9.0000000
#                           prob = 0.1818182
#
#Estimation Method:        mvue for 'prob'
#
#Data:                     dat, size.vec
#
#Sample Size:              3
#-----

# Clean up
#-----
rm(dat)

```

enorm

Estimate Parameters of a Normal (Gaussian) Distribution

Description

Estimate the mean and standard deviation parameters of a [normal \(Gaussian\) distribution](#), and optionally construct a confidence interval for the mean or the variance.

Usage

```

enorm(x, method = "mvue", ci = FALSE, ci.type = "two-sided",
      ci.method = "exact", conf.level = 0.95, ci.param = "mean")

```

Arguments

<code>x</code>	numeric vector of observations.
<code>method</code>	character string specifying the method of estimation. Possible values are "mvue" (minimum variance unbiased; the default), and "mle/mme" (maximum likelihood/method of moments). See the DETAILS section for more information on these estimation methods.
<code>ci</code>	logical scalar indicating whether to compute a confidence interval for the mean or variance. The default value is FALSE.
<code>ci.type</code>	character string indicating what kind of confidence interval to compute. The possible values are "two-sided" (the default), "lower", and "upper". This argument is ignored if <code>ci=FALSE</code> .
<code>ci.method</code>	character string indicating what method to use to construct the confidence interval for the mean or variance. The only possible value is "exact" (the default). See the DETAILS section for more information. This argument is ignored if <code>ci=FALSE</code> .

<code>conf.level</code>	a scalar between 0 and 1 indicating the confidence level of the confidence interval. The default value is <code>conf.level=0.95</code> . This argument is ignored if <code>ci=FALSE</code> .
<code>ci.param</code>	character string indicating which parameter to create a confidence interval for. The possible values are <code>ci.param="mean"</code> (the default) and <code>ci.param="variance"</code> . This argument is ignored if <code>ci=FALSE</code> .

Details

If `x` contains any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

Let $\underline{x} = (x_1, x_2, \dots, x_n)$ be a vector of n observations from an [normal \(Gaussian\) distribution](#) with parameters $\text{mean}=\mu$ and $\text{sd}=\sigma$.

Estimation

Minimum Variance Unbiased Estimation (method="mvue")

The minimum variance unbiased estimators (mvue's) of the mean and variance are:

$$\hat{\mu}_{mvue} = \bar{x} = \frac{1}{n} \sum_{i=1}^n x_i \quad (1)$$

$$\hat{\sigma}_{mvue}^2 = s^2 = \frac{1}{n-1} \sum_{i=1}^n (x_i - \bar{x})^2 \quad (2)$$

(Johnson et al., 1994; Forbes et al., 2011). Note that when `method="mvue"`, the estimated standard deviation is the square root of the mvue of the variance, but is not itself an mvue.

Maximum Likelihood/Method of Moments Estimation (method="mle/mme")

The maximum likelihood estimator (mle) and method of moments estimator (mme) of the mean are both the same as the mvue of the mean given in equation (1) above. The mle and mme of the variance is given by:

$$\hat{\sigma}_{mle}^2 = s_m^2 = \frac{n-1}{n} s^2 = \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2 \quad (3)$$

When `method="mle/mme"`, the estimated standard deviation is the square root of the mle of the variance, and is itself an mle.

Confidence Intervals

Confidence Interval for the Mean (ci.param="mean")

When `ci=TRUE` and `ci.param = "mean"`, the usual confidence interval for μ is constructed as follows. If `ci.type="two-sided"`, a the $(1 - \alpha)100\%$ confidence interval for μ is given by:

$$\left[\hat{\mu} - t(n-1, 1 - \alpha/2) \frac{\hat{\sigma}}{\sqrt{n}}, \hat{\mu} + t(n-1, 1 - \alpha/2) \frac{\hat{\sigma}}{\sqrt{n}} \right] \quad (4)$$

where $t(\nu, p)$ is the p 'th quantile of [Student's t-distribution](#) with ν degrees of freedom (Zar, 2010; Gilbert, 1987; Ott, 1995; Helsel and Hirsch, 1992).

If `ci.type="lower"`, the $(1 - \alpha)100\%$ confidence interval for μ is given by:

$$\left[\hat{\mu} - t(n-1, 1 - \alpha) \frac{\hat{\sigma}}{\sqrt{n}}, \infty \right] \quad (5)$$

and if `ci.type="upper"`, the confidence interval is given by:

$$[-\infty, \hat{\mu} + t(n-1, 1-\alpha/2) \frac{\hat{\sigma}}{\sqrt{n}}] \quad (6)$$

Confidence Interval for the Variance (`ci.param="variance"`)

When `ci=TRUE` and `ci.param = "variance"`, the usual confidence interval for σ^2 is constructed as follows. A two-sided $(1-\alpha)100\%$ confidence interval for σ^2 is given by:

$$\left[\frac{(n-1)s^2}{\chi_{n-1,1-\alpha/2}^2}, \frac{(n-1)s^2}{\chi_{n-1,\alpha/2}^2} \right] \quad (7)$$

Similarly, a one-sided upper $(1-\alpha)100\%$ confidence interval for the population variance is given by:

$$\left[0, \frac{(n-1)s^2}{\chi_{n-1,\alpha}^2} \right] \quad (8)$$

and a one-sided lower $(1-\alpha)100\%$ confidence interval for the population variance is given by:

$$\left[\frac{(n-1)s^2}{\chi_{n-1,1-\alpha}^2}, \infty \right] \quad (9)$$

(van Belle et al., 2004; Zar, 2010).

Value

a list of class "estimate" containing the estimated parameters and other information. See [estimate.object](#) for details.

Note

The normal and lognormal distribution are probably the two most frequently used distributions to model environmental data. In order to make any kind of probability statement about a normally-distributed population (of chemical concentrations for example), you have to first estimate the mean and standard deviation (the population parameters) of the distribution. Once you estimate these parameters, it is often useful to characterize the uncertainty in the estimate of the mean or variance. This is done with confidence intervals.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Berthouex, P.M., and L.C. Brown. (2002). *Statistics for Environmental Engineers*. Second Edition. Lewis Publishers, Boca Raton, FL.
- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.
- Gilbert, R.O. (1987). *Statistical Methods for Environmental Pollution Monitoring*. Van Nostrand Reinhold, New York, NY.
- Helsel, D.R., and R.M. Hirsch. (1992). *Statistical Methods in Water Resources Research*. Elsevier, New York, NY, Chapter 7.
- Johnson, N. L., S. Kotz, and N. Balakrishnan. (1994). *Continuous Univariate Distributions, Volume 1*. Second Edition. John Wiley and Sons, New York.

Millard, S.P., and N.K. Neerchal. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, FL.

Ott, W.R. (1995). *Environmental Statistics and Data Analysis*. Lewis Publishers, Boca Raton, FL.

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.

van Belle, G., L.D. Fisher, Heagerty, P.J., and Lumley, T. (2004). *Biostatistics: A Methodology for the Health Sciences, 2nd Edition*. John Wiley & Sons, New York.

Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ.

See Also

[Normal](#).

Examples

```
# Generate 20 observations from a N(3, 2) distribution, then estimate
# the parameters and create a 95% confidence interval for the mean.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
dat <- rnorm(20, mean = 3, sd = 2)
enorm(dat, ci = TRUE)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Normal
#
#Estimated Parameter(s):      mean = 2.861160
#                               sd   = 1.180226
#
#Estimation Method:           mvue
#
#Data:                         dat
#
#Sample Size:                  20
#
#Confidence Interval for:      mean
#
#Confidence Interval Method:   Exact
#
#Confidence Interval Type:     two-sided
#
#Confidence Level:             95%
#
#Confidence Interval:          LCL = 2.308798
#                               UCL = 3.413523
#
#-----

# Using the same data, construct an upper 90% confidence interval for
# the variance.
```

```

enorm(dat, ci = TRUE, ci.type = "upper", ci.param = "variance")$interval

#Confidence Interval for:      variance
#
#Confidence Interval Method:   Exact
#
#Confidence Interval Type:    upper
#
#Confidence Level:            95%
#
#Confidence Interval:         LCL = 0.000000
#                             UCL = 2.615963

#-----

# Clean up
#-----
rm(dat)

#-----

# Using the Reference area TcCB data in the data frame EPA.94b.tccb.df,
# estimate the mean and standard deviation of the log-transformed data,
# and construct a 95% confidence interval for the mean.

with(EPA.94b.tccb.df, enorm(log(TcCB[Area == "Reference"]), ci = TRUE))

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:         Normal
#
#Estimated Parameter(s):      mean = -0.6195712
#                             sd   = 0.4679530
#
#Estimation Method:           mvue
#
#Data:                        log(TcCB[Area == "Reference"])
#
#Sample Size:                 47
#
#Confidence Interval for:     mean
#
#Confidence Interval Method:   Exact
#
#Confidence Interval Type:    two-sided
#
#Confidence Level:            95%
#
#Confidence Interval:         LCL = -0.7569673
#                             UCL = -0.4821751

```

Description

Estimate the mean and standard deviation of a [normal \(Gaussian\) distribution](#) given a sample of data that has been subjected to Type I censoring, and optionally construct a confidence interval for the mean.

Usage

```
enormCensored(x, censored, method = "mle", censoring.side = "left",
  ci = FALSE, ci.method = "profile.likelihood", ci.type = "two-sided",
  conf.level = 0.95, n.bootstraps = 1000, use.acc.con = FALSE,
  pivot.statistic = "z", nmc = 1000, seed = NULL, ...)
```

Arguments

x	numeric vector of observations. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
censored	numeric or logical vector indicating which values of x are censored. This must be the same length as x. If the mode of censored is "logical", TRUE values correspond to elements of x that are censored, and FALSE values correspond to elements of x that are not censored. If the mode of censored is "numeric", it must contain only 1's and 0's; 1 corresponds to TRUE and 0 corresponds to FALSE. Missing (NA) values are allowed but will be removed.
method	character string specifying the method of estimation. For singly censored data, the possible values are: "mle" (maximum likelihood; the default), "bcmle" (bias-corrected maximum likelihood), "qq.reg" (quantile-quantile regression), "qq.reg.w.cen.level" (quantile-quantile regression including the censoring level), "impute.w.qq.reg" (moment estimation based on imputation using the qq.reg method), "impute.w.qq.reg.w.cen.level" (moment estimation based on imputation using the qq.reg.w.cen.level method), "impute.w.mle" (moment estimation based on imputation using the mle), "iterative.impute.w.q" (moment estimation based on iterative imputation using the qq.reg method), "m.est" (robust M-estimation), and "half.cen.level" (moment estimation based on setting the censored observations to half the censoring level). For multiply censored data, the possible values are: "mle" (maximum likelihood; the default), "qq.reg" (quantile-quantile regression), "impute.w.qq.reg" (moment estimation based on imputation using the qq.reg method), and "half.cen.level" (moment estimation based on setting the censored observations to half the censoring level). See the DETAILS section for more information.
censoring.side	character string indicating on which side the censoring occurs. The possible values are "left" (the default) and "right".
ci	logical scalar indicating whether to compute a confidence interval for the mean or variance. The default value is ci=FALSE.
ci.method	character string indicating what method to use to construct the confidence interval for the mean. The possible values are "profile.likelihood" (profile likelihood; the default), "normal.approx" (normal approximation), "normal.approx.w.cov" (normal approximation taking into account the covariance between the estimated mean and standard deviation; only available for singly censored data), "gpc" (generalized pivotal quantity), and "bootstrap" (based on bootstrapping). See the DETAILS section for more information. This argument is ignored if ci=FALSE.

<code>ci.type</code>	character string indicating what kind of confidence interval to compute. The possible values are "two-sided" (the default), "lower", and "upper". This argument is ignored if <code>ci=FALSE</code> .
<code>conf.level</code>	a scalar between 0 and 1 indicating the confidence level of the confidence interval. The default value is <code>conf.level=0.95</code> . This argument is ignored if <code>ci=FALSE</code> .
<code>n.bootstraps</code>	numeric scalar indicating how many bootstraps to use to construct the confidence interval for the mean when <code>ci.type="bootstrap"</code> . This argument is ignored if <code>ci=FALSE</code> and/or <code>ci.method</code> does not equal "bootstrap".
<code>use.acc.con</code>	logical scalar indicating whether to use the acceleration constant when computing the bias-corrected bootstrap confidence interval (see the DETAILS section). The default value is <code>FALSE</code> . This argument is ignored if <code>ci=FALSE</code> and/or <code>ci.method</code> does not equal "bootstrap".
<code>pivot.statistic</code>	character string indicating which pivot statistic to use in the construction of the confidence interval for the mean when <code>ci.method="normal.approx"</code> or <code>ci.method="normal.approx.w.cov"</code> (see the DETAILS section). The possible values are <code>pivot.statistic="z"</code> (the default) and <code>pivot.statistic="t"</code> . When <code>pivot.statistic="t"</code> you may supply the argument <code>ci.sample.size</code> (see below). The argument <code>pivot.statistic</code> is ignored if <code>ci=FALSE</code> .
<code>nmc</code>	numeric scalar indicating the number of Monte Carlo simulations to run when <code>ci.method="gpq"</code> . The default is <code>nmc=1000</code> . This argument is ignored if <code>ci=FALSE</code> .
<code>seed</code>	integer supplied to the function <code>set.seed</code> and used when <code>ci.method="bootstrap"</code> or <code>ci.method="gpq"</code> . The default value is <code>seed=NULL</code> , in which case the current value of <code>.Random.seed</code> is used. This argument is ignored when <code>ci=FALSE</code> .
<code>...</code>	additional arguments to pass to other functions. <ul style="list-style-type: none"> <code>prob.method</code>. Character string indicating what method to use to compute the plotting positions (empirical probabilities) when method is one of "qq.reg", "qq.reg.w.cen.level", "impute.w.qq.reg", "impute.w.qq.reg.w.cen.level", "impute.w.mle", or "iterative.impute.w.qq.reg". Possible values are "kaplan-meier" (product-limit method of Kaplan and Meier (1958)), "nelson" (hazard plotting method of Nelson (1972)), "michael-schucany" (generalization of the product-limit method due to Michael and Schucany (1986)), and "hirsch-stedinger" (generalization of the product-limit method due to Hirsch and Stedinger (1987)). The default value is <code>prob.method="michael-schucany"</code>. The "nelson" method is only available for <code>censoring.side="right"</code>. See the DETAILS section and the help file for <code>ppointsCensored</code> for more information. <code>plot.pos.con</code>. Numeric scalar between 0 and 1 containing the value of the plotting position constant to use when method is one of "qq.reg", "qq.reg.w.cen.level", "impute.w.qq.reg", "impute.w.qq.reg.w.cen.level", "impute.w.mle", or "iterative.impute.w.qq.reg". The default value is <code>plot.pos.con=0.375</code>. See the DETAILS section and the help file for <code>ppointsCensored</code> for more information. <code>ci.sample.size</code>. Numeric scalar indicating what sample size to assume to construct the confidence interval for the mean if <code>pivot.statistic="t"</code> and <code>ci.method="normal.approx"</code> or <code>ci.method="normal.approx.w.cov"</code>. When method equals "mle" or "bcmle", the default value is the expected number of uncensored observations, otherwise it is the observed number of uncensored observations.

- `lb.impute`. Numeric scalar indicating the lower bound for imputed observations when method is one of "impute.w.qq.reg", "impute.w.qq.reg.w.cen.level", "impute.w.mle", or "iterative.impute.w.qq.reg". Imputed values smaller than this value will be set to this value. The default is `lb.impute=-Inf`.
- `ub.impute`. Numeric scalar indicating the upper bound for imputed observations when method is one of "impute.w.qq.reg", "impute.w.qq.reg.w.cen.level", "impute.w.mle", or "iterative.impute.w.qq.reg". Imputed values larger than this value will be set to this value. The default is `ub.impute=Inf`.
- `convergence`. Character string indicating the kind of convergence criterion when method="iterative.impute.w.qq.reg". The possible values are "relative" (the default) and "absolute". See the DETAILS section for more information.
- `tol`. Numeric scalar indicating the convergence tolerance when method="iterative.impute.w". The default value is `tol=1e-6`. If convergence="relative", then the relative difference in the old and new estimates of the mean and the relative difference in the old and new estimates of the standard deviation must be less than `tol` for convergence to be achieved. If convergence="absolute", then the absolute difference in the old and new estimates of the mean and the absolute difference in the old and new estimates of the standard deviation must be less than `tol` for convergence to be achieved.
- `max.iter`. Numeric scalar indicating the maximum number of iterations when method="iterative.impute.w.qq.reg".
- `t.df`. Numeric scalar greater than or equal to 1 that determines the robustness and efficiency properties of the estimator when method="m.est". The default value is `t.df=3`.

Details

If `x` or `censored` contain any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

Let \underline{x} denote a vector of N observations from a [normal distribution](#) with mean μ and standard deviation σ . Assume n ($0 < n < N$) of these observations are known and c ($c = N - n$) of these observations are all censored below (left-censored) or all censored above (right-censored) at k fixed censoring levels

$$T_1, T_2, \dots, T_k; k \geq 1 \quad (1)$$

For the case when $k \geq 2$, the data are said to be Type I ***multiply censored***. For the case when $k = 1$, set $T = T_1$. If the data are left-censored and all n known observations are greater than or equal to T , or if the data are right-censored and all n known observations are less than or equal to T , then the data are said to be Type I ***singly censored*** (Nelson, 1982, p.7), otherwise they are considered to be Type I multiply censored.

Let c_j denote the number of observations censored below or above censoring level T_j for $j = 1, 2, \dots, k$, so that

$$\sum_{i=1}^k c_j = c \quad (2)$$

Let $x_{(1)}, x_{(2)}, \dots, x_{(N)}$ denote the "ordered" observations, where now "observation" means either the actual observation (for uncensored observations) or the censoring level (for censored observations). For right-censored data, if a censored observation has the same value as an uncensored one, the uncensored observation should be placed first. For left-censored data, if a censored observation has the same value as an uncensored one, the censored observation should be placed first.

Note that in this case the quantity $x_{(i)}$ does not necessarily represent the i 'th "largest" observation from the (unknown) complete sample.

Finally, let Ω (omega) denote the set of n subscripts in the "ordered" sample that correspond to uncensored observations.

ESTIMATION

Estimation Methods for Multiply and Singly Censored Data

The following methods are available for multiply and singly censored data.

Maximum Likelihood Estimation (method="mle")

For Type I left censored data, the likelihood function is given by:

$$L(\mu, \sigma | \underline{x}) = \binom{N}{c_1 c_2 \dots c_k n} \prod_{j=1}^k [F(T_j)]^{c_j} \prod_{i \in \Omega} f[x_{(i)}] \quad (3)$$

where f and F denote the probability density function (pdf) and cumulative distribution function (cdf) of the population. That is,

$$f(t) = \phi\left(\frac{t - \mu}{\sigma}\right) \quad (4)$$

$$F(t) = \Phi\left(\frac{t - \mu}{\sigma}\right) \quad (5)$$

where ϕ and Φ denote the pdf and cdf of the standard normal distribution, respectively (Cohen, 1963; 1991, pp.6, 50). For left singly censored data, Equation (3) simplifies to:

$$L(\mu, \sigma | \underline{x}) = \binom{N}{c} [F(T)]^c \prod_{i=c+1}^n f[x_{(i)}] \quad (6)$$

Similarly, for Type I right censored data, the likelihood function is given by:

$$L(\mu, \sigma | \underline{x}) = \binom{N}{c_1 c_2 \dots c_k n} \prod_{j=1}^k [1 - F(T_j)]^{c_j} \prod_{i \in \Omega} f[x_{(i)}] \quad (7)$$

and for right singly censored data this simplifies to:

$$L(\mu, \sigma | \underline{x}) = \binom{N}{c} [1 - F(T)]^c \prod_{i=1}^n f[x_{(i)}] \quad (8)$$

The maximum likelihood estimators are computed by maximizing the likelihood function. For right-censored data, Cohen (1963; 1991, pp.50-51) shows that taking partial derivatives of the log-likelihood function with respect to μ and σ and setting these to 0 produces the following two simultaneous equations:

$$\bar{x} - \mu = -\sigma \sum_{j=1}^k \left(\frac{c_j}{n}\right) Q_j \quad (9)$$

$$s^2 + (\bar{x} - \mu)^2 = \sigma^2 \left[1 - \sum_{j=1}^k \zeta_j \left(\frac{c_j}{n}\right) Q_j\right] \quad (10)$$

where

$$\bar{x} = \frac{1}{n} \sum_{i \in \Omega} x_{(i)} \quad (11)$$

$$s^2 = \frac{1}{n} \sum_{i \in \Omega} (x_{(i)} - \bar{x})^2 \quad (12)$$

$$Q_j = Q(\zeta_j) \quad (13)$$

$$\zeta_j = \frac{T_j - \mu}{\sigma} \quad (14)$$

$$Q(t) = \frac{\phi(t)}{1 - \Phi(t)} \quad (15)$$

Note that the quantity defined in Equation (11) is simply the mean of the uncensored observations, the quantity defined in Equation (12) is simply the method of moments estimator of variance based on the uncensored observations, and the function $Q(\cdot)$ defined in Equation (15) is the hazard function for the standard normal distribution.

For left-censored data, Equations (9) and (10) stay the same, except ζ is replaced with $-\zeta$.

The function `enormCensored` computes the maximum likelihood estimators by solving Equations (9) and (10) and uses the quantile-quantile regression estimators (see below) as initial values.

Quantile-Quantile Regression (method="qq.reg")

This method is sometimes called the *probability plot method* (Nelson, 1982, Chapter 3; Gilbert, 1987, pp.134-136; Helsel and Hirsch, 1992, p. 361), and more recently also called *parametric regression on order statistics* or *ROS* (USEPA, 2009; Helsel, 2012). In the case of no censoring, it is well known (e.g., Nelson, 1982, p.113; Cleveland, 1993, p.31) that for the standard normal (Gaussian) quantile-quantile plot (i.e., the plot of the sorted observations (empirical quantiles) versus standard normal quantiles; see `qqPlot`), the intercept and slope of the fitted least-squares line estimate the mean and standard deviation, respectively. Specifically, the estimates of μ and σ are found by computing the least-squares estimates in the following model:

$$x_{(i)} = \mu + \sigma \Phi^{-1}(p_i) + \epsilon_i, \quad i = 1, 2, \dots, N \quad (16)$$

where

$$p_i = \frac{i - a}{N - 2a + 1} \quad (17)$$

denotes the plotting position associated with the i 'th largest value, a is a constant such that $0 \leq a \leq 1$ (the plotting position constant), and Φ denotes the cumulative distribution function (cdf) of the standard normal distribution. The default value of a is 0.375 (see below).

This method can be adapted to the case of left (right) singly censored data as follows. Plot the n uncensored observations against the n largest (smallest) normal quantiles, where the normal quantiles are computed based on a sample size of N , fit the least-squares line to this plot, and estimate the mean and standard deviation from the intercept and slope, respectively. That is, use Equations (16) and (17), but for right singly censored data use $i = 1, 2, \dots, n$, and for left singly censored data use $i = (c + 1), (c + 2), \dots, N$.

The argument `plot.pos.con` (see the entry for `...` in the ARGUMENTS section above) determines the value of the plotting positions computed in Equation (18). The default value is `plot.pos.con=0.375`. See the help file for `qqPlot` for more information.

This method is discussed by Haas and Scheff (1990). In the context of lognormal data, Travis and Land (1990) suggest exponentiating the predicted 50'th percentile from this fit to estimate the geometric mean (i.e., the median of the lognormal distribution).

This method is easily extended to multiply censored data. Equation (16) becomes

$$x_{(i)} = \mu + \sigma \Phi^{-1}(p_i) + \epsilon_i, \quad i \in \Omega \quad (18)$$

where Ω denotes the set of n subscripts associated with the uncensored observations in the ordered sample. The plotting positions are computed by calling the **EnvStats** function `ppointsCensored`. The argument `prob.method` determines the method of computing the plotting positions (default is `prob.method="hirsch-stedinger"`), and the argument `plot.pos.con` determines the plotting position constant (default is `plot.pos.con=0.375`). (See the entry for ... in the ARGUMENTS section above.) Both Helsel (2012) and USEPA (2009) also use the Hirsch-Stedinger probability method but set the plotting position constant to 0.

Imputation Using Quantile-Quantile Regression (`method="impute.w.qq.reg"`)

This method is also called *robust ROS* (USEPA, 2009; Helsel, 2012). It involves using the quantile-quantile regression method (`method="qq.reg"`) to fit a regression line (and thus initially estimate the mean and standard deviation), and then imputing the values of the censored observations by predicting them from the regression equation. The final estimates of the mean and standard deviation are then computed using the usual formulas (see [enorm](#)) based on the observed and imputed values.

The imputed values are computed as:

$$\hat{x}_{(i)} = \hat{\mu}_{qqreg} + \hat{\sigma}_{qqreg} \Phi^{-1}(p_i), \quad i \notin \Omega \quad (19)$$

See the help file for `ppointsCensored` for information on how the plotting positions for the censored observations are computed.

The argument `prob.method` determines the method of computing the plotting positions (default is `prob.method="hirsch-stedinger"`), and the argument `plot.pos.con` determines the plotting position constant (default is `plot.pos.con=0.375`). (See the entry for ... in the ARGUMENTS section above.) Both Helsel (2012) and USEPA (2009) also use the Hirsch-Stedinger probability method but set the plotting position constant to 0.

The arguments `lb.impute` and `ub.impute` determine the lower and upper bounds for the imputed values. Imputed values smaller than `lb.impute` are set to this value. Imputed values larger than `ub.impute` are set to this value. The default values are `lb.impute=-Inf` and `ub.impute=Inf`. See the entry for ... in the ARGUMENTS section above.

For singly censored data, this is the NR method of Gilliom and Helsel (1986, p. 137). In the context of lognormal data, this method is discussed by Hashimoto and Trussell (1983), Gilliom and Helsel (1986), and El-Shaarawi (1989), and is referred to as the LR or Log-Probability Method.

For multiply censored data, this method was developed in the context of lognormal data by Helsel and Cohn (1988) using the formulas for plotting positions given in Hirsch and Stedinger (1987) and Weibull plotting positions (i.e., `prob.method="hirsch-stedinger"` and `plot.pos.con=0`).

Setting Censored Observations to Half the Censoring Level (`method="half.cen.level"`)

This method is applicable only to left censored data that is bounded below by 0. This method involves simply replacing all the censored observations with half their detection limit, and then computing the mean and standard deviation with the usual formulas (see [enorm](#)).

This method is included only to allow comparison of this method to other methods. **Setting left-censored observations to half the censoring level is not recommended.**

For singly censored data, this method is discussed by Gleit (1985), Haas and Scheff (1990), and El-Shaarawi and Esterby (1992). El-Shaarawi and Esterby (1992) show that these estimators are biased and inconsistent (i.e., the bias remains even as the sample size increases).

For multiply censored data, this method was studied by Helsel and Cohn (1988).

Estimation Methods for Singly Censored Data

The following methods are available only for singly censored data.

Bias-Corrected Maximum Likelihood Estimation (method="bcmle")

The maximum likelihood estimates of μ and σ are biased. The bias tends to 0 as the sample size increases, but it can be considerable for small sample sizes, especially in the case of a large percentage of censored observations (Saw, 1961b). Schmee et al. (1985) note that bias and variances of the mle's are of the order $1/N$ (see for example, Bain and Engelhardt, 1991), and that for 90% censoring the bias is negligible if N is at least 100. (For less intense censoring, even fewer observations are needed.)

The exact bias of each estimator is extremely difficult to compute. Saw (1961b), however, derived the first-order term (i.e., the term of order $1/N$) in the bias of the mle's of μ and σ and proposed bias-corrected mle's. His bias-corrected estimators were derived for the case of Type II singly censored data. Schneider (1986, p.110) and Haas and Scheff (1990), however, state that this bias correction should reduce the bias of the estimators in the case of Type I censoring as well.

Based on the tables of bias-correction terms given in Saw (1961b), Schneider (1986, pp.107-110) performed a least-squares fit to produce the following computational formulas for right-censored data:

$$B_{\mu} = -\exp[2.692 - 5.493 \frac{n}{N+1}] \quad (20)$$

$$B_{\sigma} = -[0.312 + 0.859 \frac{n}{N+1}]^{-2} \quad (21)$$

$$\hat{\mu}_{bcmle} = \hat{\mu}_{mle} - \frac{\hat{\sigma}_{mle}}{N+1} B_{\mu} \quad (22)$$

$$\hat{\sigma}_{bcmle} = \hat{\sigma}_{mle} - \frac{\hat{\sigma}_{mle}}{N+1} B_{\sigma} \quad (23)$$

For left-censored data, Equation (22) becomes:

$$\hat{\mu}_{bcmle} = \hat{\mu}_{mle} + \frac{\hat{\sigma}_{mle}}{N+1} B_{\mu} \quad (22)$$

Quantile-Quantile Regression Including the Censoring Level (method="qq.reg.w.cen.level")

This is a modification of the quantile-quantile regression method and was proposed by El-Shaarawi (1989) in the context of lognormal data. El-Shaarawi's idea is to include the censoring level and an associated plotting position, along with the uncensored observations and their associated plotting positions, in order to include information about the value of the censoring level T .

For left singly censored data, the modification involves adding the point $[\Phi^{-1}(p_c), T]$ to the plot before fitting the least-squares line. For right singly censored data, the point $[\Phi^{-1}(p_{n+1}), T]$ is added to the plot before fitting the least-squares line.

El-Shaarawi (1989) also proposed replacing the estimated normal quantiles with the exact expected values of normal order statistics, and using the values in their variance-covariance matrix to perform a weighted least least-squared regression. These last two modifications are not incorporated here.

Imputation Using Quantile-Quantile Regression Including the Censoring Level (method="impute.w.qq.reg.w.cen.level")

This is exactly the same method as imputation using quantile-quantile regression (method="impute.w.qq.reg"), except that the quantile-quantile regression including the censoring level method (method="qq.reg.w.cen.level") is used to fit the regression line. In the context of lognormal data, this method is discussed by El-Shaarawi (1989), which he denotes as the Modified LR Method.

Imputation Using Maximum Likelihood (method="impute.w.mle")

This is exactly the same method as imputation with quantile-quantile regression (method="impute.w.qq.reg"),

except that the maximum likelihood method (method="mle") is used to compute the initial estimates of the mean and standard deviation. In the context of lognormal data, this method is discussed by El-Shaarawi (1989), which he denotes as the Modified Maximum Likelihood Method.

Iterative Imputation Using Quantile-Quantile Regression (method="iterative.impute.w.qq.reg")

This method is similar to the imputation with quantile-quantile regression method (method="impute.w.qq.reg"), but iterates until the estimates of the mean and standard deviation converge. The algorithm is:

1. Compute the initial estimates of μ and σ using the "impute.w.qq.reg" method. (Actually, any suitable estimates will do.)
2. Using the current values of μ and σ and Equation (19), compute new imputed values of the censored observations.
3. Use the new imputed values along with the uncensored observations to compute new estimates of μ and σ based on the usual formulas (see [enorm](#)).
4. Repeat Steps 2 and 3 until the estimates converge (the convergence criterion is determined by the arguments tol and convergence; see the entry for ... in the ARGUMENTS section above).

This method is discussed by Gleit (1985), which he denotes as "Fill-In with Expected Values".

M-Estimators (method="m.est")

This method was contributed by Leo R. Korn (Korn and Tyler, 2001). This method finds location and scale estimates that are consistent at the normal model and robust to deviations from the normal model, including both outliers on the right and outliers on the left above and below the limit of detection. The estimates are found by solving the simultaneous equations:

$$\sum_{i=1}^c h_{\nu}\left(\frac{T-\mu}{\sigma}\right) + \sum_{i=c+1}^N \psi_{\nu}\left(\frac{x_i-\mu}{\sigma}\right) = 0 \quad (23)$$

$$\sum_{i=1}^c \lambda_{\nu}\left(\frac{T-\mu}{\sigma}\right) + \sum_{i=c+1}^N \chi_{\nu}\left(\frac{x_i-\mu}{\sigma}\right) = 0 \quad (24)$$

where

$$H_{\nu}(r) = -\log[F_{\nu}(r)] \quad (25)$$

$$h_{\nu}(r) = \frac{d}{dr} H_{\nu}(r) = H'_{\nu}(r) \quad (26)$$

$$\rho_{\nu}(r) = -\log[f_{\nu}(r)] \quad (27)$$

$$\psi_{\nu}(r) = \frac{d}{dr} \rho_{\nu}(r) = \rho'_{\nu}(r) \quad (28)$$

$$\lambda_{\nu}(r) = r h_{\nu}(r) \quad (29)$$

$$\chi_{\nu}(r) = r \psi_{\nu}(r) - 1 \quad (30)$$

and f_{ν} and F_{ν} denote the probability density function (pdf) and cumulative distribution function (cdf) of [Student's t-distribution](#) with ν degrees of freedom.

This results in an M-estimating equation based on the t-density function (Korn and Tyler., 2001). Since the t-density has heavier tails than the normal density, this M-estimator will tend to down-weight values that are far away from the center of the data. When censoring is present, neither the location nor the scale estimates are consistent at the normal model. A computational correction is performed that converts the above M-estimator to another M-estimator that is consistent at the normal model, even under censoring.

The degrees of freedom parameter ν is set by the argument `t.df` and may be viewed as a tuning parameter that will determine the robustness and efficiency properties. When `t.df` is large, the estimator is similar to the usual mle and the output will then be very close to that when `method="mle"`. As `t.df` decreases, the efficiency will decline and the outlier rejection property will increase in strength. Choosing `t.df=3` (the default) provides a good combination of efficiency and robustness. A reasonable strategy is to transform the data so that they are approximately symmetric (often the log transformation for environmental data is appropriate) and then apply the M-estimator using `t.df=3`.

CONFIDENCE INTERVALS

This section explains how confidence intervals for the mean μ are computed.

Likelihood Profile (`ci.method="profile.likelihood"`)

This method was proposed by Cox (1970, p.88), and Venzon and Moolgavkar (1988) introduced an efficient method of computation. This method is also discussed by Stryhn and Christensen (2003) and Royston (2007). The idea behind this method is to invert the likelihood-ratio test to obtain a confidence interval for the mean μ while treating the standard deviation σ as a nuisance parameter. Equation (3) above shows the form of the likelihood function $L(\mu, \sigma | \underline{x})$ for multiply left-censored data, and Equation (7) shows the function for multiply right-censored data.

Following Stryhn and Christensen (2003), denote the maximum likelihood estimates of the mean and standard deviation by (μ^*, σ^*) . The likelihood ratio test statistic (G^2) of the hypothesis $H_0 : \mu = \mu_0$ (where μ_0 is a fixed value) equals the drop in $2\log(L)$ between the “full” model and the reduced model with μ fixed at μ_0 , i.e.,

$$G^2 = 2\{\log[L(\mu^*, \sigma^*)] - \log[L(\mu_0, \sigma_0^*)]\} \quad (30)$$

where σ_0^* is the maximum likelihood estimate of σ for the reduced model (i.e., when $\mu = \mu_0$). Under the null hypothesis, the test statistic G^2 follows a [chi-squared distribution](#) with 1 degree of freedom.

Alternatively, we may express the test statistic in terms of the profile likelihood function L_1 for the mean μ , which is obtained from the usual likelihood function by maximizing over the parameter σ , i.e.,

$$L_1(\mu) = \max_{\sigma} L(\mu, \sigma) \quad (31)$$

Then we have

$$G^2 = 2\{\log[L_1(\mu^*)] - \log[L_1(\mu_0)]\} \quad (32)$$

A two-sided $(1 - \alpha)100\%$ confidence interval for the mean μ consists of all values of μ_0 for which the test is not significant at level *alpha*:

$$\mu_0 : G^2 \leq \chi_{1,1-\alpha}^2 \quad (33)$$

where $\chi_{\nu,p}^2$ denotes the p 'th quantile of the [chi-squared distribution](#) with ν degrees of freedom.

Normal Approximation (`ci.method="normal.approx"`)

This method constructs approximate $(1 - \alpha)100\%$ confidence intervals for μ based on the assumption that the estimator of μ is approximately normally distributed. That is, a two-sided $(1 - \alpha)100\%$ confidence interval for μ is constructed as:

$$[\hat{\mu} - t_{1-\alpha/2, m-1} \hat{\sigma}_{\hat{\mu}}, \hat{\mu} + t_{1-\alpha/2, m-1} \hat{\sigma}_{\hat{\mu}}] \quad (34)$$

where $\hat{\mu}$ denotes the estimate of μ , $\hat{\sigma}_{\hat{\mu}}$ denotes the estimated asymptotic standard deviation of the estimator of μ , m denotes the assumed sample size for the confidence interval, and $t_{p,\nu}$ denotes the

p 'th quantile of [Student's t-distribution](#) with ν degrees of freedom. One-sided confidence intervals are computed in a similar fashion.

The argument `ci.sample.size` determines the value of m (see the entry for ... in the ARGUMENTS section above). When `method` equals "mle" or "bcmle", the default value is the expected number of uncensored observations, otherwise it is the observed number of uncensored observations. This is simply an ad-hoc method of constructing confidence intervals and is not based on any published theoretical results.

When `pivot.statistic="z"`, the p 'th quantile from the [standard normal distribution](#) is used in place of the p 'th quantile from Student's t-distribution.

Approximate Confidence Interval Based on Maximum Likelihood Estimators

When `method="mle"`, the standard deviation of the mle of μ is estimated based on the inverse of the Fisher Information matrix. The estimated variance-covariance matrix for the estimates of μ and σ are based on the observed information matrix, formulas for which are given in Cohen (1991).

Approximate Confidence Interval Based on Bias-Corrected Maximum Likelihood Estimators

When `method="bcmle"` (available only for singly censored data), the same procedures are used to construct the confidence interval as for `method="mle"`. The true variance of the bias-corrected mle of μ is necessarily larger than the variance of the mle of μ (although the differences in the variances goes to 0 as the sample size gets large). Hence this method of constructing a confidence interval leads to intervals that are too short for small sample sizes, but these intervals should be better centered about the true value of μ .

Approximate Confidence Interval Based on Other Estimators

When `method` is some value other than "mle", the standard deviation of the estimated mean is approximated by

$$\hat{\sigma}_{\hat{\mu}} = \frac{\hat{\sigma}}{\sqrt{m}} \quad (35)$$

where, as already noted, m denotes the assumed sample size. This is simply an ad-hoc method of constructing confidence intervals and is not based on any published theoretical results.

Normal Approximation Using Covariance (`ci.method="normal.approx.w.cov"`) This method is only available for singly censored data and only applicable when `method="mle"` or `method="bcmle"`. It was proposed by Schneider (1986, pp. 191-193) for the case of Type II censoring, but is applicable to any situation where the estimated mean and standard deviation are consistent estimators and are correlated. In particular, the mle's of μ and σ are correlated under Type I censoring as well.

Schneider's idea is to determine two positive quantities z_1, z_2 such that

$$Pr(\hat{\mu} + z_1 \hat{\sigma} < \mu) = \frac{\alpha}{2} \quad (36)$$

$$Pr(\hat{\mu} - z_2 \hat{\sigma} > \mu) = \frac{\alpha}{2} \quad (37)$$

so that

$$[\hat{\mu} - z_2 \hat{\sigma}, \hat{\mu} + z_1 \hat{\sigma}] \quad (38)$$

is a $(1 - \alpha)100\%$ confidence interval for μ .

For cases where the estimators of μ and σ are independent (e.g., complete samples), it is well known that setting

$$z_1 = z_2 = \frac{t_{1-\alpha/2, N}}{\sqrt{N}} \quad (39)$$

yields an exact confidence interval and setting

$$z_1 = z_2 = \frac{z_{1-\alpha/2}}{\sqrt{N}} \quad (40)$$

where z_p denotes the p 'th quantile of the standard normal distribution yields an approximate confidence interval that is asymptotically correct.

For the general case, Schneider (1986) considers the random variable

$$W(z) = \hat{\mu} + z\hat{\sigma} \quad (41)$$

and provides formulas for z_1 and z_2 .

Note that the resulting confidence interval for the mean is not symmetric about the estimated mean. Also note that the quantity m is a random variable for Type I censoring, while Schneider (1986) assumed it to be fixed since he derived the result for Type II censoring (in which case $m = n$).

Bootstrap and Bias-Corrected Bootstrap Approximation (ci.method="bootstrap")

The bootstrap is a nonparametric method of estimating the distribution (and associated distribution parameters and quantiles) of a sample statistic, regardless of the distribution of the population from which the sample was drawn. The bootstrap was introduced by Efron (1979) and a general reference is Efron and Tibshirani (1993).

In the context of deriving an approximate $(1 - \alpha)100\%$ confidence interval for the population mean μ , the bootstrap can be broken down into the following steps:

1. Create a bootstrap sample by taking a random sample of size N from the observations in \underline{x} , where sampling is done with replacement. Note that because sampling is done with replacement, the same element of \underline{x} can appear more than once in the bootstrap sample. Thus, the bootstrap sample will usually not look exactly like the original sample (e.g., the number of censored observations in the bootstrap sample will often differ from the number of censored observations in the original sample).
2. Estimate μ based on the bootstrap sample created in Step 1, using the same method that was used to estimate μ using the original observations in \underline{x} . Because the bootstrap sample usually does not match the original sample, the estimate of μ based on the bootstrap sample will usually differ from the original estimate based on \underline{x} .
3. Repeat Steps 1 and 2 B times, where B is some large number. For the function `enormCensored`, the number of bootstraps B is determined by the argument `n.bootstraps` (see the section ARGUMENTS above). The default value of `n.bootstraps` is 1000.
4. Use the B estimated values of μ to compute the empirical cumulative distribution function of this estimator of μ (see `ecdfPlot`), and then create a confidence interval for μ based on this estimated cdf.

The percentile interval (Efron and Tibshirani, 1993, p.170) is computed as:

$$[\hat{G}^{-1}(\frac{\alpha}{2}), \hat{G}^{-1}(\frac{1-\alpha}{2})] \quad (42)$$

where $\hat{G}(t)$ denotes the empirical cdf evaluated at t and thus $\hat{G}^{-1}(p)$ denotes the p 'th empirical quantile, that is, the p 'th quantile associated with the empirical cdf. The function `enormCensored` calls the R function `quantile` to compute the empirical quantiles used in Equation (42).

The percentile method bootstrap confidence interval is only first-order accurate (Efron and Tibshirani, 1993, pp.187-188), meaning that the probability that the confidence interval will contain the true value of μ can be off by k/\sqrt{N} , where k is some constant. Efron and Tibshirani (1993, pp.184-188) proposed a bias-corrected and accelerated interval that is second-order accurate, meaning that the probability that the confidence interval will contain the true value of μ may be off by k/N instead of k/\sqrt{N} . The bias-corrected and accelerated confidence interval is computed as:

$$[\hat{G}^{-1}(\alpha_1), \hat{G}^{-1}(\alpha_2)] \quad (43)$$

where

$$\alpha_1 = \Phi\left[\hat{z}_0 + \frac{\hat{z}_0 + z_{\alpha/2}}{1 - \hat{a}(z_0 + z_{\alpha/2})}\right] \quad (44)$$

$$\alpha_2 = \Phi\left[\hat{z}_0 + \frac{\hat{z}_0 + z_{1-\alpha/2}}{1 - \hat{a}(z_0 + z_{1-\alpha/2})}\right] \quad (45)$$

$$\hat{z}_0 = \Phi^{-1}[\hat{G}(\hat{\mu})] \quad (46)$$

$$\hat{a} = \frac{\sum_{i=1}^N (\hat{\mu}_{(\cdot)} - \hat{\mu}_{(i)})^3}{6[\sum_{i=1}^N (\hat{\mu}_{(\cdot)} - \hat{\mu}_{(i)})^2]^{3/2}} \quad (47)$$

where the quantity $\hat{\mu}_{(i)}$ denotes the estimate of μ using all the values in \underline{x} except the i 'th one, and

$$\hat{\mu}_{(\cdot)} = \frac{1}{N} \sum_{i=1}^N \mu_{(i)} \quad (48)$$

The constant \hat{z}_0 incorporates the bias correction, and the constant \hat{a} is the acceleration constant. The term “acceleration” refers to the rate of change of the standard error of the estimate of μ with respect to the true value of μ (Efron and Tibshirani, 1993, p.186). For a normal (Gaussian) distribution, the standard error of the estimate of μ does not depend on the value of μ , hence the acceleration constant is not really necessary. The argument `use.acc.con` (see the section `ARGUMENTS` above) determines whether the acceleration constant \hat{a} is computed as in Equation (47) (`use.acc.con=TRUE`) or is set to 0 (`use.acc.con=FALSE`). The default value is `use.acc.con=FALSE`.

When `ci.method="bootstrap"`, the function `enormCensored` computes both the percentile method and bias-corrected bootstrap confidence intervals.

This method of constructing confidence intervals for censored data was studied by Shumway et al. (1989).

Generalized Pivotal Quantity (`ci.method="gpq"`)

This method was introduced by Schmee et al. (1985) and is discussed by Krishnamoorthy and Mathew (2009). The idea is essentially to use a parametric bootstrap to estimate the correct pivotal quantities z_1 and z_2 in Equation (38) above. For singly censored data, these quantities are computed as follows:

1. Generate a random sample of N observations from a standard normal (i.e., $N(0,1)$) distribution and let $z_{(1)}, z_{(2)}, \dots, z_{(N)}$ denote the ordered (sorted) observations.
2. Set the smallest c observations to be censored.
3. Compute the estimates of μ and σ using the method specified by the `method` argument, and denote these estimates as $\hat{\mu}^*, \hat{\sigma}^*$.
4. Compute the t-like pivotal quantity $\hat{t} = \hat{\mu}^* / \hat{\sigma}^*$.
5. Repeat steps 1-4 `nmc` times to produce an empirical distribution of the t-like pivotal quantity.

The function `enormCensored` calls the function `gpqCiNormSinglyCensored` to generate the distribution of pivotal quantities in the case of singly censored data. A two-sided $(1-\alpha)100\%$ confidence interval for μ is then computed as:

$$[\hat{\mu} - \hat{t}_{1-(\alpha/2)}\hat{\sigma}, \hat{\mu} - \hat{t}_{\alpha/2}\hat{\sigma}] \quad (49)$$

where \hat{t}_p denotes the p 'th empirical quantile of the `nmc` generated \hat{t} values.

Schmee et al. (1985) derived this method in the context of Type II singly censored data (for which these limits are exact within Monte Carlo error), but state that according to Regal (1982) this method

produces confidence intervals that are close approximations to the correct limits for Type I censored data.

For multiply censored data, this method has been extended as follows. The algorithm stays the same, except that Step 2 becomes:

2. Set the i 'th ordered generated observation to be censored or not censored according to whether the i 'th observed observation in the original data is censored or not censored.

The function `enormCensored` calls the function `gpqCiNormMultiplyCensored` to generate the distribution of pivotal quantities in the case of multiply censored data.

Value

a list of class "estimateCensored" containing the estimated parameters and other information. See `estimateCensored.object` for details.

Note

A sample of data contains censored observations if some of the observations are reported only as being below or above some censoring level. In environmental data analysis, Type I left-censored data sets are common, with values being reported as "less than the detection limit" (e.g., Helsel, 2012). Data sets with only one censoring level are called *singly censored*; data sets with multiple censoring levels are called *multiply* or *progressively censored*.

Statistical methods for dealing with censored data sets have a long history in the field of survival analysis and life testing. More recently, researchers in the environmental field have proposed alternative methods of computing estimates and confidence intervals in addition to the classical ones such as maximum likelihood estimation.

Helsel (2012, Chapter 6) gives an excellent review of past studies of the properties of various estimators based on censored environmental data.

In practice, it is better to use a confidence interval for the mean or a joint confidence region for the mean and standard deviation, rather than rely on a single point-estimate of the mean. Since confidence intervals and regions depend on the properties of the estimators for both the mean and standard deviation, the results of studies that simply evaluated the performance of the mean and standard deviation separately cannot be readily extrapolated to predict the performance of various methods of constructing confidence intervals and regions. Furthermore, for several of the methods that have been proposed to estimate the mean based on type I left-censored data, standard errors of the estimates are not available, hence it is not possible to construct confidence intervals (El-Shaarawi and Dolan, 1989).

Few studies have been done to evaluate the performance of methods for constructing confidence intervals for the mean or joint confidence regions for the mean and standard deviation when data are subjected to single or multiple censoring.

Schmee et al. (1985) studied Type II censoring for a normal distribution and noted that the bias and variances of the maximum likelihood estimators are of the order $1/N$, and that the bias is negligible for $N = 100$ and as much as 90% censoring. (If the proportion of censored observations is less than 90%, the bias becomes negligible for smaller sample sizes.) For small samples with moderate to high censoring, however, the bias of the mle's causes confidence intervals based on them using a normal approximation (e.g., `method="mle"` and `ci.method="normal.approx"`) to be too short. Schmee et al. (1985) provide tables for exact confidence intervals for sample sizes up to $N = 100$ that were created based on Monte Carlo simulation. Schmee et al. (1985) state that these tables should work well for Type I censored data as well.

Shumway et al. (1989) evaluated the coverage of 90% confidence intervals for the mean based on using a Box-Cox transformation to induce normality, computing the mle's based on the normal distribution, then computing the mean in the original scale. They considered three methods of constructing confidence intervals: the delta method, the bootstrap, and the bias-corrected bootstrap. Shumway et al. (1989) used three parent distributions in their study: Normal(3,1), the square of this distribution, and the exponentiation of this distribution (i.e., a lognormal distribution). Based on sample sizes of 10 and 50 with a censoring level at the 10'th or 20'th percentile, Shumway et al. (1989) found that the delta method performed quite well and was superior to the bootstrap method. Millard et al. (2014; in preparation) show that the coverage of profile likelihood method is excellent.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Bain, L.J., and M. Engelhardt. (1991). *Statistical Analysis of Reliability and Life-Testing Models*. Marcel Dekker, New York, 496pp.
- Cohen, A.C. (1959). Simplified Estimators for the Normal Distribution When Samples are Singly Censored or Truncated. *Technometrics* **1**(3), 217–237.
- Cohen, A.C. (1963). Progressively Censored Samples in Life Testing. *Technometrics* **5**, 327–339
- Cohen, A.C. (1991). *Truncated and Censored Samples*. Marcel Dekker, New York, New York, 312pp.
- Cox, D.R. (1970). *Analysis of Binary Data*. Chapman & Hall, London. 142pp.
- Efron, B. (1979). Bootstrap Methods: Another Look at the Jackknife. *The Annals of Statistics* **7**, 1–26.
- Efron, B., and R.J. Tibshirani. (1993). *An Introduction to the Bootstrap*. Chapman and Hall, New York, 436pp.
- El-Shaarawi, A.H. (1989). Inferences About the Mean from Censored Water Quality Data. *Water Resources Research* **25**(4) 685–690.
- El-Shaarawi, A.H., and D.M. Dolan. (1989). Maximum Likelihood Estimation of Water Quality Concentrations from Censored Data. *Canadian Journal of Fisheries and Aquatic Sciences* **46**, 1033–1039.
- El-Shaarawi, A.H., and S.R. Esterby. (1992). Replacement of Censored Observations by a Constant: An Evaluation. *Water Research* **26**(6), 835–844.
- El-Shaarawi, A.H., and A. Naderi. (1991). Statistical Inference from Multiply Censored Environmental Data. *Environmental Monitoring and Assessment* **17**, 339–347.
- Gibbons, R.D., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*, Second Edition. John Wiley & Sons, Hoboken.
- Gilliom, R.J., and D.R. Helsel. (1986). Estimation of Distributional Parameters for Censored Trace Level Water Quality Data: 1. Estimation Techniques. *Water Resources Research* **22**, 135–146.
- Gleit, A. (1985). Estimation for Small Normal Data Sets with Detection Limits. *Environmental Science and Technology* **19**, 1201–1206.
- Haas, C.N., and P.A. Scheff. (1990). Estimation of Averages in Truncated Samples. *Environmental Science and Technology* **24**(6), 912–919.
- Hashimoto, L.K., and R.R. Trussell. (1983). Evaluating Water Quality Data Near the Detection Limit. Paper presented at the Advanced Technology Conference, American Water Works Association, Las Vegas, Nevada, June 5-9, 1983.

- Helsel, D.R. (1990). Less than Obvious: Statistical Treatment of Data Below the Detection Limit. *Environmental Science and Technology* **24**(12), 1766–1774.
- Helsel, D.R. (2012). *Statistics for Censored Environmental Data Using Minitab and R, Second Edition*. John Wiley & Sons, Hoboken, New Jersey.
- Helsel, D.R., and T.A. Cohn. (1988). Estimation of Descriptive Statistics for Multiply Censored Water Quality Data. *Water Resources Research* **24**(12), 1997–2004.
- Hirsch, R.M., and J.R. Stedinger. (1987). Plotting Positions for Historical Floods and Their Precision. *Water Resources Research* **23**(4), 715–727.
- Korn, L.R., and D.E. Tyler. (2001). Robust Estimation for Chemical Concentration Data Subject to Detection Limits. In Fernholz, L., S. Morgenthaler, and W. Stahel, eds. *Statistics in Genetics and in the Environmental Sciences*. Birkhauser Verlag, Basel, pp.41–63.
- Krishnamoorthy K., and T. Mathew. (2009). *Statistical Tolerance Regions: Theory, Applications, and Computation*. John Wiley and Sons, Hoboken.
- Michael, J.R., and W.R. Schucany. (1986). Analysis of Data from Censored Samples. In D'Agostino, R.B., and M.A. Stephens, eds. *Goodness-of Fit Techniques*. Marcel Dekker, New York, 560pp, Chapter 11, 461–496.
- Millard, S.P., P. Dixon, and N.K. Neerchal. (2014; in preparation). *Environmental Statistics with R*. CRC Press, Boca Raton, Florida.
- Nelson, W. (1982). *Applied Life Data Analysis*. John Wiley and Sons, New York, 634pp.
- Newman, M.C., P.M. Dixon, B.B. Looney, and J.E. Pinder. (1989). Estimating Mean and Variance for Environmental Samples with Below Detection Limit Observations. *Water Resources Bulletin* **25**(4), 905–916.
- Pettitt, A. N. (1983). Re-Weighted Least Squares Estimation with Censored and Grouped Data: An Application of the EM Algorithm. *Journal of the Royal Statistical Society, Series B* **47**, 253–260.
- Regal, R. (1982). Applying Order Statistic Censored Normal Confidence Intervals to Time Censored Data. Unpublished manuscript, University of Minnesota, Duluth, Department of Mathematical Sciences.
- Royston, P. (2007). Profile Likelihood for Estimation and Confidence Intervals. *The Stata Journal* **7**(3), pp. 376–387.
- Saw, J.G. (1961b). The Bias of the Maximum Likelihood Estimators of Location and Scale Parameters Given a Type II Censored Normal Sample. *Biometrika* **48**, 448–451.
- Schmee, J., D.Gladstein, and W. Nelson. (1985). Confidence Limits for Parameters of a Normal Distribution from Singly Censored Samples, Using Maximum Likelihood. *Technometrics* **27**(2) 119–128.
- Schneider, H. (1986). *Truncated and Censored Samples from Normal Populations*. Marcel Dekker, New York, New York, 273pp.
- Shumway, R.H., A.S. Azari, and P. Johnson. (1989). Estimating Mean Concentrations Under Transformations for Environmental Data With Detection Limits. *Technometrics* **31**(3), 347–356.
- Stryhn, H., and J. Christensen. (2003). *Confidence Intervals by the Profile Likelihood Method, with Applications in Veterinary Epidemiology*. Contributed paper at ISVEE X (November 2003, Chile). <http://people.upei.ca/hstryhn/stryhn208.pdf>.
- Travis, C.C., and M.L. Land. (1990). Estimating the Mean of Data Sets with Nondetectable Values. *Environmental Science and Technology* **24**, 961–962.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. Chapter 15.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

Venzon, D.J., and S.H. Moolgavkar. (1988). A Method for Computing Profile-Likelihood-Based Confidence Intervals. *Journal of the Royal Statistical Society, Series C (Applied Statistics)* **37**(1), pp. 87–94.

See Also

[Normal](#), [enorm](#), [estimateCensored.object](#).

Examples

```
# Chapter 15 of USEPA (2009) gives several examples of estimating the mean
# and standard deviation of a lognormal distribution on the log-scale using
# manganese concentrations (ppb) in groundwater at five background wells.
# In EnvStats these data are stored in the data frame
# EPA.09.Ex.15.1.manganese.df.

# Here we will estimate the mean and standard deviation using the MLE,
# Q-Q regression (also called parametric regression on order statistics
# or ROS; e.g., USEPA, 2009 and Helsel, 2012), and imputation with Q-Q
# regression (also called robust ROS).

# We will log-transform the original observations and then call
# enormCensored. Alternatively, we could have more simply called
# elnormCensored.

# First look at the data:
#-----

EPA.09.Ex.15.1.manganese.df

#   Sample  Well Manganese.Orig.ppb Manganese.ppb Censored
#1      1 Well.1             <5           5.0      TRUE
#2      2 Well.1            12.1          12.1     FALSE
#3      3 Well.1            16.9          16.9     FALSE
#...
#23     3 Well.5             3.3           3.3     FALSE
#24     4 Well.5             8.4           8.4     FALSE
#25     5 Well.5             <2           2.0      TRUE

longToWide(EPA.09.Ex.15.1.manganese.df,
  "Manganese.Orig.ppb", "Sample", "Well",
  paste.row.name = TRUE)

#           Well.1 Well.2 Well.3 Well.4 Well.5
#Sample.1      <5      <5      <5      6.3      17.9
#Sample.2     12.1      7.7      5.3     11.9     22.7
#Sample.3     16.9     53.6     12.6      10      3.3
#Sample.4     21.6      9.5    106.3      <2      8.4
#Sample.5      <2     45.9     34.5     77.2      <2

# Now estimate the mean and standard deviation on the log-scale
# using the MLE:
```

```

#-----

with(EPA.09.Ex.15.1.manganese.df,
     enormCensored(log(Manganese.ppb), Censored))

#Results of Distribution Parameter Estimation
#Based on Type I Censored Data
#-----
#
#Assumed Distribution:          Normal
#
#Censoring Side:                left
#
#Censoring Level(s):           0.6931472 1.6094379
#
#Estimated Parameter(s):       mean = 2.215905
#                               sd   = 1.356291
#
#Estimation Method:            MLE
#
#Data:                          log(Manganese.ppb)
#
#Censoring Variable:           Censored
#
#Sample Size:                  25
#
#Percent Censored:              24%

# Now compare the MLE with the estimators based on
# Q-Q regression and imputation with Q-Q regression
#-----

with(EPA.09.Ex.15.1.manganese.df,
     enormCensored(log(Manganese.ppb), Censored))$parameters
#   mean      sd
#2.215905 1.356291

with(EPA.09.Ex.15.1.manganese.df,
     enormCensored(log(Manganese.ppb), Censored,
                    method = "qq.reg"))$parameters
#   mean      sd
#2.293742 1.283635

with(EPA.09.Ex.15.1.manganese.df,
     enormCensored(log(Manganese.ppb), Censored,
                    method = "impute.w.qq.reg"))$parameters
#   mean      sd
#2.298656 1.238104

#-----

# The method used to estimate quantiles for a Q-Q plot is
# determined by the argument prob.method. For the functions
# enormCensored and elnormCensored, for any estimation
# method that involves Q-Q regression, the default value of
# prob.method is "hirsch-stedinger" and the default value for the
# plotting position constant is plot.pos.con=0.375.

```



```

# Both Helsel (2012) and USEPA (2009) also use the Hirsch-Stedinger
# probability method but set the plotting position constant to 0.

with(EPA.09.Ex.15.1.manganese.df,
     enormCensored(log(Manganese.ppb), Censored,
                   method = "impute.w.qq.reg", plot.pos.con = 0))$parameters
#   mean      sd
#2.277175 1.261431

#-----

# Using the same data as above, compute a confidence interval
# for the mean on the log-scale using the profile-likelihood
# method.

with(EPA.09.Ex.15.1.manganese.df,
     enormCensored(log(Manganese.ppb), Censored, ci = TRUE))

#Results of Distribution Parameter Estimation
#Based on Type I Censored Data
#-----
#
#Assumed Distribution:      Normal
#
#Censoring Side:           left
#
#Censoring Level(s):       0.6931472 1.6094379
#
#Estimated Parameter(s):   mean = 2.215905
#                           sd   = 1.356291
#
#Estimation Method:        MLE
#
#Data:                     log(Manganese.ppb)
#
#Censoring Variable:       Censored
#
#Sample Size:              25
#
#Percent Censored:         24%
#
#Confidence Interval for:   mean
#
#Confidence Interval Method: Profile Likelihood
#
#Confidence Interval Type:  two-sided
#
#Confidence Level:         95%
#
#Confidence Interval:       LCL = 1.595062
#                           UCL = 2.771197

```

Description

Estimate the mean, standard deviation (left-censored data only), and standard error of the mean nonparametrically given a sample of data from a non-negative-valued distribution that has been subjected to Type I censoring, and optionally construct a confidence interval for the mean.

Usage

```
enparCensored(x, censored, censoring.side = "left", correct.se = FALSE,
  ci = FALSE, ci.method = "normal.approx", ci.type = "two-sided",
  conf.level = 0.95, pivot.statistic = "z", ci.sample.size = NULL,
  n.bootstraps = 1000, use.acc.con = FALSE)
```

Arguments

<code>x</code>	numeric vector of observations. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
<code>censored</code>	numeric or logical vector indicating which values of <code>x</code> are censored. This must be the same length as <code>x</code> . If the mode of <code>censored</code> is "logical", TRUE values correspond to elements of <code>x</code> that are censored, and FALSE values correspond to elements of <code>x</code> that are not censored. If the mode of <code>censored</code> is "numeric", it must contain only 1's and 0's; 1 corresponds to TRUE and 0 corresponds to FALSE. Missing (NA) values are allowed but will be removed.
<code>censoring.side</code>	character string indicating on which side the censoring occurs. The possible values are "left" (the default) and "right".
<code>correct.se</code>	logical scalar indicating whether to multiply the estimated standard error by a factor to correct for bias. The default value is <code>correct.se=FALSE</code> . See the DETAILS section below.
<code>ci</code>	logical scalar indicating whether to compute a confidence interval for the mean or variance. The default value is <code>ci=FALSE</code> .
<code>ci.method</code>	character string indicating what method to use to construct the confidence interval for the mean. The possible values are "normal.approx" (normal approximation; the default), and "bootstrap" (based on bootstrapping). See the DETAILS section for more information. This argument is ignored if <code>ci=FALSE</code> .
<code>ci.type</code>	character string indicating what kind of confidence interval to compute. The possible values are "two-sided" (the default), "lower", and "upper". This argument is ignored if <code>ci=FALSE</code> .
<code>conf.level</code>	a scalar between 0 and 1 indicating the confidence level of the confidence interval. The default value is <code>conf.level=0.95</code> . This argument is ignored if <code>ci=FALSE</code> .
<code>pivot.statistic</code>	character string indicating which statistic to use for the confidence interval for the mean when <code>ci.method="normal.approx"</code> . Possible values are "z" (confidence interval based on the z-statistic; the default), and "t" (confidence interval based on the t-statistic). When <code>pivot.statistic="t"</code> you may supply the argument <code>ci.sample.size</code> (see below). This argument is ignored if <code>ci=FALSE</code> .
<code>ci.sample.size</code>	numeric scalar or a NULL object indicating what sample size to assume when computing the confidence interval for the mean when <code>ci.method="normal.approx"</code> and <code>pivot.statistic="t"</code> . The default value is <code>ci.sample.size=NULL</code> , in which case <code>ci.sample.size</code> is equal to the number of uncensored observations. This argument is ignored if <code>ci=FALSE</code> .

n.bootstraps	numeric scalar indicating how many bootstraps to use to construct the confidence interval for the mean when ci.type="bootstrap". This argument is ignored if ci=FALSE and/or ci.method does not equal "bootstrap".
use.acc.con	logical scalar indicating whether to use the acceleration constant when computing the bias-corrected bootstrap confidence interval (see the DETAILS section). The default value is FALSE. This argument is ignored if ci=FALSE and/or ci.method does not equal "bootstrap".

Details

Let $\underline{x} = x_1, x_2, \dots, x_N$ denote a vector of N observations from some non-negative-valued continuous distribution with mean μ and standard deviation σ . Assume n ($0 < n < N$) of these observations are known and c ($c = N - n$) of these observations are all censored below (left-censored) or all censored above (right-censored) at k fixed censoring levels

$$T_1, T_2, \dots, T_k; \quad k \geq 1 \quad (1)$$

Let $t_{(1)}, t_{(2)}, \dots, t_{(p)}$ denote the p ordered *distinct* uncensored observations, where $p \leq n$. Let m_i denote the number of uncensored observations equal to $t_{(i)}$, let n_i denote the number of uncensored observations less than or equal to $t_{(i)}$, and let r_i denote the total number of observations (number of uncensored observations and number of censoring levels for censored observations) less than or equal to $t_{(i)}$.

Estimation

When the Kaplan-Meier estimator is used to construct the survival function, it can be shown that the area under this curve is an estimator of the mean of the distribution, and can be as efficient or more efficient than parametric estimators of the mean (Meier, 2004; Helsel, 2012; Lee and Wang, 2003). Let $\hat{F}(t)$ denote the Kaplan-Meier estimator of the empirical cumulative distribution function (ecdf) evaluated at t , and let $\hat{S}(t) = 1 - \hat{F}(t)$ denote the estimated survival function evaluated at t . (See the help files for [ecdfPlotCensored](#) and [qqPlotCensored](#) for an explanation of how the Kaplan-Meier estimator of the ecdf is computed.)

For left-censored data, the estimated mean, standard deviation, and standard error of the mean are given by:

$$\hat{\mu} = \sum_{i=1}^p t_{(i)} \{ \hat{F}[t_{(i)}] - \hat{F}[t_{(i-1)}] \} \quad (2)$$

$$\hat{\sigma} = \left[\sum_{i=1}^p (t_{(i)} - \hat{\mu})^2 \{ \hat{F}[t_{(i)}] - \hat{F}[t_{(i-1)}] \} \right]^{1/2} \quad (3)$$

$$\hat{\sigma}_{\hat{\mu}} = \left[\sum_{i=1}^{p-1} a_i^2 \frac{m_{i+1}}{n_{i+1}(n_{i+1} - m_{i+1})} \right]^{1/2} \quad (4)$$

where

$$t_{(0)} = 0; \hat{F}(t_{(0)}) = 0 \quad (5)$$

by definition, and

$$a_i = \sum_{j=1}^i (t_{(j+1)} - t_{(j)}) \hat{F}[t_{(j)}]; \quad i = 1, 2, \dots, p-1 \quad (6)$$

(USEPA, 2009, p.15-10; Singh et al., 2010, p.110; Gillespie et al., 2010; Beal, 2010). Kaplan and Meier suggest a bias correction for the estimated standard error of the mean (Lee and Wang, 2003,

p.75). When `correct.se=TRUE`, Equation (4) becomes:

$$\hat{\sigma}_{\hat{\mu}} = \left[\frac{n}{n-1} \sum_{i=1}^{p-1} a_i^2 \frac{m_{i+1}}{n_{i+1}(n_{i+1} - m_{i+1})} \right]^{1/2} \quad (7)$$

For right censored data, the equations are:

$$\hat{\mu} = \sum_{i=1}^p (t_{(i)} - t_{(i-1)}) \hat{S}[t_{(i-1)}] \quad (8)$$

$$\hat{\sigma}_{\hat{\mu}} = \left[\sum_{j=1}^{p-1} \frac{A_j^2}{(N - r_j)(N - r_j + 1)} \right]^{1/2} \quad (9)$$

where A_j denotes the area under the estimated survival curve to the right of $t_{(j)}$:

$$A_j = \sum_{i=j}^p i = j^{p-1} \hat{S}[t_{(j)}] (t_{(i+1)} - t_{(i)}) \quad (10)$$

(Lee and Wang, 2003, pp.74–75; Helsel, 2012, pp.70-76). When `correct.se=TRUE`, Equation (9) becomes:

$$\hat{\sigma}_{\hat{\mu}} = \left[\frac{n}{n-1} \sum_{j=1}^{p-1} \frac{A_j^2}{(N - r_j)(N - r_j + 1)} \right]^{1/2} \quad (11)$$

(Lee and Wang, 2003, p.75).

Confidence Intervals

This section explains how confidence intervals for the mean μ are computed.

Normal Approximation (`ci.method="normal.approx"`)

This method constructs approximate $(1 - \alpha)100\%$ confidence intervals for μ based on the assumption that the estimator of μ is approximately normally distributed. That is, a two-sided $(1 - \alpha)100\%$ confidence interval for μ is constructed as:

$$[\hat{\mu} - t_{1-\alpha/2, m-1} \hat{\sigma}_{\hat{\mu}}, \hat{\mu} + t_{1-\alpha/2, m-1} \hat{\sigma}_{\hat{\mu}}] \quad (12)$$

where $\hat{\mu}$ denotes the estimate of μ , $\hat{\sigma}_{\hat{\mu}}$ denotes the estimated asymptotic standard deviation of the estimator of μ , m denotes the assumed sample size for the confidence interval, and $t_{p, \nu}$ denotes the p 'th quantile of [Student's t-distribution](#) with ν degrees of freedom. One-sided confidence intervals are computed in a similar fashion.

The argument `ci.sample.size` determines the value of m . By default, it is the observed number of uncensored observations. This is simply an ad-hoc method of constructing confidence intervals and is not based on any published theoretical results.

When `pivot.statistic="z"`, the p 'th quantile from the [standard normal distribution](#) is used in place of the p 'th quantile from Student's t-distribution.

Bootstrap and Bias-Corrected Bootstrap Approximation (`ci.method="bootstrap"`)

The bootstrap is a nonparametric method of estimating the distribution (and associated distribution parameters and quantiles) of a sample statistic, regardless of the distribution of the population from which the sample was drawn. The bootstrap was introduced by Efron (1979) and a general reference is Efron and Tibshirani (1993).

In the context of deriving an approximate $(1 - \alpha)100\%$ confidence interval for the population mean μ , the bootstrap can be broken down into the following steps:

1. Create a bootstrap sample by taking a random sample of size N from the observations in \underline{x} , where sampling is done with replacement. Note that because sampling is done with replacement, the same element of \underline{x} can appear more than once in the bootstrap sample. Thus, the bootstrap sample will usually not look exactly like the original sample (e.g., the number of censored observations in the bootstrap sample will often differ from the number of censored observations in the original sample).
2. Estimate μ based on the bootstrap sample created in Step 1, using the same method that was used to estimate μ using the original observations in \underline{x} . Because the bootstrap sample usually does not match the original sample, the estimate of μ based on the bootstrap sample will usually differ from the original estimate based on \underline{x} .
3. Repeat Steps 1 and 2 B times, where B is some large number. For the function `enparCensored`, the number of bootstraps B is determined by the argument `n.bootstraps` (see the section ARGUMENTS above). The default value of `n.bootstraps` is 1000.
4. Use the B estimated values of μ to compute the empirical cumulative distribution function of this estimator of μ (see `ecdfPlot`), and then create a confidence interval for μ based on this estimated cdf.

The percentile interval (Efron and Tibshirani, 1993, p.170) is computed as:

$$[\hat{G}^{-1}(\frac{\alpha}{2}), \hat{G}^{-1}(\frac{1-\alpha}{2})] \quad (13)$$

where $\hat{G}(t)$ denotes the empirical cdf evaluated at t and thus $\hat{G}^{-1}(p)$ denotes the p 'th empirical quantile, that is, the p 'th quantile associated with the empirical cdf. The function `enparCensored` calls the R function `quantile` to compute the empirical quantiles used in equation (13).

The percentile method bootstrap confidence interval is only first-order accurate (Efron and Tibshirani, 1993, pp.187-188), meaning that the probability that the confidence interval will contain the true value of μ can be off by k/\sqrt{N} , where k is some constant. Efron and Tibshirani (1993, pp.184-188) proposed a bias-corrected and accelerated interval that is second-order accurate, meaning that the probability that the confidence interval will contain the true value of μ may be off by k/N instead of k/\sqrt{N} . The bias-corrected and accelerated confidence interval is computed as:

$$[\hat{G}^{-1}(\alpha_1), \hat{G}^{-1}(\alpha_2)] \quad (14)$$

where

$$\alpha_1 = \Phi[\hat{z}_0 + \frac{\hat{z}_0 + z_{\alpha/2}}{1 - \hat{a}(z_0 + z_{\alpha/2})}] \quad (15)$$

$$\alpha_2 = \Phi[\hat{z}_0 + \frac{\hat{z}_0 + z_{1-\alpha/2}}{1 - \hat{a}(z_0 + z_{1-\alpha/2})}] \quad (16)$$

$$\hat{z}_0 = \Phi^{-1}[\hat{G}(\hat{\mu})] \quad (17)$$

$$\hat{a} = \frac{\sum_{i=1}^N (\hat{\mu}_{(\cdot)} - \hat{\mu}_{(i)})^3}{6[\sum_{i=1}^N (\hat{\mu}_{(\cdot)} - \hat{\mu}_{(i)})^2]^{3/2}} \quad (18)$$

where the quantity $\hat{\mu}_{(i)}$ denotes the estimate of μ using all the values in \underline{x} except the i 'th one, and

$$\hat{\mu}_{(\cdot)} = \frac{1}{N} \sum_{i=1}^N \hat{\mu}_{(i)} \quad (19)$$

The constant \hat{z}_0 incorporates the bias correction, and the constant \hat{a} is the acceleration constant. The term “acceleration” refers to the rate of change of the standard error of the estimate of μ with respect to the true value of μ (Efron and Tibshirani, 1993, p.186). For a normal (Gaussian) distribution, the

standard error of the estimate of μ does not depend on the value of μ , hence the acceleration constant is not really necessary. The argument `use.acc.con` (see the section ARGUMENTS above) determines whether the acceleration constant \hat{a} is computed as in Equation (18) (`use.acc.con=TRUE`) or is set to 0 (`use.acc.con=FALSE`). The default value is `use.acc.con=FALSE`.

When `ci.method="bootstrap"`, the function `enparCensored` computes both the percentile method and bias-corrected bootstrap confidence intervals.

Value

a list of class "estimateCensored" containing the estimated parameters and other information. See [estimateCensored.object](#) for details.

Note

A sample of data contains censored observations if some of the observations are reported only as being below or above some censoring level. In environmental data analysis, Type I left-censored data sets are common, with values being reported as "less than the detection limit" (e.g., Helsel, 2012). Data sets with only one censoring level are called *singly censored*; data sets with multiple censoring levels are called *multiply* or *progressively censored*.

Statistical methods for dealing with censored data sets have a long history in the field of survival analysis and life testing. More recently, researchers in the environmental field have proposed alternative methods of computing estimates and confidence intervals in addition to the classical ones such as maximum likelihood estimation.

Helsel (2012, Chapter 6) gives an excellent review of past studies of the properties of various estimators based on censored environmental data.

In practice, it is better to use a confidence interval for the mean or a joint confidence region for the mean and standard deviation, rather than rely on a single point-estimate of the mean. Since confidence intervals and regions depend on the properties of the estimators for both the mean and standard deviation, the results of studies that simply evaluated the performance of the mean and standard deviation separately cannot be readily extrapolated to predict the performance of various methods of constructing confidence intervals and regions. Furthermore, for several of the methods that have been proposed to estimate the mean based on type I left-censored data, standard errors of the estimates are not available, hence it is not possible to construct confidence intervals (El-Shaarawi and Dolan, 1989).

Few studies have been done to evaluate the performance of methods for constructing confidence intervals for the mean or joint confidence regions for the mean and standard deviation when data are subjected to single or multiple censoring.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Beal, D. (2010). *A Macro for Calculating Summary Statistics on Left Censored Environmental Data Using the Kaplan-Meier Method*. Paper SDA-09, presented at Southeast SAS Users Group 2010, September 26-28, Savannah, GA.
- Efron, B. (1979). Bootstrap Methods: Another Look at the Jackknife. *The Annals of Statistics* **7**, 1–26.
- Efron, B., and R.J. Tibshirani. (1993). *An Introduction to the Bootstrap*. Chapman and Hall, New York, 436pp.

- El-Shaarawi, A.H., and D.M. Dolan. (1989). Maximum Likelihood Estimation of Water Quality Concentrations from Censored Data. *Canadian Journal of Fisheries and Aquatic Sciences* **46**, 1033–1039.
- Gillespie, B.W., Q. Chen, H. Reichert, A. Franzblau, E. Hedgeman, J. Lepkowski, P. Adriaens, A. Demond, W. Luksemburg, and D.H. Garabrant. (2010). Estimating Population Distributions When Some Data Are Below a Limit of Detection by Using a Reverse Kaplan-Meier Estimator. *Epidemiology* **21**(4), S64–S70.
- Helsel, D.R. (2012). *Statistics for Censored Environmental Data Using Minitab and R, Second Edition*. John Wiley & Sons, Hoboken, New Jersey.
- Irwin, J.O. (1949). The Standard Error of an Estimate of Expectation of Life, with Special Reference to Expectation of Tumourless Life in Experiments with Mice. *Journal of Hygiene* **47**, 188–189.
- Kaplan, E.L., and P. Meier. (1958). Nonparametric Estimation From Incomplete Observations. *Journal of the American Statistical Association* **53**, 457–481.
- Lee, E.T., and J.W. Wang. (2003). *Statistical Methods for Survival Data Analysis, Third Edition*. John Wiley & Sons, Hoboken, New Jersey, 513pp.
- Meier, P., T. Karrison, R. Chappell, and H. Xie. (2004). The Price of Kaplan-Meier. *Journal of the American Statistical Association* **99**(467), 890–896.
- Nelson, W. (1982). *Applied Life Data Analysis*. John Wiley and Sons, New York, 634pp.
- Singh, A., N. Armbya, and A. Singh. (2010). *ProUCL Version 4.1.00 Technical Guide (Draft)*. EPA/600/R-07/041, May 2010. Office of Research and Development, U.S. Environmental Protection Agency, Washington, D.C.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.
- USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

See Also

[ecdfPlotCensored](#), [qqPlotCensored](#), [estimateCensored.object](#).

Examples

```
# Example 15-1 of USEPA (2009, page 15-10) gives an example of
# estimating the mean and standard deviation nonparametrically
# using the Kaplan-Meier estimators based on censored manganese
# concentrations (ppb) in groundwater collected at 5 monitoring
# wells. The data for this example are stored in
# EPA.09.Ex.15.1.manganese.df.

# First look at the data:
#-----

EPA.09.Ex.15.1.manganese.df

#   Sample  Well Manganese.Orig.ppb Manganese.ppb Censored
#1      1 Well.1             <5             5.0      TRUE
#2      2 Well.1             12.1            12.1     FALSE
#3      3 Well.1             16.9            16.9     FALSE
```

```
#...
#23      3 Well.5              3.3          3.3  FALSE
#24      4 Well.5              8.4          8.4  FALSE
#25      5 Well.5              <2          2.0   TRUE
```

```
longToWide(EPA.09.Ex.15.1.manganese.df,
  "Manganese.Orig.ppb", "Sample", "Well",
  paste.row.name = TRUE)
```

```
#      Well.1 Well.2 Well.3 Well.4 Well.5
#Sample.1    <5    <5    <5    6.3   17.9
#Sample.2   12.1    7.7    5.3   11.9   22.7
#Sample.3   16.9   53.6   12.6    10    3.3
#Sample.4   21.6    9.5  106.3    <2    8.4
#Sample.5    <2   45.9   34.5   77.2    <2
```

```
#-----
```

```
# Following Example 15-1 in USEPA (2009, p.15-10),
# estimate the log-scale mean and standard deviation
# nonparametrically using the Kaplan-Meier method
```

```
#-----
```

```
with(EPA.09.Ex.15.1.manganese.df,
  enparCensored(log(Manganese.ppb), Censored, ci = TRUE))
```

```
#Results of Distribution Parameter Estimation
```

```
#Based on Type I Censored Data
```

```
#-----
```

```
#
#Assumed Distribution:          None
#
#Censoring Side:               left
#
#Censoring Level(s):           0.6931472 1.6094379
#
#Estimated Parameter(s):       mean    = 2.4144518
#                               sd      = 1.0478026
#                               se.mean = 0.2095605
#
#Estimation Method:            Kaplan-Meier
#
#Data:                         log(Manganese.ppb)
#
#Censoring Variable:           Censored
#
#Sample Size:                  25
#
#Percent Censored:             24%
#
#Confidence Interval for:      mean
#
#Confidence Interval Method:    Normal Approximation
#
#Confidence Interval Type:      two-sided
#
#Confidence Level:             95%
#
```



```

#Confidence Interval:          LCL = 2.003721
#                             UCL = 2.825183

#-----

# Of course, if you are estimating the mean nonparametrically,
# there is no reason to use a log transformation to attempt
# to induce normality.

# Now estimate the mean and standard deviation on the
# original scale nonparametrically using the
# Kaplan-Meier method.
#-----

with(EPA.09.Ex.15.1.manganese.df,
     enparCensored(Manganese.ppb, Censored, ci = TRUE))

#Results of Distribution Parameter Estimation
#Based on Type I Censored Data
#-----
#
#Assumed Distribution:          None
#
#Censoring Side:                left
#
#Censoring Level(s):           2 5
#
#Estimated Parameter(s):       mean    = 20.140000
#                               sd      = 25.129918
#                               se.mean =  5.025984
#
#Estimation Method:            Kaplan-Meier
#
#Data:                          Manganese.ppb
#
#Censoring Variable:           Censored
#
#Sample Size:                   25
#
#Percent Censored:              24%
#
#Confidence Interval for:       mean
#
#Confidence Interval Method:    Normal Approximation
#
#Confidence Interval Type:      two-sided
#
#Confidence Level:              95%
#
#Confidence Interval:          LCL = 10.28925
#                               UCL = 29.99075

```

Description

Daily measurements of ozone concentration, wind speed, temperature, and solar radiation in New York City for 153 consecutive days between May 1 and September 30, 1973.

Usage

```
Environmental.df  
Air.df
```

Format

The data frame `Environmental.df` has 153 observations on the following 4 variables.

ozone Average ozone concentration (of hourly measurements) of in parts per billion.

radiation Solar radiation (from 08:00 to 12:00) in langleys.

temperature Maximum daily temperature in degrees Fahrenheit.

wind Average wind speed (at 07:00 and 10:00) in miles per hour.

Row names are the dates the data were collected.

The data frame `Air.df` is the same as `Environmental.df` except that the column ozone is the cube root of average ozone concentration.

Details

Data on ozone (ppb), solar radiation (langleys), temperature (degrees Fahrenheit), and wind speed (mph) for 153 consecutive days between May 1 and September 30, 1973. These data are a superset of the data contained in the data frame `environmental` in the package **lattice**.

Source

Chambers et al. (1983), pp. 347-349.

References

Chambers, J.M., W.S. Cleveland, B. Kleiner, and P.A. Tukey. (1983). *Graphical Methods for Data Analysis*. Duxbury Press, Boston, MA, 395pp.

Cleveland, W.S. (1993). *Visualizing Data*. Hobart Press, Summit, New Jersey, 360pp.

Cleveland, W.S. (1994). *The Elements of Graphing Data*. Revised Edition. Hobart Press, Summit, New Jersey, 297pp.

Examples

```
# Scatterplot matrix  
pairs(Environmental.df)  
  
pairs(Air.df)  
  
# Time series plot for ozone  
attach(Environmental.df)  
dates <- as.Date(row.names(Environmental.df), format = "%m/%d/%Y")  
plot(dates, ozone, type = "l",  
      xlab = "Time (Year = 1973)", ylab = "Ozone (ppb)",
```

```
main = "Time Series Plot of Daily Ozone Measures")
detach("Environmental.df")
rm(dates)
```

EPA.02d.Ex.2.ug.per.L.vec*Concentrations in Exhibit 2 of 2002d USEPA Guidance Document*

Description

Concentrations (*mug*/L) from an exposure unit.

Usage

```
data(EPA.02d.Ex.2.ug.per.L.vec)
```

Format

a numeric vector of concentrations (*mug*/L)

Source

USEPA. (2002d). *Calculating Upper Confidence Limits for Exposure Point Concentrations at Hazardous Waste Sites*. OSWER 9285.6-10, December 2002. Office of Emergency and Remedial Response, U.S. Environmental Protection Agency, Washington, D.C., p. 9.

EPA.02d.Ex.4.mg.per.kg.vec*Concentrations in Exhibit 4 of 2002d USEPA Guidance Document*

Description

Concentrations (mg/kg) from an exposure unit.

Usage

```
data(EPA.02d.Ex.4.mg.per.kg.vec)
```

Format

a numeric vector of concentrations (mg/kg)

Source

USEPA. (2002d). *Calculating Upper Confidence Limits for Exposure Point Concentrations at Hazardous Waste Sites*. OSWER 9285.6-10, December 2002. Office of Emergency and Remedial Response, U.S. Environmental Protection Agency, Washington, D.C., p. 11.

EPA.02d.Ex.6.mg.per.kg.vec

Concentrations in Exhibit 6 of 2002d USEPA Guidance Document

Description

Concentrations (mg/kg) from an exposure unit.

Usage

data(EPA.02d.Ex.6.mg.per.kg.vec)

Format

a numeric vector of concentrations (mg/kg)

Source

USEPA. (2002d). *Calculating Upper Confidence Limits for Exposure Point Concentrations at Hazardous Waste Sites*. OSWER 9285.6-10, December 2002. Office of Emergency and Remedial Response, U.S. Environmental Protection Agency, Washington, D.C., p. 13.

EPA.02d.Ex.9.mg.per.L.vec

Concentrations in Exhibit 9 of 2002d USEPA Guidance Document

Description

Concentrations (mg/L) from an exposure unit.

Usage

data(EPA.02d.Ex.9.mg.per.L.vec)

Format

a numeric vector of concentrations (mg/L)

Source

USEPA. (2002d). *Calculating Upper Confidence Limits for Exposure Point Concentrations at Hazardous Waste Sites*. OSWER 9285.6-10, December 2002. Office of Emergency and Remedial Response, U.S. Environmental Protection Agency, Washington, D.C., p. 16.

EPA.09.Ex.10.1.nickel.df*Nickel Concentrations from Example 10-1 of 2009 USEPA Guidance Document*

Description

Nickel concentrations (ppb) from four wells (five observations per year for each well). The Guidance Document has the label “Year” instead of “Well”; corrected in Errata.

Usage

EPA.09.Ex.10.1.nickel.df

Format

A data frame with 20 observations on the following 3 variables.

Month a numeric vector indicating the month the sample was taken

Well a factor indicating the well number

Nickel.ppb a numeric vector of nickel concentrations (ppb)

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery, Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C., p.10-12.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.11.1.arsenic.df*Arsenic Concentrations from Example 11-1 of 2009 USEPA Guidance Document*

Description

Arsenic concentrations (ppb) at six wells (four observations per well).

Usage

EPA.09.Ex.11.1.arsenic.df

Format

A data frame with 24 observations on the following 3 variables.

Arsenic.ppb a numeric vector of arsenic concentrations (ppb)

Month a factor indicating the month of collection

Well a factor indicating the well number

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.11-3.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.12.1.ccl4.df

Carbon Tetrachloride Concentrations from Example 12-1 of 2009 USEPA Guidance Document

Description

Carbon tetrachloride (CCL4) concentrations (ppb) at five background wells (four measures at each well).

Usage

EPA.09.Ex.12.1.ccl4.df

Format

A data frame with 20 observations on the following 2 variables.

Well a factor indicating the well number

CCL4.ppb a numeric vector of CCL4 concentrations (ppb)

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.12-3.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.12.4.naphthalene.df

Naphthalene Concentrations from Example 12-4 of 2009 USEPA Guidance Document

Description

Naphthalene concentrations (ppb) at five background wells (five quarterly measures at each well).

Usage

EPA.09.Ex.12.4.naphthalene.df

Format

A data frame with 25 observations on the following 3 variables.

Quarter a numeric vector indicating the quarter the sample was taken

Well a factor indicating the well number

Naphthalene.ppb a numeric vector of naphthalene concentrations (ppb)

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.12-12.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.13.1.iron.df

Iron Concentrations from Example 13-1 of 2009 USEPA Guidance Document

Description

Dissolved iron (Fe) concentrations (ppm) at six upgradient wells (four quarterly measures at each well).

Usage

EPA.09.Ex.13.1.iron.df

Format

A data frame with 24 observations on the following 4 variables.

Month a numeric vector indicating the month the sample was taken

Year a numeric vector indicating the year the sample was taken

Well a factor indicating the well number

Iron.ppm a numeric vector of iron concentrations (ppm)

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.13-3.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.14.1.manganese.df

Manganese Concentrations from Example 14-1 of 2009 USEPA Guidance Document

Description

Manganese concentrations (ppm) at four background wells (eight quarterly measures at each well).

Usage

EPA.09.Ex.14.1.manganese.df

Format

A data frame with 32 observations on the following 3 variables.

Quarter a numeric vector indicating the quarter the sample was taken

Well a factor indicating the well number

Manganese.ppm a numeric vector of manganese concentrations (ppm)

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.14-5.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.14.3.alkalinity.df

Alkalinity Measures from Example 14-3 of 2009 USEPA Guidance Document

Description

Alkalinity measures (mg/L) collected from leachate at a solid waste landfill during a four and a half year period.

Usage

EPA.09.Ex.14.3.alkalinity.df

Format

A data frame with 54 observations on the following 2 variables.

Date a Date object indicating the date of collection

Alkalinity.mg.per.L a numeric vector of alkalinity measures (mg/L)

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.14-14.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.14.4.arsenic.df

Arsenic Concentrations from Example 14-4 of 2009 USEPA Guidance Document

Description

Sixteen quarterly measures of arsenic concentrations (ppb).

Usage

EPA.09.Ex.14.4.arsenic.df

Format

A data frame with 16 observations on the following 4 variables.

Sample.Date a factor indicating the month and year of collection

Month a factor indicating the month of collection

Year a factor indicating the year of collection

Arsenic.ppb a numeric vector of arsenic concentrations (ppb)

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.14-18.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.14.8.df

Analyte Concentrations from Example 14-8 of 2009 USEPA Guidance Document

Description

Monthly unadjusted and adjusted analyte concentrations over a 3-year period. Adjusted concentrations are computed by subtracting the monthly mean and adding the overall mean.

Usage

EPA.09.Ex.14.8.df

Format

A data frame with 36 observations on the following 4 variables.

Month a factor indicating the month of collection

Year a numeric vector indicating the year of collection

Unadj.Conc a numeric vector of unadjusted concentrations

Adj.Conc a numeric vector adjusted concentrations

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.14-32.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.15.1.manganese.df

Manganese Concentrations from Example 15-1 of 2009 USEPA Guidance Document

Description

Manganese concentrations (ppb) at five background wells (five measures at each well).

Usage

EPA.09.Ex.15.1.manganese.df

Format

A data frame with 25 observations on the following 5 variables.

Sample a numeric vector indicating the sample number (1-5)

Well a factor indicating the well number

Manganese.Orig.ppb a character vector of the original manganese concentrations (ppb)

Manganese.ppb a numeric vector of manganese concentrations with non-detects coded to their detection limit

Censored a logical vector indicating which observations are censored

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.15-10.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.16.1.sulfate.df

Sulfate Concentrations from Example 16-1 of 2009 USEPA Guidance Document

Description

Sulfate concentrations (ppm) at one background well and one downgradient well (eight quarterly measures at each well).

Usage

EPA.09.Ex.16.1.sulfate.df

Format

A data frame with 16 observations on the following 4 variables.

Month a factor indicating the month of collection

Year a factor indicating the year of collection

Well.type a factor indicating the well type (background vs. downgradient)

Sulfate.ppm a numeric vector of sulfate concentrations (ppm)

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.16-6.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.16.2.benzene.df

Benzene Concentrations from Example 16-2 of 2009 USEPA Guidance Document

Description

Benzene concentrations (ppb) at one background and one downgradient well (eight monthly measures at each well).

Usage

EPA.09.Ex.16.2.benzene.df

Format

A data frame with 16 observations on the following 3 variables.

Month a factor indicating the month of collection

Well.type a factor indicating the well type (background vs. downgradient)

Benzene.ppb a numeric vector of benzene concentrations (ppb)

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.16-9.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.16.4.copper.df

Copper Concentrations from Example 16-4 of 2009 USEPA Guidance Document

Description

Copper concentrations (ppb) at two background wells and one compliance well (six measures at each well).

Usage

EPA.09.Ex.16.4.copper.df

Format

A data frame with 18 observations on the following 4 variables.

Month a factor indicating the month of collection

Well a factor indicating the well number

Well.type a factor indicating the well type (background vs. compliance)

Copper.ppb a numeric vector of copper concentrations (ppb)

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.16-19.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.16.5.PCE.df *Tetrachloroethylene Concentrations from Example 16-5 of 2009 USEPA Guidance Document*

Description

Tetrachloroethylene (PCE) concentrations (ppb) at one background well and one compliance well.

Usage

EPA.09.Ex.16.5.PCE.df

Format

A data frame with 14 observations on the following 4 variables.

Well.type a factor with levels Background Compliance

PCE.Orig.ppb a character vector of original PCE concentrations (ppb)

PCE.ppb a numeric vector of PCE concentrations (ppb) with nondetects set to their detection limit

Censored a logical vector indicating which observations are censored

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.16-22.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.17.1.loglead.df

*Log-transformed Lead Concentrations from Example 17-1 of 2009
USEPA Guidance Document*

Description

Log-transformed lead concentrations (ppb) at two background and four compliance wells (four quarterly measures at each well).

Usage

EPA.09.Ex.17.1.loglead.df

Format

A data frame with 24 observations on the following 4 variables.

Month a factor indicating the month of collection; 1 = Jan, 2 = Apr, 3 = Jul, 4 = Oct

Well a factor indicating the well number

Well.type a factor indicating the well type (background vs. compliance)

LogLead a numeric vector of log-transformed lead concentrations (ppb)

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.17-7.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.17.2.toluene.df

Toluene Concentrations from Example 17-2 of 2009 USEPA Guidance Document

Description

Toluene concentrations (ppb) at two background and three compliance wells (five monthly measures at each well).

Usage

EPA.09.Ex.17.2.toluene.df

Format

A data frame with 25 observations on the following 6 variables.

Month a factor indicating the month of collection

Well a factor indicating the well number

Well.type a factor indicating the well type (background vs. compliance)

Toluene.ppb.orig a character vector of original toluene concentrations (ppb)

Toluene.ppb a numeric vector of toluene concentrations (ppb) with nondetects set to their detection limit

Censored a logical vector indicating which observations are censored

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.17-13.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.17.3.chrysene.df

Chrysene Concentrations from Example 17-3 of 2009 USEPA Guidance Document

Description

Chrysene concentrations (ppb) at two background and three compliance wells (four monthly measures at each well).

Usage

EPA.09.Ex.17.3.chrysene.df

Format

A data frame with 20 observations on the following 4 variables.

Month a factor indicating the month of collection

Well a factor indicating the well number

Well.type a factor indicating the well type (background vs. compliance)

Chrysene.ppb a numeric vector of chrysene concentrations (ppb)

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.17-17.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.17.3.log.chrysene.df

*Log-transformed Chrysene Concentrations from Example 17-3 of
2009 USEPA Guidance Document*

Description

Log-transformed chrysene concentrations (ppb) at two background and three compliance wells (four monthly measures at each well).

Usage

EPA.09.Ex.17.3.log.chrysene.df

Format

A data frame with 20 observations on the following 4 variables.

Month a factor indicating the month of collection

Well a factor indicating the well number

Well.type a factor indicating the well type (background vs. compliance)

Log.Chrysene.ppb a numeric vector of log-transformed chrysene concentrations (ppb)

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.17-18.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.17.4.copper.df

Copper Concentrations from Example 17-4 of 2009 USEPA Guidance Document

Description

Copper concentrations (ppb) at three background and two compliance wells (eight monthly measures at the background wells, four monthly measures at the compliance wells).

Usage

EPA.09.Ex.17.4.copper.df

Format

A data frame with 40 observations on the following 6 variables.

Month a factor indicating the month of collection

Well a factor indicating the well number

Well.type a factor indicating the well type (background vs. compliance)

Copper.ppb.orig a character vector of original copper concentrations (ppb)

Copper.ppb a numeric vector of copper concentrations with nondetects set to their detection limit

Censored a logical vector indicating which observations are censored

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.17-21.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.17.5.chloride.df

Chloride Concentrations from Example 17-5 of 2009 USEPA Guidance Document

Description

Chloride concentrations (ppm) collected over a five-year period at a solid waste landfill.

Usage

EPA.09.Ex.17.5.chloride.df

Format

A data frame with 19 observations on the following 4 variables.

Date a Date object indicating the date of collection

Chloride.ppm a numeric vector of chloride concentrations (ppm)

Elapsed.Days a numeric vector indicating the number of days since January 1, 2002

Residuals a numeric vector of residuals from a linear regression trend fit

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.17-26.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.17.6.sulfate.df

Sulfate Concentrations from Example 17-6 of 2009 USEPA Guidance Document

Description

Sulfate concentrations (ppm) collected over several years. The date of collection is simply indicated by month and year of collection. The column Date is a Date object where the day of the month has been arbitrarily set to 1.

Usage

EPA.09.Ex.17.6.sulfate.df

Format

A data frame with 23 observations on the following 6 variables.

Sample.No a numeric vector indicating the sample number

Year a numeric vector indicating the year of collection

Month a numeric vector indicating the month of collection

Sampling.Date a numeric vector indicating the year and month of collection

Date a Date object indicating the date of collection, where the day of the month is arbitrarily set to 1

Sulfate.ppm a numeric vector of sulfate concentrations (ppm)

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.17-33.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.17.7.sodium.df

Sodium Concentrations from Example 17-7 of 2009 USEPA Guidance Document

Description

Sodium concentrations (ppm) collected over several years. The sample dates are recorded as the year of collection (2-digit format) plus a fractional part indicating when during the year the sample was collected.

Usage

EPA.09.Ex.17.7.sodium.df

Format

A data frame with 10 observations on the following 2 variables.

Year a numeric vector indicating the year of collection (a fractional number)

Sodium.ppm a numeric vector of sodium concentrations (ppm)

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.17-36.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.18.1.arsenic.df

Arsenic Concentrations from Example 18-1 of 2009 USEPA Guidance Document

Description

Arsenic concentrations (ppb) in a single well at a solid waste landfill. Four observations per year over four years. Years 1-3 are the background period and Year 4 is the compliance period.

Usage

EPA.09.Ex.18.1.arsenic.df

Format

A data frame with 16 observations on the following 3 variables.

Year a factor indicating the year of collection

Sampling.Period a factor indicating the sampling period (background vs. compliance)

Arsenic.ppb a numeric vector of arsenic concentrations (ppb)

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.18-10.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.18.2.chrysene.df

Chrysene Concentrations from Example 18-2 of 2009 USEPA Guidance Document

Description

Chrysene concentrations (ppb) at two background wells and one compliance well (four monthly measures at each well).

Usage

EPA.09.Ex.18.2.chrysene.df

Format

A data frame with 12 observations on the following 4 variables.

Month a factor indicating the month of collection

Well a factor indicating the well number

Well.type a factor indicating the well type (background vs. compliance)

Chrysene.ppb a numeric vector of chrysene concentrations (ppb)

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.18-15.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.18.3.TCE.df *Trichloroethylene Concentrations from Example 18-3 of 2009 USEPA Guidance Document*

Description

Trichloroethylene (TCE) concentrations (ppb) at three background wells and one compliance well. Six monthly measures at each background well, three monthly measures at the compliance well.

Usage

EPA.09.Ex.18.3.TCE.df

Format

A data frame with 24 observations on the following 6 variables.

Month a factor indicating the month of collection

Well a factor indicating the well number

Well.type a factor indicating the well type (background vs. compliance)

TCE.ppb.orig a character vector of original TCE concentrations (ppb)

TCE.ppb a numeric vector of TCE concentrations (ppb) with nondetects set to their detection limit

Censored a logical vector indicating which observations are censored

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.18-19.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.18.4.xylene.df

Xylene Concentrations from Example 18-4 of 2009 USEPA Guidance Document

Description

Xylene concentrations (ppb) at three background wells and one compliance well. Eight monthly measures at each compliance well; three monthly measures at the compliance well.

Usage

EPA.09.Ex.18.4.xylene.df

Format

A data frame with 32 observations on the following 6 variables.

Month a factor indicating the month of collection

Well a factor indicating the well number

Well.type a factor indicating the well type (background vs. compliance)

Xylene.ppb.orig a character vector of original xylene concentrations (ppb)

Xylene.ppb a numeric vector of xylene concentrations (ppb) with nondetects set to their detection limit

Censored a logical vector indicating which observations are censored

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.18-22.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.19.1.sulfate.df

Sulfate Concentrations from Example 19-1 of 2009 USEPA Guidance Document

Description

Sulfate concentrations (mg/L) at four background wells.

Usage

EPA.09.Ex.19.1.sulfate.df

Format

A data frame with 25 observations on the following 7 variables.

Well a factor indicating the well number

Month a numeric vector indicating the month of collection

Day a numeric vector indicating the day of the month of collection

Year a numeric vector indicating the year of collection

Date a Date object indicating the date of collection

Sulfate.mg.per.l a numeric vector of sulfate concentrations (mg/L)

log.Sulfate.mg.per.l a numeric vector of log-transformed sulfate concentrations (mg/L)

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.19-17.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.19.2.chloride.df

Chloride Concentrations from Example 19-2 of 2009 USEPA Guidance Document

Description

Chloride concentrations (mg/L) at 10 compliance wells at a solid waste landfill. One year of quarterly measures at each well.

Usage

EPA.09.Ex.19.2.chloride.df

Format

A data frame with 40 observations on the following 2 variables.

Well a factor indicating the well number

Chloride.mg.per.l a numeric vector of chloride concentrations (mg/L)

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.19-19.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.19.5.mercury.df

Mercury Concentrations from Example 19-5 of 2009 USEPA Guidance Document

Description

Mercury concentrations (ppb) at four background and two compliance wells.

Usage

EPA.09.Ex.19.5.mercury.df

Format

A data frame with 36 observations on the following 6 variables.

Event a factor indicating the time of collection

Well a factor indicating the well number

Well.type a factor indicating the well type (background vs. compliance)

Mercury.ppb.orig a character vector of original mercury concentrations (ppb)

Mercury.ppb a numeric vector of mercury concentrations (ppb) with nondetects set to their detection limit

Censored a logical vector indicating which observations are censored

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.19-33.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.20.1.nickel.df

Nickel Concentrations from Example 20-1 of 2009 USEPA Guidance Document

Description

Nickel concentrations (ppb) at a single well. Eight monthly measures during the background period and eight monthly measures during the compliance period.

Usage

EPA.09.Ex.20.1.nickel.df

Format

A data frame with 16 observations on the following 4 variables.

Month a factor indicating the month of collection

Year a factor indicating the year of collection

Period a factor indicating the period (baseline vs. compliance)

Nickel.ppb a numeric vector of nickel concentrations (ppb)

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.20-4.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.21.1.aldicarb.df

Aldicarb Concentrations from Example 21-1 of 2009 USEPA Guidance Document

Description

Aldicarb concentrations (ppb) at three compliance wells (four monthly measures at each well).

Usage

EPA.09.Ex.21.1.aldicarb.df

Format

A data frame with 12 observations on the following 3 variables.

Month a factor indicating the month of collection

Well a factor indicating the well number

Aldicarb.ppb a numeric vector of aldicarb concentrations (ppb)

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.21-4.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.21.2.benzene.df

Benzene Concentrations from Example 21-2 of 2009 USEPA Guidance Document

Description

Benzene concentrations (ppb) collected at a landfill that previously handled smelter waste and is now undergoing remediation efforts.

Usage

EPA.09.Ex.21.2.benzene.df

Format

A data frame with 8 observations on the following 4 variables.

Month a numeric vector indicating the month of collection

Benzene.ppb.orig a character vector of original benzene concentrations (ppb)

Benzene.ppb a numeric vector of benzene concentrations (ppb) with nondetects set to their detection limit

Censored a logical vector indicating which observations are censored

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.21-7.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.21.5.beryllium.df

Beryllium Concentrations from Example 21-5 of 2009 USEPA Guidance Document

Description

Beryllium concentrations (ppb) at one well (four years of quarterly measures).

Usage

data(EPA.09.Ex.21.5.beryllium.df)

Format

A data frame with 16 observations on the following 3 variables.

Year a factor indicating the year of collection

Quarter a factor indicating the quarter of collection

Beryllium.ppb a numeric vector of beryllium concentrations (ppb)

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.21-18.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.21.6.nitrate.df

Nitrate Concentrations from Example 21-6 of 2009 USEPA Guidance Document

Description

Nitrate concentrations (mg/L) at a well used for drinking water.

Usage

EPA.09.Ex.21.6.nitrate.df

Format

A data frame with 12 observations on the following 5 variables.

Sampling.Date a character vector indicating the sampling date

Date a Date object indicating the sampling date

Nitrate.mg.per.l.orig a character vector of original nitrate concentrations (mg/L)

Nitrate.mg.per.l a numeric vector of nitrate concentrations (mg/L) with nondetects set to their detection limit

Censored a logical vector indicating which observations are censored

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.21-22.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.21.7.TCE.df *Trichloroethylene Concentrations from Example 21-7 of 2009 USEPA Guidance Document*

Description

Trichloroethylene (TCE) concentrations (ppb) at a site undergoing remediation.

Usage

EPA.09.Ex.21.7.TCE.df

Format

A data frame with 10 observations on the following 2 variables.

Month a numeric vector indicating the month of collection

TCE.ppb a numeric vector of TCE concentrations (ppb)

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.21-26.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.22.1.VC.df *Vinyl Chloride Concentrations from Example 22-1 of 2009 USEPA Guidance Document*

Description

Vinyl Chloride (VC) concentrations (ppb) during detection monitoring for two compliance wells. Four years of quarterly measures at each well. Compliance monitoring began with Year 2 of the sampling record.

Usage

EPA.09.Ex.22.1.VC.df

Format

A data frame with 32 observations on the following 5 variables.

Year a factor indicating the year of collection

Quarter a factor indicating the quarter of collection

Period a factor indicating the period (background vs. compliance)

Well a factor indicating the well number

VC.ppb a numeric vector of VC concentrations (ppb)

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.22-6.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.22.2.Specific.Conductance.df

Specific Conductance from Example 22-2 of 2009 USEPA Guidance Document

Description

Specific conductance (μmho) collected over several years at two wells at a hazardous waste facility.

Usage

EPA.09.Ex.22.2.Specific.Conductance.df

Format

A data frame with 43 observations on the following 3 variables.

Well a factor indicating the well number

Date a Date object indicating the date of collection

Specific.Conductance.umho a numeric vector of specific conductance (μmho)

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.22-11.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.6.3.sulfate.df

Sulfate Concentrations from Example 6-3 of 2009 USEPA Guidance Document

Description

Sulfate concentrations (ppm) at two background wells (five quarterly measures at each well).

Usage

EPA.09.Ex.6.3.sulfate.df

Format

A data frame with 10 observations on the following 4 variables.

Month a numeric vector indicating the month the observations was taken

Year a numeric vector indicating the year the observation was taken

Well a factor indicating the well number

Sulfate.ppm a numeric vector of sulfate concentrations (ppm)

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.6-20.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Ex.7.1.arsenic.df

Arsenic concentrations from Example 7.1 of 2009 USEPA Guidance Document

Description

Arsenic concentrations ($\mu\text{g/L}$) at a single well, consisting of: 8 historical observations, 4 future observations for Case 1, and 4 future observations for Case 2.

Usage

EPA.09.Ex.7.1.arsenic.df

Format

A data frame with 16 observations on the following 2 variables.

Data.Source a factor with levels Historical, Case.1, Case.2

Arsenic.ug.per.l a numeric vector of arsenic concentrations ($\mu\text{g/L}$)

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.7-26.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Table.9.1.TCE.df

Trichloroethene concentrations in Table 9.1 of 2009 USEPA Guidance Document

Description

Time series of trichloroethene (TCE) concentrations (mg/L) taken at 2 separate wells. Some observations are annotated with a data qualifier of U (nondetect) or J (estimated detected concentration).

Usage

EPA.09.Table.9.1.TCE.df

Format

A data frame with 30 observations on the following 5 variables.

Date.Collectected a factor indicating the date of collection

Date a Date object indicating the date of collection

Well a factor indicating the well number

TCE.mg.per.L a numeric vector indicating the TCE concnetrations (mg/L)

Data.Qualifier a factor indicating the data qualifier

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.9-3.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Table.9.3.df *Arsenic, Mercury and Strontium Concentrations in Table 9-3 of 2009 USEPA Guidance Document*

Description

Arsenic, mercury, and strontium concentrations (mg/L) from a single well collected approximately quarterly. Nondetects are indicated by the data qualifier U.

Usage

EPA.09.Table.9.3.df

Format

A data frame with 15 observations on the following 8 variables.

Date.Collecte d a factor indicating the date of collection

Date a Date object indicating the date of collection

Arsenic.mg.per.L a numeric vector of arsenic concentrations (mg/L)

Arsenic.Data.Qualifier a factor indicating the data qualifier for arsenic

Mercury.mg.per.L a numeric vector of mercury concentrations (mg/L)

Mercury.Data.Qualifier a factor indicating the data qualifier for mercury

Strontium.mg.per.L a numeric vector of strontium concentrations

Strontium.Data.Qualifier a factor indicating the data qualifier for strontium

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.9-13.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.09.Table.9.4.nickel.vec *Nickel Concentrations in Table 9-4 of 2009 USEPA Guidance Document*

Description

Nickel concentrations (ppb) from a single well.

Usage

EPA.09.Table.9.4.nickel.vec

Format

a numeric vector of nickel concentrations (ppb)

Source

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.9-18.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

EPA.89b.aldicarb1.df *Aldicarb Concentrations from 1989 USEPA Guidance Document*

Description

Aldicarb concentrations (ppb) at three compliance wells (four monthly samples at each well).

Usage

EPA.89b.aldicarb1.df

Format

A data frame with 12 observations on the following 3 variables.

Aldicarb Aldicarb concentrations (ppb)

Month a factor indicating the month of collection

Well a factor indicating the well number

Source

USEPA. (1989b). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities, Interim Final Guidance*. EPA/530-SW-89-026. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C. p.6-4.

EPA.89b.aldicarb2.df *Aldicarb Concentrations from 1989 USEPA Guidance Document*

Description

Aldicarb concentrations (ppm) at three compliance wells (four monthly samples at each well).

Usage

EPA.89b.aldicarb2.df

Format

A data frame with 12 observations on the following 3 variables.

Aldicarb Aldicarb concentrations (ppm)

Month a factor indicating the month of collection

Well a factor indicating the well number

Source

USEPA. (1989b). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities, Interim Final Guidance*. EPA/530-SW-89-026. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C. p.6-13.

EPA.89b.benzene.df *Benzene Concentrations from 1989 USEPA Guidance Document*

Description

Benzene concentrations (ppm) at one background and five compliance wells (four monthly samples for each well).

Usage

EPA.89b.benzene.df

Format

A data frame with 24 observations on the following 6 variables.

Benzene.orig a character vector of the original observations

Benzene a numeric vector with <1 observations coded as 1

Censored a logical vector indicating which observations are censored

Month a factor indicating the month of collection

Well a factor indicating the well number

Well.type a factor indicating the well type (background vs. compliance)

Source

USEPA. (1989b). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities, Interim Final Guidance*. EPA/530-SW-89-026. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C. p.5-18.

EPA.89b.cadmium.df	<i>Cadmium Concentrations from 1989 USEPA Guidance Document</i>
--------------------	---

Description

Cadmium concentrations (mg/L) at one set of background and one set of compliance wells. Non-detects reported as "BDL". Detection limit not given.

Usage

EPA.89b.cadmium.df

Format

A data frame with 88 observations on the following 4 variables.

Cadmium.orig a character vector of the original cadmium observations (mg/L)

Cadmium a numeric vector with BDL coded as 0

Censored a logical vector indicating which observations are censored

Well.type a factor indicating the well type (background vs. compliance)

Source

USEPA. (1989b). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities, Interim Final Guidance*. EPA/530-SW-89-026. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C. p.8-6.

EPA.89b.chlordane1.df	<i>Chlordane Concentrations from 1989 USEPA Guidance Document</i>
-----------------------	---

Description

Chlordane concentrations (ppm) in 24 water samples. Two possible phases: dissolved (18 observations) and immiscible (6 observations).

Usage

EPA.89b.chlordane1.df

Format

A data frame with 24 observations on the following 2 variables.

Chlordane Chlordane concentrations (ppm)

Phase a factor indicating the phase (dissolved vs. immiscible)

Source

USEPA. (1989b). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities, Interim Final Guidance*. EPA/530-SW-89-026. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C. p.4-8.

EPA.89b.chlordane2.df *Chlordane Concentrations from 1989 USEPA Guidance Document*

Description

Chlordane concentrations (ppb) at one background and one compliance well. Observations taken during four separate months over two years. Four replicates taken for each “month/year/well type” combination.

Usage

data(EPA.89b.chlordane2.df)

Format

A data frame with 32 observations on the following 5 variables.

Chlordane Chlordane concentration (ppb)

Month a factor indicating the month of collection

Year a numeric vector indicating the year of collection (85 or 86)

Replicate a factor indicating the replicate number

Well.type a factor indicating the well type (background vs. compliance)

Source

USEPA. (1989b). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities, Interim Final Guidance*. EPA/530-SW-89-026. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C. p.5-27.

EPA.89b.edb.df *EDB Concentrations from 1989 USEPA Guidance Document*

Description

EDB concentrations (ppb) at three compliance wells (four monthly samples at each well).

Usage

EPA.89b.edb.df

Format

A data frame with 12 observations on the following 3 variables.

EDB EDB concentrations (ppb)

Month a factor indicating the month of collection

Well a factor indicating the well number

Source

USEPA. (1989b). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities, Interim Final Guidance*. EPA/530-SW-89-026. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C. p.6-6.

EPA.89b.lead.df*Lead Concentrations from 1989 USEPA Guidance Document*

Description

Lead concentrations (ppm) at two background and four compliance wells (four monthly samples for each well).

Usage

EPA.89b.lead.df

Format

A data frame with 24 observations on the following 4 variables.

Lead Lead concentrations (ppm)

Month a factor indicating the month of collection

Well a factor indicating the well number

Well.type a factor indicating the well type (background vs. compliance)

Source

USEPA. (1989b). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities, Interim Final Guidance*. EPA/530-SW-89-026. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C. p.5-23.

EPA.89b.loglead.df	<i>Log-transformed Lead Concentrations from 1989 USEPA Guidance Document</i>
--------------------	--

Description

Log-transformed lead concentrations ($\mu\text{g/L}$) at two background and four compliance wells (four monthly samples for each well).

Usage

EPA.89b.loglead.df

Format

A data frame with 24 observations on the following 4 variables.

LogLead Natural logarithm of lead concentrations ($\mu\text{g/L}$)

Month a factor indicating the month of collection

Well a factor indicating the well number

Well.type a factor indicating the well type (background vs. compliance)

Source

USEPA. (1989b). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities, Interim Final Guidance*. EPA/530-SW-89-026. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C. p.5-11.

EPA.89b.manganese.df	<i>Manganese Concentrations from 1989 USEPA Guidance Document</i>
----------------------	---

Description

Manganese concentrations at six monitoring wells (four monthly samples for each well).

Usage

EPA.89b.manganese.df

Format

A data frame with 24 observations on the following 3 variables.

Manganese Manganese concentrations

Month a factor indicating the month of collection

Well a factor indicating the well number

Source

USEPA. (1989b). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities, Interim Final Guidance*. EPA/530-SW-89-026. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C. p.4-19.

EPA.89b.sulfate.df	<i>Sulfate Concentrations from 1989 USEPA Guidance Document</i>
--------------------	---

Description

Sulfate concentrations (mg/L). Nondetects reported as <1450.

Usage

```
data(EPA.89b.sulfate.df)
```

Format

A data frame with 24 observations on the following 3 variables.

Sulfate.orig a character vector of original sulfate concentration (mg/L)

Sulfate a numeric vector of sulfate concentrations with <1450 coded as 1450

Censored a logical vector indicating which observations are censored

Source

USEPA. (1989b). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities, Interim Final Guidance*. EPA/530-SW-89-026. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C. p.8-9.

EPA.89b.t29.df	<i>T-29 Concentrations from 1989 USEPA Guidance Document</i>
----------------	--

Description

T-29 concentrations (ppm) at two compliance wells (four monthly samples at each well, four replicates within each month). Detection limit is not given.

Usage

```
EPA.89b.t29.df
```

Format

A data frame with 32 observations on the following 6 variables.

T29.orig a character vector of the original T-29 concentrations (ppm)

T29 a numeric vector of T-29 concentrations with <? coded as 0

Censored a logical vector indicating which observations are censored

Month a factor indicating the month of collection

Replicate a factor indicating the replicate number

Well a factor indicating the well number

Source

USEPA. (1989b). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities, Interim Final Guidance*. EPA/530-SW-89-026. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C. p.6-10.

EPA.89b.toc.vec	<i>Total Organic Carbon Concentrations from 1989 USEPA Guidance Document</i>
-----------------	--

Description

Numeric vector containing total organic carbon (TOC) concentrations (mg/L).

Usage

EPA.89b.toc.vec

Format

A numeric vector with 19 elements containing TOC concentrations (mg/L).

Source

USEPA. (1989b). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities, Interim Final Guidance*. EPA/530-SW-89-026. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C. p.8-13.

EPA.92c.arsenic1.df	<i>Arsenic Concentrations from 1992 USEPA Guidance Document</i>
---------------------	---

Description

Arsenic concentrations (ppm) at six monitoring wells (four monthly samples for each well).

Usage

EPA.92c.arsenic1.df

Format

A data frame with 24 observations on the following 3 variables.

Arsenic Arsenic concentrations (ppm)

Month a factor indicating the month of collection

Well a factor indicating the well number

Source

USEPA. (1992c). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities: Addendum to Interim Final Guidance*. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C. p.21.

EPA.92c.arsenic2.df *Arsenic Concentrations from 1992 USEPA Guidance Document*

Description

Arsenic concentrations (ppb) at three background wells and one compliance well (six monthly samples for each well; first four missing at compliance well). Nondetects reported as <5.

Usage

EPA.92c.arsenic2.df

Format

A data frame with 24 observations on the following 6 variables.

Arsenic.orig a character vector of original arsenic concentrations (ppb)

Arsenic a numeric vector of arsenic concentrations with <5 coded as 5

Censored a logical vector indicating which observations are censored

Month a factor indicating the month of collection

Well a factor indicating the well number

Well.type a factor indicating the well type (background vs. compliance)

Source

USEPA. (1992c). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities: Addendum to Interim Final Guidance*. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C. p.60.

EPA.92c.arsenic3.df *Arsenic Concentrations from 1992 USEPA Guidance Document*

Description

Arsenic concentrations at one background and one compliance monitoring well. Three years of observations for background well, two years of observations for compliance well, four samples per year for each well.

Usage

EPA.92c.arsenic3.df

Format

A data frame with 20 observations on the following 3 variables.

Arsenic a numeric vector of arsenic concentrations

Year a factor indicating the year of collection

Well.type a factor indicating the well type (background vs. compliance)

Source

USEPA. (1992c). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities: Addendum to Interim Final Guidance*. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C.

EPA.92c.benzene1.df *Benzene Concentrations from 1992 USEPA Guidance Document*

Description

Benzene concentrations (ppb) at six background wells (six monthly samples for each well). Non-detects reported as <2.

Usage

EPA.92c.benzene1.df

Format

A data frame with 36 observations on the following 5 variables.

Benzene.orig a character vector of original benzene concentrations (ppb)

Benzene a numeric vector of benzene concentrations with <2 coded as 2

Censored a logical vector indicating which observations are censored

Month a factor indicating the month of collection

Well a factor indicating the well number

Source

USEPA. (1992c). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities: Addendum to Interim Final Guidance*. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C. p.36.

EPA.92c.benzene2.df *Benzene Concentrations from 1992 USEPA Guidance Document*

Description

Benzene concentrations (ppb) at one background and one compliance well. Four observations per month for each well. Background well sampled in months 1,2, and 3; compliance well sampled in months 4 and 5.

Usage

EPA.92c.benzene2.df

Format

A data frame with 20 observations on the following 3 variables.

Benzene a numeric vector of benzene concentrations (ppb)

Month a factor indicating the month of collection

Well.type a factor indicating the well type (background vs. compliance)

Source

USEPA. (1992c). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities: Addendum to Interim Final Guidance*. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C. p.56.

EPA.92c.cc14.df

Carbon Tetrachloride Concentrations from 1992 USEPA Guidance Document

Description

Carbon tetrachloride (CCL4) concentrations (ppb) at five wells (four monthly samples at each well).

Usage

EPA.92c.cc14.df

Format

A data frame with 20 observations on the following 3 variables.

CCL4 a numeric vector of carbon tetrachloride concentrations (ppb)

Month a factor indicating the month of collection

Well a factor indicating the well number

Source

USEPA. (1992c). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities: Addendum to Interim Final Guidance*. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C. p.80.

EPA.92c.chrysene.df	<i>Chrysene Concentrations from 1992 USEPA Guidance Document</i>
---------------------	--

Description

Chrysene concentrations (ppb) at five compliance wells (four monthly samples for each well).

Usage

EPA.92c.chrysene.df

Format

A data frame with 20 observations on the following 3 variables.

Chrysene a numeric vector of chrysene concentrations (ppb)

Month a factor indicating the month of collection

Well a factor indicating the well number

Source

USEPA. (1992c). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities: Addendum to Interim Final Guidance*. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C. p.52.

EPA.92c.copper1.df	<i>Copper Concentrations from 1992 USEPA Guidance Document</i>
--------------------	--

Description

Copper concentrations (ppb) at two background and one compliance wells (six monthly samples for each well).

Usage

EPA.92c.copper1.df

Format

A data frame with 18 observations on the following 4 variables.

Copper a numeric vector of copper concentrations (ppb)

Month a factor indicating the month of collection

Well a factor indicating the well number

Well.type a factor indicating the well type (background vs. compliance)

Source

USEPA. (1992c). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities: Addendum to Interim Final Guidance*. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C. p.47.

EPA.92c.copper2.df	<i>Copper Concentrations from 1992 USEPA Guidance Document</i>
--------------------	--

Description

Copper concentrations (ppb) at three background and two compliance wells (eight monthly samples for each well; first four missing at compliance wells). Nondetects reported as <5.

Usage

EPA.92c.copper2.df

Format

A data frame with 40 observations on the following 6 variables.

Copper.orig a character vector of original copper concentrations (ppb)

Copper a numeric vector of copper concentrations with <5 coded as 5

Censored a logical vector indicating which observations are censored

Month a factor indicating the month of collection

Well a factor indicating the well number

Well.type a factor indicating the well type (background vs. compliance)

Source

USEPA. (1992c). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities: Addendum to Interim Final Guidance*. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C. p.55.

EPA.92c.lognickel1.df	<i>Log-transformed Nickel Concentrations from 1992 USEPA Guidance Document</i>
-----------------------	--

Description

Log-transformed nickel concentrations (ppb) at four monitoring wells (five monthly samples for each well).

Usage

EPA.92c.lognickel1.df

Format

A data frame with 20 observations on the following 3 variables.

LogNickel a numeric vector of log-transformed nickel concentrations (ppb)

Month a factor indicating the month of collection

Well a factor indicating the well number

Source

USEPA. (1992c). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities: Addendum to Interim Final Guidance*. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C. p.15.

EPA.92c.nickel1.df	<i>Nickel Concentrations from 1992 USEPA Guidance Document</i>
--------------------	--

Description

Nickel concentrations (ppb) at four monitoring wells (five monthly samples for each well).

Usage

EPA.92c.nickel1.df

Format

A data frame with 20 observations on the following 3 variables.

Nickel a numeric vector of nickel concentrations (ppb)

Month a factor indicating the month of collection

Well a factor indicating the well number

Source

USEPA. (1992c). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities: Addendum to Interim Final Guidance*. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C. p.7.

EPA.92c.nickel2.df	<i>Nickel Concentrations from 1992 USEPA Guidance Document</i>
--------------------	--

Description

Nickel concentrations (ppb) at a monitoring well (eight months of samples, two samples for each sampling occasion).

Usage

EPA.92c.nickel2.df

Format

A data frame with 16 observations on the following 3 variables.

Nickel a numeric vector of nickel concentrations (ppb)

Month a factor indicating the month of collection

Sample a factor indicating the sample (replicate) number

Source

USEPA. (1992c). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities: Addendum to Interim Final Guidance*. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C. p.78.

EPA.92c.toluene.df	<i>Toluene Concentrations from 1992 USEPA Guidance Document</i>
--------------------	---

Description

Toluene concentrations (ppb) at two background and three compliance wells (five monthly samples at each well). Nondetects reported as <5.

Usage

EPA.92c.toluene.df

Format

A data frame with 25 observations on the following 6 variables.

Toluene.orig a character vector of original toluene concentrations (ppb)

Toluene a numeric vector of toluene concentrations with <5 coded as 5

Censored a logical vector indicating which observations are censored

Month a factor indicating the month of collection

Well a factor indicating the well number

Well.type a factor indicating the well type (background vs. compliance)

Source

USEPA. (1992c). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities: Addendum to Interim Final Guidance*. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C. p.43.

EPA.92c.zinc.df	<i>Zinc Concentrations from 1992 USEPA Guidance Document</i>
-----------------	--

Description

Zinc concentrations (ppb) at five background wells (eight samples for each well). Nondetects reported as <7.

Usage

EPA.92c.zinc.df

Format

A data frame with 40 observations on the following 5 variables.

Zinc.orig a character vector of original zinc concentrations (ppb)

Zinc a numeric vector of zinc concentrations with <7 coded as 7

Censored a logical vector indicating which observations are censored

Sample a factor indicating the sample number

Well a factor indicating the well number

Source

USEPA. (1992c). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities: Addendum to Interim Final Guidance*. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C. p.30.

EPA.92d.chromium.df	<i>Chromium Concentrations from 1992 USEPA Guidance Document</i>
---------------------	--

Description

Chromium concentrations (mg/kg) in soil samples collected randomly over a Superfund site.

Usage

EPA.92d.chromium.df

Format

A data frame with 15 observations on the following variable.

Cr a numeric vector of chromium concentrations (mg/kg)

Source

USEPA. (1992d). *Supplemental Guidance to RAGS: Calculating the Concentration Term*. Publication 9285.7-081, May 1992. Intermittent Bulletin, Volume 1, Number 1. Office of Emergency and Remedial Response, Hazardous Site Evaluation Division, OS-230. Office of Solid Waste and Emergency Response, U.S. Environmental Protection Agency, Washington, D.C.

EPA.92d.chromium.vec *Chromium Concentrations from 1992 USEPA Guidance Document*

Description

Chromium concentrations (mg/kg) in soil samples collected randomly over a Superfund site.

Usage

EPA.92d.chromium.vec

Format

A numeric vector with 15 observations.

Source

USEPA. (1992d). *Supplemental Guidance to RAGS: Calculating the Concentration Term*. Publication 9285.7-081, May 1992. Intermittent Bulletin, Volume 1, Number 1. Office of Emergency and Remedial Response, Hazardous Site Evaluation Division, OS-230. Office of Solid Waste and Emergency Response, U.S. Environmental Protection Agency, Washington, D.C.

EPA.94b.lead.df *Lead Concentrations from 1994 USEPA Guidance Document*

Description

Lead concentrations (mg/Kg) in soil samples at a reference area and a cleanup area. Nondetects reported as <39. There are 14 observations for each area.

Usage

EPA.94b.lead.df

Format

A data frame with 28 observations on the following 4 variables.

Lead.orig a character vector of original lead concentrations (mg/Kg)

Lead a numeric vector of lead concentrations with <39 coded as 39

Censored a logical vector indicating which observations are censored

Area a factor indicating the area (cleanup vs. reference)

Source

USEPA. (1994b). *Statistical Methods for Evaluating the Attainment of Cleanup Standards, Volume 3: Reference-Based Standards for Soils and Solid Media*. EPA/230-R-94-004. Office of Policy, Planning, and Evaluation, U.S. Environmental Protection Agency, Washington, D.C. pp.6.20–6.21.

EPA.94b.tccb.df	<i>1,2,3,4-Tetrachlorobenzene Concentrations from 1994 USEPA Guidance Document</i>
-----------------	--

Description

1,2,3,4-Tetrachlorobenzene (TcCB) concentrations (ppb) in soil samples at a reference area and a cleanup area. There are 47 observations for the reference area and 77 for the cleanup area. There is only one nondetect in the dataset (it's in the cleanup area), and it is reported as ND. Here it is assumed the nondetect is less than the smallest reported value, which is 0.09 ppb. Note that on page 6.23 of USEPA (1994b), a value of 25.5 for the Cleanup Unit was erroneously omitted.

Usage

EPA.94b.tccb.df

Format

A data frame with 124 observations on the following 4 variables.

TcCB.orig a character vector with the original tetrachlorobenzene concentrations (ppb)

TcCB a numeric vector of tetrachlorobenzene with <0.99 coded as 0.99

Censored a logical vector indicating which observations are censored

Area a factor indicating the area (cleanup vs. reference)

Source

USEPA. (1994b). *Statistical Methods for Evaluating the Attainment of Cleanup Standards, Volume 3: Reference-Based Standards for Soils and Solid Media*. EPA/230-R-94-004. Office of Policy, Planning, and Evaluation, U.S. Environmental Protection Agency, Washington, D.C. pp.6.22-6.25.

EPA.97.cadmium.111.df	<i>Calibration Data for Cadmium at Mass 111</i>
-----------------------	---

Description

Calibration data for cadmium at mass 111 (ng/L; method 1638 ICPMS) that appeared in Gibbons et al. (1997b) and were provided to them by the U.S. EPA.

Usage

EPA.97.cadmium.111.df

Format

A data frame with 35 observations on the following 2 variables.

Cadmium Observed concentration of cadmium (ng/L)

Spike "True" concentration of cadmium taken from a standard (ng/L)

Source

Gibbons, R.D., D.E. Coleman, and R.F. Maddalone. (1997b). Response to Comment on "An Alternative Minimum Level Definition for Analytical Quantification". *Environmental Science and Technology*, **31**(12), 3729–3731.

 epareto

Estimate Parameters of a Pareto Distribution

Description

Estimate the location and shape parameters of a [Pareto distribution](#).

Usage

```
epareto(x, method = "mle", plot.pos.con = 0.375)
```

Arguments

x	numeric vector of observations.
method	character string specifying the method of estimation. Possible values are "mle" (maximum likelihood; the default), and "lse" (least-squares). See the DETAILS section for more information on these estimation methods.
plot.pos.con	numeric scalar between 0 and 1 containing the value of the plotting position constant used to construct the values of the empirical cdf. The default value is plot.pos.con=0.375. This argument is used only when method="lse".

Details

If x contains any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

Let $\underline{x} = (x_1, x_2, \dots, x_n)$ be a vector of n observations from a [Pareto distribution](#) with parameters location= η and shape= θ .

Maximum Likelihood Estimation (method="mle")

The maximum likelihood estimators (mle's) of η and θ are given by (Evans et al., 1993; p.122; Johnson et al., 1994, p.581):

$$\hat{\eta}_{mle} = x_{(1)} \quad (1)$$

$$\hat{\theta}_{mle} = n \left[\sum_{i=1}^n \log\left(\frac{x_i}{\hat{\eta}_{mle}}\right) \right]^{-1} \quad (2)$$

where $x_{(1)}$ denotes the first order statistic (i.e., the minimum value).

Least-Squares Estimation (method="lse")

The least-squares estimators (lse's) of η and θ are derived as follows. Let X denote a [Pareto](#) random variable with parameters location= η and shape= θ . It can be shown that

$$\log[1 - F(x)] = \theta \log(\eta) - \theta \log(x) \quad (3)$$

where F denotes the cumulative distribution function of X . Set

$$y_i = \log[1 - \hat{F}(x_i)] \quad (4)$$

$$z_i = \log(x_i) \quad (5)$$

where $\hat{F}(x)$ denotes the empirical cumulative distribution function evaluated at x . The least-squares estimates of η and θ are obtained by solving the regression equation

$$y_i = \beta_0 + \beta_1 z_i \quad (6)$$

and setting

$$\hat{\theta}_{lse} = -\hat{\beta}_1 \quad (7)$$

$$\hat{\eta}_{lse} = \exp\left(\frac{\hat{\beta}_0}{\hat{\theta}_{lse}}\right) \quad (8)$$

(Johnson et al., 1994, p.580).

Value

a list of class "estimate" containing the estimated parameters and other information. See [estimate.object](#) for details.

Note

The Pareto distribution is named after Vilfredo Pareto (1848-1923), a professor of economics. It is derived from Pareto's law, which states that the number of persons N having income $\geq x$ is given by:

$$N = Ax^{-\theta}$$

where θ denotes Pareto's constant and is the shape parameter for the probability distribution.

The Pareto distribution takes values on the positive real line. All values must be larger than the "location" parameter η , which is really a threshold parameter. There are three kinds of Pareto distributions. The one described here is the Pareto distribution of the first kind. Stable Pareto distributions have $0 < \theta < 2$. Note that the r 'th moment only exists if $r < \theta$.

The Pareto distribution is related to the [exponential distribution](#) and [logistic distribution](#) as follows. Let X denote a Pareto random variable with location= η and shape= θ . Then $\log(X/\eta)$ has an exponential distribution with parameter rate= θ , and $-\log\{[(X/\eta)^\theta] - 1\}$ has a logistic distribution with parameters location=0 and scale=1.

The Pareto distribution has a very long right-hand tail. It is often applied in the study of socioeconomic data, including the distribution of income, firm size, population, and stock price fluctuations.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.
- Johnson, N. L., S. Kotz, and N. Balakrishnan. (1994). *Continuous Univariate Distributions, Volume 1*. Second Edition. John Wiley and Sons, New York.

See Also

[Pareto](#).

Examples

```

# Generate 30 observations from a Pareto distribution with parameters
# location=1 and shape=1 then estimate the parameters.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
dat <- rpareto(30, location = 1, shape = 1)
epareto(dat)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Pareto
#
#Estimated Parameter(s):      location = 1.009046
#                             shape    = 1.079850
#
#Estimation Method:           mle
#
#Data:                        dat
#
#Sample Size:                  30

#-----

# Compare the results of using the least-squares estimators:

epareto(dat, method="lse")$parameters
#location    shape
#1.085924 1.144180

#-----

# Clean up
#-----

rm(dat)

```

epdfPlot

*Plot Empirical Probability Density Function***Description**

Produces an empirical probability density function plot.

Usage

```

epdfPlot(x, discrete = FALSE, density.arg.list = NULL, plot.it = TRUE,
  add = FALSE, epdf.col = "black", epdf.lwd = 3 * par("cex"), epdf.lty = 1,
  curve.fill = FALSE, curve.fill.col = "cyan", ...,
  type = ifelse(discrete, "h", "l"), main = NULL, xlab = NULL, ylab = NULL,
  xlim = NULL, ylim = NULL)

```

Arguments

<code>x</code>	numeric vector of observations. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
<code>discrete</code>	logical scalar indicating whether the assumed parent distribution of <code>x</code> is discrete (<code>discrete=TRUE</code>) or continuous (<code>discrete=FALSE</code> ; the default).
<code>density.arg.list</code>	list with arguments to the density function. The default value is <code>density.arg.list=NULL</code> . This argument is ignored if <code>discrete=TRUE</code> .
<code>plot.it</code>	logical scalar indicating whether to produce a plot or add to the current plot (see <code>add</code>) on the current graphics device. The default value is <code>plot.it=TRUE</code> .
<code>add</code>	logical scalar indicating whether to add the empirical pdf to the current plot (<code>add=TRUE</code>) or generate a new plot (<code>add=FALSE</code> ; the default). This argument is ignored if <code>plot.it=FALSE</code> .
<code>epdf.col</code>	a numeric scalar or character string determining the color of the empirical pdf line or points. The default value is <code>epdf.col="black"</code> . See the entry for <code>col</code> in the help file for par for more information.
<code>epdf.lwd</code>	a numeric scalar determining the width of the empirical pdf line. The default value is <code>epdf.lwd=3*par("cex")</code> . See the entry for <code>lwd</code> in the help file for par for more information.
<code>epdf.lty</code>	a numeric scalar determining the line type of the empirical pdf line. The default value is <code>ecdf.lty=1</code> . See the entry for <code>lty</code> in the help file for par for more information.
<code>curve.fill</code>	a logical scalar indicating whether to fill in the area below the empirical pdf curve with the color specified by <code>curve.fill.col</code> . The default value is <code>curve.fill=FALSE</code> .
<code>curve.fill.col</code>	a numeric scalar or character string indicating what color to use to fill in the area below the empirical pdf curve. The default value is <code>curve.fill.col="cyan"</code> . This argument is ignored if <code>curve.fill=FALSE</code> .
<code>type, main, xlab, ylab, xlim, ylim, ...</code>	additional graphical parameters (see lines and par). In particular, the argument <code>type</code> specifies the kind of line type. By default, the function <code>epdfPlot</code> plots histogram-like vertical lines (<code>type="h"</code>) when <code>discrete=TRUE</code> , and plots a straight line between points (<code>type="l"</code>) when <code>discrete=FALSE</code> . The user may override these defaults by supplying the graphics parameter <code>type</code> (<code>type="h"</code> for histogram-like vertical lines, <code>type="l"</code> for linear interpolation, <code>type="p"</code> for points only, etc.).

Details

When a distribution is discrete and can only take on a finite number of values, the empirical pdf plot is the same as the standard relative frequency histogram; that is, each bar of the histogram represents the proportion of the sample equal to that particular number (or category). When a distribution is continuous, the function `epdfPlot` calls the R function [density](#) to compute the estimated probability density at a number of evenly spaced points between the minimum and maximum values.

Value

`epdfPlot` invisibly returns a list with the following components:

<code>x</code>	numeric vector of ordered quantiles.
<code>f.x</code>	numeric vector of the associated estimated values of the pdf.

Note

An *empirical probability density function (epdf) plot* is a graphical tool that can be used in conjunction with other graphical tools such as [histograms](#) and [boxplots](#) to assess the characteristics of a set of data.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Chambers, J.M., W.S. Cleveland, B. Kleiner, and P.A. Tukey. (1983). *Graphical Methods for Data Analysis*. Duxbury Press, Boston, MA.

See the REFERENCES section in the help file for [density](#).

See Also

[Empirical](#), [pdfPlot](#), [ecdfPlot](#), [cdfPlot](#), [cdfCompare](#), [qqPlot](#).

Examples

```
# Using Reference Area TcCB data in EPA.94b.tccb.df,
# create a histogram of the log-transformed observations,
# then superimpose the empirical pdf plot.

windows()
log.TcCB <- with(EPA.94b.tccb.df, log(TcCB[Area == "Reference"]))

hist(log.TcCB, freq = FALSE, xlim = c(-2, 1),
     col = "cyan", xlab = "log [ TcCB (ppb) ]",
     ylab = "Relative Frequency",
     main = "Reference Area TcCB with Empirical PDF")

epdfPlot(log.TcCB, add = TRUE)

#=====

# Generate 20 observations from a Poisson distribution with
# parameter lambda = 10, and plot the empirical PDF.

set.seed(875)
x <- rpois(20, lambda = 10)
windows()
epdfPlot(x, discrete = TRUE)

#=====

# Clean up
#-----
rm(log.TcCB, x)
graphics.off()
```

epois

*Estimate Parameter of a Poisson Distribution***Description**

Estimate the mean of a [Poisson distribution](#), and optionally construct a confidence interval for the mean.

Usage

```
epois(x, method = "mle/mme/mvue", ci = FALSE, ci.type = "two-sided",
      ci.method = "exact", conf.level = 0.95)
```

Arguments

<code>x</code>	numeric vector of observations.
<code>method</code>	character string specifying the method of estimation. Currently the only possible value is "mle/mme/mvue" (maximum likelihood/method of moments/minimum variance unbiased; the default). See the DETAILS section for more information.
<code>ci</code>	logical scalar indicating whether to compute a confidence interval for the location or scale parameter. The default value is FALSE.
<code>ci.type</code>	character string indicating what kind of confidence interval to compute. The possible values are "two-sided" (the default), "lower", and "upper". This argument is ignored if ci=FALSE.
<code>ci.method</code>	character string indicating what method to use to construct the confidence interval for the location or scale parameter. Possible values are "exact" (the default), "pearson.hartley.approx" (Pearson-Hartley approximation), and "normal.approx" (normal approximation). See the DETAILS section for more information. This argument is ignored if ci=FALSE.
<code>conf.level</code>	a scalar between 0 and 1 indicating the confidence level of the confidence interval. The default value is conf.level=0.95. This argument is ignored if ci=FALSE.

Details

If `x` contains any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

Let $\underline{x} = (x_1, x_2, \dots, x_n)$ be a vector of n observations from a [Poisson distribution](#) with parameter $\text{lambda}=\lambda$. It can be shown (e.g., Forbes et al., 2009) that if y is defined as:

$$y = \sum_{i=1}^n x_i \quad (1)$$

then y is an observation from a Poisson distribution with parameter $\text{lambda}=n\lambda$.

Estimation

The maximum likelihood, method of moments, and minimum variance unbiased estimator (mle/mme/mvue) of λ is given by:

$$\hat{\lambda} = \bar{x} \quad (2)$$

where

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i = \frac{y}{n} \quad (3)$$

Confidence Intervals

There are three possible ways to construct a confidence interval for λ : based on the exact distribution of the estimator of λ (`ci.type="exact"`), based on an approximation of Pearson and Hartley (`ci.type="pearson.hartley.approx"`), or based on the normal approximation (`ci.type="normal.approx"`).

Exact Confidence Interval (`ci.method="exact"`)

If `ci.type="two-sided"`, an exact $(1 - \alpha)100\%$ confidence interval for λ can be constructed as $[LCL, UCL]$, where the confidence limits are computed such that:

$$Pr[Y \geq y | \lambda = LCL] = \frac{\alpha}{2} \quad (4)$$

$$Pr[Y \leq y | \lambda = UCL] = \frac{\alpha}{2} \quad (5)$$

where y is defined in equation (1) and Y denotes a Poisson random variable with parameter $\lambda = n\lambda$.

If `ci.type="lower"`, $\alpha/2$ is replaced with α in equation (4) and UCL is set to ∞ .

If `ci.type="upper"`, $\alpha/2$ is replaced with α in equation (5) and LCL is set to 0.

Note that an exact upper confidence bound can be computed even when all observations are 0.

Pearson-Hartley Approximation (`ci.method="pearson.hartley.approx"`)

For a two-sided $(1 - \alpha)100\%$ confidence interval for λ , the Pearson and Hartley approximation (Zar, 2010, p.587; Pearson and Hartley, 1970, p.81) is given by:

$$\left[\frac{\chi_{2n\bar{x}, \alpha/2}^2}{2n}, \frac{\chi_{2n\bar{x}+2, 1-\alpha/2}^2}{2n} \right] \quad (6)$$

where $\chi_{\nu, p}^2$ denotes the p 'th quantile of the [chi-square distribution](#) with ν degrees of freedom. One-sided confidence intervals are computed in a similar fashion.

Normal Approximation (`ci.method="normal.approx"`) An approximate $(1 - \alpha)100\%$ confidence interval for λ can be constructed assuming the distribution of the estimator of λ is approximately normally distributed. A two-sided confidence interval is constructed as:

$$[\hat{\lambda} - z_{1-\alpha/2} \hat{\sigma}_{\hat{\lambda}}, \hat{\lambda} + z_{1-\alpha/2} \hat{\sigma}_{\hat{\lambda}}] \quad (7)$$

where z_p is the p 'th quantile of the standard normal distribution, and the quantity

$$\hat{\sigma}_{\hat{\lambda}} = \sqrt{\hat{\lambda}/n} \quad (8)$$

denotes the estimated asymptotic standard deviation of the estimator of λ .

One-sided confidence intervals are constructed in a similar manner.

Value

a list of class "estimate" containing the estimated parameters and other information. See [estimate.object](#) for details.

Note

The [Poisson distribution](#) is named after Poisson, who derived this distribution as the limiting distribution of the [binomial distribution](#) with parameters $\text{size}=N$ and $\text{prob}=p$, where N tends to infinity, p tends to 0, and Np stays constant.

In this context, the Poisson distribution was used by Bortkiewicz (1898) to model the number of deaths (per annum) from kicks by horses in Prussian Army Corps. In this case, p , the probability of death from this cause, was small, but the number of soldiers exposed to this risk, N , was large.

The Poisson distribution has been applied in a variety of fields, including quality control (modeling number of defects produced in a process), ecology (number of organisms per unit area), and queueing theory. Gibbons (1987b) used the Poisson distribution to model the number of detected compounds per scan of the 32 volatile organic priority pollutants (VOC), and also to model the distribution of chemical concentration (in ppb).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.
- Gibbons, R.D. (1987b). Statistical Models for the Analysis of Volatile Organic Compounds in Waste Disposal Sites. *Ground Water* **25**, 572-580.
- Gibbons, R.D., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*, Second Edition. John Wiley & Sons, Hoboken.
- Johnson, N. L., S. Kotz, and A. Kemp. (1992). *Univariate Discrete Distributions*. Second Edition. John Wiley and Sons, New York, Chapter 4.
- Pearson, E.S., and H.O. Hartley, eds. (1970). *Biometrika Tables for Statisticians, Volume 1*. Cambridge University Press, New York, p.81.
- Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ, pp. 585–586.

See Also

[Poisson](#).

Examples

```
# Generate 20 observations from a Poisson distribution with parameter
# lambda=2, then estimate the parameter and construct a 90% confidence
# interval.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
dat <- rpois(20, lambda = 2)
epois(dat, ci = TRUE, conf.level = 0.9)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Poisson
#
```

```

#Estimated Parameter(s):      lambda = 1.8
#
#Estimation Method:           mle/mme/mvue
#
#Data:                        dat
#
#Sample Size:                 20
#
#Confidence Interval for:     lambda
#
#Confidence Interval Method:   exact
#
#Confidence Interval Type:     two-sided
#
#Confidence Level:            90%
#
#Confidence Interval:         LCL = 1.336558
#                               UCL = 2.377037

#-----

# Compare the different ways of constructing confidence intervals for
# lambda using the same data as in the previous example:

epois(dat, ci = TRUE, ci.method = "pearson",
      conf.level = 0.9)$interval$limits
#      LCL      UCL
#1.336558 2.377037

epois(dat, ci = TRUE, ci.method = "normal.approx",
      conf.level = 0.9)$interval$limits
#      LCL      UCL
#1.306544 2.293456

#-----

# Clean up
#-----

rm(dat)

```

epoisCensored	<i>Estimate Mean of a Poisson Distribution Based on Type I Censored Data</i>
---------------	--

Description

Estimate the mean of a [Poisson distribution](#) given a sample of data that has been subjected to Type I censoring, and optionally construct a confidence interval for the mean.

Usage

```

epoisCensored(x, censored, method = "mle", censoring.side = "left",
  ci = FALSE, ci.method = "profile.likelihood", ci.type = "two-sided",

```

```
conf.level = 0.95, n.bootstraps = 1000, use.acc.con = FALSE,
pivot.statistic = "z", ci.sample.size = sum(!censored))
```

Arguments

<code>x</code>	numeric vector of observations. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
<code>censored</code>	numeric or logical vector indicating which values of <code>x</code> are censored. This must be the same length as <code>x</code> . If the mode of <code>censored</code> is "logical", TRUE values correspond to elements of <code>x</code> that are censored, and FALSE values correspond to elements of <code>x</code> that are not censored. If the mode of <code>censored</code> is "numeric", it must contain only 1's and 0's; 1 corresponds to TRUE and 0 corresponds to FALSE. Missing (NA) values are allowed but will be removed.
<code>method</code>	character string specifying the method of estimation. The possible values are: "mle" (maximum likelihood; the default), and "half.cen.level" (moment estimation based on setting the censored observations to half the censoring level).
<code>censoring.side</code>	character string indicating on which side the censoring occurs. The possible values are "left" (the default) and "right".
<code>ci</code>	logical scalar indicating whether to compute a confidence interval for the mean or variance. The default value is <code>ci=FALSE</code> .
<code>ci.method</code>	character string indicating what method to use to construct the confidence interval for the mean. The possible values are "profile.likelihood" (profile likelihood; the default), "normal.approx" (normal approximation), and "bootstrap" (based on bootstrapping). See the DETAILS section for more information. This argument is ignored if <code>ci=FALSE</code> .
<code>ci.type</code>	character string indicating what kind of confidence interval to compute. The possible values are "two-sided" (the default), "lower", and "upper". This argument is ignored if <code>ci=FALSE</code> .
<code>conf.level</code>	a scalar between 0 and 1 indicating the confidence level of the confidence interval. The default value is <code>conf.level=0.95</code> . This argument is ignored if <code>ci=FALSE</code> .
<code>n.bootstraps</code>	numeric scalar indicating how many bootstraps to use to construct the confidence interval for the mean when <code>ci.type="bootstrap"</code> . This argument is ignored if <code>ci=FALSE</code> and/or <code>ci.method</code> does not equal "bootstrap".
<code>use.acc.con</code>	logical scalar indicating whether to use the acceleration constant when computing the bias-corrected bootstrap confidence interval (see the DETAILS section). The default value is FALSE. This argument is ignored if <code>ci=FALSE</code> and/or <code>ci.method</code> does not equal "bootstrap".
<code>pivot.statistic</code>	character string indicating which pivot statistic to use in the construction of the confidence interval for the mean when <code>ci.method="normal.approx"</code> (see the DETAILS section). The possible values are <code>pivot.statistic="z"</code> (the default) and <code>pivot.statistic="t"</code> . When <code>pivot.statistic="t"</code> you may supply the argument <code>ci.sample.size</code> (see below). The argument <code>pivot.statistic</code> is ignored if <code>ci=FALSE</code> .
<code>ci.sample.size</code>	numeric scalar indicating what sample size to assume to construct the confidence interval for the mean if <code>pivot.statistic="t"</code> and <code>ci.method="normal.approx"</code> . The default value is the number of uncensored observations.

Details

If x or $censored$ contain any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

Let \underline{x} denote a vector of N observations from a [Poisson distribution](#) with mean λ .

Assume n ($0 < n < N$) of these observations are known and c ($c = N - n$) of these observations are all censored below (left-censored) or all censored above (right-censored) at k fixed censoring levels

$$T_1, T_2, \dots, T_k; \quad k \geq 1 \quad (1)$$

For the case when $k \geq 2$, the data are said to be Type I **multiply censored**. For the case when $k = 1$, set $T = T_1$. If the data are left-censored and all n known observations are greater than or equal to T , or if the data are right-censored and all n known observations are less than or equal to T , then the data are said to be Type I **singly censored** (Nelson, 1982, p.7), otherwise they are considered to be Type I multiply censored.

Let c_j denote the number of observations censored below or above censoring level T_j for $j = 1, 2, \dots, k$, so that

$$\sum_{j=1}^k c_j = c \quad (2)$$

Let $x_{(1)}, x_{(2)}, \dots, x_{(N)}$ denote the “ordered” observations, where now “observation” means either the actual observation (for uncensored observations) or the censoring level (for censored observations). For right-censored data, if a censored observation has the same value as an uncensored one, the uncensored observation should be placed first. For left-censored data, if a censored observation has the same value as an uncensored one, the censored observation should be placed first.

Note that in this case the quantity $x_{(i)}$ does not necessarily represent the i ’th “largest” observation from the (unknown) complete sample.

Finally, let Ω (omega) denote the set of n subscripts in the “ordered” sample that correspond to uncensored observations.

Estimation

Maximum Likelihood Estimation (method="mle")

For Type I left censored data, the likelihood function is given by:

$$L(\lambda|\underline{x}) = \binom{N}{c_1 c_2 \dots c_k n} \prod_{j=1}^k [F(T_j)]^{c_j} \prod_{i \in \Omega} f[x_{(i)}] \quad (3)$$

where f and F denote the probability density function (pdf) and cumulative distribution function (cdf) of the population (Cohen, 1963; Cohen, 1991, pp.6, 50). That is,

$$f(t) = \frac{e^{-\lambda} \lambda^t}{t!}, \quad x = 0, 1, 2, \dots \quad (4)$$

$$F(t) = \sum_{i=0}^t f(i) = \sum_{i=0}^t \frac{e^{-\lambda} \lambda^i}{i!} \quad (5)$$

(Johnson et al., 1992, p.151). For left singly censored data, equation (3) simplifies to:

$$L(\lambda|\underline{x}) = \binom{N}{c} [F(T)]^c \prod_{i=c+1}^n f[x_{(i)}] \quad (6)$$

Similarly, for Type I right censored data, the likelihood function is given by:

$$L(\lambda|\underline{x}) = \binom{N}{c_1 c_2 \dots c_k n} \prod_{j=1}^k [1 - F(T_j)]^{c_j} \prod_{i \in \Omega} f[x_{(i)}] \quad (7)$$

and for right singly censored data this simplifies to:

$$L(\lambda|\underline{x}) = \binom{N}{c} [1 - F(T)]^c \prod_{i=1}^n f[x_{(i)}] \quad (8)$$

The maximum likelihood estimators are computed by maximizing the likelihood function. For right-censored data, taking the derivative of the log-likelihood function with respect to λ and setting this to 0 produces the following equation:

$$\bar{x} = \lambda \left\{ 1 - \sum_{j=1}^K \frac{c_j}{n} \left[\frac{f(T_j)}{1 - F(T_j)} \right] \right\} \quad (9)$$

where

$$\bar{x} = \frac{1}{n} \sum_{i \in \Omega} x_i \quad (10)$$

Note that the quantity defined in equation (10) is simply the mean of the uncensored observations.

For left-censored data, taking the derivative of the log-likelihood function with respect to λ and setting this to 0 produces the following equation:

$$\bar{x} = \lambda \left\{ 1 + \sum_{j=1}^K \frac{c_j}{n} \left[\frac{f(T_j - 1)}{F(T_j - 1)} \right] \right\} \quad (11)$$

The function `epoisCensored` computes the maximum likelihood estimator of λ by solving Equation (9) (right-censored data) or Equation (11) (left-censored data); it uses the sample mean of the uncensored observations as the initial value.

Setting Censored Observations to Half the Censoring Level (`method="half.cen.level"`)

This method is applicable only to left censored data. This method involves simply replacing all the censored observations with half their detection limit, and then computing the mean and standard deviation with the usual formulas (see [epois](#)).

This method is included only to allow comparison of this method to other methods. ***Setting left-censored observations to half the censoring level is not recommended.***

Confidence Intervals

This section explains how confidence intervals for the mean λ are computed.

Likelihood Profile (`ci.method="profile.likelihood"`)

This method was proposed by Cox (1970, p.88), and Venzon and Moolgavkar (1988) introduced an efficient method of computation. This method is also discussed by Stryhn and Christensen (2003) and Royston (2007). The idea behind this method is to invert the likelihood-ratio test to obtain a confidence interval for the mean λ . Equation (3) above shows the form of the likelihood function $L(\lambda|\underline{x})$ for multiply left-censored data, and Equation (7) shows the function for multiply right-censored data.

Following Stryhn and Christensen (2003), denote the maximum likelihood estimate of the mean by λ^* . The likelihood ratio test statistic (G^2) of the hypothesis $H_0 : \lambda = \lambda_0$ (where λ_0 is a fixed value) equals the drop in $2\log(L)$ between the “full” model and the reduced model with λ fixed at μ_0 , i.e.,

$$G^2 = 2\{\log[L(\lambda^*)] - \log[L(\lambda_0)]\} \quad (11)$$

. Under the null hypothesis, the test statistic G^2 follows a [chi-squared distribution](#) with 1 degree of freedom.

A two-sided $(1 - \alpha)100\%$ confidence interval for the mean λ consists of all values of λ_0 for which the test is not significant at level *alpha*:

$$\lambda_0 : G^2 \leq \chi_{1,1-\alpha}^2 \quad (12)$$

where $\chi_{\nu,p}^2$ denotes the p 'th quantile of the [chi-squared distribution](#) with ν degrees of freedom.

Normal Approximation (ci.method="normal.approx")

This method constructs approximate $(1 - \alpha)100\%$ confidence intervals for λ based on the assumption that the estimator of λ is approximately normally distributed. That is, a two-sided $(1 - \alpha)100\%$ confidence interval for λ is constructed as:

$$[\hat{\lambda} - t_{1-\alpha/2,m-1}\hat{\sigma}_{\hat{\lambda}}, \hat{\lambda} + t_{1-\alpha/2,m-1}\hat{\sigma}_{\hat{\lambda}}] \quad (13)$$

where $\hat{\lambda}$ denotes the estimate of λ , $\hat{\sigma}_{\hat{\lambda}}$ denotes the estimated asymptotic standard deviation of the estimator of λ , m denotes the assumed sample size for the confidence interval, and $t_{p,\nu}$ denotes the p 'th quantile of [Student's t-distribution](#) with ν degrees of freedom. One-sided confidence intervals are computed in a similar fashion.

The argument ci.sample.size determines the value of m and by default is equal to the number of uncensored observations. This is simply an ad-hoc method of constructing confidence intervals and is not based on any published theoretical results.

When pivot.statistic="z", the p 'th quantile from the [standard normal distribution](#) is used in place of the p 'th quantile from Student's t-distribution.

When λ is estimated with the maximum likelihood estimator (method="mle"), the variance of $\hat{\lambda}$ is estimated based on the inverse of the Fisher Information matrix. When λ is estimated using the half-censoring-level method (method="half.cen.level"), the variance of $\hat{\lambda}$ is estimated as:

$$\hat{\sigma}_{\hat{\lambda}}^2 = \frac{\hat{\lambda}}{m} \quad (14)$$

where m denotes the assumed sample size (see above).

Bootstrap and Bias-Corrected Bootstrap Approximation (ci.method="bootstrap")

The bootstrap is a nonparametric method of estimating the distribution (and associated distribution parameters and quantiles) of a sample statistic, regardless of the distribution of the population from which the sample was drawn. The bootstrap was introduced by Efron (1979) and a general reference is Efron and Tibshirani (1993).

In the context of deriving an approximate $(1 - \alpha)100\%$ confidence interval for the population mean λ , the bootstrap can be broken down into the following steps:

1. Create a bootstrap sample by taking a random sample of size N from the observations in \underline{x} , where sampling is done with replacement. Note that because sampling is done with replacement, the same element of \underline{x} can appear more than once in the bootstrap sample. Thus, the bootstrap sample will usually not look exactly like the original sample (e.g., the number of censored observations in the bootstrap sample will often differ from the number of censored observations in the original sample).

2. Estimate λ based on the bootstrap sample created in Step 1, using the same method that was used to estimate λ using the original observations in \underline{x} . Because the bootstrap sample usually does not match the original sample, the estimate of λ based on the bootstrap sample will usually differ from the original estimate based on \underline{x} .
3. Repeat Steps 1 and 2 B times, where B is some large number. For the function `epoisCensored`, the number of bootstraps B is determined by the argument `n.bootstraps` (see the section ARGUMENTS above). The default value of `n.bootstraps` is 1000.
4. Use the B estimated values of λ to compute the empirical cumulative distribution function of this estimator of λ (see `ecdfPlot`), and then create a confidence interval for λ based on this estimated cdf.

The percentile interval (Efron and Tibshirani, 1993, p.170) is computed as:

$$[\hat{G}^{-1}(\frac{\alpha}{2}), \hat{G}^{-1}(\frac{1-\alpha}{2})] \quad (15)$$

where $\hat{G}(t)$ denotes the empirical cdf evaluated at t and thus $\hat{G}^{-1}(p)$ denotes the p 'th empirical quantile, that is, the p 'th quantile associated with the empirical cdf. The function `epoisCensored` calls the R function `quantile` to compute the empirical quantiles used in Equation (15).

The percentile method bootstrap confidence interval is only first-order accurate (Efron and Tibshirani, 1993, pp.187-188), meaning that the probability that the confidence interval will contain the true value of λ can be off by k/\sqrt{N} , where k is some constant. Efron and Tibshirani (1993, pp.184-188) proposed a bias-corrected and accelerated interval that is second-order accurate, meaning that the probability that the confidence interval will contain the true value of λ may be off by k/N instead of k/\sqrt{N} . The bias-corrected and accelerated confidence interval is computed as:

$$[\hat{G}^{-1}(\alpha_1), \hat{G}^{-1}(\alpha_2)] \quad (16)$$

where

$$\alpha_1 = \Phi[\hat{z}_0 + \frac{\hat{z}_0 + z_{\alpha/2}}{1 - \hat{a}(z_0 + z_{\alpha/2})}] \quad (17)$$

$$\alpha_2 = \Phi[\hat{z}_0 + \frac{\hat{z}_0 + z_{1-\alpha/2}}{1 - \hat{a}(z_0 + z_{1-\alpha/2})}] \quad (18)$$

$$\hat{z}_0 = \Phi^{-1}[\hat{G}(\hat{\lambda})] \quad (19)$$

$$\hat{a} = \frac{\sum_{i=1}^N (\hat{\lambda}_{(\cdot)} - \hat{\lambda}_{(i)})^3}{6[\sum_{i=1}^N (\hat{\lambda}_{(\cdot)} - \hat{\lambda}_{(i)})^2]^{3/2}} \quad (20)$$

where the quantity $\hat{\lambda}_{(i)}$ denotes the estimate of λ using all the values in \underline{x} except the i 'th one, and

$$\hat{\lambda}_{(\cdot)} = \frac{1}{N} \sum_{i=1}^N \hat{\lambda}_{(i)} \quad (21)$$

The constant \hat{z}_0 incorporates the bias correction, and the constant \hat{a} is the acceleration constant. The term “acceleration” refers to the rate of change of the standard error of the estimate of λ with respect to the true value of λ (Efron and Tibshirani, 1993, p.186). For a normal (Gaussian) distribution, the standard error of the estimate of λ does not depend on the value of λ , hence the acceleration constant is not really necessary. The argument `use.acc.con` (see the section ARGUMENTS above) determines whether the acceleration constant \hat{a} is computed as in equation (20) (`use.acc.con=TRUE`) or is set to 0 (`use.acc.con=FALSE`). The default value is `use.acc.con=FALSE`.

When `ci.method="bootstrap"`, the function `epoisCensored` computes both the percentile method and bias-corrected bootstrap confidence intervals.

Value

a list of class "estimateCensored" containing the estimated parameters and other information. See [estimateCensored.object](#) for details.

Note

A sample of data contains censored observations if some of the observations are reported only as being below or above some censoring level. In environmental data analysis, Type I left-censored data sets are common, with values being reported as "less than the detection limit" (e.g., Helsel, 2012). Data sets with only one censoring level are called *singly censored*; data sets with multiple censoring levels are called *multiply* or *progressively censored*.

Statistical methods for dealing with censored data sets have a long history in the field of survival analysis and life testing. More recently, researchers in the environmental field have proposed alternative methods of computing estimates and confidence intervals in addition to the classical ones such as maximum likelihood estimation. Helsel (2012, Chapter 6) gives an excellent review of past studies of the properties of various estimators for parameters of a normal or lognormal distribution based on censored environmental data.

In practice, it is better to use a confidence interval for the mean or a joint confidence region for the mean and standard deviation (or coefficient of variation), rather than rely on a single point-estimate of the mean. Few studies have been done to evaluate the performance of methods for constructing confidence intervals for the mean or joint confidence regions for the mean and coefficient of variation of a Poisson distribution when data are subjected to single or multiple censoring.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Cohen, A.C. (1991). *Truncated and Censored Samples*. Marcel Dekker, New York, New York, 312pp.
- Cox, D.R. (1970). *Analysis of Binary Data*. Chapman & Hall, London. 142pp.
- Efron, B. (1979). Bootstrap Methods: Another Look at the Jackknife. *The Annals of Statistics* **7**, 1–26.
- Efron, B., and R.J. Tibshirani. (1993). *An Introduction to the Bootstrap*. Chapman and Hall, New York, 436pp.
- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions, Fourth Edition*. John Wiley and Sons, Hoboken, NJ.
- Helsel, D.R. (2012). *Statistics for Censored Environmental Data Using Minitab and R, Second Edition*. John Wiley & Sons, Hoboken, New Jersey.
- Johnson, N. L., S. Kotz, and A. Kemp. (1992). *Univariate Discrete Distributions, Second Edition*. John Wiley and Sons, New York, Chapter 4.
- Millard, S.P., P. Dixon, and N.K. Neerchal. (2014; in preparation). *Environmental Statistics with R*. CRC Press, Boca Raton, Florida.
- Nelson, W. (1982). *Applied Life Data Analysis*. John Wiley and Sons, New York, 634pp.
- Royston, P. (2007). Profile Likelihood for Estimation and Confidence Intervals. *The Stata Journal* **7**(3), pp. 376–387.

Stryhn, H., and J. Christensen. (2003). *Confidence Intervals by the Profile Likelihood Method, with Applications in Veterinary Epidemiology*. Contributed paper at ISVEE X (November 2003, Chile). <http://people.upei.ca/hstryhn/stryhn208.pdf>.

Venzon, D.J., and S.H. Moolgavkar. (1988). A Method for Computing Profile-Likelihood-Based Confidence Intervals. *Journal of the Royal Statistical Society, Series C (Applied Statistics)* **37**(1), pp. 87–94.

See Also

Poisson, `epois`, `estimateCensored.object`.

Examples

```
# Generate 20 observations from a Poisson distribution with
# parameter lambda=10, and censor the values less than 10.
# Then generate 20 more observations from the same distribution
# and censor the values less than 20. Then estimate the mean
# using the maximum likelihood method.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(300)
dat.1 <- rpois(20, lambda=10)
censored.1 <- dat.1 < 10
dat.1[censored.1] <- 10

dat.2 <- rpois(20, lambda=10)
censored.2 <- dat.2 < 20
dat.2[censored.2] <- 20

dat <- c(dat.1, dat.2)
censored <- c(censored.1, censored.2)

epoisCensored(dat, censored, ci = TRUE)

#Results of Distribution Parameter Estimation
#Based on Type I Censored Data
#-----
#
#Assumed Distribution:      Poisson
#
#Censoring Side:           left
#
#Censoring Level(s):       10 20
#
#Estimated Parameter(s):   lambda = 11.05402
#
#Estimation Method:        MLE
#
#Data:                     dat
#
#Censoring Variable:       censored
#
#Sample Size:              40
#
#Percent Censored:         65%
#
```

```
#Confidence Interval for:      lambda
#
#Confidence Interval Method:   Profile Likelihood
#
#Confidence Interval Type:    two-sided
#
#Confidence Level:           95%
#
#Confidence Interval:        LCL =  9.842894
#                             UCL = 12.846484

#-----

# Clean up
#-----
rm(dat.1, censored.1, dat.2, censored.2, dat, censored)
```

eqbeta	<i>Estimate Quantiles of a Beta Distribution</i>
--------	--

Description

Estimate quantiles of a [beta distribution](#).

Usage

```
eqbeta(x, p = 0.5, method = "mle", digits = 0)
```

Arguments

x	a numeric vector of observations, or an object resulting from a call to an estimating function that assumes a beta distribution (e.g., ebeta). If x is a numeric vector, missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
p	numeric vector of probabilities for which quantiles will be estimated. All values of p must be between 0 and 1. The default value is p=0.5.
method	character string specifying the method to use to estimate the shape and scale parameters of the distribution. The possible values are "mle" (maximum likelihood; the default), "mme" (method of moments), and "mmue" (method of moments based on the unbiased estimator of variance). See the DETAILS section of the help file for ebeta for more information.
digits	an integer indicating the number of decimal places to round to when printing out the value of 100*p. The default value is digits=0.

Details

The function eqbeta returns estimated quantiles as well as estimates of the shape1 and shape2 parameters.

Quantiles are estimated by 1) estimating the shape1 and shape2 parameters by calling [ebeta](#), and then 2) calling the function [qbeta](#) and using the estimated values for shape1 and shape2.

Value

If `x` is a numeric vector, `eqbeta` returns a list of class "estimate" containing the estimated quantile(s) and other information. See [estimate.object](#) for details.

If `x` is the result of calling an estimation function, `eqbeta` returns a list whose class is the same as `x`. The list contains the same components as `x`, as well as components called `quantiles` and `quantile.method`.

Note

The beta distribution takes real values between 0 and 1. Special cases of the beta are the [Uniform](#)[0,1] when `shape1=1` and `shape2=1`, and the arcsin distribution when `shape1=0.5` and `shape2=0.5`. The arcsin distribution appears in the theory of random walks. The beta distribution is used in Bayesian analyses as a conjugate to the binomial distribution.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.

Johnson, N. L., S. Kotz, and N. Balakrishnan. (1995). *Continuous Univariate Distributions, Volume 2*. Second Edition. John Wiley and Sons, New York.

See Also

[ebeta](#), [Beta](#), [estimate.object](#).

Examples

```
# Generate 20 observations from a beta distribution with parameters
# shape1=2 and shape2=4, then estimate the parameters via
# maximum likelihood and estimate the 90'th percentile.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
dat <- rbeta(20, shape1 = 2, shape2 = 4)
eqbeta(dat, p = 0.9)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Beta
#
#Estimated Parameter(s):      shape1 =  5.392221
#                              shape2 = 11.823233
#
#Estimation Method:           mle
#
#Estimated Quantile(s):       90'th %ile = 0.4592796
#
#Quantile Estimation Method:   Quantile(s) Based on
#                              mle Estimators
```

```
#
#Data:                                dat
#
#Sample Size:                          20

#-----
# Clean up

rm(dat)
```

eqbinom

Estimate Quantiles of a Binomial Distribution

Description

Estimate quantiles of a [binomial distribution](#).

Usage

```
eqbinom(x, size = NULL, p = 0.5, method = "mle/mme/mvue", digits = 0)
```

Arguments

x	numeric or logical vector of observations, or an object resulting from a call to an estimating function that assumes a binomial distribution (e.g., ebinom). If x is a vector of observations, then when size is not supplied, x must be a numeric vector of 0s (“failures”) and 1s (“successes”), or else a logical vector of FALSE values (“failures”) and TRUE values (“successes”). When size is supplied, x must be a non-negative integer containing the number of “successes” out of the number of trials indicated by size. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
size	positive integer indicating the of number of trials; size must be at least as large as the value of x.
p	numeric vector of probabilities for which quantiles will be estimated. All values of p must be between 0 and 1. The default value is p=0.5.
method	character string specifying the method of estimation. The only possible value is "mle/mme/mvue" (maximum likelihood, method of moments, and minimum variance unbiased). See the DETAILS section of the help file for ebinom for more information.
digits	an integer indicating the number of decimal places to round to when printing out the value of 100*p. The default value is digits=0.

Details

The function eqbinom returns estimated quantiles as well as estimates of the prob parameter.

Quantiles are estimated by 1) estimating the prob parameter by calling [ebinom](#), and then 2) calling the function [qbinom](#) and using the estimated value for prob.

Value

If x is a numeric vector, eqbinom returns a list of class "estimate" containing the estimated quantile(s) and other information. See [estimate.object](#) for details.

If x is the result of calling an estimation function, eqbinom returns a list whose class is the same as x . The list contains the same components as x , as well as components called quantiles and quantile.method.

Note

The binomial distribution is used to model processes with binary (Yes-No, Success-Failure, Heads-Tails, etc.) outcomes. It is assumed that the outcome of any one trial is independent of any other trial, and that the probability of "success", p , is the same on each trial. A binomial discrete random variable X is the number of "successes" in n independent trials. A special case of the binomial distribution occurs when $n = 1$, in which case X is also called a Bernoulli random variable.

In the context of environmental statistics, the binomial distribution is sometimes used to model the proportion of times a chemical concentration exceeds a set standard in a given period of time (e.g., Gilbert, 1987, p.143). The binomial distribution is also used to compute an upper bound on the overall Type I error rate for deciding whether a facility or location is in compliance with some set standard. Assume the null hypothesis is that the facility is in compliance. If a test of hypothesis is conducted periodically over time to test compliance and/or several tests are performed during each time period, and the facility or location is always in compliance, and each single test has a Type I error rate of α , and the result of each test is independent of the result of any other test (usually not a reasonable assumption), then the number of times the facility is declared out of compliance when in fact it is in compliance is a binomial random variable with probability of "success" $p = \alpha$ being the probability of being declared out of compliance (see USEPA, 2009).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Agresti, A., and B.A. Coull. (1998). Approximate is Better than "Exact" for Interval Estimation of Binomial Proportions. *The American Statistician*, **52**(2), 119–126.
- Agresti, A., and B. Caffo. (2000). Simple and Effective Confidence Intervals for Proportions and Differences of Proportions Result from Adding Two Successes and Two Failures. *The American Statistician*, **54**(4), 280–288.
- Berthouex, P.M., and L.C. Brown. (1994). *Statistics for Environmental Engineers*. Lewis Publishers, Boca Raton, FL, Chapters 2 and 15.
- Cochran, W.G. (1977). *Sampling Techniques*. John Wiley and Sons, New York, Chapter 3.
- Fisher, R.A., and F. Yates. (1963). *Statistical Tables for Biological, Agricultural, and Medical Research*. 6th edition. Hafner, New York, 146pp.
- Fleiss, J. L. (1981). *Statistical Methods for Rates and Proportions*. Second Edition. John Wiley and Sons, New York, Chapters 1-2.
- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.
- Gilbert, R.O. (1987). *Statistical Methods for Environmental Pollution Monitoring*. Van Nostrand Reinhold, New York, NY, Chapter 11.
- Johnson, N. L., S. Kotz, and A.W. Kemp. (1992). *Univariate Discrete Distributions*. Second Edition. John Wiley and Sons, New York, Chapter 3.

Millard, S.P., and Neerchal, N.K. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, Florida.

Newcombe, R.G. (1998a). Two-Sided Confidence Intervals for the Single Proportion: Comparison of Seven Methods. *Statistics in Medicine*, **17**, 857–872.

Ott, W.R. (1995). *Environmental Statistics and Data Analysis*. Lewis Publishers, Boca Raton, FL, Chapter 4.

USEPA. (1989b). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities, Interim Final Guidance*. EPA/530-SW-89-026. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C.

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.6-38.

Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ, Chapter 24.

See Also

[ebinom](#), [Binomial](#), [estimate.object](#).

Examples

```
# Generate 20 observations from a binomial distribution with
# parameters size=1 and prob=0.2, then estimate the 'prob'
# parameter and the 90'th percentile.
# (Note: the call to set.seed simply allows you to reproduce this example.

set.seed(251)
dat <- rbinom(20, size = 1, prob = 0.2)
eqbinom(dat, p = 0.9)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Binomial
#
#Estimated Parameter(s):      size = 20.0
#                             prob =  0.1
#
#Estimation Method:           mle/mme/mvue for 'prob'
#
#Estimated Quantile(s):       90'th %ile = 4
#
#Quantile Estimation Method:   Quantile(s) Based on
#                             mle/mme/mvue for 'prob' Estimators
#
#Data:                         dat
#
#Sample Size:                  20
#
#
#-----
```

```
# Clean up
rm(dat)
```

eqevd

Estimate Quantiles of an Extreme Value (Gumbel) Distribution

Description

Estimate quantiles of an [extreme value distribution](#).

Usage

```
eqevd(x, p = 0.5, method = "mle", pwme.method = "unbiased",
      plot.pos.cons = c(a = 0.35, b = 0), digits = 0)
```

Arguments

- | | |
|---------------|---|
| x | a numeric vector of observations, or an object resulting from a call to an estimating function that assumes an extreme value distribution (e.g., eevd). If x is a numeric vector, missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed. |
| p | numeric vector of probabilities for which quantiles will be estimated. All values of p must be between 0 and 1. The default value is p=0.5. |
| method | character string specifying the method to use to estimate the location and scale parameters. Possible values are "mle" (maximum likelihood; the default), "mme" (methods of moments), "mmue" (method of moments based on the unbiased estimator of variance), and "pwme" (probability-weighted moments). See the DETAILS section of the help file for eevd for more information on these estimation methods. |
| pwme.method | character string specifying what method to use to compute the probability-weighted moments when method="pwme". The possible values are "unbiased" (method based on the U-statistic; the default), or "plotting.position" (method based on the plotting position formula). See the DETAILS section of the help file for eevd for more information. This argument is ignored if method is not equal to "pwme". |
| plot.pos.cons | numeric vector of length 2 specifying the constants used in the formula for the plotting positions when method="pwme" and pwme.method="plotting.position". The default value is plot.pos.cons=c(a=0.35, b=0). If this vector has a names attribute with the value c("a", "b") or c("b", "a"), then the elements will be matched by name in the formula for computing the plotting positions. Otherwise, the first element is mapped to the name "a" and the second element to the name "b". See the DETAILS section of the help file for eevd for more information. This argument is ignored if method is not equal to "pwme" or if pwme.method="unbiased". |
| digits | an integer indicating the number of decimal places to round to when printing out the value of 100*p. The default value is digits=0. |

Details

The function `eqevd` returns estimated quantiles as well as estimates of the location and scale parameters.

Quantiles are estimated by 1) estimating the location and scale parameters by calling `eevd`, and then 2) calling the function `qevd` and using the estimated values for location and scale.

Value

If `x` is a numeric vector, `eqevd` returns a list of class "estimate" containing the estimated quantile(s) and other information. See `estimate.object` for details.

If `x` is the result of calling an estimation function, `eqevd` returns a list whose class is the same as `x`. The list contains the same components as `x`, as well as components called `quantiles` and `quantile.method`.

Note

There are three families of extreme value distributions. The one described here is the [Type I, also called the Gumbel extreme value distribution or simply Gumbel distribution](#). The name "extreme value" comes from the fact that this distribution is the limiting distribution (as n approaches infinity) of the greatest value among n independent random variables each having the same continuous distribution.

The Gumbel extreme value distribution is related to the [exponential distribution](#) as follows. Let Y be an [exponential random variable](#) with parameter $\text{rate}=\lambda$. Then $X = \eta - \log(Y)$ has an extreme value distribution with parameters $\text{location}=\eta$ and $\text{scale}=1/\lambda$.

The distribution described above and assumed by `eevd` is the *largest* extreme value distribution. The smallest extreme value distribution is the limiting distribution (as n approaches infinity) of the smallest value among n independent random variables each having the same continuous distribution. If X has a largest extreme value distribution with parameters $\text{location}=\eta$ and $\text{scale}=\theta$, then $Y = -X$ has a smallest extreme value distribution with parameters $\text{location}=-\eta$ and $\text{scale}=\theta$. The smallest extreme value distribution is related to the [Weibull distribution](#) as follows. Let Y be a [Weibull random variable](#) with parameters $\text{shape}=\beta$ and $\text{scale}=\alpha$. Then $X = \log(Y)$ has a smallest extreme value distribution with parameters $\text{location}=\log(\alpha)$ and $\text{scale}=1/\beta$.

The extreme value distribution has been used extensively to model the distribution of streamflow, flooding, rainfall, temperature, wind speed, and other meteorological variables, as well as material strength and life data.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Castillo, E. (1988). *Extreme Value Theory in Engineering*. Academic Press, New York, pp.184–198.
- Downton, F. (1966). Linear Estimates of Parameters in the Extreme Value Distribution. *Technometrics* **8**(1), 3–17.
- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.
- Greenwood, J.A., J.M. Landwehr, N.C. Matalas, and J.R. Wallis. (1979). Probability Weighted Moments: Definition and Relation to Parameters of Several Distributions Expressible in Inverse Form. *Water Resources Research* **15**(5), 1049–1054.

Hosking, J.R.M., J.R. Wallis, and E.F. Wood. (1985). Estimation of the Generalized Extreme-Value Distribution by the Method of Probability-Weighted Moments. *Technometrics* **27**(3), 251–261.

Johnson, N. L., S. Kotz, and N. Balakrishnan. (1995). *Continuous Univariate Distributions, Volume 2*. Second Edition. John Wiley and Sons, New York.

Landwehr, J.M., N.C. Matalas, and J.R. Wallis. (1979). Probability Weighted Moments Compared With Some Traditional Techniques in Estimating Gumbel Parameters and Quantiles. *Water Resources Research* **15**(5), 1055–1064.

Tiago de Oliveira, J. (1963). Decision Results for the Parameters of the Extreme Value (Gumbel) Distribution Based on the Mean and Standard Deviation. *Trabajos de Estadística* **14**, 61–81.

See Also

[eevd](#), [Extreme Value Distribution](#), [estimate.object](#).

Examples

```
# Generate 20 observations from an extreme value distribution with
# parameters location=2 and scale=1, then estimate the parameters
# and estimate the 90'th percentile.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
dat <- revd(20, location = 2)
eqevd(dat, p = 0.9)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:      Extreme Value
#
#Estimated Parameter(s):   location = 1.9684093
#                           scale    = 0.7481955
#
#Estimation Method:        mle
#
#Estimated Quantile(s):    90'th %ile = 3.652124
#
#Quantile Estimation Method: Quantile(s) Based on
#                           mle Estimators
#
#Data:                      dat
#
#Sample Size:               20

#-----

# Clean up
#-----
rm(dat)
```

eqexp

*Estimate Quantiles of an Exponential Distribution***Description**

Estimate quantiles of an [exponential distribution](#).

Usage

```
eqexp(x, p = 0.5, method = "mle/mme", digits = 0)
```

Arguments

- | | |
|--------|--|
| x | a numeric vector of observations, or an object resulting from a call to an estimating function that assumes an exponential distribution (e.g., eexp). If x is a numeric vector, missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed. |
| p | numeric vector of probabilities for which quantiles will be estimated. All values of p must be between 0 and 1. The default value is p=0.5. |
| method | character string specifying the method to use to estimate the rate parameter. Currently the only possible value is "mle/mme" (maximum likelihood/method of moments; the default). See the DETAILS section of the help file for eexp for more information. |
| digits | an integer indicating the number of decimal places to round to when printing out the value of 100*p. The default value is digits=0. |

Details

The function eqexp returns estimated quantiles as well as the estimate of the rate parameter.

Quantiles are estimated by 1) estimating the rate parameter by calling [eexp](#), and then 2) calling the function [qexp](#) and using the estimated value for rate.

Value

If x is a numeric vector, eqexp returns a list of class "estimate" containing the estimated quantile(s) and other information. See [estimate.object](#) for details.

If x is the result of calling an estimation function, eqexp returns a list whose class is the same as x. The list contains the same components as x, as well as components called quantiles and quantile.method.

Note

The [exponential distribution](#) is a special case of the [gamma distribution](#), and takes on positive real values. A major use of the exponential distribution is in life testing where it is used to model the lifetime of a product, part, person, etc.

The exponential distribution is the only continuous distribution with a "lack of memory" property. That is, if the lifetime of a part follows the exponential distribution, then the distribution of the time until failure is the same as the distribution of the time until failure given that the part has survived to time t .

The exponential distribution is related to the double exponential (also called Laplace) distribution, and to the [extreme value distribution](#).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.

Johnson, N. L., S. Kotz, and N. Balakrishnan. (1994). *Continuous Univariate Distributions, Volume 1*. Second Edition. John Wiley and Sons, New York.

See Also

[eexp](#), [Exponential](#), [estimate.object](#).

Examples

```
# Generate 20 observations from an exponential distribution with parameter
# rate=2, then estimate the parameter and estimate the 90th percentile.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
dat <- rexp(20, rate = 2)
eqexp(dat, p = 0.9)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:           Exponential
#
#Estimated Parameter(s):       rate = 2.260587
#
#Estimation Method:            mle/mme
#
#Estimated Quantile(s):        90'th %ile = 1.018578
#
#Quantile Estimation Method:    Quantile(s) Based on
#                               mle/mme Estimators
#
#Data:                          dat
#
#Sample Size:                   20
#
#-----

# Clean up
#-----
rm(dat)
```

eqgamma

*Estimate Quantiles of a Gamma Distribution***Description**

Estimate quantiles of a [gamma distribution](#), and optionally construct a confidence interval for a quantile.

Usage

```
eqgamma(x, p = 0.5, method = "mle", ci = FALSE,
        ci.type = "two-sided", conf.level = 0.95,
        normal.approx.transform = "kulkarni.powar", digits = 0)
```

```
eqgammaAlt(x, p = 0.5, method = "mle", ci = FALSE,
           ci.type = "two-sided", conf.level = 0.95,
           normal.approx.transform = "kulkarni.powar", digits = 0)
```

Arguments

- | | |
|-------------------------|---|
| x | a numeric vector of observations, or an object resulting from a call to an estimating function that assumes a gamma distribution (e.g., egamma or egammaAlt). If <code>ci=TRUE</code> then x must be a numeric vector of observations. If x is a numeric vector, missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed. |
| p | numeric vector of probabilities for which quantiles will be estimated. All values of p must be between 0 and 1. When <code>ci=TRUE</code> , p must be a scalar. The default value is <code>p=0.5</code> . |
| method | character string specifying the method to use to estimate the shape and scale parameters of the distribution. The possible values are "mle" (maximum likelihood; the default), "bcmle" (bias-corrected mle), "mme" (method of moments), and "mmue" (method of moments based on the unbiased estimator of variance). See the DETAILS section of the help file for egamma for more information. |
| ci | logical scalar indicating whether to compute a confidence interval for the quantile. The default value is <code>ci=FALSE</code> . |
| ci.type | character string indicating what kind of confidence interval for the quantile to compute. The possible values are "two-sided" (the default), "lower", and "upper". This argument is ignored if <code>ci=FALSE</code> . |
| conf.level | a scalar between 0 and 1 indicating the confidence level of the confidence interval. The default value is <code>conf.level=0.95</code> . This argument is ignored if <code>ci=FALSE</code> . |
| normal.approx.transform | character string indicating which power transformation to use. Possible values are "kulkarni.powar" (the default), "cube.root", and "fourth.root". See the DETAILS section for more information. This argument is ignored if <code>ci=FALSE</code> . |
| digits | an integer indicating the number of decimal places to round to when printing out the value of <code>100*p</code> . The default value is <code>digits=0</code> . |

Details

The function `eqgamma` returns estimated quantiles as well as estimates of the shape and scale parameters. The function `eqgammaAlt` returns estimated quantiles as well as estimates of the mean and coefficient of variation.

Quantiles are estimated by 1) estimating the shape and scale parameters by calling `egamma`, and then 2) calling the function `qgamma` and using the estimated values for shape and scale.

The confidence interval for a quantile is computed by:

1. using a power transformation on the original data to induce approximate normality,
2. using `eqnorm` to compute the confidence interval, and then
3. back-transforming the interval to create a confidence interval on the original scale.

This is similar to what is done to create tolerance intervals for a gamma distribution (Krishnamoorthy et al., 2008), and there is a one-to-one relationship between confidence intervals for a quantile and tolerance intervals (see the DETAILS section of the help file for `eqnorm`). The value `normal.approx.transform="cube.root"` uses the cube root transformation suggested by Wilson and Hilferty (1931) and used by Krishnamoorthy et al. (2008) and Singh et al. (2010b), and the value `normal.approx.transform="fourth.root"` uses the fourth root transformation suggested by Hawkins and Wixley (1986) and used by Singh et al. (2010b). The default value `normal.approx.transform="kulkarni.powar"` uses the “Optimum Power Normal Approximation Method” of Kulkarni and Powar (2010). The “optimum” power r is determined by:

$$\begin{aligned} r &= -0.0705 - 0.178 \text{ shape} + 0.475 \sqrt{\text{shape}} && \text{if } \text{shape} \leq 1.5 \\ r &= 0.246 && \text{if } \text{shape} > 1.5 \end{aligned}$$

where *shape* denotes the estimate of the shape parameter. Although Kulkarni and Powar (2010) use the maximum likelihood estimate of shape to determine the power r , for the functions `eqgamma` and `eqgammaAlt` the power r is based on whatever estimate of shape is used (e.g., `method="mle"`, `method="bcml"`, etc.).

Value

If x is a numeric vector, `eqgamma` and `eqgammaAlt` return a list of class “estimate” containing the estimated quantile(s) and other information. See `estimate.object` for details.

If x is the result of calling an estimation function, `eqgamma` and `eqgammaAlt` return a list whose class is the same as x . The list contains the same components as x , as well as components called `quantiles` and `quantile.method`. In addition, if `ci=TRUE`, the returned list contains a component called `interval` containing the confidence interval information. If x already has a component called `interval`, this component is replaced with the confidence interval information.

Note

The gamma distribution takes values on the positive real line. Special cases of the gamma are the [exponential](#) distribution and the [chi-square distributions](#). Applications of the gamma include life testing, statistical ecology, queuing theory, inventory control, and precipitation processes. A gamma distribution starts to resemble a normal distribution as the shape parameter tends to infinity.

Some EPA guidance documents (e.g., Singh et al., 2002; Singh et al., 2010a,b) strongly recommend against using a lognormal model for environmental data and recommend trying a gamma distribution instead.

Percentiles are sometimes used in environmental standards and regulations. For example, Berthouex and Brown (2002, p.71) note that England has water quality limits based on the 90th and 95th percentiles of monitoring data not exceeding specified levels. They also note that the U.S. EPA has specifications for air quality monitoring, aquatic standards on toxic chemicals, and maximum daily limits for industrial effluents that are all based on percentiles. Given the importance of these quantities, it is essential to characterize the amount of uncertainty associated with the estimates of these quantities. This is done with confidence intervals.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Berthouex, P.M., and L.C. Brown. (2002). *Statistics for Environmental Engineers*. Lewis Publishers, Boca Raton.
- Conover, W.J. (1980). *Practical Nonparametric Statistics*. Second Edition. John Wiley and Sons, New York.
- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.
- Gibbons, R.D., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*, Second Edition. John Wiley & Sons, Hoboken.
- Hawkins, D. M., and R.A.J. Wixley. (1986). A Note on the Transformation of Chi-Squared Variables to Normality. *The American Statistician*, **40**, 296–298.
- Johnson, N.L., S. Kotz, and N. Balakrishnan. (1994). *Continuous Univariate Distributions, Volume 1*. Second Edition. John Wiley and Sons, New York, Chapter 17.
- Krishnamoorthy K., T. Mathew, and S. Mukherjee. (2008). Normal-Based Methods for a Gamma Distribution: Prediction and Tolerance Intervals and Stress-Strength Reliability. *Technometrics*, **50**(1), 69–78.
- Krishnamoorthy K., and T. Mathew. (2009). *Statistical Tolerance Regions: Theory, Applications, and Computation*. John Wiley and Sons, Hoboken.
- Kulkarni, H.V., and S.K. Powar. (2010). A New Method for Interval Estimation of the Mean of the Gamma Distribution. *Lifetime Data Analysis*, **16**, 431–447.
- Singh, A., A.K. Singh, and R.J. Iaci. (2002). *Estimation of the Exposure Point Concentration Term Using a Gamma Distribution*. EPA/600/R-02/084. October 2002. Technology Support Center for Monitoring and Site Characterization, Office of Research and Development, Office of Solid Waste and Emergency Response, U.S. Environmental Protection Agency, Washington, D.C.
- Singh, A., R. Maichle, and N. Armbya. (2010a). *ProUCL Version 4.1.00 User Guide (Draft)*. EPA/600/R-07/041, May 2010. Office of Research and Development, U.S. Environmental Protection Agency, Washington, D.C.
- Singh, A., N. Armbya, and A. Singh. (2010b). *ProUCL Version 4.1.00 Technical Guide (Draft)*. EPA/600/R-07/041, May 2010. Office of Research and Development, U.S. Environmental Protection Agency, Washington, D.C.
- Wilson, E.B., and M.M. Hilferty. (1931). The Distribution of Chi-Squares. *Proceedings of the National Academy of Sciences*, **17**, 684–688.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

See Also

[eqgamma](#), [GammaDist](#), [estimate.object](#), [eqnorm](#), [tolIntGamma](#).

Examples

```
# Generate 20 observations from a gamma distribution with parameters
# shape=3 and scale=2, then estimate the 90th percentile and create
# a one-sided upper 95% confidence interval for that percentile.
# (Note: the call to set.seed simply allows you to reproduce this
# example.)

set.seed(250)
dat <- rgamma(20, shape = 3, scale = 2)
eqgamma(dat, p = 0.9, ci = TRUE, ci.type = "upper")

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Gamma
#
#Estimated Parameter(s):      shape = 2.203862
#                              scale = 2.174928
#
#Estimation Method:           mle
#
#Estimated Quantile(s):       90'th %ile = 9.113446
#
#Quantile Estimation Method:   Quantile(s) Based on
#                              mle Estimators
#
#Data:                         dat
#
#Sample Size:                  20
#
#Confidence Interval for:      90'th %ile
#
#Confidence Interval Method:   Exact using
#                              Kulkarni & Powar (2010)
#                              transformation to Normality
#                              based on mle of 'shape'
#
#Confidence Interval Type:     upper
#
#Confidence Level:             95%
#
#Confidence Interval:          LCL = 0.00000
#                              UCL = 13.79733
#
#-----
# Compare these results with the true 90'th percentile:

qgamma(p = 0.9, shape = 3, scale = 2)
```



```

#[1] 10.64464

#-----

# Using the same data as in the previous example, use eqgammaAlt
# to estimate the mean and cv based on the bias-corrected
# estimate of shape, and use the cube-root transformation to
# normality.

eqgammaAlt(dat, p = 0.9, method = "bcmle", ci = TRUE,
  ci.type = "upper", normal.approx.transform = "cube.root")

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Gamma
#
#Estimated Parameter(s):      mean = 4.7932408
#                               cv   = 0.7242165
#
#Estimation Method:           bcmle of 'shape'
#
#Estimated Quantile(s):       90'th %ile = 9.428
#
#Quantile Estimation Method:   Quantile(s) Based on
#                               bcmle of 'shape'
#
#Data:                         dat
#
#Sample Size:                  20
#
#Confidence Interval for:      90'th %ile
#
#Confidence Interval Method:    Exact using
#                               Wilson & Hilferty (1931) cube-root
#                               transformation to Normality
#
#Confidence Interval Type:     upper
#
#Confidence Level:             95%
#
#Confidence Interval:          LCL = 0.00000
#                               UCL = 12.89643
#
#-----

# Clean up
rm(dat)

#-----

# Example 17-3 of USEPA (2009, p. 17-17) shows how to construct a
# beta-content upper tolerance limit with 95% coverage and
# 95% confidence using chrysene data and assuming a lognormal
# distribution. Here we will use the same chrysene data but assume a
# gamma distribution.

```

```

# A beta-content upper tolerance limit with 95% coverage and
# 95% confidence is equivalent to the 95% upper confidence limit for
# the 95th percentile.

attach(EPA.09.Ex.17.3.chrysene.df)
Chrysene <- Chrysene.ppb[Well.type == "Background"]
eqgamma(Chrysene, p = 0.95, ci = TRUE, ci.type = "upper")

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Gamma
#
#Estimated Parameter(s):      shape = 2.806929
#                              scale = 5.286026
#
#Estimation Method:           mle
#
#Estimated Quantile(s):       95'th %ile = 31.74348
#
#Quantile Estimation Method:   Quantile(s) Based on
#                              mle Estimators
#
#Data:                        Chrysene
#
#Sample Size:                 8
#
#Confidence Interval for:     95'th %ile
#
#Confidence Interval Method:   Exact using
#                              Kulkarni & Powar (2010)
#                              transformation to Normality
#                              based on mle of 'shape'
#
#Confidence Interval Type:    upper
#
#Confidence Level:           95%
#
#Confidence Interval:         LCL = 0.00000
#                              UCL = 69.32425
#
#-----
# Clean up

rm(Chrysene)
detach("EPA.09.Ex.17.3.chrysene.df")

```

eqgeom

Estimate Quantiles of a Geometric Distribution

Description

Estimate quantiles of a [geometric distribution](#).

Usage

```
eqgeom(x, p = 0.5, method = "mle/mme", digits = 0)
```

Arguments

x	a numeric vector of observations, or an object resulting from a call to an estimating function that assumes a geometric distribution (e.g., eqgeom). If x is a numeric vector, missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
p	numeric vector of probabilities for which quantiles will be estimated. All values of p must be between 0 and 1. The default value is p=0.5.
method	character string specifying the method to use to estimate the probability parameter. Possible values are "mle/mme" (maximum likelihood and method of moments; the default) and "mvue" (minimum variance unbiased). You cannot use method="mvue" if length(x)=1. See the DETAILS section of the help file for eqgeom for more information on these estimation methods.
digits	an integer indicating the number of decimal places to round to when printing out the value of 100*p. The default value is digits=0.

Details

The function `eqgeom` returns estimated quantiles as well as the estimate of the rate parameter.

Quantiles are estimated by 1) estimating the probability parameter by calling [eqgeom](#), and then 2) calling the function [qgeom](#) and using the estimated value for the probability parameter.

Value

If x is a numeric vector, `eqgeom` returns a list of class "estimate" containing the estimated quantile(s) and other information. See [estimate.object](#) for details.

If x is the result of calling an estimation function, `eqgeom` returns a list whose class is the same as x. The list contains the same components as x, as well as components called `quantiles` and `quantile.method`.

Note

The [geometric distribution](#) with parameter $\text{prob}=p$ is a special case of the [negative binomial distribution](#) with parameters $\text{size}=1$ and $\text{prob}=p$.

The negative binomial distribution has its roots in a gambling game where participants would bet on the number of tosses of a coin necessary to achieve a fixed number of heads. The negative binomial distribution has been applied in a wide variety of fields, including accident statistics, birth-and-death processes, and modeling spatial distributions of biological organisms.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.
- Johnson, N. L., S. Kotz, and A. Kemp. (1992). *Univariate Discrete Distributions*. Second Edition. John Wiley and Sons, New York, Chapter 5.

See Also

[egeom](#), [Geometric](#), [enbinom](#), [NegBinomial](#), [estimate.object](#).

Examples

```
# Generate an observation from a geometric distribution with parameter
# prob=0.2, then estimate the parameter prob and the 90'th percentile.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
dat <- rgeom(1, prob = 0.2)
dat
#[1] 4

eqgeom(dat, p = 0.9)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Geometric
#
#Estimated Parameter(s):      prob = 0.2
#
#Estimation Method:           mle/mme
#
#Estimated Quantile(s):       90'th %ile = 10
#
#Quantile Estimation Method:   Quantile(s) Based on
#                               mle/mme Estimators
#
#Data:                         dat
#
#Sample Size:                  1
#-----

# Clean up
#-----
rm(dat)
```

eqgevd

Estimate Quantiles of a Generalized Extreme Value Distribution

Description

Estimate quantiles of a [generalized extreme value distribution](#).

Usage

```
eqgevd(x, p = 0.5, method = "mle", pwme.method = "unbiased",
       tsoe.method = "med", plot.pos.cons = c(a = 0.35, b = 0), digits = 0)
```

Arguments

<code>x</code>	a numeric vector of observations, or an object resulting from a call to an estimating function that assumes a generalized extreme value distribution (e.g., egevd). If <code>x</code> is a numeric vector, missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
<code>p</code>	numeric vector of probabilities for which quantiles will be estimated. All values of <code>p</code> must be between 0 and 1. The default value is <code>p=0.5</code> .
<code>method</code>	character string specifying the method to use to estimate the location, scale, and threshold parameters. Possible values are "mle" (maximum likelihood; the default), "pwme" (probability-weighted moments), and "tsoe" (two-stage order-statistics estimator of Castillo and Hadi (1994)). See the DETAILS section of the help file for egevd for more information on these estimation methods.
<code>pwme.method</code>	character string specifying what method to use to compute the probability-weighted moments when <code>method="pwme"</code> . The possible values are "ubiased" (method based on the U-statistic; the default), or "plotting.position" (method based on the plotting position formula). See the DETAILS section of the help file for egevd for more information. This argument is ignored if <code>method</code> is not equal to "pwme".
<code>tsoe.method</code>	character string specifying the robust function to apply in the second stage of the two-stage order-statistics estimator when <code>method="tsoe"</code> . Possible values are "med" (median; the default), and "lms" (least median of squares). See the DETAILS section of the help file for egevd for more information on these estimation methods. This argument is ignored if <code>method</code> is not equal to "tsoe".
<code>plot.pos.cons</code>	numeric vector of length 2 specifying the constants used in the formula for the plotting positions when <code>method="pwme"</code> and <code>pwme.method="plotting.position"</code> . The default value is <code>plot.pos.cons=c(a=0.35, b=0)</code> . If this vector has a names attribute with the value <code>c("a", "b")</code> or <code>c("b", "a")</code> , then the elements will be matched by name in the formula for computing the plotting positions. Otherwise, the first element is mapped to the name "a" and the second element to the name "b". See the DETAILS section of the help file for egevd for more information. This argument is used only if <code>method="tsoe"</code> , or if both <code>method="pwme"</code> and <code>pwme.method="plotting.position"</code> .
<code>digits</code>	an integer indicating the number of decimal places to round to when printing out the value of <code>100*p</code> . The default value is <code>digits=0</code> .

Details

The function `eqgevd` returns estimated quantiles as well as estimates of the location, scale and threshold parameters.

Quantiles are estimated by 1) estimating the location, scale, and threshold parameters by calling [egevd](#), and then 2) calling the function [qgevd](#) and using the estimated values for location, scale, and threshold.

Value

If `x` is a numeric vector, `eqgevd` returns a list of class "estimate" containing the estimated quantile(s) and other information. See [estimate.object](#) for details.

If `x` is the result of calling an estimation function, `eqgevd` returns a list whose class is the same as `x`. The list contains the same components as `x`, as well as components called `quantiles` and `quantile.method`.

Note

Two-parameter [extreme value distributions](#) (EVD) have been applied extensively since the 1930's to several fields of study, including the distributions of hydrological and meteorological variables, human lifetimes, and strength of materials. The three-parameter [generalized extreme value distribution](#) (GEVD) was introduced by Jenkinson (1955) to model annual maximum and minimum values of meteorological events. Since then, it has been used extensively in the hydrological and meteorological fields.

The three families of EVDs are all special kinds of GEVDs. When the shape parameter $\kappa = 0$, the GEVD reduces to the Type I extreme value (Gumbel) distribution. (The function [zTestGevdShape](#) allows you to test the null hypothesis $H_0 : \kappa = 0$.) When $\kappa > 0$, the GEVD is the same as the Type II extreme value distribution, and when $\kappa < 0$ it is the same as the Type III extreme value distribution.

Hosking et al. (1985) compare the asymptotic and small-sample statistical properties of the PWME with the MLE and Jenkinson's (1969) method of sextiles. Castillo and Hadi (1994) compare the small-sample statistical properties of the MLE, PWME, and TSOE. Hosking and Wallis (1995) compare the small-sample properties of unbiased L -moment estimators vs. plotting-position L -moment estimators. (PWMEs can be written as linear combinations of L -moments and thus have equivalent statistical properties.) Hosking and Wallis (1995) conclude that unbiased estimators should be used for almost all applications.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Castillo, E., and A. Hadi. (1994). Parameter and Quantile Estimation for the Generalized Extreme-Value Distribution. *Environmetrics* **5**, 417–432.
- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.
- Greenwood, J.A., J.M. Landwehr, N.C. Matalas, and J.R. Wallis. (1979). Probability Weighted Moments: Definition and Relation to Parameters of Several Distributions Expressible in Inverse Form. *Water Resources Research* **15**(5), 1049–1054.
- Hosking, J.R.M. (1984). Testing Whether the Shape Parameter is Zero in the Generalized Extreme-Value Distribution. *Biometrika* **71**(2), 367–374.
- Hosking, J.R.M. (1985). Algorithm AS 215: Maximum-Likelihood Estimation of the Parameters of the Generalized Extreme-Value Distribution. *Applied Statistics* **34**(3), 301–310.
- Hosking, J.R.M., J.R. Wallis, and E.F. Wood. (1985). Estimation of the Generalized Extreme-Value Distribution by the Method of Probability-Weighted Moments. *Technometrics* **27**(3), 251–261.
- Jenkinson, A.F. (1969). *Statistics of Extremes*. *Technical Note 98*, World Meteorological Office, Geneva.
- Johnson, N. L., S. Kotz, and N. Balakrishnan. (1995). *Continuous Univariate Distributions, Volume 2*. Second Edition. John Wiley and Sons, New York.
- Landwehr, J.M., N.C. Matalas, and J.R. Wallis. (1979). Probability Weighted Moments Compared With Some Traditional Techniques in Estimating Gumbel Parameters and Quantiles. *Water Resources Research* **15**(5), 1055–1064.
- Macleod, A.J. (1989). Remark AS R76: A Remark on Algorithm AS 215: Maximum Likelihood Estimation of the Parameters of the Generalized Extreme-Value Distribution. *Applied Statistics* **38**(1), 198–199.

Prescott, P., and A.T. Walden. (1980). Maximum Likelihood Estimation of the Parameters of the Generalized Extreme-Value Distribution. *Biometrika* **67**(3), 723–724.

Prescott, P., and A.T. Walden. (1983). Maximum Likelihood Estimation of the Three-Parameter Generalized Extreme-Value Distribution from Censored Samples. *Journal of Statistical Computing and Simulation* **16**, 241–250.

See Also

[egevd](#), [Generalized Extreme Value Distribution](#), [Extreme Value Distribution](#), [eevd](#), [estimate.object](#).

Examples

```
# Generate 20 observations from a generalized extreme value distribution
# with parameters location=2, scale=1, and shape=0.2, then compute the
# MLEs of location, shape, and threshold, and estimate the 90th percentile.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(498)
dat <- rgevd(20, location = 2, scale = 1, shape = 0.2)
eqgevd(dat, p = 0.9)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Generalized Extreme Value
#
#Estimated Parameter(s):      location = 1.6144631
#                               scale    = 0.9867007
#                               shape    = 0.2632493
#
#Estimation Method:           mle
#
#Estimated Quantile(s):       90'th %ile = 3.289912
#
#Quantile Estimation Method:   Quantile(s) Based on
#                               mle Estimators
#
#Data:                         dat
#
#Sample Size:                  20

#-----

# Clean up
#-----
rm(dat)
```

eqhyper

Estimate Quantiles of a Hypergeometric Distribution

Description

Estimate quantiles of a [hypergeometric distribution](#).

Usage

```
eqhyper(x, m = NULL, total = NULL, k = NULL, p = 0.5, method = "mle", digits = 0)
```

Arguments

<code>x</code>	non-negative integer indicating the number of white balls out of a sample of size k drawn without replacement from the urn, or an object resulting from a call to an estimating function that assumes a hypergeometric distribution (e.g., ehyper). Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed.
<code>m</code>	non-negative integer indicating the number of white balls in the urn. You must supply <code>m</code> or <code>total</code> , but not both. Missing values (NAs) are not allowed.
<code>total</code>	positive integer indicating the total number of balls in the urn (i.e., $m+n$). You must supply <code>m</code> or <code>total</code> , but not both. Missing values (NAs) are not allowed.
<code>k</code>	positive integer indicating the number of balls drawn without replacement from the urn. Missing values (NAs) are not allowed.
<code>p</code>	numeric vector of probabilities for which quantiles will be estimated. All values of <code>p</code> must be between 0 and 1. The default value is <code>p=0.5</code> .
<code>method</code>	character string specifying the method of estimating the parameters of the hypergeometric distribution. Possible values are "mle" (maximum likelihood; the default) and "mvue" (minimum variance unbiased). The mvue method is only available when you are estimating m (i.e., when you supply the argument <code>total</code>). See the DETAILS section of the help file for ehyper for more information on these estimation methods.
<code>digits</code>	an integer indicating the number of decimal places to round to when printing out the value of $100 \cdot p$. The default value is <code>digits=0</code> .

Details

The function `eqhyper` returns estimated quantiles as well as estimates of the hypergeometric distribution parameters.

Quantiles are estimated by 1) estimating the distribution parameters by calling [ehyper](#), and then 2) calling the function [qhyper](#) and using the estimated values for the distribution parameters.

Value

If `x` is a numeric vector, `eqhyper` returns a list of class "estimate" containing the estimated quantile(s) and other information. See [estimate.object](#) for details.

If `x` is the result of calling an estimation function, `eqhyper` returns a list whose class is the same as `x`. The list contains the same components as `x`, as well as components called `quantiles` and `quantile.method`.

Note

The [hypergeometric distribution](#) can be described by an urn model with M white balls and N black balls. If K balls are drawn *with* replacement, then the number of white balls in the sample of size K follows a [binomial distribution](#) with parameters `size=K` and `prob=M/(M+N)`. If K balls are drawn *without* replacement, then the number of white balls in the sample of size K follows a [hypergeometric distribution](#) with parameters `m=M`, `n=N`, and `k=K`.

The name “hypergeometric” comes from the fact that the probabilities associated with this distribution can be written as successive terms in the expansion of a function of a Gaussian hypergeometric series.

The hypergeometric distribution is applied in a variety of fields, including quality control and estimation of animal population size. It is also the distribution used to compute probabilities for [Fishers’s exact test](#) for a 2x2 contingency table.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.

Johnson, N. L., S. Kotz, and A. Kemp. (1992). *Univariate Discrete Distributions*. Second Edition. John Wiley and Sons, New York, Chapter 6.

See Also

[ehyper](#), [Hypergeometric](#), [estimate.object](#).

Examples

```
# Generate an observation from a hypergeometric distribution with
# parameters m=10, n=30, and k=5, then estimate the parameter m, and
# the 80'th percentile.
# Note: the call to set.seed simply allows you to reproduce this example.
# Also, the only parameter actually estimated is m; once m is estimated,
# n is computed by subtracting the estimated value of m (8 in this example)
# from the given of value of m+n (40 in this example). The parameters
# n and k are shown in the output in order to provide information on
# all of the parameters associated with the hypergeometric distribution.
```

```
set.seed(250)
dat <- rhyper(nn = 1, m = 10, n = 30, k = 5)
dat
#[1] 1
```

```
eqhyper(dat, total = 40, k = 5, p = 0.8)
```

```
#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Hypergeometric
#
#Estimated Parameter(s):      m =  8
#                               n = 32
#                               k =  5
#
#Estimation Method:           mle for 'm'
#
#Estimated Quantile(s):       80'th %ile = 2
#
#Quantile Estimation Method:   Quantile(s) Based on
```

```
#                               mle for 'm' Estimators
#
#Data:                           dat
#
#Sample Size:                     1

#-----

# Clean up
#-----
rm(dat)
```

eqlnorm

Estimate Quantiles of a Lognormal Distribution

Description

Estimate quantiles of a [lognormal distribution](#), and optionally construct a confidence interval for a quantile.

Usage

```
eqlnorm(x, p = 0.5, method = "qmle", ci = FALSE,
        ci.method = "exact", ci.type = "two-sided", conf.level = 0.95,
        digits = 0)
```

Arguments

x	a numeric vector of positive observations, or an object resulting from a call to an estimating function that assumes a lognormal distribution (i.e., elnorm , elnormSinglyCensored , or elnormMultiplyCensored). You <i>cannot</i> use objects resulting from a call to estimating functions that use the alternative parameterization such as elnormAlt . If x is a numeric vector, missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
p	numeric vector of probabilities for which quantiles will be estimated. All values of p must be between 0 and 1. When ci=TRUE, p must be a scalar. The default value is p=0.5.
method	character string indicating what method to use to estimate the quantile(s). The possible values are "qmle" (quasi maximum likelihood; the default) and "mvue" (minimum variance unbiased). The method "mvue" is available only when p=0.5 (i.e., when you are estimating the median). See the DETAILS section for more information.
ci	logical scalar indicating whether to compute a confidence interval for the quantile. The default value is ci=FALSE.
ci.method	character string indicating what method to use to construct the confidence interval for the quantile. The possible values are "exact" (exact method; the default) and "normal.approx" (normal approximation). See the DETAILS section for more information.
ci.type	character string indicating what kind of confidence interval for the quantile to compute. The possible values are "two-sided" (the default), "lower", and "upper". This argument is ignored if ci=FALSE.

<code>conf.level</code>	a scalar between 0 and 1 indicating the confidence level of the confidence interval. The default value is <code>conf.level=0.95</code> . This argument is ignored if <code>ci=FALSE</code> .
<code>digits</code>	an integer indicating the number of decimal places to round to when printing out the value of $100 \cdot p$. The default value is <code>digits=0</code> .

Details

If x contains any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

Quantiles and their associated confidence intervals are constructed by calling the function `eqnorm` using the log-transformed data and then exponentiating the quantiles and confidence limits.

In the special case when $p=0.5$ and `method="mvue"`, the estimated median is computed using the method given in Gilbert (1987, p.172) and Bradu and Mundlak (1970).

Value

If x is a numeric vector, `eqlnorm` returns a list of class "estimate" containing the estimated quantile(s) and other information. See `estimate.object` for details.

If x is the result of calling an estimation function, `eqlnorm` returns a list whose class is the same as x . The list contains the same components as x , as well as components called `quantiles` and `quantile.method`. In addition, if `ci=TRUE`, the returned list contains a component called `interval` containing the confidence interval information. If x already has a component called `interval`, this component is replaced with the confidence interval information.

Note

Percentiles are sometimes used in environmental standards and regulations. For example, Berthouex and Brown (2002, p.71) note that England has water quality limits based on the 90th and 95th percentiles of monitoring data not exceeding specified levels. They also note that the U.S. EPA has specifications for air quality monitoring, aquatic standards on toxic chemicals, and maximum daily limits for industrial effluents that are all based on percentiles. Given the importance of these quantities, it is essential to characterize the amount of uncertainty associated with the estimates of these quantities. This is done with confidence intervals.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Berthouex, P.M., and L.C. Brown. (2002). *Statistics for Environmental Engineers*. Lewis Publishers, Boca Raton.
- Bradru, D., and Y. Mundlak. (1970). Estimation in Lognormal Linear Models. *Journal of the American Statistical Association* **65**, 198-211.
- Conover, W.J. (1980). *Practical Nonparametric Statistics*. Second Edition. John Wiley and Sons, New York.
- Gibbons, R.D., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*, Second Edition. John Wiley & Sons, Hoboken.
- Gilbert, R.O. (1987). *Statistical Methods for Environmental Pollution Monitoring*. Van Nostrand Reinhold, New York.

- Helsel, D.R., and R.M. Hirsch. (1992). *Statistical Methods in Water Resources Research*. Elsevier, New York, NY, pp.88-90.
- Johnson, N.L., and B.L. Welch. (1940). Applications of the Non-Central t-Distribution. *Biometrika* **31**, 362-389.
- Millard, S.P., and Neerchal, N.K. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, Florida.
- Owen, D.B. (1962). *Handbook of Statistical Tables*. Addison-Wesley, Reading, MA.
- Stedinger, J. (1983). Confidence Intervals for Design Events. *Journal of Hydraulic Engineering* **109**(1), 13-27.
- Stedinger, J.R., R.M. Vogel, and E. Foufoula-Georgiou. (1993). Frequency Analysis of Extreme Events. In: Maidment, D.R., ed. *Handbook of Hydrology*. McGraw-Hill, New York, Chapter 18, pp.29-30.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.
- USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

See Also

[eqnorm](#), [Lognormal](#), [elnorm](#), [estimate.object](#).

Examples

```
# Generate 20 observations from a lognormal distribution with
# parameters meanlog=3 and sdlog=0.5, then estimate the 90th
# percentile and create a one-sided upper 95% confidence interval
# for that percentile.
# (Note: the call to set.seed simply allows you to reproduce this
# example.)

set.seed(47)
dat <- rlnorm(20, meanlog = 3, sdlog = 0.5)
eqlnorm(dat, p = 0.9, ci = TRUE, ci.type = "upper")

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Lognormal
#
#Estimated Parameter(s):      meanlog = 2.9482139
#                               sdlog   = 0.4553215
#
#Estimation Method:           mvue
#
#Estimated Quantile(s):       90'th %ile = 34.18312
#
#Quantile Estimation Method:   qmle
#
#Data:                         dat
#
```

```

#Sample Size:                20
#
#Confidence Interval for:    90'th %ile
#
#Confidence Interval Method: Exact
#
#Confidence Interval Type:   upper
#
#Confidence Level:          95%
#
#Confidence Interval:       LCL =  0.00000
#                           UCL = 45.84008

#-----
# Compare these results with the true 90'th percentile:

qlnorm(p = 0.9, meanlog = 3, sdlog = 0.5)
#[1] 38.1214

#-----

# Clean up
rm(dat)

#-----

# Example 17-3 of USEPA (2009, p. 17-17) shows how to construct a
# beta-content upper tolerance limit with 95% coverage and 95%
# confidence using chrysene data and assuming a lognormal
# distribution.

# A beta-content upper tolerance limit with 95% coverage and 95%
# confidence is equivalent to the 95% upper confidence limit for the
# 95th percentile.

attach(EPA.09.Ex.17.3.chrysene.df)
Chrysene <- Chrysene.ppb[Well.type == "Background"]
eqlnorm(Chrysene, p = 0.95, ci = TRUE, ci.type = "upper")

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:       Lognormal
#
#Estimated Parameter(s):    meanlog = 2.5085773
#                           sdlog   = 0.6279479
#
#Estimation Method:         mvue
#
#Estimated Quantile(s):     95'th %ile = 34.51727
#
#Quantile Estimation Method: qmle
#
#Data:                      Chrysene
#
#Sample Size:               8
#

```

```
#Confidence Interval for:      95'th %ile
#
#Confidence Interval Method:   Exact
#
#Confidence Interval Type:    upper
#
#Confidence Level:           95%
#
#Confidence Interval:        LCL =  0.0000
#                             UCL = 90.9247

#-----
# Clean up

rm(Chrysene)
detach("EPA.09.Ex.17.3.chrysene.df")
```

eqlnorm3	<i>Estimate Quantiles of a Three-Parameter Lognormal Distribution</i>
----------	---

Description

Estimate quantiles of a [three-parameter lognormal distribution](#).

Usage

```
eqlnorm3(x, p = 0.5, method = "lmle", digits = 0)
```

Arguments

x	a numeric vector of observations, or an object resulting from a call to an estimating function that assumes a three-parameter lognormal distribution (e.g., elnorm3). If x is a numeric vector, missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
p	numeric vector of probabilities for which quantiles will be estimated. All values of p must be between 0 and 1. When ci=TRUE, p must be a scalar. The default value is p=0.5.
method	character string specifying the method of estimating the distribution parameters. Possible values are "lmle" (local maximum likelihood; the default), "mme" (method of moments), "mmue" (method of moments using an unbaised estimate of variance), "mmme" (modified method of moments due to Cohen and Whitten (1980)), "zero.skew" (zero-skewness estimator due to Griffiths (1980)), and "royston.skew" (estimator based on Royston's (1992b) index of skewness). See the DETAILS section of the help file for elnorm3 for more information on these estimation methods.
digits	an integer indicating the number of decimal places to round to when printing out the value of 100*p. The default value is digits=0.

Details

If x contains any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

Quantiles are estimated by 1) estimating the distribution parameters by calling `elnorm3`, and then 2) calling the function `qlnorm3` and using the estimated distribution parameters.

Value

If x is a numeric vector, `eqlnorm3` returns a list of class "estimate" containing the estimated quantile(s) and other information. See `estimate.object` for details.

If x is the result of calling an estimation function, `eqlnorm3` returns a list whose class is the same as x . The list contains the same components as x , as well as components called `quantiles` and `quantile.method`.

Note

The problem of estimating the parameters of a three-parameter lognormal distribution has been extensively discussed by Aitchison and Brown (1957, Chapter 6), Calitz (1973), Cohen (1951), Cohen (1988), Cohen and Whitten (1980), Cohen et al. (1985), Griffiths (1980), Harter and Moore (1966), Hill (1963), and Royston (1992b). Stedinger (1980) and Hoshi et al. (1984) discuss fitting the three-parameter lognormal distribution to hydrologic data.

The global maximum likelihood estimates are inadmissible. In the past, several researchers have found that the local maximum likelihood estimates (lmle's) occasionally fail because of convergence problems, but they were not using the likelihood profile and reparameterization of Griffiths (1980). Cohen (1988) recommends the modified methods of moments estimators over lmle's because they are easy to compute, they are unbiased with respect to μ and σ^2 (the mean and standard deviation on the log-scale), their variances are minimal or near minimal, and they do not suffer from regularity problems.

Because the distribution of the lmle of the threshold parameter γ is far from normal for moderate sample sizes (Griffiths, 1980), it is questionable whether confidence intervals for γ or the median based on asymptotic variances and covariances will perform well. Cohen and Whitten (1980) and Cohen et al. (1985), however, found that the asymptotic variances and covariances are reasonably close to corresponding simulated variances and covariances for the modified method of moments estimators (`method="mmme"`). In a simulation study (5000 monte carlo trials), Royston (1992b) found that the coverage of confidence intervals for γ based on the likelihood profile (`ci.method="likelihood.profile"`) was very close the nominal level (94.1% for a nominal level of 95%), although not symmetric. Royston (1992b) also found that the coverage of confidence intervals for γ based on the skewness method (`ci.method="skewness"`) was also very close (95.4%) and symmetric.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Aitchison, J., and J.A.C. Brown (1957). *The Lognormal Distribution (with special references to its uses in economics)*. Cambridge University Press, London, Chapter 5.
- Calitz, F. (1973). Maximum Likelihood Estimation of the Parameters of the Three-Parameter Lognormal Distribution—a Reconsideration. *Australian Journal of Statistics* **15**(3), 185–190.

- Cohen, A.C. (1951). Estimating Parameters of Logarithmic-Normal Distributions by Maximum Likelihood. *Journal of the American Statistical Association* **46**, 206–212.
- Cohen, A.C. (1988). Three-Parameter Estimation. In Crow, E.L., and K. Shimizu, eds. *Lognormal Distributions: Theory and Applications*. Marcel Dekker, New York, Chapter 4.
- Cohen, A.C., and B.J. Whitten. (1980). Estimation in the Three-Parameter Lognormal Distribution. *Journal of the American Statistical Association* **75**, 399–404.
- Cohen, A.C., B.J. Whitten, and Y. Ding. (1985). Modified Moment Estimation for the Three-Parameter Lognormal Distribution. *Journal of Quality Technology* **17**, 92–99.
- Crow, E.L., and K. Shimizu. (1988). *Lognormal Distributions: Theory and Applications*. Marcel Dekker, New York, Chapter 2.
- Griffiths, D.A. (1980). Interval Estimation for the Three-Parameter Lognormal Distribution via the Likelihood Function. *Applied Statistics* **29**, 58–68.
- Harter, H.L., and A.H. Moore. (1966). Local-Maximum-Likelihood Estimation of the Parameters of Three-Parameter Lognormal Populations from Complete and Censored Samples. *Journal of the American Statistical Association* **61**, 842–851.
- Heyde, C.C. (1963). On a Property of the Lognormal Distribution. *Journal of the Royal Statistical Society, Series B* **25**, 392–393.
- Hill, .B.M. (1963). The Three-Parameter Lognormal Distribution and Bayesian Analysis of a Point-Source Epidemic. *Journal of the American Statistical Association* **58**, 72–84.
- Hoshi, K., J.R. Stedinger, and J. Burges. (1984). Estimation of Log-Normal Quantiles: Monte Carlo Results and First-Order Approximations. *Journal of Hydrology* **71**, 1–30.
- Johnson, N. L., S. Kotz, and N. Balakrishnan. (1994). *Continuous Univariate Distributions, Volume 1*. Second Edition. John Wiley and Sons, New York.
- Royston, J.P. (1992b). Estimation, Reference Ranges and Goodness of Fit for the Three-Parameter Log-Normal Distribution. *Statistics in Medicine* **11**, 897–912.
- Stedinger, J.R. (1980). Fitting Lognormal Distributions to Hydrologic Data. *Water Resources Research* **16**(3), 481–490.

See Also

[elnorm3](#), [Lognormal3](#), [Lognormal](#), [LognormalAlt](#), [Normal](#).

Examples

```
# Generate 20 observations from a 3-parameter lognormal distribution
# with parameters meanlog=1.5, sdlog=1, and threshold=10, then use
# Cohen and Whitten's (1980) modified moments estimators to estimate
# the parameters, and estimate the 90th percentile.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
dat <- rlnorm3(20, meanlog = 1.5, sdlog = 1, threshold = 10)
eqlnorm3(dat, method = "mmme", p = 0.9)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          3-Parameter Lognormal
#
#Estimated Parameter(s):      meanlog    = 1.5206664
```



```

#                      sdlog      = 0.5330974
#                      threshold = 9.6620403
#
#Estimation Method:      mmme
#
#Estimated Quantile(s):  90'th %ile = 18.72194
#
#Quantile Estimation Method: Quantile(s) Based on
#                          mmme Estimators
#
#Data:                  dat
#
#Sample Size:           20

# Clean up
#-----
rm(dat)

```

eqlogis

Estimate Quantiles of a Logistic Distribution

Description

Estimate quantiles of a [logistic distribution](#).

Usage

```
eqlogis(x, p = 0.5, method = "mle", digits = 0)
```

Arguments

x	a numeric vector of observations, or an object resulting from a call to an estimating function that assumes a logistic distribution (e.g., elogis). If x is a numeric vector, missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
p	numeric vector of probabilities for which quantiles will be estimated. All values of p must be between 0 and 1. The default value is p=0.5.
method	character string specifying the method to use to estimate the distribution parameters. Possible values are "mle" (maximum likelihood; the default), "mme" (methods of moments), and "mmue" (method of moments based on the unbiased estimator of variance). See the DETAILS section of the help file for elogis for more information.
digits	an integer indicating the number of decimal places to round to when printing out the value of 100*p. The default value is digits=0.

Details

The function eqlogis returns estimated quantiles as well as estimates of the location and scale parameters.

Quantiles are estimated by 1) estimating the location and scale parameters by calling [elogis](#), and then 2) calling the function [qlogis](#) and using the estimated values for location and scale.

Value

If `x` is a numeric vector, `eqlogis` returns a list of class "estimate" containing the estimated quantile(s) and other information. See [estimate.object](#) for details.

If `x` is the result of calling an estimation function, `eqlogis` returns a list whose class is the same as `x`. The list contains the same components as `x`, as well as components called `quantiles` and `quantile.method`.

Note

The [logistic distribution](#) is defined on the real line and is unimodal and symmetric about its location parameter (the mean). It has longer tails than a normal (Gaussian) distribution. It is used to model growth curves and bioassay data.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.

Johnson, N. L., S. Kotz, and N. Balakrishnan. (1995). *Continuous Univariate Distributions, Volume 2*. Second Edition. John Wiley and Sons, New York.

See Also

[elogis](#), [Logistic](#), [estimate.object](#).

Examples

```
# Generate 20 observations from a logistic distribution with
# parameters location=0 and scale=1, then estimate the parameters
# and estimate the 90th percentile.
# (Note: the call to set.seed simply allows you to reproduce this example.)
```

```
set.seed(250)
dat <- rlogis(20)
eqlogis(dat, p = 0.9)
```

```
#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Logistic
#
#Estimated Parameter(s):      location = -0.2181845
#                               scale    =  0.8152793
#
#Estimation Method:           mle
#
#Estimated Quantile(s):       90'th %ile = 1.573167
#
#Quantile Estimation Method:   Quantile(s) Based on
#                               mle Estimators
#
```

```
#Data:                                dat
#
#Sample Size:                          20

#-----
# Clean up

rm(dat)
```

eqnbinom

*Estimate Quantiles of a Negative Binomial Distribution***Description**

Estimate quantiles of a [negative binomial distribution](#).

Usage

```
eqnbinom(x, size = NULL, p = 0.5, method = "mle/mme", digits = 0)
```

Arguments

x	vector of non-negative integers indicating the number of trials that took place <i>before</i> size “successes” occurred (the total number of trials that took place is $x+1$), or an object resulting from a call to an estimating function that assumes a negative binomial distribution (e.g., enbinom). If x is a vector of non-negative integers, then missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed. If $\text{length}(x)=n$ and n is greater than 1, it is assumed that x represents observations from n separate negative binomial experiments that all had the same probability of success (prob), but possibly different values of size.
size	vector of positive integers indicating the number of “successes” that must be observed before the trials are stopped. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed. The length of size must be 1 or else the same length as x.
p	numeric vector of probabilities for which quantiles will be estimated. All values of p must be between 0 and 1. The default value is $p=0.5$.
method	character string specifying the method of estimating the probability parameter. Possible values are “mle/mme” (maximum likelihood and method of moments; the default) and “mvue” (minimum variance unbiased). You cannot use <code>method=“mvue”</code> if the sum of the elements in size is 1. See the DETAILS section of the help file for enbinom for more information on these estimation methods.
digits	an integer indicating the number of decimal places to round to when printing out the value of $100 \times p$. The default value is <code>digits=0</code> .

Details

The function eqnbinom returns estimated quantiles as well as estimates of the prob parameter.

Quantiles are estimated by 1) estimating the prob parameter by calling [enbinom](#), and then 2) calling the function [qnbinom](#) and using the estimated value for prob.

Value

If `x` is a numeric vector, `eqnbinom` returns a list of class "estimate" containing the estimated quantile(s) and other information. See `estimate.object` for details.

If `x` is the result of calling an estimation function, `eqnbinom` returns a list whose class is the same as `x`. The list contains the same components as `x`, as well as components called `quantiles` and `quantile.method`.

Note

The [negative binomial distribution](#) has its roots in a gambling game where participants would bet on the number of tosses of a coin necessary to achieve a fixed number of heads. The negative binomial distribution has been applied in a wide variety of fields, including accident statistics, birth-and-death processes, and modeling spatial distributions of biological organisms.

The [geometric distribution](#) with parameter $\text{prob}=p$ is a special case of the negative binomial distribution with parameters $\text{size}=1$ and $\text{prob}=p$.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.

Johnson, N. L., S. Kotz, and A. Kemp. (1992). *Univariate Discrete Distributions*. Second Edition. John Wiley and Sons, New York, Chapter 5.

See Also

[enbinom](#), [NegBinomial](#), [egeom](#), [Geometric](#), [estimate.object](#).

Examples

```
# Generate an observation from a negative binomial distribution with
# parameters size=2 and prob=0.2, then estimate the parameter prob
# and the 90th percentile.
# Note: the call to set.seed simply allows you to reproduce this example.
# Also, the only parameter that is estimated is prob; the parameter
# size is supplied in the call to enbinom. The parameter size is printed in
# order to show all of the parameters associated with the distribution.

set.seed(250)
dat <- rnbinom(1, size = 2, prob = 0.2)
dat
#[1] 5

eqnbinom(dat, size = 2, p = 0.9)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Negative Binomial
#
#Estimated Parameter(s):      size = 2.0000000
```

```

#                                prob = 0.2857143
#
#Estimation Method:              mle/mme for 'prob'
#
#Estimated Quantile(s):         90'th %ile = 11
#
#Quantile Estimation Method:     Quantile(s) Based on
#                                mle/mme for 'prob' Estimators
#
#Data:                           dat, 2
#
#Sample Size:                    1

#-----
# Clean up

rm(dat)

```

eqnorm

Estimate Quantiles of a Normal Distribution

Description

Estimate quantiles of a [normal distribution](#), and optionally construct a confidence interval for a quantile.

Usage

```
eqnorm(x, p = 0.5, method = "qmle", ci = FALSE,
       ci.method = "exact", ci.type = "two-sided", conf.level = 0.95,
       digits = 0, warn = TRUE)
```

Arguments

x	a numeric vector of observations, or an object resulting from a call to an estimating function that assumes a normal (Gaussian) distribution (i.e., enorm , enormSinglyCensored , or enormMultiplyCensored). If x is a numeric vector, missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
p	numeric vector of probabilities for which quantiles will be estimated. All values of p must be between 0 and 1. When ci=TRUE, p must be a scalar. The default value is p=0.5.
method	character string indicating what method to use to estimate the quantile(s). Currently the only possible value is method="qmle" (quasi maximum likelihood). See the DETAILS section for more information.
ci	logical scalar indicating whether to compute a confidence interval for the quantile. The default value is ci=FALSE.
ci.method	character string indicating what method to use to construct the confidence interval for the quantile. The possible values are "exact" (exact method; the default) and "normal.approx" (normal approximation). See the DETAILS section for more information.

<code>ci.type</code>	character string indicating what kind of confidence interval for the quantile to compute. The possible values are "two-sided" (the default), "lower", and "upper". This argument is ignored if <code>ci=FALSE</code> .
<code>conf.level</code>	a scalar between 0 and 1 indicating the confidence level of the confidence interval. The default value is <code>conf.level=0.95</code> . This argument is ignored if <code>ci=FALSE</code> .
<code>digits</code>	an integer indicating the number of decimal places to round to when printing out the value of $100 \times p$. The default value is <code>digits=0</code> .
<code>warn</code>	logical scalar indicating whether to warn in the case when <code>ci=TRUE</code> , <code>ci.method="exact"</code> , and the supplied object <code>x</code> is of class "estimate" but did not use <code>method="mvue"</code> for estimation.

Details

If `x` contains any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

Quantiles are estimated by 1) estimating the mean and standard deviation parameters by calling `enorm` with `method="mvue"`, and then 2) calling the function `qnorm` and using the estimated values for mean and standard deviation. This estimator of the p 'th quantile is sometimes called the quasi-maximum likelihood estimator (qml; Cohn et al., 1989) because if the maximum likelihood estimator of standard deviation were used in place of the minimum variance unbiased one, then this estimator of the quantile would be the mle of the p 'th quantile.

When `ci=TRUE` and `ci.method="exact"`, the confidence interval for a quantile is computed by using the relationship between a confidence interval for a quantile and a tolerance interval. Specifically, it can be shown (e.g., Conover, 1980, pp.119-121) that an upper confidence interval for the p 'th quantile with confidence level $100(1 - \alpha)\%$ is equivalent to an upper β -content tolerance interval with coverage $100p\%$ and confidence level $100(1 - \alpha)\%$. Also, a lower confidence interval for the p 'th quantile with confidence level $100(1 - \alpha)\%$ is equivalent to a lower β -content tolerance interval with coverage $100(1 - p)\%$ and confidence level $100(1 - \alpha)\%$. See the help file for `tolIntNorm` for information on tolerance intervals for a normal distribution.

When `ci=TRUE` and `ci.method="normal.approx"`, the confidence interval for a quantile is computed by assuming the estimated quantile has an approximately normal distribution and using the asymptotic variance to construct the confidence interval (see Stedinger, 1983; Stedinger et al., 1993).

Value

If `x` is a numeric vector, `eqnorm` returns a list of class "estimate" containing the estimated quantile(s) and other information. See `estimate.object` for details.

If `x` is the result of calling an estimation function, `eqnorm` returns a list whose class is the same as `x`. The list contains the same components as `x`, as well as components called `quantiles` and `quantile.method`. In addition, if `ci=TRUE`, the returned list contains a component called `interval` containing the confidence interval information. If `x` already has a component called `interval`, this component is replaced with the confidence interval information.

Note

Percentiles are sometimes used in environmental standards and regulations. For example, Berthouex and Brown (2002, p.71) note that England has water quality limits based on the 90th and 95th percentiles of monitoring data not exceeding specified levels. They also note that the U.S. EPA has specifications for air quality monitoring, aquatic standards on toxic chemicals, and maximum

daily limits for industrial effluents that are all based on percentiles. Given the importance of these quantities, it is essential to characterize the amount of uncertainty associated with the estimates of these quantities. This is done with confidence intervals.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Berthouex, P.M., and L.C. Brown. (2002). *Statistics for Environmental Engineers*. Lewis Publishers, Boca Raton.
- Conover, W.J. (1980). *Practical Nonparametric Statistics*. Second Edition. John Wiley and Sons, New York.
- Gibbons, R.D., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*, Second Edition. John Wiley & Sons, Hoboken.
- Gilbert, R.O. (1987). *Statistical Methods for Environmental Pollution Monitoring*. Van Nostrand Reinhold, New York, NY, pp.132-136.
- Helsel, D.R., and R.M. Hirsch. (1992). *Statistical Methods in Water Resources Research*. Elsevier, New York, NY, pp.88-90.
- Johnson, N.L., and B.L. Welch. (1940). Applications of the Non-Central t-Distribution. *Biometrika* **31**, 362-389.
- Millard, S.P., and Neerchal, N.K. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, Florida.
- Owen, D.B. (1962). *Handbook of Statistical Tables*. Addison-Wesley, Reading, MA.
- Stedinger, J. (1983). Confidence Intervals for Design Events. *Journal of Hydraulic Engineering* **109**(1), 13-27.
- Stedinger, J.R., R.M. Vogel, and E. Foufoula-Georgiou. (1993). Frequency Analysis of Extreme Events. In: Maidment, D.R., ed. *Handbook of Hydrology*. McGraw-Hill, New York, Chapter 18, pp.29-30.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.
- USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

See Also

[enorm](#), [tolIntNorm](#), [Normal](#), [estimate.object](#).

Examples

```
# Generate 20 observations from a normal distribution with
# parameters mean=10 and sd=2, then estimate the 90th
# percentile and create a one-sided upper 95% confidence interval
# for that percentile.
# (Note: the call to set.seed simply allows you to reproduce this
# example.)
```

```

set.seed(47)
dat <- rnorm(20, mean = 10, sd = 2)
eqnorm(dat, p = 0.9, ci = TRUE, ci.type = "upper")

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Normal
#
#Estimated Parameter(s):       mean = 9.792856
#                               sd   = 1.821286
#
#Estimation Method:            mvue
#
#Estimated Quantile(s):        90'th %ile = 12.12693
#
#Quantile Estimation Method:    qmle
#
#Data:                          dat
#
#Sample Size:                   20
#
#Confidence Interval for:       90'th %ile
#
#Confidence Interval Method:     Exact
#
#Confidence Interval Type:       upper
#
#Confidence Level:              95%
#
#Confidence Interval:           LCL =      -Inf
#                               UCL = 13.30064
#
#-----
# Compare these results with the true 90'th percentile:

qnorm(p = 0.9, mean = 10, sd = 2)
#[1] 12.56310

#-----

# Clean up
rm(dat)

#-----

# Example 17-3 of USEPA (2009, p. 17-17) shows how to construct a
# beta-content upper tolerance limit with 95% coverage and 95%
# confidence using chrysene data and assuming a lognormal
# distribution.

# A beta-content upper tolerance limit with 95% coverage and 95%
# confidence is equivalent to the 95% upper confidence limit for the
# 95th percentile.

# Here we will construct a 95% upper confidence limit for the 95th

```



```

# percentile based on the log-transformed data, then exponentiate the
# result to get the confidence limit on the original scale. Note that
# it is easier to just use the function eqlnorm with the original data
# to achieve the same result.

attach(EPA.09.Ex.17.3.chrysene.df)
log.Chrysene <- log(Chrysene.ppb[Well.type == "Background"])
eqnorm(log.Chrysene, p = 0.95, ci = TRUE, ci.type = "upper")

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Normal
#
#Estimated Parameter(s):      mean = 2.5085773
#                              sd   = 0.6279479
#
#Estimation Method:           mvue
#
#Estimated Quantile(s):       95'th %ile = 3.54146
#
#Quantile Estimation Method:   qmle
#
#Data:                         log.Chrysene
#
#Sample Size:                  8
#
#Confidence Interval for:      95'th %ile
#
#Confidence Interval Method:   Exact
#
#Confidence Interval Type:     upper
#
#Confidence Level:             95%
#
#Confidence Interval:          LCL =      -Inf
#                              UCL = 4.510032

exp(4.510032)
#[1] 90.92473

#-----
# Clean up

rm(log.Chrysene)
detach("EPA.09.Ex.17.3.chrysene.df")

```

Description

Estimate quantiles of a distribution, and optionally create confidence intervals for them, without making any assumptions about the form of the distribution.

Usage

```
eqnpar(x, p = 0.5, ci = FALSE, lcl.rank = NULL, ucl.rank = NULL,
       lb = -Inf, ub = Inf, ci.type = "two-sided",
       ci.method = "exact", approx.conf.level = 0.95, digits = 0)
```

Arguments

<code>x</code>	a numeric vector of observations. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
<code>p</code>	numeric vector of probabilities for which quantiles will be estimated. All values of <code>p</code> must be between 0 and 1. When <code>ci=TRUE</code> , <code>p</code> must be a scalar. The default value is <code>p=0.5</code> .
<code>ci</code>	logical scalar indicating whether to compute a confidence interval for the quantile. The default value is <code>ci=FALSE</code> .
<code>lcl.rank</code> , <code>ucl.rank</code>	positive integers indicating the ranks of the order statistics that are used for the lower and upper bounds of the confidence interval for the specified quantile. Both arguments must be integers between 1 and the number of non-missing values in <code>x</code> , and <code>lcl.rank</code> must be strictly less than <code>ucl.rank</code> . Setting values for <code>lcl.rank</code> and/or <code>ucl.rank</code> allows the user to bypass the automatic selection of order statistics. By default the value of these arguments is <code>NULL</code> , in which case order statistics are chosen based on the value of <code>ci.type</code> and <code>ci.method</code> . If only <code>lcl.rank</code> is supplied, a lower confidence interval is constructed. If only <code>ucl.rank</code> is supplied, an upper confidence interval is constructed. If both <code>lcl.rank</code> and <code>ucl.rank</code> are supplied, a two-sided confidence interval is constructed. These arguments are ignored if <code>ci=FALSE</code> .
<code>lb</code> , <code>ub</code>	scalars indicating lower and upper bounds on the distribution. By default, <code>lb=-Inf</code> and <code>ub=Inf</code> . If you are constructing a confidence interval for a quantile from a distribution that you know has a lower bound other than <code>-Inf</code> (e.g., 0), set <code>lb</code> to this value. Similarly, if you know the distribution has an upper bound other than <code>Inf</code> , set <code>ub</code> to this value. These arguments are ignored if <code>ci=FALSE</code> .
<code>ci.type</code>	character string indicating what kind of confidence interval to compute. The possible values are "two-sided" (the default), "lower", and "upper". This argument is ignored if <code>ci=FALSE</code> , or <code>lcl.rank</code> and/or <code>ucl.rank</code> are supplied.
<code>ci.method</code>	character string indicating the method to use to construct the confidence interval. The possible values are "exact" (the default) and "normal.approx". See the DETAILS section for more information on these methods. This argument is ignored if <code>ci=FALSE</code> , or <code>lcl.rank</code> and/or <code>ucl.rank</code> are supplied.
<code>approx.conf.level</code>	a scalar between 0 and 1 indicating the desired confidence level of the confidence interval. The default value is 0.95. The true confidence level usually will not be exactly equal to <code>approx.conf.level</code> (see DETAILS). This argument is ignored if <code>ci=FALSE</code> , or <code>lcl.rank</code> and/or <code>ucl.rank</code> are supplied.
<code>digits</code>	an integer indicating the number of decimal places to round to when printing out the value of <code>100*p</code> . The default value is <code>digits=0</code> .

Details

If `x` contains any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

Estimation

The function eqnpar calls the R function `quantile` to estimate quantiles.

Confidence Intervals

Let x_1, x_2, \dots, x_n denote a sample of n independent and identically distributed random variables from some arbitrary distribution. Furthermore, let $x_{(i)}$ denote the i 'th order statistic for these n random variables. That is,

$$x_{(1)} \leq x_{(2)} \leq \dots \leq x_{(n)} \quad (1)$$

Finally, let x_p denote the p 'th quantile of the distribution, that is:

$$Pr(X < x_p) \leq p \quad (2)$$

$$Pr(X \leq x_p) \geq p \quad (3)$$

It can be shown (e.g., Conover, 1980, pp. 114-116) that for the i 'th order statistic:

$$Pr[x_p < x_{(i)}] = F_{B(n,p)}[i-1]; \quad i = 1, 2, \dots, n \quad (4)$$

where $F_{B(n,p)}[y]$ denotes the cumulative distribution function of a [binomial random variable](#) with parameters `size=n` and `prob=p` evaluated at y . This fact is used to construct exact confidence intervals for quantiles (see below).

Two-Sided Confidence Interval (`ci.type="two-sided"`)

A two-sided nonparametric confidence interval for the p 'th quantile is constructed as:

$$[x_{(r)}, x_{(s)}] \quad (5)$$

where

$$1 \leq r \leq (n-1) \quad (6)$$

$$2 \leq s \leq n \quad (7)$$

$$r < s \quad (8)$$

Note that the argument `lcl.rank` corresponds to r , and the argument `ucl.rank` corresponds to s .

This confidence interval has an associated confidence level that is at least as large as:

$$F_{B(n,p)}[s-1] - F_{B(n,p)}[r-1] \quad (9)$$

for a discrete distribution and exactly equal to this value for a continuous distribution. This is because:

$$\begin{aligned} & Pr[x_{(r)} \leq x_p \leq x_{(s)}] \\ &= Pr[x_p \leq x_{(s)}] - Pr[x_p < x_{(r)}] \\ &= Pr[x_p < x_{(s)}] + Pr[x_p = x_{(s)}] - Pr[x_p < x_{(r)}] \\ &\geq Pr[x_p < x_{(s)}] - Pr[x_p < x_{(r)}] \\ &= F_{B(n,p)}[s-1] - F_{B(n,p)}[r-1] \quad (10) \end{aligned}$$

Exact Method (`ci.method="exact"`)

When `lcl.rank` (r) and `ucl.rank` (s) are not supplied by the user, and `ci.method="exact"`, r and

s are chosen such that r is the smallest integer satisfying equation (11) below, and s is initially set to the largest integer satisfying equation (12) below:

$$F_{B(n,p)}[r-1] \geq \frac{\alpha}{2} \quad (11)$$

$$F_{B(n,p)}[s-1] \leq \frac{1-\alpha}{2} \quad (12)$$

where $\alpha = 1 - \text{approx.conf.level}$. Let r^* and s^* denote the values of r and s chosen in this manner. If

$$F_{B(n,p)}[s^*] - F_{B(n,p)}[r^* - 1] \leq 1 - \alpha \quad (13)$$

then s is set to $s^* + 1$.

Approximate Method (ci.method="approx")

When lcl.rank (r) and ucl.rank (s) are not supplied by the user and ci.method="normal.approx", r and s are chosen such that

$$r = np - h \quad (14)$$

$$s = np + h \quad (15)$$

$$h = t_{n-1, 1-\alpha/2} \sqrt{np(1-p)} \quad (16)$$

where $t_{\nu,q}$ denotes the q 'th quantile of [Student's t-distribution](#) with ν degrees of freedom, and then r and s are rounded to the nearest integer.

One-Sided Lower Confidence Interval (ci.type="lower")

A one-sided lower nonparametric confidence interval for the p 'th quantile is constructed as:

$$[x_{(r)}, ub] \quad (17)$$

where ub denotes the value of the ub argument (the user-supplied upper bound).

Exact Method (ci.method="exact")

When lcl.rank (r) is not supplied by the user, and ci.method="exact", r is chosen such that it is the smallest integer satisfying the following equation:

$$F_{B(n,p)}[r-1] \geq \alpha \quad (18)$$

where $\alpha = 1 - \text{approx.conf.level}$.

Approximate Method (ci.method="approx")

When lcl.rank (r) is not supplied by the user and ci.method="normal.approx", r is chosen such that

$$r = np - t_{n-1, 1-\alpha} \sqrt{np(1-p)} \quad (19)$$

and then r is rounded to the nearest integer.

One-Sided Upper Confidence Interval (ci.type="upper")

A one-sided upper nonparametric confidence interval for the p 'th quantile is constructed as:

$$[lb, x_{(s)}] \quad (20)$$

where lb denotes the value of the lb argument (the user-supplied lower bound).

Exact Method (ci.method="exact")

When ucl.rank (s) is not supplied by the user, and ci.method="exact", s is chosen such that it is the largest integer satisfying the following equation:

$$F_{B(n,p)}[s-1] \leq 1-\alpha \quad (21)$$

where $\alpha = 1 - \text{approx.conf.level}$.

Approximate Method (ci.method="approx")

When ucl.rank (s) is not supplied by the user and ci.method="normal.approx", s is chosen such that

$$s = np + t_{n-1, 1-\alpha} \sqrt{np(1-p)} \quad (22)$$

and then s is rounded to the nearest integer.

Note on Value of Confidence Level

Because of the discrete nature of order statistics, no matter what the underlying distribution, the value of the confidence level associated with the nonparametric confidence interval for the p 'th quantile will usually not be precisely equal to approx.conf.level.

Value

a list of class "estimate" containing the estimated quantile(s) and other information. See [estimate.object](#) for details.

Note

Percentiles are sometimes used in environmental standards and regulations. For example, Berthouex and Brown (2002, p.71) note that England has water quality limits based on the 90th and 95th percentiles of monitoring data not exceeding specified levels. They also note that the U.S. EPA has specifications for air quality monitoring, aquatic standards on toxic chemicals, and maximum daily limits for industrial effluents that are all based on percentiles. Given the importance of these quantities, it is essential to characterize the amount of uncertainty associated with the estimates of these quantities. This is done with confidence intervals.

It can be shown (e.g., Conover, 1980, pp.119-121) that an upper confidence interval for the p 'th quantile with confidence level $100(1-\alpha)\%$ is equivalent to an upper β -content tolerance interval with coverage $100p\%$ and confidence level $100(1-\alpha)\%$. Also, a lower confidence interval for the p 'th quantile with confidence level $100(1-\alpha)\%$ is equivalent to a lower β -content tolerance interval with coverage $100(1-p)\%$ and confidence level $100(1-\alpha)\%$. See the help file for [tolIntNpar](#) for more information on nonparametric tolerance intervals.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Berthouex, P.M., and L.C. Brown. (2002). *Statistics for Environmental Engineers*. Lewis Publishers, Boca Raton.
- Conover, W.J. (1980). *Practical Nonparametric Statistics*. Second Edition. John Wiley and Sons, New York.
- Gilbert, R.O. (1987). *Statistical Methods for Environmental Pollution Monitoring*. Van Nostrand Reinhold, New York, NY, pp.132-136.

Helsel, D.R., and R.M. Hirsch. (1992). *Statistical Methods in Water Resources Research*. Elsevier, New York, NY, pp.88-90.

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.

Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ.

See Also

[quantile](#), [tolIntNpar](#), [Estimating Distribution Quantiles](#), [Tolerance Intervals](#), [estimate.object](#).

Examples

```
# Generate 20 observations from a cauchy distribution with parameters
# location=0, scale=1. The true 75th percentile of this distribution is 1.
# Use eqnpar to estimate the 75th percentile and construct a 90% confidence interval.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
dat <- rcauchy(20, location = 0, scale = 1)
eqnpar(dat, p = 0.75, ci = TRUE, approx.conf.level = 0.9)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          None
#
#Estimated Quantile(s):        75'th %ile = 1.524903
#
#Quantile Estimation Method:    Nonparametric
#
#Data:                          dat
#
#Sample Size:                  20
#
#Confidence Interval for:       75'th %ile
#
#Confidence Interval Method:    exact
#
#Confidence Interval Type:      two-sided
#
#Confidence Level:              87.38755%
#
#Confidence Limit Rank(s):      13 19
#
#Confidence Interval:           LCL = 1.018038
#                               UCL = 2.215660
#
#-----

# In the above example, the true confidence level is 87% instead of 90%.
# Let's try to construct a confidence interval with a confidence level that is
# at least 90% by supplying our own indices for the order statistics to use for
# the confidence limits. In the above example, the 13'th and 19'th order statistics
# are used to construct the confidence interval. Let's try the 12'th and 19'th:
```

```
eqnpar(dat, p = 0.75, ci = TRUE, lcl.rank = 12, ucl.rank = 19)
```

```
#Results of Distribution Parameter Estimation
```

```
#-----
```

```
#
```

```
#Assumed Distribution:      None
```

```
#
```

```
#Estimated Quantile(s):    75'th %ile = 1.524903
```

```
#
```

```
#Quantile Estimation Method: Nonparametric
```

```
#
```

```
#Data:                     dat
```

```
#
```

```
#Sample Size:              20
```

```
#
```

```
#Confidence Interval for:  75'th %ile
```

```
#
```

```
#Confidence Interval Method: exact
```

```
#
```

```
#Confidence Interval Type: two-sided
```

```
#
```

```
#Confidence Level:         93.47622%
```

```
#
```

```
#Confidence Limit Rank(s): 12 19
```

```
#
```

```
#Confidence Interval:      LCL = 0.7494692
```

```
#
```

```
UCL = 2.215660
```

```
#-----
```

```
# Clean up
```

```
rm(dat)
```

```
#-----
```

```
# Modify Example 17-4 on page 17-21 of USEPA (2009). This example uses
# copper concentrations (ppb) from 3 background wells to set an upper
# limit for 2 compliance wells. Here we will compute an upper 95% confidence interval for
# the 95'th percentile of the distribution of copper concentrations in the background wells.
# The data are stored in EPA.92c.copper2.df. Note that even though these data are
# Type I left singly censored, it is still possible to compute an estimate of the
# 95'th percentile.
```

```
EPA.92c.copper2.df
```

```
#   Copper.orig Copper Censored Month Well Well.type
#1      <5      5.0      TRUE      1      1 Background
#2      <5      5.0      TRUE      2      1 Background
#3      7.5      7.5     FALSE      3      1 Background
#...
#9      9.2      9.2     FALSE      1      2 Background
#10     <5      5.0      TRUE      2      2 Background
#11     <5      5.0      TRUE      3      2 Background
#...
#17     <5      5.0      TRUE      1      3 Background
#18     5.4      5.4     FALSE      2      3 Background
#19     6.7      6.7     FALSE      3      3 Background
#...
```

```
#29      6.2    6.2    FALSE    5    4 Compliance
#30      <5    5.0     TRUE     6    4 Compliance
#31      7.8    7.8    FALSE     7    4 Compliance
#...
#38      <5    5.0     TRUE     6    5 Compliance
#39      5.6    5.6    FALSE     7    5 Compliance
#40      <5    5.0     TRUE     8    5 Compliance
```

```
with(EPA.92c.copper2.df,
     eqnpar(Copper[Well.type=="Background"], p = 0.95, ci = TRUE, lb = 0,
           ci.type = "upper", approx.conf.level = 0.95))
```

```
#Results of Distribution Parameter Estimation
```

```
#-----
```

```
#
#Assumed Distribution:          None
#
#Estimated Quantile(s):        95'th %ile = 7.925
#
#Quantile Estimation Method:   Nonparametric
#
#Data:                          Copper[Well.type == "Background"]
#
#Sample Size:                   24
#
#Confidence Interval for:       95'th %ile
#
#Confidence Interval Method:    exact
#
#Confidence Interval Type:      upper
#
#Confidence Level:              70.8011%
#
#Confidence Limit Rank(s):      24
#
#Confidence Interval:           LCL = 0.0
#                               UCL = 9.2
```

```
#-----
```

```
# For the above example, the true confidence level is 71% instead of 95%.
# This is a function of the small sample size. In fact, as Example 17-4 shows, the
# largest quantile for which you can construct a nonparametric confidence interval that
# will have associated confidence level of 95% is the 88'th percentile:
```

```
with(EPA.92c.copper2.df,
     eqnpar(Copper[Well.type=="Background"], p = 0.88, ci = TRUE,
           ucl.rank = 24, lb = 0, ci.type = "upper", approx.conf.level = 0.95))
```

```
#Results of Distribution Parameter Estimation
```

```
#-----
```

```
#
#Assumed Distribution:          None
#
#Estimated Quantile(s):        88'th %ile = 6.892
#
```



```

#Quantile Estimation Method:      Nonparametric
#
#Data:                            Copper[Well.type == "Background"]
#
#Sample Size:                     24
#
#Confidence Interval for:         88'th %ile
#
#Confidence Interval Method:      exact
#
#Confidence Interval Type:        upper
#
#Confidence Level:                95.3486%
#
#Confidence Limit Rank(s):        24
#
#Confidence Interval:             LCL = 0.0
#                                 UCL = 9.2

```

eqpareto

Estimate Quantiles of a Pareto Distribution

Description

Estimate quantiles of a [Pareto distribution](#).

Usage

```
eqpareto(x, p = 0.5, method = "mle", plot.pos.con = 0.375, digits = 0)
```

Arguments

<code>x</code>	a numeric vector of observations, or an object resulting from a call to an estimating function that assumes a Pareto distribution (e.g., epareto). If <code>x</code> is a numeric vector, missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
<code>p</code>	numeric vector of probabilities for which quantiles will be estimated. All values of <code>p</code> must be between 0 and 1. The default value is <code>p=0.5</code> .
<code>method</code>	character string specifying the method of estimating the distribution parameters. Possible values are "mle" (maximum likelihood; the default), and "lse" (least-squares). See the DETAILS section of the help file for epareto for more information on these estimation methods.
<code>plot.pos.con</code>	numeric scalar between 0 and 1 containing the value of the plotting position constant used to construct the values of the empirical cdf. The default value is <code>plot.pos.con=0.375</code> . This argument is used only when <code>method="lse"</code> .
<code>digits</code>	an integer indicating the number of decimal places to round to when printing out the value of <code>100*p</code> . The default value is <code>digits=0</code> .

Details

The function `eqpareto` returns estimated quantiles as well as estimates of the location and scale parameters.

Quantiles are estimated by 1) estimating the location and scale parameters by calling `epareto`, and then 2) calling the function `qpareto` and using the estimated values for location and scale.

Value

If `x` is a numeric vector, `eqpareto` returns a list of class "estimate" containing the estimated quantile(s) and other information. See `estimate.object` for details.

If `x` is the result of calling an estimation function, `eqpareto` returns a list whose class is the same as `x`. The list contains the same components as `x`, as well as components called `quantiles` and `quantile.method`.

Note

The Pareto distribution is named after Vilfredo Pareto (1848-1923), a professor of economics. It is derived from Pareto's law, which states that the number of persons N having income $\geq x$ is given by:

$$N = Ax^{-\theta}$$

where θ denotes Pareto's constant and is the shape parameter for the probability distribution.

The Pareto distribution takes values on the positive real line. All values must be larger than the "location" parameter η , which is really a threshold parameter. There are three kinds of Pareto distributions. The one described here is the Pareto distribution of the first kind. Stable Pareto distributions have $0 < \theta < 2$. Note that the r 'th moment only exists if $r < \theta$.

The Pareto distribution is related to the [exponential distribution](#) and [logistic distribution](#) as follows. Let X denote a Pareto random variable with `location`= η and `shape`= θ . Then $\log(X/\eta)$ has an exponential distribution with parameter `rate`= θ , and $-\log\{[(X/\eta)^\theta] - 1\}$ has a logistic distribution with parameters `location`=0 and `scale`=1.

The Pareto distribution has a very long right-hand tail. It is often applied in the study of socioeconomic data, including the distribution of income, firm size, population, and stock price fluctuations.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.

Johnson, N. L., S. Kotz, and N. Balakrishnan. (1994). *Continuous Univariate Distributions, Volume 1*. Second Edition. John Wiley and Sons, New York.

See Also

[epareto](#), [Pareto](#), [estimate.object](#).

Examples

```
# Generate 30 observations from a Pareto distribution with
# parameters location=1 and shape=1 then estimate the parameters
# and the 90'th percentile.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
dat <- rpareto(30, location = 1, shape = 1)
eqpareto(dat, p = 0.9)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:      Pareto
#
#Estimated Parameter(s):   location = 1.009046
#                           shape    = 1.079850
#
#Estimation Method:        mle
#
#Estimated Quantile(s):    90'th %ile = 8.510708
#
#Quantile Estimation Method: Quantile(s) Based on
#                           mle Estimators
#
#Data:                     dat
#
#Sample Size:              30

#-----

# Clean up
#-----
rm(dat)
```

eqpois

Estimate Quantiles of a Poisson Distribution

Description

Estimate quantiles of an [Poisson distribution](#), and optionally construct a confidence interval for a quantile.

Usage

```
eqpois(x, p = 0.5, method = "mle/mme/mvue", ci = FALSE, ci.method = "exact",
       ci.type = "two-sided", conf.level = 0.95, digits = 0)
```

Arguments

x a numeric vector of observations, or an object resulting from a call to an estimating function that assumes an Poisson distribution (e.g., [epois](#)). If ci=TRUE then x must be a numeric vector of observations. If x is a numeric vector, missing

	(NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
p	numeric vector of probabilities for which quantiles will be estimated. All values of p must be between 0 and 1. When ci=TRUE, p must be a scalar. The default value is p=0.5.
method	character string specifying the method to use to estimate the mean. Currently the only possible value is "mle/mme/mvue" (maximum likelihood/method of moments/minimum variance unbiased; the default). See the DETAILS section of the help file for eqpois for more information.
ci	logical scalar indicating whether to compute a confidence interval for the specified quantile. The default value is ci=FALSE.
ci.method	character string indicating what method to use to construct the confidence interval for the quantile. The only possible value is "exact" (exact method; the default). See the DETAILS section for more information.
ci.type	character string indicating what kind of confidence interval to compute. The possible values are "two-sided" (the default), "lower", and "upper". This argument is ignored if ci=FALSE.
conf.level	a scalar between 0 and 1 indicating the confidence level of the confidence interval. The default value is conf.level=0.95. This argument is ignored if ci=FALSE.
digits	an integer indicating the number of decimal places to round to when printing out the value of 100*p. The default value is digits=0.

Details

The function `eqpois` returns estimated quantiles as well as the estimate of the mean parameter.

Estimation

Let X denote a [Poisson random variable](#) with parameter λ . Let $x_{p|\lambda}$ denote the p 'th quantile of the distribution. That is,

$$Pr(X < x_{p|\lambda}) \leq p \leq Pr(X \leq x_{p|\lambda}) \quad (1)$$

Note that due to the discrete nature of the Poisson distribution, there will be several values of p associated with one value of X . For example, for $\lambda = 2$, the value 1 is the p 'th quantile for any value of p between 0.14 and 0.406.

Let \underline{x} denote a vector of n observations from a Poisson distribution with parameter λ . The p 'th quantile is estimated as the p 'th quantile from a Poisson distribution assuming the true value of λ is equal to the estimated value of λ . That is:

$$\hat{x}_{p|\lambda} = x_{p|\lambda=\hat{\lambda}} \quad (2)$$

where

$$\hat{\lambda} = \bar{x} = \frac{1}{n} \sum_{i=1}^n x_i \quad (3)$$

Because the estimator in equation (3) is the maximum likelihood estimator of λ (see the help file for [eqpois](#)), the estimated quantile is the maximum likelihood estimator.

Quantiles are estimated by 1) estimating the mean parameter by calling [eqpois](#), and then 2) calling the function [qpois](#) and using the estimated value for the mean parameter.

Confidence Intervals

It can be shown (e.g., Conover, 1980, pp.119-121) that an upper confidence interval for the p 'th quantile with confidence level $100(1 - \alpha)\%$ is equivalent to an upper β -content tolerance interval with coverage $100p\%$ and confidence level $100(1 - \alpha)\%$. Also, a lower confidence interval for the p 'th quantile with confidence level $100(1 - \alpha)\%$ is equivalent to a lower β -content tolerance interval with coverage $100(1 - p)\%$ and confidence level $100(1 - \alpha)\%$.

Thus, based on the theory of tolerance intervals for a Poisson distribution (see [tolIntPois](#)), if `ci.type="upper"`, a one-sided upper $100(1 - \alpha)\%$ confidence interval for the p 'th quantile is constructed as:

$$[0, x_{p|\lambda=UCL}] \quad (4)$$

where UCL denotes the upper $100(1 - \alpha)\%$ confidence limit for λ (see the help file for [epois](#) for information on how UCL is computed).

Similarly, if `ci.type="lower"`, a one-sided lower $100(1 - \alpha)\%$ confidence interval for the p 'th quantile is constructed as:

$$[x_{p|\lambda=LCL}, \infty] \quad (5)$$

where LCL denotes the lower $100(1 - \alpha)\%$ confidence limit for λ (see the help file for [epois](#) for information on how LCL is computed).

Finally, if `ci.type="two-sided"`, a two-sided $100(1 - \alpha)\%$ confidence interval for the p 'th quantile is constructed as:

$$[x_{p|\lambda=LCL}, x_{p|\lambda=UCL}] \quad (6)$$

where LCL and UCL denote the two-sided lower and upper $100(1 - \alpha)\%$ confidence limits for λ (see the help file for [epois](#) for information on how LCL and UCL are computed).

Value

If `x` is a numeric vector, `eqpois` returns a list of class "estimate" containing the estimated quantile(s) and other information. See [estimate.object](#) for details.

If `x` is the result of calling an estimation function, `eqpois` returns a list whose class is the same as `x`. The list contains the same components as `x`, as well as components called `quantiles` and `quantile.method`.

Note

Percentiles are sometimes used in environmental standards and regulations. For example, Berthouex and Brown (2002, p.71) state:

The U.S. EPA has specifications for air quality monitoring that are, in effect, percentile limitations. ... The U.S. EPA has provided guidance for setting aquatic standards on toxic chemicals that require estimating 99th percentiles and using this statistic to make important decisions about monitoring and compliance. They have also used the 99th percentile to establish maximum daily limits for industrial effluents (e.g., pulp and paper).

Given the importance of these quantities, it is essential to characterize the amount of uncertainty associated with the estimates of these quantities. This is done with confidence intervals.

The [Poisson distribution](#) is named after Poisson, who derived this distribution as the limiting distribution of the [binomial distribution](#) with parameters `size=N` and `prob=p`, where N tends to infinity, p tends to 0, and Np stays constant.

In this context, the Poisson distribution was used by Bortkiewicz (1898) to model the number of deaths (per annum) from kicks by horses in Prussian Army Corps. In this case, p , the probability of death from this cause, was small, but the number of soldiers exposed to this risk, N , was large.

The Poisson distribution has been applied in a variety of fields, including quality control (modeling number of defects produced in a process), ecology (number of organisms per unit area), and queueing theory. Gibbons (1987b) used the Poisson distribution to model the number of detected compounds per scan of the 32 volatile organic priority pollutants (VOC), and also to model the distribution of chemical concentration (in ppb).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Berthouex, P.M., and L.C. Brown. (2002). *Statistics for Environmental Engineers*. Second Edition. Lewis Publishers, Boca Raton, FL.
- Berthouex, P.M., and I. Hau. (1991). Difficulties Related to Using Extreme Percentiles for Water Quality Regulations. *Research Journal of the Water Pollution Control Federation* **63**(6), 873–879.
- Conover, W.J. (1980). *Practical Nonparametric Statistics*. Second Edition. John Wiley and Sons, New York, Chapter 3.
- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.
- Gibbons, R.D. (1987b). Statistical Models for the Analysis of Volatile Organic Compounds in Waste Disposal Sites. *Ground Water* **25**, 572–580.
- Gibbons, R.D., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*, Second Edition. John Wiley & Sons, Hoboken.
- Johnson, N. L., S. Kotz, and A. Kemp. (1992). *Univariate Discrete Distributions*. Second Edition. John Wiley and Sons, New York, Chapter 4.
- Pearson, E.S., and H.O. Hartley, eds. (1970). *Biometrika Tables for Statisticians, Volume 1*. Cambridge University Press, New York, p.81.
- Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ.

See Also

[epois](#), [Poisson](#), [codeestimate.object](#).

Examples

```
# Generate 20 observations from a Poisson distribution with parameter
# lambda=2. The true 90'th percentile of this distribution is 4 (actually,
# 4 is the p'th quantile for any value of p between 0.86 and 0.947).
# Here we will use eqpois to estimate the 90'th percentile and construct a
# two-sided 95% confidence interval for this percentile.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
dat <- rpois(20, lambda = 2)
eqpois(dat, p = 0.9, ci = TRUE)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Poisson
#
```

```
#Estimated Parameter(s):      lambda = 1.8
#
#Estimation Method:           mle/mme/mvue
#
#Estimated Quantile(s):       90'th %ile = 4
#
#Quantile Estimation Method:   mle
#
#Data:                         dat
#
#Sample Size:                  20
#
#Confidence Interval for:      90'th %ile
#
#Confidence Interval Method:   Exact
#
#Confidence Interval Type:     two-sided
#
#Confidence Level:             95%
#
#Confidence Interval:          LCL = 3
#                               UCL = 5

# Clean up
#-----
rm(dat)
```

equnif

Estimate Quantiles of a Uniform Distribution

Description

Estimate quantiles of a [uniform distribution](#).

Usage

```
equnif(x, p = 0.5, method = "mle", digits = 0)
```

Arguments

<code>x</code>	a numeric vector of observations, or an object resulting from a call to an estimating function that assumes a uniform distribution (e.g., equnif). If <code>x</code> is a numeric vector, missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
<code>p</code>	numeric vector of probabilities for which quantiles will be estimated. All values of <code>p</code> must be between 0 and 1. The default value is <code>p=0.5</code> .
<code>method</code>	character string specifying the method of estimating the distribution parameters. The possible values are "mle" (maximum likelihood; the default), "mme" (method of moments), and "mmue" (method of moments based on the unbiased estimator of variance). See the DETAILS section of the help file for equnif for more information on these estimation methods.
<code>digits</code>	an integer indicating the number of decimal places to round to when printing out the value of <code>100*p</code> . The default value is <code>digits=0</code> .

Details

The function `equnif` returns estimated quantiles as well as estimates of the location and scale parameters.

Quantiles are estimated by 1) estimating the location and scale parameters by calling `eunif`, and then 2) calling the function `qunif` and using the estimated values for location and scale.

Value

If `x` is a numeric vector, `equnif` returns a list of class "estimate" containing the estimated quantile(s) and other information. See `estimate.object` for details.

If `x` is the result of calling an estimation function, `equnif` returns a list whose class is the same as `x`. The list contains the same components as `x`, as well as components called `quantiles` and `quantile.method`.

Note

The [uniform distribution](#) (also called the rectangular distribution) with parameters `min` and `max` takes on values on the real line between `min` and `max` with equal probability. It has been used to represent the distribution of round-off errors in tabulated values. Another important application is that the distribution of the cumulative distribution function (cdf) of any kind of continuous random variable follows a uniform distribution with parameters `min=0` and `max=1`.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.

Johnson, N. L., S. Kotz, and N. Balakrishnan. (1995). *Continuous Univariate Distributions, Volume 2*. Second Edition. John Wiley and Sons, New York.

See Also

[eunif](#), [Uniform](#), [estimate.object](#).

Examples

```
# Generate 20 observations from a uniform distribution with parameters
# min=-2 and max=3, then estimate the parameters via maximum likelihood
# and estimate the 90th percentile.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
dat <- runif(20, min = -2, max = 3)
equnif(dat, p = 0.9)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Uniform
#
#Estimated Parameter(s):      min = -1.574529
```



```

#                               max = 2.837006
#
#Estimation Method:             mle
#
#Estimated Quantile(s):         90'th %ile = 2.395852
#
#Quantile Estimation Method:    Quantile(s) Based on
#                               mle Estimators
#
#Data:                          dat
#
#Sample Size:                   20

#-----
# Clean up

rm(dat)

```

eqweibull

Estimate Quantiles of a Weibull Distribution

Description

Estimate quantiles of a [Weibull distribution](#).

Usage

```
eqweibull(x, p = 0.5, method = "mle", digits = 0)
```

Arguments

x	a numeric vector of observations, or an object resulting from a call to an estimating function that assumes a Weibull distribution (e.g., eweibull). If x is a numeric vector, missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
p	numeric vector of probabilities for which quantiles will be estimated. All values of p must be between 0 and 1. The default value is p=0.5.
method	character string specifying the method of estimating the distribution parameters. Possible values are "mle" (maximum likelihood; the default), "mme" (methods of moments), and "mmue" (method of moments based on the unbiased estimator of variance). See the DETAILS section of the help file for eweibull for more information.
digits	an integer indicating the number of decimal places to round to when printing out the value of 100*p. The default value is digits=0.

Details

The function eqweibull returns estimated quantiles as well as estimates of the shape and scale parameters.

Quantiles are estimated by 1) estimating the shape and scale parameters by calling [eweibull](#), and then 2) calling the function [qweibull](#) and using the estimated values for shape and scale.

Value

If x is a numeric vector, `eqweibull` returns a list of class "estimate" containing the estimated quantile(s) and other information. See `estimate.object` for details.

If x is the result of calling an estimation function, `eqweibull` returns a list whose class is the same as x . The list contains the same components as x , as well as components called `quantiles` and `quantile.method`.

Note

The [Weibull distribution](#) is named after the Swedish physicist Waloddi Weibull, who used this distribution to model breaking strengths of materials. The Weibull distribution has been extensively applied in the fields of reliability and quality control.

The [exponential distribution](#) is a special case of the Weibull distribution: a Weibull random variable with parameters $\text{shape}=1$ and $\text{scale}=\beta$ is equivalent to an exponential random variable with parameter $\text{rate}=1/\beta$.

The Weibull distribution is related to the [Type I extreme value \(Gumbel\) distribution](#) as follows: if X is a random variable from a Weibull distribution with parameters $\text{shape}=\alpha$ and $\text{scale}=\beta$, then

$$Y = -\log(X) \quad (10)$$

is a random variable from an extreme value distribution with parameters $\text{location}=-\log(\beta)$ and $\text{scale}=1/\alpha$.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.

Johnson, N. L., S. Kotz, and N. Balakrishnan. (1994). *Continuous Univariate Distributions, Volume 1*. Second Edition. John Wiley and Sons, New York.

See Also

[eweibull](#), [Weibull](#), [Exponential](#), [EVD](#), [estimate.object](#).

Examples

```
# Generate 20 observations from a Weibull distribution with parameters
# shape=2 and scale=3, then estimate the parameters via maximum likelihood,
# and estimate the 90'th percentile.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
dat <- rweibull(20, shape = 2, scale = 3)
eqweibull(dat, p = 0.9)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Weibull
```

```

#
#Estimated Parameter(s):      shape = 2.673098
#                               scale = 3.047762
#
#Estimation Method:           mle
#
#Estimated Quantile(s):       90'th %ile = 4.163755
#
#Quantile Estimation Method:   Quantile(s) Based on
#                               mle Estimators
#
#Data:                         dat
#
#Sample Size:                  20

#-----

# Clean up
#-----
rm(dat)

```

eqzmlnorm

Estimate Quantiles of a Zero-Modified Lognormal (Delta) Distribution

Description

Estimate quantiles of a [zero-modified lognormal distribution](#) or a [zero-modified lognormal distribution \(alternative parameterization\)](#).

Usage

```
eqzmlnorm(x, p = 0.5, method = "mvue", digits = 0)
```

```
eqzmlnormAlt(x, p = 0.5, method = "mvue", digits = 0)
```

Arguments

- | | |
|--------|---|
| x | <p>a numeric vector of positive observations, or an object resulting from a call to an estimating function that assumes a zero-modified lognormal distribution.</p> <p>For eqzmlnorm, if x is an object, it must be the result of calling ezmlnorm, not ezmlnormAlt.</p> <p>For eqzmlnormAlt, if x is an object, it must be the result of calling ezmlnormAlt, not ezmlnorm.</p> <p>If x is a numeric vector, missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.</p> |
| p | <p>numeric vector of probabilities for which quantiles will be estimated. All values of p must be between 0 and 1. When ci=TRUE, p must be a scalar. The default value is p=0.5.</p> |
| method | <p>character string specifying the method of estimation. The only possible value is "mvue" (minimum variance unbiased; the default). See the DETAILS section of the help file for ezmlnorm for more information.</p> |

`digits` an integer indicating the number of decimal places to round to when printing out the value of $100 \times p$. The default value is `digits=0`.

Details

The functions `eqzmlnorm` and `eqzmlnormAlt` return estimated quantiles as well as estimates of the distribution parameters.

Quantiles are estimated by 1) estimating the distribution parameters by calling `ezmlnorm` or `ezmlnormAlt`, and then 2) calling the function `qzmlnorm` or `qzmlnormAlt` and using the estimated distribution parameters.

Value

If `x` is a numeric vector, `eqzmlnorm` and `eqzmlnormAlt` return a list of class "estimate" containing the estimated quantile(s) and other information. See `estimate.object` for details.

If `x` is the result of calling an estimation function, `eqzmlnorm` and `eqzmlnormAlt` return a list whose class is the same as `x`. The list contains the same components as `x`, as well as components called `quantiles` and `quantile.method`.

Note

The zero-modified lognormal (delta) distribution is sometimes used to model chemical concentrations for which some observations are reported as "Below Detection Limit" (the nondetects are assumed equal to 0). See, for example, Gilliom and Helsel (1986), Owen and DeRouen (1980), and Gibbons et al. (2009, Chapter 12). USEPA (2009, Chapter 15) recommends this strategy only in specific situations, and Helsel (2012, Chapter 1) strongly discourages this approach to dealing with non-detects.

A variation of the zero-modified lognormal (delta) distribution is the [zero-modified normal distribution](#), in which a normal distribution is mixed with a positive probability mass at 0.

One way to try to assess whether a zero-modified lognormal (delta), zero-modified normal, censored normal, or censored lognormal is the best model for the data is to construct both censored and detects-only probability plots (see [qqPlotCensored](#)).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Aitchison, J. (1955). On the Distribution of a Positive Random Variable Having a Discrete Probability Mass at the Origin. *Journal of the American Statistical Association* **50**, 901–908.
- Aitchison, J., and J.A.C. Brown (1957). *The Lognormal Distribution (with special reference to its uses in economics)*. Cambridge University Press, London. pp.94-99.
- Crow, E.L., and K. Shimizu. (1988). *Lognormal Distributions: Theory and Applications*. Marcel Dekker, New York, pp.47–51.
- Gibbons, RD., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*. Second Edition. John Wiley and Sons, Hoboken, NJ.
- Gilliom, R.J., and D.R. Helsel. (1986). Estimation of Distributional Parameters for Censored Trace Level Water Quality Data: 1. Estimation Techniques. *Water Resources Research* **22**, 135–146.
- Helsel, D.R. (2012). *Statistics for Censored Environmental Data Using Minitab and R*. Second Edition. John Wiley and Sons, Hoboken, NJ, Chapter 1.

Johnson, N. L., S. Kotz, and A.W. Kemp. (1992). *Univariate Discrete Distributions*. Second Edition. John Wiley and Sons, New York, p.312.

Owen, W., and T. DeRouen. (1980). Estimation of the Mean for Lognormal Data Containing Zeros and Left-Censored Values, with Applications to the Measurement of Worker Exposure to Air Contaminants. *Biometrics* **36**, 707–719.

USEPA (1992c). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities: Addendum to Interim Final Guidance*. Office of Solid Waste, Permits and State Programs Division, US Environmental Protection Agency, Washington, D.C.

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.

See Also

[ezmlnorm](#), [Zero-Modified Lognormal](#), [ezmlnormAlt](#), [Zero-Modified Lognormal \(Alternative Parameterization\)](#), [Zero-Modified Normal](#), [Lognormal](#).

Examples

```
# Generate 100 observations from a zero-modified lognormal (delta)
# distribution with mean=2, cv=1, and p.zero=0.5, then estimate the
# parameters and also the 80'th and 90'th percentiles.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
dat <- rzmlnormAlt(100, mean = 2, cv = 1, p.zero = 0.5)
eqzmlnormAlt(dat, p = c(0.8, 0.9))

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Zero-Modified Lognormal (Delta)
#
#Estimated Parameter(s):      mean          = 1.9604561
#                               cv             = 0.9169411
#                               p.zero         = 0.4500000
#                               mean.zmlnorm   = 1.0782508
#                               cv.zmlnorm     = 1.5307175
#
#Estimation Method:           mvue
#
#Estimated Quantile(s):       80'th %ile = 1.897451
#                               90'th %ile = 2.937976
#
#Quantile Estimation Method:   Quantile(s) Based on
#                               mvue Estimators
#
#Data:                         dat
#
#Sample Size:                  100

#-----

# Compare the estimated quatiles with the true quantiles
```

```
qzmlnormAlt(mean = 2, cv = 1, p.zero = 0.5, p = c(0.8, 0.9))
#[1] 1.746299 2.849858

#-----

# Clean up
rm(dat)
```

eqzmnorm

*Estimate Quantiles of a Zero-Modified Normal Distribution***Description**

Estimate quantiles of a [zero-modified normal distribution](#).

Usage

```
eqzmnorm(x, p = 0.5, method = "mvue", digits = 0)
```

Arguments

<code>x</code>	a numeric vector of observations, or an object resulting from a call to an estimating function that assumes a zero-modified normal distribution (e.g., ezmnorm). If <code>x</code> is a numeric vector, missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
<code>p</code>	numeric vector of probabilities for which quantiles will be estimated. All values of <code>p</code> must be between 0 and 1. The default value is <code>p=0.5</code> .
<code>method</code>	character string specifying the method of estimating the distribution parameters. Currently, the only possible value is "mvue" (minimum variance unbiased; the default). See the DETAILS section of the help file for ezmnorm for more information.
<code>digits</code>	an integer indicating the number of decimal places to round to when printing out the value of $100 \cdot p$. The default value is <code>digits=0</code> .

Details

The function `eqzmnorm` returns estimated quantiles as well as estimates of the distribution parameters.

Quantiles are estimated by 1) estimating the distribution parameters by calling [ezmnorm](#), and then 2) calling the function [qzmnorm](#) and using the estimated values for the distribution parameters.

Value

If `x` is a numeric vector, `eqzmnorm` returns a list of class "estimate" containing the estimated quantile(s) and other information. See [estimate.object](#) for details.

If `x` is the result of calling an estimation function, `eqzmnorm` returns a list whose class is the same as `x`. The list contains the same components as `x`, as well as components called `quantiles` and `quantile.method`.


```

#                                sd.zmnorm   = 2.465829
#
#Estimation Method:              mvue
#
#Estimated Quantile(s):          80'th %ile = 4.706298
#                                90'th %ile = 5.779250
#
#Quantile Estimation Method:     Quantile(s) Based on
#                                mvue Estimators
#
#Data:                           dat
#
#Sample Size:                    100

#-----

# Compare the estimated quantiles with the true quantiles

qzmnorm(mean = 4, sd = 2, p.zero = 0.5, p = c(0.8, 0.9))
#[1] 4.506694 5.683242

#-----

# Clean up
rm(dat)

```

estimate.object	S3 Class "estimate"
-----------------	---------------------

Description

Objects of S3 class "estimate" are returned by any of the **EnvStats** functions that estimate the parameters or quantiles of a probability distribution and optionally construct confidence, prediction, or tolerance intervals based on a sample of data assumed to come from that distribution.

Details

Objects of S3 class "estimate" are lists that contain information about the estimated distribution parameters, quantiles, and intervals. The names of the **EnvStats** functions that produce objects of class "estimate" have the following forms:

Form of Function Name	Result
<i>eabb</i>	Parameter Estimation
<i>eqabb</i>	Quantile Estimation
<i>predIntAbb</i>	Prediction Interval
<i>tolIntAbb</i>	Tolerance Interval

where *abb* denotes the abbreviation of the name of a probability distribution (see the help file for [Distribution.df](#) for a list of available probability distributions and their abbreviations), and *Abb* denotes the same thing as *abb* except the first letter of the abbreviation for the probability distribution is capitalized.

See the help files [Estimating Distribution Parameters](#) and [Estimating Distribution Quantiles](#) for lists of functions that estimate distribution parameters and quantiles. See the help files [Prediction Intervals](#) and [Tolerance Intervals](#) for lists of functions that create prediction and tolerance intervals.

For example:

- The function [enorm](#) returns an object of class "estimate" (a list) with information about the estimated mean and standard deviation of the assumed normal (Gaussian) distribution, as well as an optional confidence interval for the mean.
- The function [eqnorm](#) returns a list of class "estimate" with information about the estimated mean and standard deviation of the assumed normal distribution, the estimated user-specified quantile(s), and an optional confidence interval for a single quantile.
- The function [predIntNorm](#) returns a list of class "estimate" with information about the estimated mean and standard deviation of the assumed normal distribution, along with a prediction interval for a user-specified number of future observations (or means, medians, or sums).
- The function [tolIntNorm](#) returns a list of class "estimate" with information about the estimated mean and standard deviation of the assumed normal distribution, along with a tolerance interval.

Value

Required Components

The following components must be included in a legitimate list of class "estimate".

distribution	character string indicating the name of the assumed distribution (this equals "Nonparametric") for nonparametric procedures).
sample.size	numeric scalar indicating the sample size used to estimate the parameters or quantiles.
data.name	character string indicating the name of the data object used to compute the estimated parameters or quantiles.
bad.obs	numeric scalar indicating the number of missing (NA), undefined (NaN) and/or infinite (Inf, -Inf) values that were removed from the data object prior to performing the estimation.

Optional Components

The following components may optionally be included in a legitimate list of class "estimate".

parameters	(parametric estimation only) a numeric vector with a names attribute containing the names and values of the estimated distribution parameters.
n.param.est	(parametric estimation only) a scalar indicating the number of distribution parameters estimated.
method	(parametric estimation only) a character string indicating the method used to compute the estimated parameters.
quantiles	a numeric vector of estimated quantiles.
quantile.method	a character string indicating the method of quantile estimation.
interval	a list of class "intervalEstimate" containing information on a confidence, tolerance, or prediction interval.

All lists of class "intervalEstimate" contain the following component:

name	a character string indicating the kind of interval. Possible values are: "Confidence", "Tolerance", or "Prediction".
------	--

The number and names of the other components in a list of class "intervalEstimate" depends on the kind of interval it is. These components may include:

parameter	a character string indicating the parameter for which the interval is constructed (e.g., "mean", "95'th %ile", etc.).
limits	a numeric vector containing the lower and upper bounds of the interval.
type	the type of interval (i.e., "two-sided", "lower", or "upper").
method	the method used to construct the interval (e.g., "normal.approx").
conf.level	the confidence level associated with the interval.
sample.size	the sample size associated with the interval.
dof	(parametric intervals only) the degrees of freedom associated with the interval.
limit.ranks	(nonparametric intervals only) the rank(s) of the order statistic(s) used to construct the interval.
m	(prediction intervals only) the total number of future observations (n.mean=1, n.median=1, or n.sum=1) or averages (n.mean>1), medians (n.median>1), or sums (n.sum>1).
k	(prediction intervals only) the minimum number of future observations (n.mean=1, n.median=1, or n.sum=1), or averages (n.mean>1), medians (n.median>1) or sums (n.sum>1) out of the total m that the interval should contain.
n.mean	(prediction intervals only) the sample size associated with the future averages that should be contained in the interval.
n.median	(prediction intervals only) the sample size associated with the future medians that should be contained in the interval.
n.sum	(Poisson prediction intervals only) the sample size associated with the future sums that should be contained in the interval.
rule	(simultaneous prediction intervals only) the rule used to construct the simultaneous prediction interval.
delta.over.sigma	(simultaneous prediction intervals only) numeric scalar indicating the ratio Δ/σ . The quantity Δ (delta) denotes the difference between the mean of the population that was sampled to construct the prediction interval, and the mean of the population that will be sampled to produce the future observations. The quantity σ (sigma) denotes the population standard deviation for both populations.

Methods

Generic functions that have methods for objects of class "estimate" include:

[print](#).

Note

Since objects of class "estimate" are lists, you may extract their components with the \$ and [[operators.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

See Also

[Estimating Distribution Parameters](#), [Estimating Distribution Quantiles](#), [Distribution.df](#), [Prediction Intervals](#), [Tolerance Intervals](#), [estimateCensored.object](#).

Examples

```
# Create an object of class "estimate", then print it out.
# (Note: the call to set.seed simply allows you to reproduce
# this example.)

set.seed(250)

dat <- rnorm(20, mean = 3, sd = 2)

estimate.obj <- enorm(dat, ci = TRUE)

mode(estimate.obj)
#[1] "list"

class(estimate.obj)
#[1] "estimate"

names(estimate.obj)
#[1] "distribution" "sample.size" "parameters"
#[4] "n.param.est"  "method"         "data.name"
#[7] "bad.obs"      "interval"

names(estimate.obj$interval)
#[1] "name"          "parameter"      "limits"
#[4] "type"          "method"         "conf.level"
#[7] "sample.size"  "dof"

estimate.obj

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Normal
#
#Estimated Parameter(s):      mean = 2.861160
#                               sd   = 1.180226
#
#Estimation Method:           mvue
#
#Data:                         dat
#
#Sample Size:                  20
#
#Confidence Interval for:      mean
#
#Confidence Interval Method:   Exact
#
```

```

#Confidence Interval Type:      two-sided
#
#Confidence Level:             95%
#
#Confidence Interval:          LCL = 2.308798
#                               UCL = 3.413523

#-----

# Extract the confidence limits for the mean

estimate.obj$interval$limits
#      LCL      UCL
#2.308798 3.413523

#-----

# Clean up

rm(dat, estimate.obj)

```

```
estimateCensored.object
```

```
S3 Class "estimateCensored"
```

Description

Objects of S3 class "estimateCensored" are returned by any of the **EnvStats** functions that estimate the parameters or quantiles of a probability distribution and optionally construct confidence, prediction, or tolerance intervals based on a sample of *censored* data assumed to come from that distribution.

Details

Objects of S3 class "estimateCensored" are lists that contain information about the estimated distribution parameters, quantiles, and intervals. The names of the **EnvStats** functions that produce objects of class "estimateCensored" have the following forms:

Form of Function Name	Result
<i>eabb</i> Censored	Parameter Estimation
<i>eqabb</i> Censored	Quantile Estimation
<i>predIntAbb</i> Censored	Prediction Interval
<i>tolIntAbb</i> Censored	Tolerance Interval

where *abb* denotes the abbreviation of the name of a probability distribution (see the help file for [Distribution.df](#) for a list of available probability distributions and their abbreviations), and *Abb* denotes the same thing as *abb* except the first letter of the abbreviation for the probability distribution is capitalized.

See the sections **Estimating Distribution Parameters**, **Estimating Distribution Quantiles**, and **Prediction and Tolerance Intervals** in the help file [EnvStats Functions for Censored Data](#) for a list of functions that estimate distribution parameters, estimate distribution quantiles, create prediction

intervals, or create tolerance intervals using censored data.

For example:

- The function `enormCensored` returns an object of class "estimateCensored" (a list) with information about the estimated mean and standard deviation of the assumed normal (Gaussian) distribution, information about the amount and side of censoring, and also an optional confidence interval for the mean.
- The function `eqnormCensored` returns a list of class "estimateCensored" with information about the estimated mean and standard deviation of the assumed normal distribution, information about the amount and side of censoring, the estimated user-specified quantile(s), and an optional confidence interval for a single quantile.
- The function `tolIntNormCensored` returns a list of class "estimateCensored" with information about the estimated mean and standard deviation of the assumed normal distribution, information about the amount and side of censoring, and the computed tolerance interval.

Value

Required Components

The following components must be included in a legitimate list of class "estimateCensored".

distribution	character string indicating the name of the assumed distribution (this equals "Nonparametric") for nonparametric procedures).
sample.size	numeric scalar indicating the sample size used to estimate the parameters or quantiles.
censoring.side	character string indicating whether the data are left- or right-censored.
censoring.levels	numeric scalar or vector indicating the censoring level(s).
percent.censored	numeric scalar indicating the percent of non-missing observations that are censored.
data.name	character string indicating the name of the data object used to compute the estimateCensored parameters or quantiles.
censoring.name	character string indicating the name of the data object used to identify which values are censored.
bad.obs	numeric scalar indicating the number of missing (NA), undefined (NaN) and/or infinite (Inf, -Inf) values that were removed from the data object prior to performing the estimation.

Optional Components

The following components may optionally be included in a legitimate list of class "estimateCensored".

parameters	(parametric estimation only) a numeric vector with a names attribute containing the names and values of the estimateCensored distribution parameters.
n.param.est	(parametric estimation only) a scalar indicating the number of distribution parameters estimateCensored.
method	(parametric estimation only) a character string indicating the method used to compute the estimateCensored parameters.
quantiles	a numeric vector of estimateCensored quantiles.

quantile.method	a character string indicating the method of quantile estimation.
interval	a list of class "intervalEstimate" containing information on a confidence, tolerance, or prediction interval.

All lists of class "intervalEstimateCensored" contain the following component:

name	a character string indicating the kind of interval. Possible values are: "Confidence", "Tolerance", or "Prediction".
------	--

The number and names of the other components in a list of class "intervalEstimate" depends on the kind of interval it is. These components may include:

parameter	a character string indicating the parameter for which the interval is constructed (e.g., "mean", "95'th %ile", etc.).
limits	a numeric vector containing the lower and upper bounds of the interval.
type	the type of interval (i.e., "two-sided", "lower", or "upper").
method	the method used to construct the interval (e.g., "normal.approx").
conf.level	the confidence level associated with the interval.
sample.size	the sample size associated with the interval.
dof	(parametric intervals only) the degrees of freedom associated with the interval.
limit.ranks	(nonparametric intervals only) the rank(s) of the order statistic(s) used to construct the interval.
m	(prediction intervals only) the total number of future observations (n.mean=1, n.median=1, or n.sum=1) or averages (n.mean>1), medians (n.median>1), or sums (n.sum>1).
k	(prediction intervals only) the minimum number of future observations (n.mean=1, n.median=1, or n.sum=1), or averages (n.mean>1), medians (n.median>1) or sums (n.sum>1) out of the total m that the interval should contain.
n.mean	(prediction intervals only) the sample size associated with the future averages that should be contained in the interval.
n.median	(prediction intervals only) the sample size associated with the future medians that should be contained in the interval.
n.sum	(Poisson prediction intervals only) the sample size associated with the future sums that should be contained in the interval.
rule	(simultaneous prediction intervals only) the rule used to construct the simultaneous prediction interval.
delta.over.sigma	(simultaneous prediction intervals only) numeric scalar indicating the ratio Δ/σ . The quantity Δ (delta) denotes the difference between the mean of the population that was sampled to construct the prediction interval, and the mean of the population that will be sampled to produce the future observations. The quantity σ (sigma) denotes the population standard deviation for both populations.

Methods

Generic functions that have methods for objects of class "estimateCensored" include:

[print](#).

Note

Since objects of class "estimateCensored" are lists, you may extract their components with the \$ and [[operators.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

See Also

[EnvStats Functions for Censored Data](#), [Distribution.df](#), [estimate.object](#).

Examples

```
# Create an object of class "estimateCensored", then print it out.
# (Note: the call to set.seed simply allows you to reproduce
# this example.)

set.seed(250)

dat <- rnorm(20, mean = 100, sd = 20)
censored <- dat < 90
dat[censored] <- 90

estimateCensored.obj <- enormCensored(dat, censored, ci = TRUE)

mode(estimateCensored.obj)
#[1] "list"

class(estimateCensored.obj)
#[1] "estimateCensored"

names(estimateCensored.obj)
# [1] "distribution"      "sample.size"      "censoring.side"   "censoring.levels"
# [5] "percent.censored"  "parameters"       "n.param.est"     "method"
# [9] "data.name"        "censoring.name"   "bad.obs"         "interval"
#[13] "var.cov.params"

names(estimateCensored.obj$interval)
#[1] "name"      "parameter" "limits"    "type"      "method"    "conf.level"

estimateCensored.obj

#Results of Distribution Parameter Estimation
#Based on Type I Censored Data
#-----
#
#Assumed Distribution:      Normal
#
#Censoring Side:           left
#
#Censoring Level(s):      90
#
#Estimated Parameter(s):   mean = 96.52796
#                           sd   = 14.62275
```

```

#
#Estimation Method:          MLE
#
#Data:                       dat
#
#Censoring Variable:         censored
#
#Sample Size:                20
#
#Percent Censored:           25%
#
#Confidence Interval for:     mean
#
#Confidence Interval Method:  Profile Likelihood
#
#Confidence Interval Type:    two-sided
#
#Confidence Level:           95%
#
#Confidence Interval:         LCL = 88.82415
#                               UCL = 103.27604

#-----

# Extract the confidence limits for the mean

estimateCensored.obj$interval$limits
#      LCL      UCL
# 91.7801 103.7839

#-----

# Clean up

rm(dat, censored, estimateCensored.obj)

```

EulersConstant

Euler's Constant

Description

Explanation of Euler's Constant.

Details

Euler's Constant, here denoted ϵ , is a real-valued number that can be defined in several ways. Johnson et al. (1992, p. 5) use the definition:

$$\epsilon = \lim_{n \rightarrow \infty} [1 + \frac{1}{2} + \frac{1}{3} + \dots + \frac{1}{n} - \log(n)]$$

and note that it can also be expressed as

$$\epsilon = -\Psi(1)$$

where $\Psi()$ is the [digamma function](#) (Johnson et al., 1992, p.8).

The value of Euler's Constant, to 10 decimal places, is 0.5772156649.

The expression for the mean of a [Type I extreme value \(Gumbel\) distribution](#) involves Euler's constant; hence Euler's constant is used to compute the method of moments estimators for this distribution (see [eevd](#)).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Johnson, N. L., S. Kotz, and A.W. Kemp. (1992). *Univariate Discrete Distributions*. Second Edition. John Wiley and Sons, New York, pp.4-8.

Johnson, N. L., S. Kotz, and N. Balakrishnan. (1995). *Continuous Univariate Distributions, Volume 2*. Second Edition. John Wiley and Sons, New York.

See Also

[Extreme Value Distribution](#), [eevd](#).

eunif

Estimate Parameters of a Uniform Distribution

Description

Estimate the minimum and maximum parameters of a [uniform distribution](#).

Usage

```
eunif(x, method = "mle")
```

Arguments

x	numeric vector of observations. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
method	character string specifying the method of estimation. The possible values are "mle" (maximum likelihood; the default), "mme" (method of moments), and "mmue" (method of moments based on the unbiased estimator of variance). See the DETAILS section for more information on these estimation methods.

Details

If x contains any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

Let $\underline{x} = (x_1, x_2, \dots, x_n)$ be a vector of n observations from an [uniform distribution](#) with parameters $\min=a$ and $\max=b$. Also, let $x_{(i)}$ denote the i 'th order statistic.

Estimation

Maximum Likelihood Estimation (method="mle")

The maximum likelihood estimators (mle's) of a and b are given by (Johnson et al, 1995, p.286):

$$\hat{a}_{mle} = x_{(1)} \quad (1)$$

$$\hat{b}_{mle} = x_{(n)} \quad (2)$$

Method of Moments Estimation (method="mme")

The method of moments estimators (mme's) of a and b are given by (Forbes et al., 2011):

$$\hat{a}_{mme} = \bar{x} - \sqrt{3}s_m \quad (3)$$

$$\hat{b}_{mme} = \bar{x} + \sqrt{3}s_m \quad (4)$$

where

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i \quad (5)$$

$$s_m^2 = \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2 \quad (6)$$

Method of Moments Estimation Based on the Unbiased Estimator of Variance (method="mmue")

The method of moments estimators based on the unbiased estimator of variance are exactly the same as the method of moments estimators given in equations (3-6) above, except that the method of moments estimator of variance in equation (6) is replaced with the unbiased estimator of variance:

$$\hat{a}_{mmue} = \bar{x} - \sqrt{3}s \quad (7)$$

$$\hat{b}_{mmue} = \bar{x} + \sqrt{3}s \quad (8)$$

where

$$s^2 = \frac{1}{n-1} \sum_{i=1}^n (x_i - \bar{x})^2 \quad (9)$$

Value

a list of class "estimate" containing the estimated parameters and other information. See [estimate.object](#) for details.

Note

The [uniform distribution](#) (also called the rectangular distribution) with parameters `min` and `max` takes on values on the real line between `min` and `max` with equal probability. It has been used to represent the distribution of round-off errors in tabulated values. Another important application is that the distribution of the cumulative distribution function (cdf) of any kind of continuous random variable follows a uniform distribution with parameters `min=0` and `max=1`.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.

Johnson, N. L., S. Kotz, and N. Balakrishnan. (1995). *Continuous Univariate Distributions, Volume 2*. Second Edition. John Wiley and Sons, New York.

See Also

[Uniform](#), [estimate.object](#).

Examples

```
# Generate 20 observations from a uniform distribution with parameters
# min=-2 and max=3, then estimate the parameters via maximum likelihood.
# (Note: the call to set.seed simply allows you to reproduce this example.)
```

```
set.seed(250)
dat <- runif(20, min = -2, max = 3)
eunif(dat)
```

```
#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Uniform
#
#Estimated Parameter(s):      min = -1.574529
#                               max =  2.837006
#
#Estimation Method:           mle
#
#Data:                         dat
#
#Sample Size:                  20
```

```
#-----
```

```
# Compare the three methods of estimation:
```

```
eunif(dat, method = "mle")$parameters
#      min      max
#-1.574529  2.837006
```

```
eunif(dat, method = "mme")$parameters
#      min      max
#-1.988462  2.650737
```

```
eunif(dat, method = "mmue")$parameters
#      min      max
#-2.048721  2.710996
```

```
#-----
```

```
# Clean up
#-----
```

rm(dat)

EVD

The Extreme Value (Gumbel) Distribution

Description

Density, distribution function, quantile function, and random generation for the (largest) extreme value distribution.

Usage

```
devd(x, location = 0, scale = 1)
pevd(q, location = 0, scale = 1)
qevd(p, location = 0, scale = 1)
revd(n, location = 0, scale = 1)
```

Arguments

x	vector of quantiles.
q	vector of quantiles.
p	vector of probabilities between 0 and 1.
n	sample size. If length(n) is larger than 1, then length(n) random values are returned.
location	vector of location parameters.
scale	vector of positive scale parameters.

Details

Let X be an extreme value random variable with parameters $\text{location}=\eta$ and $\text{scale}=\theta$. The density function of X is given by:

$$f(x; \eta, \theta) = \frac{1}{\theta} e^{-(x-\eta)/\theta} \exp[-e^{-(x-\eta)/\theta}]$$

where $-\infty < x, \eta < \infty$ and $\theta > 0$.

The cumulative distribution function of X is given by:

$$F(x; \eta, \theta) = \exp[-e^{-(x-\eta)/\theta}]$$

The p^{th} quantile of X is given by:

$$x_p = \eta - \theta \log[-\log(p)]$$

The mode, mean, variance, skew, and kurtosis of X are given by:

$$\text{Mode}(X) = \eta$$

$$E(X) = \eta + \epsilon \theta$$

$$\text{Var}(X) = \theta^2 \pi^2 / 6$$

$$\text{Skew}(X) = \sqrt{\beta_1} = 1.139547$$

$$\text{Kurtosis}(X) = \beta_2 = 5.4$$

where ϵ denotes [Euler's constant](#), which is equivalent to `-digamma(1)`.

Value

density (devd), probability (pevd), quantile (qevd), or random sample (revd) for the extreme value distribution with location parameter(s) determined by location and scale parameter(s) determined by scale.

Note

There are three families of extreme value distributions. The one described here is the Type I, also called the Gumbel extreme value distribution or simply Gumbel distribution. The name “extreme value” comes from the fact that this distribution is the limiting distribution (as n approaches infinity) of the greatest value among n independent random variables each having the same continuous distribution.

The Gumbel extreme value distribution is related to the [exponential distribution](#) as follows. Let Y be an [exponential](#) random variable with parameter $\text{rate}=\lambda$. Then $X = \eta - \log(Y)$ has an extreme value distribution with parameters $\text{location}=\eta$ and $\text{scale}=1/\lambda$.

The distribution described above and used by devd, pevd, qevd, and revd is the *largest* extreme value distribution. The smallest extreme value distribution is the limiting distribution (as n approaches infinity) of the smallest value among n independent random variables each having the same continuous distribution. If X has a largest extreme value distribution with parameters $\text{location}=\eta$ and $\text{scale}=\theta$, then $Y = -X$ has a smallest extreme value distribution with parameters $\text{location}=-\eta$ and $\text{scale}=\theta$. The smallest extreme value distribution is related to the [Weibull distribution](#) as follows. Let Y be a [Weibull random variable](#) with parameters $\text{shape}=\beta$ and $\text{scale}=\alpha$. Then $X = \log(Y)$ has a smallest extreme value distribution with parameters $\text{location}=\log(\alpha)$ and $\text{scale}=1/\beta$.

The extreme value distribution has been used extensively to model the distribution of streamflow, flooding, rainfall, temperature, wind speed, and other meteorological variables, as well as material strength and life data.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.

Johnson, N. L., S. Kotz, and N. Balakrishnan. (1995). *Continuous Univariate Distributions, Volume 2*. Second Edition. John Wiley and Sons, New York.

See Also

[eevd](#), [GEVD](#), [Probability Distributions and Random Numbers](#).

Examples

```
# Density of an extreme value distribution with location=0, scale=1,
# evaluated at 0.5:

devd(.5)
#[1] 0.3307043

#-----
```

```
# The cdf of an extreme value distribution with location=1, scale=2,
# evaluated at 0.5:

pevd(.5, 1, 2)
#[1] 0.2769203

#-----

# The 25'th percentile of an extreme value distribution with
# location=-2, scale=0.5:

qevd(.25, -2, 0.5)
#[1] -2.163317

#-----

# Random sample of 4 observations from an extreme value distribution with
# location=5, scale=2.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(20)
revd(4, 5, 2)
#[1] 9.070406 7.669139 4.511481 5.903675
```

eweibull

Estimate Parameters of a Weibull Distribution

Description

Estimate the shape and scale parameters of a [Weibull distribution](#).

Usage

```
eweibull(x, method = "mle")
```

Arguments

x	numeric vector of observations. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
method	character string specifying the method of estimation. Possible values are "mle" (maximum likelihood; the default), "mme" (methods of moments), and "mmue" (method of moments based on the unbiased estimator of variance). See the DETAILS section for more information on these estimation methods.

Details

If x contains any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

Let $\underline{x} = (x_1, x_2, \dots, x_n)$ be a vector of n observations from an [Weibull distribution](#) with parameters $\text{shape}=\alpha$ and $\text{scale}=\beta$.

Estimation

Maximum Likelihood Estimation (method="mle")

The maximum likelihood estimators (mle's) of α and β are the solutions of the simultaneous equations (Forbes et al., 2011):

$$\hat{\alpha}_{mle} = \frac{n}{\{(1/\hat{\beta}_{mle})^{\hat{\alpha}_{mle}} \sum_{i=1}^n [x_i^{\hat{\alpha}_{mle}} \log(x_i)]\} - \sum_{i=1}^n \log(x_i)} \quad (1)$$

$$\hat{\beta}_{mle} = \left[\frac{1}{n} \sum_{i=1}^n x_i^{\hat{\alpha}_{mle}} \right]^{1/\hat{\alpha}_{mle}} \quad (2)$$

Method of Moments Estimation (method="mme")

The method of moments estimator (mme) of α is computed by solving the equation:

$$\frac{s}{\bar{x}} = \left\{ \frac{\Gamma[(\hat{\alpha}_{mme} + 2)/\hat{\alpha}_{mme}]}{\{\Gamma[(\hat{\alpha}_{mme} + 1)/\hat{\alpha}_{mme}]\}^2} - 1 \right\}^{1/2} \quad (3)$$

and the method of moments estimator (mme) of β is then computed as:

$$\hat{\beta}_{mme} = \frac{\bar{x}}{\Gamma[(\hat{\alpha}_{mme} + 1)/\hat{\alpha}_{mme}]} \quad (4)$$

where

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i \quad (5)$$

$$s_m^2 = \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2 \quad (6)$$

and $\Gamma()$ denotes the [gamma function](#).

Method of Moments Estimation Based on the Unbiased Estimator of Variance (method="mmue")

The method of moments estimators based on the unbiased estimator of variance are exactly the same as the method of moments estimators given in equations (3-6) above, except that the method of moments estimator of variance in equation (6) is replaced with the unbiased estimator of variance:

$$s^2 = \frac{1}{n-1} \sum_{i=1}^n (x_i - \bar{x})^2 \quad (7)$$

Value

a list of class "estimate" containing the estimated parameters and other information. See [estimate.object](#) for details.

Note

The [Weibull distribution](#) is named after the Swedish physicist Waloddi Weibull, who used this distribution to model breaking strengths of materials. The Weibull distribution has been extensively applied in the fields of reliability and quality control.

The [exponential distribution](#) is a special case of the Weibull distribution: a Weibull random variable with parameters shape=1 and scale= β is equivalent to an exponential random variable with parameter rate= $1/\beta$.

The Weibull distribution is related to the [Type I extreme value \(Gumbel\) distribution](#) as follows: if X is a random variable from a Weibull distribution with parameters $\text{shape}=\alpha$ and $\text{scale}=\beta$, then

$$Y = -\log(X) \quad (10)$$

is a random variable from an extreme value distribution with parameters $\text{location}=-\log(\beta)$ and $\text{scale}=1/\alpha$.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.

Johnson, N. L., S. Kotz, and N. Balakrishnan. (1994). *Continuous Univariate Distributions, Volume 1*. Second Edition. John Wiley and Sons, New York.

See Also

[Weibull](#), [Exponential](#), [EVD](#), [estimate.object](#).

Examples

```
# Generate 20 observations from a Weibull distribution with parameters
# shape=2 and scale=3, then estimate the parameters via maximum likelihood.
# (Note: the call to set.seed simply allows you to reproduce this example.)
```

```
set.seed(250)
dat <- rweibull(20, shape = 2, scale = 3)
eweibull(dat)
```

```
#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:      Weibull
#
#Estimated Parameter(s):   shape = 2.673098
#                           scale = 3.047762
#
#Estimation Method:        mle
#
#Data:                     dat
#
#Sample Size:              20
```

```
#-----
```

```
# Use the same data as in previous example, and compute the method of
# moments estimators based on the unbiased estimator of variance:
```

```
eweibull(dat, method = "mmue")
```

```
#Results of Distribution Parameter Estimation
#-----
```



```
#
#Assumed Distribution:      Weibull
#
#Estimated Parameter(s):  shape = 2.528377
#                          scale = 3.052507
#
#Estimation Method:       mmue
#
#Data:                    dat
#
#Sample Size:             20

#-----

# Clean up
#-----
rm(dat)
```

ezmlnorm	<i>Estimate Parameters of a Zero-Modified Lognormal (Delta) Distribution</i>
----------	--

Description

Estimate the parameters of a [zero-modified lognormal distribution](#) or a [zero-modified lognormal distribution \(alternative parameterization\)](#), and optionally construct a confidence interval for the mean.

Usage

```
ezmlnorm(x, method = "mvue", ci = FALSE, ci.type = "two-sided",
         ci.method = "normal.approx", conf.level = 0.95)

ezmlnormAlt(x, method = "mvue", ci = FALSE, ci.type = "two-sided",
            ci.method = "normal.approx", conf.level = 0.95)
```

Arguments

x	numeric vector of observations. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
method	character string specifying the method of estimation. The only possible value is "mvue" (minimum variance unbiased; the default). See the DETAILS section for more information on this estimation method.
ci	logical scalar indicating whether to compute a confidence interval for the mean. The default value is FALSE. If ci=TRUE and there are less than three non-missing observations in x, or if all observations are zeros, a warning will be issued and no confidence interval will be computed.
ci.type	character string indicating what kind of confidence interval to compute. The possible values are "two-sided" (the default), "lower", and "upper". This argument is ignored if ci=FALSE.

<code>ci.method</code>	character string indicating what method to use to construct the confidence interval for the mean. The only possible value is "normal.approx" (the default). See the DETAILS section for more information. This argument is ignored if <code>ci=FALSE</code> .
<code>conf.level</code>	a scalar between 0 and 1 indicating the confidence level of the confidence interval. The default value is <code>conf.level=0.95</code> . This argument is ignored if <code>ci=FALSE</code> .

Details

If \mathbf{x} contains any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

Let $\underline{x} = (x_1, x_2, \dots, x_n)$ be a vector of n observations from a [zero-modified lognormal distribution](#) with parameters $\text{meanlog}=\mu$, $\text{sdlog}=\sigma$, and $\text{p.zero}=p$. Alternatively, let $\underline{x} = (x_1, x_2, \dots, x_n)$ be a vector of n observations from a [zero-modified lognormal distribution \(alternative parameterization\)](#) with parameters $\text{mean}=\theta$, $\text{cv}=\tau$, and $\text{p.zero}=p$.

Let r denote the number of observations in \underline{x} that are equal to 0, and order the observations so that x_1, x_2, \dots, x_r denote the r zero observations and $x_{r+1}, x_{r+2}, \dots, x_n$ denote the $n - r$ non-zero observations.

Note that θ is *not* the mean of the zero-modified lognormal distribution; it is the mean of the lognormal part of the distribution. Similarly, τ is *not* the coefficient of variation of the zero-modified lognormal distribution; it is the coefficient of variation of the lognormal part of the distribution.

Let γ , δ , and ϕ denote the mean, standard deviation, and coefficient of variation of the overall zero-modified lognormal (delta) distribution. Let η denote the standard deviation of the lognormal part of the distribution, so that $\eta = \theta\tau$. Aitchison (1955) shows that:

$$\gamma = (1 - p)\theta \quad (1)$$

$$\delta^2 = (1 - p)\eta^2 + p(1 - p)\theta^2 \quad (2)$$

so that

$$\phi = \frac{\delta}{\gamma} = \frac{\sqrt{\tau^2 + p}}{\sqrt{1 - p}} \quad (3)$$

Estimation

Minimum Variance Unbiased Estimation (method="mvue")

Aitchison (1955) shows that the minimum variance unbiased estimators (mvue's) of γ and δ are:

$$\hat{\gamma}_{mvue} = \begin{cases} (1 - \frac{r}{n})e^{\bar{y}}g_{n-r-1}(\frac{s^2}{2}) & \text{if } r < n - 1, \\ x_n/n & \text{if } r = n - 1, \\ 0 & \text{if } r = n \end{cases} \quad (4)$$

$$\hat{\delta}_{mvue}^2 = \begin{cases} (1 - \frac{r}{n})e^{2\bar{y}}\{g_{n-r-1}(2s^2) - \frac{n-r-1}{n-1}g_{n-r-1}[\frac{(n-r-2)s^2}{n-r-1}]\} & \text{if } r < n - 1, \\ x_n^2/n & \text{if } r = n - 1, \\ 0 & \text{if } r = n \end{cases} \quad (5)$$

where

$$y_i = \log(x_i), \quad r = r + 1, r + 2, \dots, n \quad (6)$$

$$\bar{y} = \frac{1}{n-r} \sum_{i=r+1}^n y_i \quad (7)$$

$$s^2 = \frac{1}{n-r-1} \sum_{i=r+1}^n (y_i - \bar{y})^2 \quad (8)$$

$$g_m(z) = \sum_{i=0}^{\infty} \frac{m^i(m+2i)}{m(m+2) \cdots (m+2i)} \left(\frac{m}{m+1}\right)^i \left(\frac{z^i}{i!}\right) \quad (9)$$

Note that when $r = n-1$ or $r = n$, the estimator of γ is simply the sample mean for all observations (including zero values), and the estimator for δ^2 is simply the sample variance for all observations.

The expected value and asymptotic variance of the mvue of γ are (Aitchison and Brown, 1957, p.99; Owen and DeRouen, 1980):

$$E(\hat{\gamma}_{mvue}) = \gamma \quad (10)$$

$$AVar(\hat{\gamma}_{mvue}) = \frac{1}{n} \exp(2\mu + \sigma^2)(1-p) \left(p + \frac{2\sigma^2 + \sigma^4}{2}\right) \quad (11)$$

Confidence Intervals

Based on Normal Approximation (ci.method="normal.approx")

An approximate $(1 - \alpha)100\%$ confidence interval for γ is constructed based on the assumption that the estimator of γ is approximately normally distributed. Thus, an approximate two-sided $(1 - \alpha)100\%$ confidence interval for γ is constructed as:

$$[\hat{\gamma}_{mvue} - t_{n-2, 1-\alpha/2} \hat{\sigma}_{\hat{\gamma}}, \hat{\gamma}_{mvue} + t_{n-2, 1-\alpha/2} \hat{\sigma}_{\hat{\gamma}}] \quad (12)$$

where $t_{\nu, p}$ is the p 'th quantile of [Student's t-distribution](#) with ν degrees of freedom, and the quantity $\hat{\sigma}_{\hat{\gamma}}$ is the estimated standard deviation of the mvue of γ , and is computed by replacing the values of μ , σ , and p in equation (11) above with their estimated values and taking the square root.

Note that there must be at least 3 non-missing observations ($n \geq 3$) and at least one observation must be non-zero ($r \leq n-1$) in order to construct a confidence interval.

One-sided confidence intervals are computed in a similar fashion.

Value

a list of class "estimate" containing the estimated parameters and other information. See [estimate.object](#) for details.

For the function ezmlnorm, the component called parameters is a numeric vector with the following estimated parameters:

Parameter Name	Explanation
meanlog	mean of the log of the lognormal part of the distribution.
sdlog	standard deviation of the log of the lognormal part of the distribution.
p.zero	probability that an observation will be 0.
mean.zmlnorm	mean of the overall zero-modified lognormal (delta) distribution.
sd.zmlnorm	standard deviation of the overall zero-modified lognormal (delta) distribution.

For the function `ezm1normAlt`, the component called `parameters` is a numeric vector with the following estimated parameters:

Parameter Name	Explanation
<code>mean</code>	mean of the lognormal part of the distribution.
<code>cv</code>	coefficient of variation of the lognormal part of the distribution.
<code>p.zero</code>	probability that an observation will be 0.
<code>mean.zm1norm</code>	mean of the overall zero-modified lognormal (delta) distribution.
<code>sd.zm1norm</code>	standard deviation of the overall zero-modified lognormal (delta) distribution.

Note

The zero-modified lognormal (delta) distribution is sometimes used to model chemical concentrations for which some observations are reported as “Below Detection Limit” (the nondetects are assumed equal to 0). See, for example, Gilliom and Helsel (1986), Owen and DeRouen (1980), and Gibbons et al. (2009, Chapter 12). USEPA (2009, Chapter 15) recommends this strategy only in specific situations, and Helsel (2012, Chapter 1) strongly discourages this approach to dealing with non-detects.

A variation of the zero-modified lognormal (delta) distribution is the [zero-modified normal distribution](#), in which a normal distribution is mixed with a positive probability mass at 0.

One way to try to assess whether a zero-modified lognormal (delta), zero-modified normal, censored normal, or censored lognormal is the best model for the data is to construct both censored and detects-only probability plots (see [qqPlotCensored](#)).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Aitchison, J. (1955). On the Distribution of a Positive Random Variable Having a Discrete Probability Mass at the Origin. *Journal of the American Statistical Association* **50**, 901–908.
- Aitchison, J., and J.A.C. Brown (1957). *The Lognormal Distribution (with special reference to its uses in economics)*. Cambridge University Press, London. pp.94-99.
- Crow, E.L., and K. Shimizu. (1988). *Lognormal Distributions: Theory and Applications*. Marcel Dekker, New York, pp.47–51.
- Gibbons, RD., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*. Second Edition. John Wiley and Sons, Hoboken, NJ.
- Gilliom, R.J., and D.R. Helsel. (1986). Estimation of Distributional Parameters for Censored Trace Level Water Quality Data: 1. Estimation Techniques. *Water Resources Research* **22**, 135–146.
- Helsel, D.R. (2012). *Statistics for Censored Environmental Data Using Minitab and R*. Second Edition. John Wiley and Sons, Hoboken, NJ, Chapter 1.
- Johnson, N. L., S. Kotz, and A.W. Kemp. (1992). *Univariate Discrete Distributions*. Second Edition. John Wiley and Sons, New York, p.312.
- Owen, W., and T. DeRouen. (1980). Estimation of the Mean for Lognormal Data Containing Zeros and Left-Censored Values, with Applications to the Measurement of Worker Exposure to Air Contaminants. *Biometrics* **36**, 707–719.
- USEPA (1992c). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities: Addendum to Interim Final Guidance*. Office of Solid Waste, Permits and State Programs Division, US Environmental Protection Agency, Washington, D.C.

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.

See Also

[Zero-Modified Lognormal](#), [Zero-Modified Normal](#), [Lognormal](#).

Examples

```
# Generate 100 observations from a zero-modified lognormal (delta)
# distribution with mean=2, cv=1, and p.zero=0.5, then estimate the
# parameters. According to equations (1) and (3) above, the overall mean
# is mean.zmlnorm=1 and the overall cv is cv.zmlnorm=sqrt(3).
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
dat <- rzmnormAlt(100, mean = 2, cv = 1, p.zero = 0.5)
ezmnormAlt(dat, ci = TRUE)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Zero-Modified Lognormal (Delta)
#
#Estimated Parameter(s):      mean          = 1.9604561
#                               cv            = 0.9169411
#                               p.zero         = 0.4500000
#                               mean.zmlnorm   = 1.0782508
#                               cv.zmlnorm     = 1.5307175
#
#Estimation Method:           mvue
#
#Data:                         dat
#
#Sample Size:                  100
#
#Confidence Interval for:      mean.zmlnorm
#
#Confidence Interval Method:    Normal Approximation
#                               (t Distribution)
#
#Confidence Interval Type:      two-sided
#
#Confidence Level:              95%
#
#Confidence Interval:           LCL = 0.748134
#                               UCL = 1.408368
#
#-----

# Clean up
rm(dat)
```

ezmnorm

*Estimate Parameters of a Zero-Modified Normal Distribution***Description**

Estimate the mean and standard deviation of a [zero-modified normal distribution](#), and optionally construct a confidence interval for the mean.

Usage

```
ezmnorm(x, method = "mvue", ci = FALSE, ci.type = "two-sided",
        ci.method = "normal.approx", conf.level = 0.95)
```

Arguments

<code>x</code>	numeric vector of observations.
<code>method</code>	character string specifying the method of estimation. Currently, the only possible value is "mvue" (minimum variance unbiased; the default). See the DETAILS section for more information.
<code>ci</code>	logical scalar indicating whether to compute a confidence interval for the mean. The default value is FALSE.
<code>ci.type</code>	character string indicating what kind of confidence interval to compute. The possible values are "two-sided" (the default), "lower", and "upper". This argument is ignored if ci=FALSE.
<code>ci.method</code>	character string indicating what method to use to construct the confidence interval for the mean. Currently the only possible value is "normal.approx" (the default). See the DETAILS section for more information.
<code>conf.level</code>	a scalar between 0 and 1 indicating the confidence level of the confidence interval. The default value is conf.level=0.95. This argument is ignored if ci=FALSE.

Details

If `x` contains any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

Let $\underline{x} = (x_1, x_2, \dots, x_n)$ be a vector of n observations from a [zero-modified normal distribution](#) with parameters $\text{mean}=\mu$, $\text{sd}=\sigma$, and $\text{p.zero}=p$. Let r denote the number of observations in \underline{x} that are equal to 0, and order the observations so that x_1, x_2, \dots, x_r denote the r zero observations, and $x_{r+1}, x_{r+2}, \dots, x_n$ denote the $n - r$ non-zero observations.

Note that μ is *not* the mean of the zero-modified normal distribution; it is the mean of the normal part of the distribution. Similarly, σ is *not* the standard deviation of the zero-modified normal distribution; it is the standard deviation of the normal part of the distribution.

Let γ and δ denote the mean and standard deviation of the overall zero-modified normal distribution. Aitchison (1955) shows that:

$$\gamma = (1 - p)\mu \quad (1)$$

$$\delta^2 = (1 - p)\sigma^2 + p(1 - p)\mu^2 \quad (2)$$

Estimation

Minimum Variance Unbiased Estimation (method="mvue")

Aitchison (1955) shows that the minimum variance unbiased estimators (mvue's) of γ and δ are:

$$\hat{\gamma}_{mvue} = \bar{x} \quad (3)$$

$$\hat{\delta}_{mvue}^2 = \begin{cases} \frac{n-r-1}{n-1}(s^*)^2 + \frac{r}{n}\left(\frac{n-r}{n-1}\right)(\bar{x}^*)^2 & \text{if } r < n-1, \\ x_n^2/n & \text{if } r = n-1, \\ 0 & \text{if } r = n \end{cases} \quad (4)$$

where

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i \quad (5)$$

$$\bar{x}^* = \frac{1}{n-r} \sum_{i=r+1}^n x_i \quad (6)$$

$$(s^*)^2 = \frac{1}{n-r-1} \sum_{i=r+1}^n (x_i - \bar{x}^*)^2 \quad (7)$$

Note that the quantity in equation (5) is the sample mean of all observations (including 0 values), the quantity in equation (6) is the sample mean of all non-zero observations, and the quantity in equation (7) is the sample variance of all non-zero observations. Also note that for $r = n-1$ or $r = n$, the estimator of δ^2 is the sample variance for all observations (including 0 values).

Confidence Intervals

Based on Normal Approximation (ci.method="normal.approx")

An approximate $(1-\alpha)100\%$ confidence interval for γ is constructed based on the assumption that the estimator of γ is approximately normally distributed. Aitchison (1955) shows that

$$\text{Var}(\hat{\gamma}_{mvue}) = \text{Var}(\bar{x}) = \frac{\delta^2}{n} \quad (8)$$

Thus, an approximate two-sided $(1-\alpha)100\%$ confidence interval for γ is constructed as:

$$\left[\hat{\gamma}_{mvue} - t_{n-2,1-\alpha/2} \frac{\hat{\delta}_{mvue}}{\sqrt{n}}, \hat{\gamma}_{mvue} + t_{n-2,1-\alpha/2} \frac{\hat{\delta}_{mvue}}{\sqrt{n}} \right] \quad (9)$$

where $t_{\nu,p}$ is the p 'th quantile of [Student's t-distribution](#) with ν degrees of freedom.

One-sided confidence intervals are computed in a similar fashion.

Value

a list of class "estimate" containing the estimated parameters and other information. See [estimate.object](#) for details.

The component called parameters is a numeric vector with the following estimated parameters:

Parameter Name	Explanation
mean	mean of the normal (Gaussian) part of the distribution.

sd	standard deviation of the normal (Gaussian) part of the distribution.
p.zero	probability that an observation will be 0.
mean.zmnorm	mean of the overall zero-modified normal distribution.
sd.zmnorm	standard deviation of the overall normal distribution.

Note

The [zero-modified normal distribution](#) is sometimes used to model chemical concentrations for which some observations are reported as “Below Detection Limit”. See, for example USEPA (1992c, pp.27-34). In most cases, however, the zero-modified lognormal (delta) distribution will be more appropriate, since chemical concentrations are bounded below at 0 (e.g., Gilliom and Helsel, 1986; Owen and DeRouen, 1980).

Once you estimate the parameters of the zero-modified normal distribution, it is often useful to characterize the uncertainty in the estimate of the mean. This is done with a confidence interval.

One way to try to assess whether a [zero-modified lognormal \(delta\)](#), [zero-modified normal](#), censored normal, or censored lognormal is the best model for the data is to construct both censored and detects-only probability plots (see [qqPlotCensored](#)).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Aitchison, J. (1955). On the Distribution of a Positive Random Variable Having a Discrete Probability Mass at the Origin. *Journal of the American Statistical Association* **50**, 901–908.
- Gilliom, R.J., and D.R. Helsel. (1986). Estimation of Distributional Parameters for Censored Trace Level Water Quality Data: 1. Estimation Techniques. *Water Resources Research* **22**, 135–146.
- Owen, W., and T. DeRouen. (1980). Estimation of the Mean for Lognormal Data Containing Zeros and Left-Censored Values, with Applications to the Measurement of Worker Exposure to Air Contaminants. *Biometrics* **36**, 707–719.
- USEPA (1992c). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities: Addendum to Interim Final Guidance*. Office of Solid Waste, Permits and State Programs Division, US Environmental Protection Agency, Washington, D.C.

See Also

[ZeroModifiedNormal](#), [Normal](#), [ezmlnorm](#), [ZeroModifiedLognormal](#), [estimate.object](#).

Examples

```
# Generate 100 observations from a zero-modified normal distribution
# with mean=4, sd=2, and p.zero=0.5, then estimate the parameters.
# According to equations (1) and (2) above, the overall mean is
# mean.zmnorm=2 and the overall standard deviation is sd.zmnorm=sqrt(6).
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
dat <- rzmnorm(100, mean = 4, sd = 2, p.zero = 0.5)
ezmnorm(dat, ci = TRUE)

#Results of Distribution Parameter Estimation
```



```

#-----
#
#Assumed Distribution:      Zero-Modified Normal
#
#Estimated Parameter(s):   mean      = 4.037732
#                           sd        = 1.917004
#                           p.zero    = 0.450000
#                           mean.zmnorm = 2.220753
#                           sd.zmnorm  = 2.465829
#
#Estimation Method:        mvue
#
#Data:                     dat
#
#Sample Size:              100
#
#Confidence Interval for:   mean.zmnorm
#
#Confidence Interval Method: Normal Approximation
#                           (t Distribution)
#
#Confidence Interval Type:  two-sided
#
#Confidence Level:          95%
#
#Confidence Interval:       LCL = 1.731417
#                           UCL = 2.710088

#-----

# Following Example 9 on page 34 of USEPA (1992c), compute an
# estimate of the mean of the zinc data, assuming a
# zero-modified normal distribution. The data are stored in
# EPA.92c.zinc.df.

head(EPA.92c.zinc.df)
#  Zinc.orig  Zinc Censored Sample Well
#1      <7  7.00      TRUE      1      1
#2     11.41 11.41     FALSE      2      1
#3      <7  7.00      TRUE      3      1
#4      <7  7.00      TRUE      4      1
#5      <7  7.00      TRUE      5      1
#6     10.00 10.00     FALSE      6      1

New.Zinc <- EPA.92c.zinc.df$Zinc
New.Zinc[EPA.92c.zinc.df$Censored] <- 0
ezmnorm(New.Zinc, ci = TRUE)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:      Zero-Modified Normal
#
#Estimated Parameter(s):   mean      = 11.891000
#                           sd        = 1.594523
#                           p.zero    = 0.500000
#                           mean.zmnorm = 5.945500

```

```

#                                sd.zmnorm   =  6.123235
#
#Estimation Method:              mvue
#
#Data:                           New.Zinc
#
#Sample Size:                    40
#
#Confidence Interval for:        mean.zmnorm
#
#Confidence Interval Method:     Normal Approximation
#                                (t Distribution)
#
#Confidence Interval Type:       two-sided
#
#Confidence Level:               95%
#
#Confidence Interval:            LCL = 3.985545
#                                UCL = 7.905455
#
#-----
# Clean up
rm(dat, New.Zinc)

```

FcnsByCat

EnvStats Functions Listed by Category

Description

Hyperlink list of **EnvStats** functions by category.

Details

- [Calibration](#)
- [Censored Data](#)
- [Data Transformations](#)
- [Estimating Distribution Parameters](#)
- [Estimating Distribution Quantiles](#)
- [Goodness-of-Fit Tests](#)
- [Hypothesis Tests](#)
- [Monte Carlo Simulation and Risk Assessment](#)
- [Plotting Probability Distributions](#)
- [Power and Sample Size Calculations](#)
- [Prediction Intervals](#)
- [Printing and Plotting Methods](#)
- [Probability Distributions and Random Numbers](#)
- [Summary Statistics](#)
- [Tolerance Intervals](#)
- [Trend Analysis](#)

FcnsByCatCalibration *EnvStats Functions for Calibration*

Description

The **EnvStats** functions listed below are useful for performing calibration and inverse prediction to determine the concentration of a chemical based on a machine signal.

Details

<i>Function Name</i>	<i>Description</i>
anovaPE	Compute lack-of-fit and pure error ANOVA table for a linear model.
calibrate	Fit a calibration line or curve.
detectionLimitCalibrate	Determine detection limit based on using a calibration line (or curve) and inverse regression.
inversePredictCalibrate	Predict concentration using a calibration line (or curve) and inverse regression.
pointwise	Pointwise confidence limits for predictions.
predict.lm	Predict method for linear model fits.

FcnsByCatCensoredData *EnvStats Functions for Censored Data*

Description

The **EnvStats** functions listed below are useful for dealing with Type I censored data.

Details**Data Transformations**

<i>Function Name</i>	<i>Description</i>
boxcoxCensored	Compute values of an objective for Box-Cox Power transformations, or compute optimal transformation, for Type I censored data.
print.boxcoxCensored	Print an object of class " boxcoxCensored ".
plot.boxcoxCensored	Plot an object of class " boxcoxCensored ".

Estimating Distribution Parameters

<i>Function Name</i>	<i>Description</i>
----------------------	--------------------

<code>egammaCensored</code>	Estimate shape and scale parameters for a gamma distribution based on Type I censored data.
<code>egammaAltCensored</code>	Estimate mean and CV for a gamma distribution based on Type I censored data.
<code>elnormCensored</code>	Estimate parameters for a lognormal distribution (log-scale) based on Type I censored data.
<code>elnormAltCensored</code>	Estimate parameters for a lognormal distribution (original scale) based on Type I censored data.
<code>enormCensored</code>	Estimate parameters for a Normal distribution based on Type I censored data.
<code>epoisCensored</code>	Estimate parameter for a Poisson distribution based on Type I censored data.
<code>enparCensored</code>	Estimate the mean and standard deviation nonparametrically.
<code>gpqCiNormSinglyCensored</code>	Generate the generalized pivotal quantity used to construct a confidence interval for the mean of a Normal distribution based on Type I singly censored data.
<code>gpqCiNormMultiplyCensored</code>	Generate the generalized pivotal quantity used to construct a confidence interval for the mean of a Normal distribution based on Type I multiply censored data.
<code>print.estimateCensored</code>	Print an object of class <code>"estimateCensored"</code> .

Estimating Distribution Quantiles

<i>Function Name</i>	<i>Description</i>
<code>eqlnormCensored</code>	Estimate quantiles of a Lognormal distribution (log-scale) based on Type I censored data, and optionally construct a confidence interval for a quantile.
<code>eqnormCensored</code>	Estimate quantiles of a Normal distribution based on Type I censored data, and optionally construct a confidence interval for a quantile.

All of the functions for computing quantiles (and associated confidence intervals) for complete (uncensored) data are listed in the help file [Estimating Distribution Quantiles](#). All of these functions, with the exception of `eqnpar`, will accept an object of class `"estimateCensored"`. Thus, you may estimate quantiles (and construct *approximate* confidence intervals) for any distribution for which:

1. There exists a function to estimate distribution parameters using censored data (see the section **Estimating Distribution Parameters** above).

2. There exists a function to estimate quantiles for that distribution based on complete data (see the help file [Estimating Distribution Quantiles](#)).

Nonparametric estimates of quantiles (and associated confidence intervals) can be constructed from censored data as long as the order statistics used in the results are above all left-censored observations or below all right-censored observations. See the help file for [eqnpar](#) for more information and examples.

Goodness-of-Fit Tests

<i>Function Name</i>	<i>Description</i>
gofTestCensored	Perform a goodness-of-fit test based on Type I left- or right-censored data.
print.gofCensored	Print an object of class " gofCensored ".
plot.gofCensored	Plot an object of class " gofCensored ".

Hypothesis Tests

<i>Function Name</i>	<i>Description</i>
twoSampleLinearRankTestCensored	Perform two-sample linear rank tests based on censored data.
print.htestCensored	Printing method for object of class " htestCensored ".

Plotting Probability Distributions

<i>Function Name</i>	<i>Description</i>
cdfCompareCensored	Plot two cumulative distribution functions based on Type I censored data.
ecdfPlotCensored	Plot an empirical cumulative distribution function based on Type I censored data.
ppointsCensored	Compute plotting positions for Type I censored data.
qqPlotCensored	Produce quantile-quantile (Q-Q) plots, also called probability plots, based on Type I censored data.

Prediction and Tolerance Intervals

<i>Function Name</i>	<i>Description</i>
gpqTolIntNormSinglyCensored	Generate the generalized pivotal quantity used to construct a tolerance interval for a Normal distribution based on Type I singly censored data.
gpqTolIntNormMultiplyCensored	Generate the generalized pivotal quantity used to construct a

	tolerance interval for a Normal distribution based on Type I multiply censored data.
<code>tolIntLnormCensored</code>	Tolerance interval for a lognormal distribution (log-scale) based on Type I censored data.
<code>tolIntNormCensored</code>	Tolerance interval for a Normal distribution based on Type I censored data.

All of the functions for computing prediction and tolerance intervals for complete (uncensored) data are listed in the help files [Prediction Intervals](#) and [Tolerance Intervals](#). All of these functions, with the exceptions of `predIntNpar` and `tolIntNpar`, will accept an object of class `"estimateCensored"`. Thus, you may construct *approximate* prediction or tolerance intervals for any distribution for which:

1. There exists a function to estimate distribution parameters using censored data (see the section **Estimating Distribution Parameters** above).
2. There exists a function to create a prediction or tolerance interval for that distribution based on complete data (see the help files [Prediction Intervals](#) and [Tolerance Intervals](#)).

Nonparametric prediction and tolerance intervals can be constructed from censored data as long as the order statistics used in the results are above all left-censored observations or below all right-censored observations. See the help files for `predIntNpar`, `predIntNparSimultaneous`, and `tolIntNpar` for more information and examples.

FcnsByCatDataTrans	<i>EnvStats Functions for Data Transformations</i>
--------------------	--

Description

The **EnvStats** functions listed below are useful for deciding on data transformations.

Details

<i>Function Name</i>	<i>Description</i>
<code>boxcox</code>	Compute values of an objective for Box-Cox transformations, or compute optimal transformation based on raw observations or residuals from a linear model.
<code>boxcoxTransform</code>	Apply a Box-Cox Power transformation to a set of data.
<code>plot.boxcox</code>	Plotting method for an object of class <code>"boxcox"</code> .
<code>plot.boxcoxLm</code>	Plotting method for an object of class <code>"boxcoxLm"</code> .
<code>print.boxcox</code>	Printing method for an object of class <code>"boxcox"</code> .
<code>print.boxcoxLm</code>	Printing method for an object of class <code>"boxcoxLm"</code> .

FcnsByCatEstDistParams

*EnvStats Functions for Estimating Distribution Parameters***Description**

The **EnvStats** functions listed below are useful for estimating distribution parameters and optionally constructing confidence intervals.

Details

<i>Function Name</i>	<i>Description</i>
ebeta	Estimate parameters of a Beta distribution
ebinom	Estimate parameter of a Binomial distribution
eexp	Estimate parameter of an Exponential distribution
eevd	Estimate parameters of an Extreme Value distribution
egamma	Estimate shape and scale parameters of a Gamma distribution
egammaAlt	Estimate mean and CV parameters of a Gamma distribution
egevd	Estimate parameters of a Generalized Extreme Value distribution
egeom	Estimate parameter of a Geometric distribution
ehyper	Estimate parameter of a Hypergeometric distribution
elogis	Estimate parameters of a Logistic distribution
elnorm	Estimate parameters of a Lognormal distribution (log-scale)
elnormAlt	Estimate parameters of a Lognormal distribution (original scale)
elnorm3	Estimate parameters of a Three-Parameter Lognormal distribution
enbinom	Estimate parameter of a Negative Binomial distribution
enorm	Estimate parameters of a Normal distribution
epareto	Estimate parameters of a Pareto distribution
epois	Estimate parameter of a Poisson distribution
eunif	Estimate parameters of a Uniform distribution
eweibull	Estimate parameters of a Weibull distribution
ezmlnorm	Estimate parameters of a Zero-Modified Lognormal (Delta) distribution (log-Scale)
ezmlnormAlt	Estimate parameters of a Zero-Modified Lognormal (Delta) distribution (original Scale)
ezmnorm	Estimate parameters of a Zero-Modified Normal distribution

FcnsByCatEstDistQuants

*EnvStats Functions for Estimating Distribution Quantiles***Description**

The **EnvStats** functions listed below are useful for estimating distribution quantiles and, for some functions, optionally constructing confidence intervals for a quantile.

Details

<i>Function Name</i>	<i>Description</i>
<code>eqbeta</code>	Estimate quantiles of a Beta distribution.
<code>eqbinom</code>	Estimate quantiles of a Binomial distribution.
<code>eqexp</code>	Estimate quantiles of an Exponential distribution.
<code>eqevd</code>	Estimate quantiles of an Extreme Value distribution.
<code>eqgamma</code>	Estimate quantiles of a Gamma distribution using the Shape and Scale Parameterization, and optionally construct a confidence interval for a quantile.
<code>eqgammaAlt</code>	Estimate quantiles of a Gamma distribution using the mean and CV Parameterization, and optionally construct a confidence interval for a quantile.
<code>eqgev</code>	Estimate quantiles of a Generalized Extreme Value distribution.
<code>eqgeom</code>	Estimate quantiles of a Geometric distribution.
<code>eqhyper</code>	Estimate quantiles of a Hypergeometric distribution.
<code>eqlogis</code>	Estimate quantiles of a Logistic distribution.
<code>eqlnorm</code>	Estimate quantiles of a Lognormal distribution (log-scale), and optionally construct a confidence interval for a quantile.
<code>eqlnorm3</code>	Estimate quantiles of a Three-Parameter Lognormal distribution.
<code>eqnbinom</code>	Estimate quantiles of a Negative Binomial distribution.
<code>eqnorm</code>	Estimate quantiles of a Normal distribution, and optionally construct a confidence interval for a quantile.
<code>eqpareto</code>	Estimate quantiles of a Pareto distribution.
<code>eqpois</code>	Estimate quantiles of a Poisson distribution, and optionally construct a confidence interval for a quantile.
<code>equnif</code>	Estimate quantiles of a Uniform distribution.
<code>eqweibull</code>	Estimate quantiles of a Weibull distribution.
<code>eqzmlnorm</code>	Estimate quantiles of a Zero-Modified Lognormal (Delta) distribution (log-scale).
<code>eqzmlnormAlt</code>	Estimate quantiles of a Zero-Modified Lognormal (Delta) distribution (original scale).
<code>eqzmnorm</code>	Estimate quantiles of a Zero-Modified Normal distribution.

FcnsByCatGOFTests

*EnvStats Functions for Goodness-of-Fit Tests***Description**

The **EnvStats** functions listed below are useful for performing goodness-of-fit tests for user-specified probability distributions.

Details

<i>Function Name</i>	<i>Description</i>
<code>gofTest</code>	Perform a goodness-of-fit test for a specified probability distribution. The resulting object is of class " gof " unless the test is the

	two-sample Kolmogorov-Smirnov test, in which case the resulting object is of class <code>"gofTwoSample"</code> .
<code>plot.gof</code>	S3 class method for plotting an object of class <code>"gof"</code> .
<code>print.gof</code>	S3 class method for printing an object of class <code>"gof"</code> .
<code>plot.gofTwoSample</code>	S3 class method for plotting an object of class <code>"gofTwoSample"</code> .
<code>print.gofTwoSample</code>	S3 class method for printing an object of class <code>"gofTwoSample"</code> .
<code>gofGroupTest</code>	Perform a goodness-of-fit test to determine whether data in a set of groups appear to all come from the same probability distribution (with possibly different parameters for each group). The resulting object is of class <code>"gofGroup"</code> .
<code>plot.gofGroup</code>	S3 class method for plotting an object of class <code>"gofGroup"</code> .
<code>print.gofGroup</code>	S3 class method for printing an object of class <code>"gofGroup"</code> .

FcnsByCatHypothTests *EnvStats Functions for Hypothesis Tests*

Description

The **EnvStats** functions listed below are useful for performing hypothesis tests not already built into R. See [Power and Sample Size Calculations](#) for a list of functions you can use to perform power and sample size calculations based on various hypothesis tests.

Details

For goodness-of-fit tests, see [Goodness-of-Fit Tests](#).

Function Name	Description
<code>chenTTest</code>	Chen's modified one-sided t-test for skewed distributions.
<code>kendallTrendTest</code>	Nonparametric test for monotonic trend based on Kendall's tau statistic (and optional confidence interval for slope).
<code>kendallSeasonalTrendTest</code>	Nonparametric test for monotonic trend within each season based on Kendall's tau statistic (and optional confidence interval for slope).
<code>oneSamplePermutationTest</code>	Fisher's one-sample randomization (permutation) test for location.
<code>quantileTest</code>	Two-sample rank test to detect a shift in a proportion of the "treated" population.
<code>quantileTestPValue</code>	Compute p-value associated with a specified combination of m , n , r and k for the quantile test. Useful for determining r and k for a given significance level α .
<code>serialCorrelationTest</code>	Test for the presence of serial correlation.
<code>signTest</code>	One- or paired-sample sign test on the

<code>twoSampleLinearRankTest</code>	median. Two-sample linear rank test to detect a shift in the “treated” population.
<code>twoSamplePermutationTestLocation</code>	Two-sample or paired-sample randomization (permutation) test for location.
<code>twoSamplePermutationTestProportion</code>	Randomization (permutation) test to compare two proportions (Fisher’s exact test).
<code>varTest</code>	One-sample test on variance or two-sample test to compare variances.
<code>varGroupTest</code>	Test for homogeneity of variance among two or more groups.
<code>zTestGevdShape</code>	Estimate the shape parameter of a Generalized Extreme Value distribution and test the null hypothesis that the true value is equal to 0.

FcnsByCatMCandRisk

*EnvStats Functions for Monte Carlo Simulation and Risk Assessment***Description**

The **EnvStats** functions listed below are useful for performing Monte Carlo simulations and risk assessment.

Details

<i>Function Name</i>	<i>Description</i>
<code>Empirical</code>	Empirical distribution based on a set of observations.
<code>simulateVector</code>	Simulate a vector of random numbers from a specified theoretical probability distribution or empirical probability distribution using either Latin hypercube sampling or simple random sampling.
<code>simulateMvMatrix</code>	Simulate a multivariate matrix of random numbers from specified theoretical probability distributions and/or empirical probability distributions based on a specified rank correlation matrix, using either Latin hypercube sampling or simple random sampling.

FcnsByCatPlotProbDists

*EnvStats Functions for Plotting Probability Distributions***Description**

The **EnvStats** functions listed below are useful for plotting probability distributions.

Details

<i>Function Name</i>	<i>Description</i>
<code>cdfCompare</code>	Plot two cumulative distribution functions with the same x -axis in order to compare them.
<code>cdfPlot</code>	Plot a cumulative distribution function.
<code>ecdfPlot</code>	Plot empirical cumulative distribution function.
<code>epdfPlot</code>	Plot empirical probability density function.
<code>pdfPlot</code>	Plot probability density function.
<code>qqPlot</code>	Produce a quantile-quantile (Q-Q) plot, also called a probability plot.
<code>qqPlotGestalt</code>	Plot several Q-Q plots from the same distribution in order to develop a Gestalt of Q-Q plots for that distribution.

FcnsByCatPower

*EnvStats Functions for Power and Sample Size Calculations***Description**

The **EnvStats** functions listed below are useful for power and sample size calculations.

Details**Confidence Intervals**

<i>Function Name</i>	<i>Description</i>
<code>ciBinomHalfWidth</code>	Compute the half-width of a confidence interval for a Binomial proportion or the difference between two proportions.
<code>ciBinomN</code>	Compute the sample size necessary to achieve a specified half-width of a confidence interval for a Binomial proportion or the difference between two proportions.
<code>plotCiBinomDesign</code>	Create plots for a sampling design based on a confidence interval for a Binomial proportion or the difference between two proportions.
<code>ciNormHalfWidth</code>	Compute the half-width of a confidence interval for the mean of a Normal distribution or the difference between two means.
<code>ciNormN</code>	Compute the sample size necessary to achieve a specified half-width of a confidence interval for the mean of a Normal distribution or the difference between two means.
<code>plotCiNormDesign</code>	Create plots for a sampling design based on a confidence interval for the mean of a Normal distribution or the difference between two means.
<code>plotCiNparDesign</code>	Create plots for a sampling design based on a nonparametric confidence interval for a percentile.

Hypothesis Tests

<i>Function Name</i>	<i>Description</i>
<code>aovN</code>	Compute the sample sizes necessary to achieve a

	specified power for a one-way fixed-effects analysis of variance test.
<code>aovPower</code>	Compute the power of a one-way fixed-effects analysis of variance test.
<code>plotAovDesign</code>	Create plots for a sampling design based on a one-way analysis of variance.
<code>propTestN</code>	Compute the sample size necessary to achieve a specified power for a one- or two-sample proportion test.
<code>propTestPower</code>	Compute the power of a one- or two-sample proportion test.
<code>propTestMdd</code>	Compute the minimal detectable difference associated with a one- or two-sample proportion test.
<code>plotPropTestDesign</code>	Create plots involving sample size, power, difference, and significance level for a one- or two-sample proportion test.
<code>tTestN</code>	Compute the sample size necessary to achieve a specified power for a one- or two-sample t-test.
<code>tTestPower</code>	Compute the power of a one- or two-sample t-test.
<code>tTestScaledMdd</code>	Compute the scaled minimal detectable difference associated with a one- or two-sample t-test.
<code>plotTTestDesign</code>	Create plots for a sampling design based on a one- or two-sample t-test.
<code>tTestLnormAltN</code>	Compute the sample size necessary to achieve a specified power for a one- or two-sample t-test, assuming lognormal data.
<code>tTestLnormAltPower</code>	Compute the power of a one- or two-sample t-test, assuming lognormal data.
<code>tTestLnormAltRatioOfMeans</code>	Compute the minimal or maximal detectable ratio of means associated with a one- or two-sample t-test, assuming lognormal data.
<code>plotTTestLnormAltDesign</code>	Create plots for a sampling design based on a one- or two-sample t-test, assuming lognormal data.
<code>linearTrendTestN</code>	Compute the sample size necessary to achieve a specified power for a t-test for linear trend.
<code>linearTrendTestPower</code>	Compute the power of a t-test for linear trend.
<code>linearTrendTestScaledMds</code>	Compute the scaled minimal detectable slope for a t-test for linear trend.
<code>plotLinearTrendTestDesign</code>	Create plots for a sampling design based on a t-test for linear trend.

Prediction Intervals

Normal Distribution Prediction Intervals

<i>Function Name</i>	<i>Description</i>
<code>predIntNormHalfWidth</code>	Compute the half-width of a prediction interval for a normal distribution.
<code>predIntNormK</code>	Compute the required value of K for a prediction interval for a Normal distribution.
<code>predIntNormN</code>	Compute the sample size necessary to achieve a specified half-width for a prediction interval for a Normal distribution.

<code>plotPredIntNormDesign</code>	Create plots for a sampling design based on the width of a prediction interval for a Normal distribution.
<code>predIntNormTestPower</code>	Compute the probability that at least one future observation (or mean) falls outside a prediction interval for a Normal distribution.
<code>plotPredIntNormTestPowerCurve</code>	Create plots for a sampling design based on a prediction interval for a Normal distribution.
<code>predIntNormSimultaneousTestPower</code>	Compute the probability that at least one set of future observations (or means) violates the given rule based on a simultaneous prediction interval for a Normal distribution.
<code>plotPredIntNormSimultaneousTestPowerCurve</code>	Create plots for a sampling design based on a simultaneous prediction interval for a Normal distribution.

Lognormal Distribution Prediction Intervals

<i>Function Name</i>	<i>Description</i>
<code>predIntLnormAltTestPower</code>	Compute the probability that at least one future observation (or geometric mean) falls outside a prediction interval for a lognormal distribution.
<code>plotPredIntLnormAltTestPowerCurve</code>	Create plots for a sampling design based on a prediction interval for a lognormal distribution.
<code>predIntLnormAltSimultaneousTestPower</code>	Compute the probability that at least one set of future observations (or geometric means) violates the given rule based on a simultaneous prediction interval for a lognormal distribution.
<code>plotPredIntLnormAltSimultaneousTestPowerCurve</code>	Create plots for a sampling design based on a simultaneous prediction interval for a lognormal distribution.

Nonparametric Prediction Intervals

<i>Function Name</i>	<i>Description</i>
<code>predIntNparConfLevel</code>	Compute the confidence level associated with a nonparametric prediction interval.
<code>predIntNparN</code>	Compute the required sample size to achieve a specified confidence level for a nonparametric prediction interval.
<code>plotPredIntNparDesign</code>	Create plots for a sampling design based on the confidence level and sample size of a

<code>predIntNparSimultaneousConfLevel</code>	nonparametric prediction interval. Compute the confidence level associated with a simultaneous nonparametric prediction interval.
<code>predIntNparSimultaneousN</code>	Compute the required sample size for a simultaneous nonparametric prediction interval.
<code>plotPredIntNparSimultaneousDesign</code>	Create plots for a sampling design based on a simultaneous nonparametric prediction interval.
<code>predIntNparSimultaneousTestPower</code>	Compute the probability that at least one set of future observations violates the given rule based on a nonparametric simultaneous prediction interval.
<code>plotPredIntNparSimultaneousTestPowerCurve</code>	Create plots for a sampling design based on a simultaneous nonparametric prediction interval.

Tolerance Intervals

<i>Function Name</i>	<i>Description</i>
<code>tolIntNormHalfWidth</code>	Compute the half-width of a tolerance interval for a normal distribution.
<code>tolIntNormK</code>	Compute the required value of K for a tolerance interval for a Normal distribution.
<code>tolIntNormN</code>	Compute the sample size necessary to achieve a specified half-width for a tolerance interval for a Normal distribution.
<code>plotTolIntNormDesign</code>	Create plots for a sampling design based on a tolerance interval for a Normal distribution.
<code>tolIntNparConfLevel</code>	Compute the confidence level associated with a nonparametric tolerance interval for a specified sample size and coverage.
<code>tolIntNparCoverage</code>	Compute the coverage associated with a nonparametric tolerance interval for a specified sample size and confidence level.
<code>tolIntNparN</code>	Compute the sample size required for a nonparametric tolerance interval with a specified coverage and confidence level.
<code>plotTolIntNparDesign</code>	Create plots for a sampling design based on a nonparametric tolerance interval.

Description

The **EnvStats** functions listed below are useful for computing prediction intervals and simultaneous prediction intervals. See [Power and Sample Size](#) for a list of functions useful for computing power and sample size for a design based on a prediction interval width, or a design based on a hypothesis test for future observations falling outside of a prediction interval.

Details

<i>Function Name</i>	<i>Description</i>
predIntGamma , predIntGammaAlt	Prediction interval for the next k observations or next set of k means for a Gamma distribution.
predIntGammaSimultaneous , predIntGammaAltSimultaneous	Construct a simultaneous prediction interval for the next r sampling occasions based on a Gamma distribution.
predIntLnorm , predIntLnormAlt	Prediction interval for the next k observations or geometric means from a Lognormal distribution.
predIntLnormSimultaneous , predIntLnormAltSimultaneous	Construct a simultaneous prediction interval for the next r sampling occasions based on a Lognormal distribution.
predIntNorm	Prediction interval for the next k observations or means from a Normal (Gaussian) distribution.
predIntNormK	Compute the value of K for a prediction interval for a Normal distribution.
predIntNormSimultaneous	Construct a simultaneous prediction interval for the next r sampling occasions based on a Normal distribution.
predIntNormSimultaneousK	Compute the value of K for a simultaneous prediction interval for the next r sampling occasions based on a Normal distribution.
predIntNpar	Nonparametric prediction interval for the next k of K observations.
predIntNparSimultaneous	Construct a nonparametric simultaneous prediction interval for the next r sampling occasions.
predIntPois	Prediction interval for the next k observations or sums from a Poisson distribution.

Description

The **EnvStats** functions listed below are printing and plotting methods for various S3 classes.

Details

Printing Methods

<i>Function Name</i>	<i>Description</i>
<code>print.boxcox</code>	Print an object that inherits from class <code>"boxcox"</code> .
<code>print.boxcoxCensored</code>	Print an object that inherits from class <code>"boxcoxCensored"</code> .
<code>print.boxcoxLm</code>	Print an object that inherits from class <code>"boxcoxLm"</code> .
<code>print.estimate</code>	Print an object that inherits from class <code>"estimate"</code> .
<code>print.intervalEstimate</code>	Print an object that inherits from class <code>"intervalEstimate"</code> .
<code>print.estimateCensored</code>	Print an object that inherits from class <code>"estimateCensored"</code> .
<code>print.intervalEstimateCensored</code>	Print an object that inherits from class <code>"intervalEstimateCensored"</code> .
<code>print.gof</code>	Print an object that inherits from class <code>"gof"</code> .
<code>print.gofCensored</code>	Print an object that inherits from class <code>"gofCensored"</code> .
<code>print.gofGroup</code>	Print an object that inherits from class <code>"gofGroup"</code> .
<code>print.gofTwoSample</code>	Print an object that inherits from class <code>"gofTwoSample"</code> .
<code>print.htest</code>	Print an object that inherits from class <code>"htest"</code> .
<code>print.htestCensored</code>	Print an object that inherits from class <code>"htestCensored"</code> .
<code>print.permutationTest</code>	Print an object that inherits from class <code>"permutationTest"</code> .
<code>print.summaryFull</code>	Print an object that inherits from class <code>"summaryFull"</code> .
<code>print.summaryStats</code>	Print an object that inherits from class <code>"summaryStats"</code> .

Plotting Methods

<i>Function Name</i>	<i>Description</i>
<code>plot.boxcox</code>	Plot an object that inherits from class <code>"boxcox"</code> .
<code>plot.boxcoxCensored</code>	Plot an object that inherits from class <code>"boxcoxCensored"</code> .
<code>plot.boxcoxLm</code>	Plot an object that inherits from class <code>"boxcoxLm"</code> .
<code>plot.gof</code>	Plot an object that inherits from class <code>"gof"</code> .
<code>plot.gofCensored</code>	Plot an object that inherits from class <code>"gofCensored"</code> .
<code>plot.gofGroup</code>	Plot an object that inherits from class <code>"gofGroup"</code> .
<code>plot.gofTwoSample</code>	Plot an object that inherits from class <code>"gofTwoSample"</code> .
<code>plot.permutationTest</code>	Plot an object that inherits from class <code>"permutationTest"</code> .

Description

Listed below are all of the probability distributions available in R and **EnvStats**. Distributions with a description in **bold** are new ones that are part of **EnvStats**. For each distribution, there are functions for generating: values for the probability density function, values for the cumulative distribution function, quantiles, and random numbers.

The data frame [Distribution.df](#) contains information about all of these probability distributions.

Details

<i>Distribution Abbreviation</i>	<i>Description</i>
beta	Beta distribution.
binom	Binomial distribution.
cauchy	Cauchy distribution.
chi	Chi distribution.
chisq	Chi-squared distribution.
exp	Exponential distribution.
evd	Extreme value distribution.
f	F-distribution.
gamma	Gamma distribution.
gammaAlt	Gamma distribution parameterized with mean and CV.
gevd	Generalized extreme value distribution.
geom	Geometric distribution.
hyper	Hypergeometric distribution.
logis	Logistic distribution.
lnorm	Lognormal distribution.
lnormAlt	Lognormal distribution parameterized with mean and CV.
lnormMix	Mixture of two lognormal distributions.
lnormMixAlt	Mixture of two lognormal distributions parameterized by their means and CVs.
lnorm3	Three-parameter lognormal distribution.
lnormTrunc	Truncated lognormal distribution.
lnormTruncAlt	Truncated lognormal distribution parameterized by mean and CV.
nbinom	Negative binomial distribution.
norm	Normal distribution.
normMix	Mixture of two normal distributions.
normTrunc	Truncated normal distribution.
pareto	Pareto distribution.
pois	Poisson distribution.
t	Student's t-distribution.
tri	Triangular distribution.
unif	Uniform distribution.
weibull	Weibull distribution.
wilcox	Wilcoxon rank sum distribution.
zmlnorm	Zero-modified lognormal (delta) distribution.
zmlnormAlt	Zero-modified lognormal (delta) distribution

[zmnorm](#) parameterized with mean and CV.
Zero-modified normal distribution.

FcnsByCatSumStats	<i>EnvStats Functions for Summary Statistics and Plots</i>
-------------------	--

Description

The **EnvStats** functions listed below create summary statistics and plots.

Details

Summary Statistics

R comes with several functions for computing summary statistics, including [mean](#), [var](#), [median](#), [range](#), [quantile](#), and [summary](#). The following functions in **EnvStats** complement these R functions.

<i>Function Name</i>	<i>Description</i>
cv	Coefficient of variation
geoMean	Geometric mean
geoSD	Geometric standard deviation
iqr	Interquartile range
kurtosis	Kurtosis
lMoment	<i>L</i> -moments
pwMoment	Probability-weighted moments
skewness	Skew
summaryFull	Extensive summary statistics
summaryStats	Summary statistics

Summary Plots

R comes with several functions for creating plots to summarize data, including [hist](#), [barplot](#), [boxplot](#), [dotchart](#), [stripchart](#), and numerous others.

The help file [Plotting Probability Distributions](#) lists several **EnvStats** functions useful for producing summary plots as well. In addition, the **EnvStats** function [stripChart](#) is a modification of [stripchart](#) that allows you to include summary statistics on the plot itself.

FcnsByCatTolInts	<i>EnvStats Functions for Tolerance Intervals</i>
------------------	---

Description

The **EnvStats** functions listed below are useful for computing tolerance intervals. See [Power and Sample Size](#) for a list of functions useful for computing power and sample size for a design based on a tolerance interval width.

Details

<i>Function Name</i>	<i>Description</i>
<code>tolIntGamma,</code> <code>tolIntGammaAlt</code>	Tolerance interval for a Gamma distribution.
<code>tolIntLnorm,</code> <code>tolIntLnormAlt</code>	Tolerance interval for a lognormal distribution.
<code>tolIntNorm</code> <code>tolIntNormK</code>	Tolerance interval for a Normal (Gaussian) distribution. Compute the constant K for a Normal (Gaussian) tolerance interval.
<code>tolIntNpar</code>	Nonparametric tolerance interval.
<code>tolIntPois</code>	Tolerance interval for a Poisson distribution.

FcnsByCatTrend

*EnvStats Functions for Trend Analysis***Description**

See [Hypothesis Tests](#).

GammaAlt

*The Gamma Distribution (Alternative Parameterization)***Description**

Density, distribution function, quantile function, and random generation for the gamma distribution with parameters mean and cv.

Usage

```

dgammaAlt(x, mean, cv = 1, log = FALSE)
pgammaAlt(q, mean, cv = 1, lower.tail = TRUE, log.p = FALSE)
qgammaAlt(p, mean, cv = 1, lower.tail = TRUE, log.p = FALSE)
rgammaAlt(n, mean, cv = 1)

```

Arguments

<code>x</code>	vector of quantiles.
<code>q</code>	vector of quantiles.
<code>p</code>	vector of probabilities between 0 and 1.
<code>n</code>	sample size. If <code>length(n)</code> is larger than 1, then <code>length(n)</code> random values are returned.
<code>mean</code>	vector of (positive) means of the distribution of the random variable.
<code>cv</code>	vector of (positive) coefficients of variation of the random variable.
<code>log, log.p</code>	logical; if TRUE, probabilities/densities p are returned as $\log(p)$.
<code>lower.tail</code>	logical; if TRUE (default), probabilities are $P[X \leq x]$, otherwise, $P[X > x]$.

Details

Let X be a random variable with a gamma distribution with parameters $\text{shape}=\alpha$ and $\text{scale}=\beta$. The relationship between these parameters and the mean ($\text{mean}=\mu$) and coefficient of variation ($\text{cv}=\tau$) of this distribution is given by:

$$\alpha = \tau^{-2} \quad (1)$$

$$\beta = \mu/\alpha \quad (2)$$

$$\mu = \alpha\beta \quad (3)$$

$$\tau = \alpha^{-1/2} \quad (4)$$

Thus, the functions `dgammaAlt`, `pgammaAlt`, `qgammaAlt`, and `rgammaAlt` call the R functions [dgamma](#), [pgamma](#), [qgamma](#), and [rgamma](#), respectively, using the values for the shape and scale parameters given by: `shape <- cv^-2`, `scale <- mean/shape`.

Value

`dgammaAlt` gives the density, `pgammaAlt` gives the distribution function, `qgammaAlt` gives the quantile function, and `rgammaAlt` generates random deviates.

Invalid arguments will result in return value NaN, with a warning.

Note

The gamma distribution takes values on the positive real line. Special cases of the gamma are the [exponential distribution](#) and the [chi-square distribution](#). Applications of the gamma include life testing, statistical ecology, queuing theory, inventory control and precipitation processes. A gamma distribution starts to resemble a normal distribution as the shape parameter α tends to infinity or the cv parameter τ tends to 0.

Some EPA guidance documents (e.g., Singh et al., 2002; Singh et al., 2010a,b) discourage using the assumption of a [lognormal distribution](#) for some types of environmental data and recommend instead assessing whether the data appear to fit a gamma distribution.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions, Fourth Edition*. John Wiley and Sons, Hoboken, NJ.
- Johnson, N. L., S. Kotz, and N. Balakrishnan. (1994). *Continuous Univariate Distributions, Volume 1*. Second Edition. John Wiley and Sons, New York.
- Singh, A., A.K. Singh, and R.J. Iaci. (2002). Estimation of the Exposure Point Concentration Term Using a Gamma Distribution. EPA/600/R-02/084. October 2002. Technology Support Center for Monitoring and Site Characterization, Office of Research and Development, Office of Solid Waste and Emergency Response, U.S. Environmental Protection Agency, Washington, D.C.
- Singh, A., R. Maichle, and N. Armbya. (2010a). *ProUCL Version 4.1.00 User Guide (Draft)*. EPA/600/R-07/041, May 2010. Office of Research and Development, U.S. Environmental Protection Agency, Washington, D.C.
- Singh, A., N. Armbya, and A. Singh. (2010b). *ProUCL Version 4.1.00 Technical Guide (Draft)*. EPA/600/R-07/041, May 2010. Office of Research and Development, U.S. Environmental Protection Agency, Washington, D.C.

See Also

[GammaDist](#), [egammaAlt](#), [Probability Distributions and Random Numbers](#).

Examples

```
# Density of a gamma distribution with parameters mean=10 and cv=2,
# evaluated at 7:

dgammaAlt(7, mean = 10, cv = 2)
#[1] 0.02139335

#-----

# The cdf of a gamma distribution with parameters mean=10 and cv=2,
# evaluated at 12:

pgammaAlt(12, mean = 10, cv = 2)
#[1] 0.7713307

#-----

# The 25'th percentile of a gamma distribution with parameters
# mean=10 and cv=2:

qgammaAlt(0.25, mean = 10, cv = 2)
#[1] 0.1056871

#-----

# A random sample of 4 numbers from a gamma distribution with
# parameters mean=10 and cv=2.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(10)
rgammaAlt(4, mean = 10, cv = 2)
#[1] 3.772004230 1.889028078 0.002987823 8.179824976
```

 geoMean

Geometric Mean

Description

Compute the sample geometric mean.

Usage

```
geoMean(x, na.rm = FALSE)
```

Arguments

x	numeric vector of observations.
na.rm	logical scalar indicating whether to remove missing values from x. If na.rm=FALSE (the default) and x contains missing values, then a missing value (NA) is returned. If na.rm=TRUE, missing values are removed from x prior to computing the coefficient of variation.

Details

If x contains any non-positive values (values less than or equal to 0), geoMean returns NA and issues a warning.

Let \underline{x} denote a vector of n observations from some distribution. The sample geometric mean is a measure of central tendency. It is defined as:

$$\bar{x}_G = \sqrt[n]{x_1 x_2 \dots x_n} = \left[\prod_{i=1}^n x_i \right]^{1/n} \quad (1)$$

that is, it is the n 'th root of the product of all n observations.

An equivalent way to define the geometric mean is by:

$$\bar{x}_G = \exp\left[\frac{1}{n} \sum_{i=1}^n \log(x_i)\right] = e^{\bar{y}} \quad (2)$$

where

$$\bar{y} = \frac{1}{n} \sum_{i=1}^n y_i \quad (3)$$

$$y_i = \log(x_i), \quad i = 1, 2, \dots, n \quad (4)$$

That is, the sample geometric mean is antilog of the sample mean of the log-transformed observations.

The geometric mean is only defined for positive observations. It can be shown that the geometric mean is less than or equal to the sample arithmetic mean with equality only when all of the observations are the same value.

Value

A numeric scalar – the sample geometric mean.

Note

The geometric mean is sometimes used to average ratios and percent changes (Zar, 2010). For the lognormal distribution, the geometric mean is the maximum likelihood estimator of the *median* of the distribution, although it is sometimes used incorrectly to estimate the mean of the distribution (see the NOTE section in the help file for [elnormAlt](#)).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Berthouex, P.M., and L.C. Brown. (2002). *Statistics for Environmental Engineers, Second Edition*. Lewis Publishers, Boca Raton, FL.

Gilbert, R.O. (1987). *Statistical Methods for Environmental Pollution Monitoring*. Van Nostrand Reinhold, NY.

Ott, W.R. (1995). *Environmental Statistics and Data Analysis*. Lewis Publishers, Boca Raton, FL.

Taylor, J.K. (1990). *Statistical Techniques for Data Analysis*. Lewis Publishers, Boca Raton, FL.

Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ.

See Also

[geoSD](#), [summaryFull](#), [Summary Statistics](#), [mean](#), [median](#).

Examples

```
# Generate 20 observations from a lognormal distribution with parameters
# mean=10 and cv=2, and compute the mean, median, and geometric mean.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
dat <- rlnormAlt(20, mean = 10, cv = 2)

mean(dat)
#[1] 5.339273

median(dat)
#[1] 3.692091

geoMean(dat)
#[1] 4.095127

#-----
# Clean up
rm(dat)
```

geoSD

*Geometric Standard Deviation.***Description**

Compute the sample geometric standard deviation.

Usage

```
geoSD(x, na.rm = FALSE, sqrt.unbiased = TRUE)
```

Arguments

<code>x</code>	numeric vector of observations.
<code>na.rm</code>	logical scalar indicating whether to remove missing values from <code>x</code> . If <code>na.rm=FALSE</code> (the default) and <code>x</code> contains missing values, then a missing value (NA) is returned. If <code>na.rm=TRUE</code> , missing values are removed from <code>x</code> prior to computing the coefficient of variation.
<code>sqrt.unbiased</code>	logical scalar specifying what method to use to compute the sample standard deviation of the log-transformed observations. If <code>sqrt.unbiased=TRUE</code> (the default), the square root of the unbiased estimator of variance is used, otherwise the method of moments estimator of standard deviation is used. See the DETAILS section for more information.

Details

If `x` contains any non-positive values (values less than or equal to 0), `geoMean` returns NA and issues a warning.

Let \underline{x} denote a vector of n observations from some distribution. The sample geometric standard deviation is a measure of variability. It is defined as:

$$s_G = \exp(s_y) \quad (1)$$

where

$$s_y = \left[\frac{1}{n-1} \sum_{i=1}^n (y_i - \bar{y})^2 \right]^{1/2} \quad (2)$$

$$y_i = \log(x_i), \quad i = 1, 2, \dots, n \quad (3)$$

That is, the sample geometric standard deviation is the antilog of the sample standard deviation of the log-transformed observations.

The sample standard deviation of the log-transformed observations shown in Equation (2) is the square root of the unbiased estimator of variance. (Note that this estimator of standard deviation is not an unbiased estimator.) Sometimes, the square root of the method of moments estimator of variance is used instead:

$$s_y = \left[\frac{1}{n} \sum_{i=1}^n (y_i - \bar{y})^2 \right]^{1/2} \quad (4)$$

This is the estimator used in Equation (1) when `sqrt.unbiased=FALSE`.

Value

A numeric scalar – the sample geometric standard deviation.

Note

The geometric standard deviation is only defined for positive observations. It is usually computed only for observations that are assumed to have come from a [lognormal distribution](#).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Berthouex, P.M., and L.C. Brown. (2002). *Statistics for Environmental Engineers, Second Edition*. Lewis Publishers, Boca Raton, FL.

Gilbert, R.O. (1987). *Statistical Methods for Environmental Pollution Monitoring*. Van Nostrand Reinhold, NY.

Leidel, N.A., K.A. Busch, and J.R. Lynch. (1977). *Occupational Exposure Sampling Strategy Manual*. U.S. Department of Health, Education, and Welfare, Public Health Service, Center for Disease Control, National Institute for Occupational Safety and Health, Cincinnati, Ohio 45226, January, 1977, pp.102–103.

Ott, W.R. (1995). *Environmental Statistics and Data Analysis*. Lewis Publishers, Boca Raton, FL.

Taylor, J.K. (1990). *Statistical Techniques for Data Analysis*. Lewis Publishers, Boca Raton, FL.

Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ.

See Also

[geoMean](#), [linkLognormal](#), [elnorm](#), [summaryFull](#), [Summary Statistics](#).

Examples

```
# Generate 2000 observations from a lognormal distribution with parameters
# mean=10 and cv=1, which implies the standard deviation (on the original
# scale) is 10. Compute the mean, geometric mean, standard deviation,
# and geometric standard deviation.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
dat <- rlnormAlt(2000, mean = 10, cv = 1)

mean(dat)
#[1] 10.23417

geoMean(dat)
#[1] 7.160154

sd(dat)
#[1] 9.786493

geoSD(dat)
#[1] 2.334358
```

```
#-----
# Clean up
rm(dat)
```

GEVD

The Generalized Extreme Value Distribution

Description

Density, distribution function, quantile function, and random generation for the generalized extreme value distribution.

Usage

```
dgevd(x, location = 0, scale = 1, shape = 0)
pgevd(q, location = 0, scale = 1, shape = 0)
qgevd(p, location = 0, scale = 1, shape = 0)
rgevd(n, location = 0, scale = 1, shape = 0)
```

Arguments

x	vector of quantiles.
q	vector of quantiles.
p	vector of probabilities between 0 and 1.
n	sample size. If length(n) is larger than 1, then length(n) random values are returned.
location	vector of location parameters.
scale	vector of positive scale parameters.
shape	vector of shape parameters.

Details

Let X be a generalized extreme value random variable with parameters $\text{location}=\eta$, $\text{scale}=\theta$, and $\text{shape}=\kappa$. When the shape parameter $\kappa = 0$, the generalized extreme value distribution reduces to the [extreme value distribution](#). When the shape parameter $\kappa \neq 0$, the cumulative distribution function of X is given by:

$$F(x; \eta, \theta, \kappa) = \exp\{-[1 - \kappa(x - \eta)/\theta]^{1/\kappa}\}$$

where $-\infty < \eta, \kappa < \infty$ and $\theta > 0$. When $\kappa > 0$, the range of x is:

$$-\infty < x \leq \eta + \theta/\kappa$$

and when $\kappa < 0$ the range of x is:

$$\eta + \theta/\kappa \leq x < \infty$$

The p^{th} quantile of X is given by:

$$x_p = \eta + \frac{\theta\{1 - [-\log(p)]^\kappa\}}{\kappa}$$

Value

density (ddevd), probability (pevd), quantile (qevd), or random sample (revd) for the generalized extreme value distribution with location parameter(s) determined by location, scale parameter(s) determined by scale, and shape parameter(s) determined by shape.

Note

Two-parameter [extreme value distributions \(EVD\)](#) have been applied extensively since the 1930's to several fields of study, including the distributions of hydrological and meteorological variables, human lifetimes, and strength of materials. The three-parameter generalized extreme value distribution (GEVD) was introduced by Jenkinson (1955) to model annual maximum and minimum values of meteorological events. Since then, it has been used extensively in the hydrological and meteorological fields.

The three families of EVDs are all special kinds of GEVDs. When the shape parameter $\kappa = 0$, the GEVD reduces to the [Type I extreme value \(Gumbel\) distribution](#). (The function [zTestGevdShape](#) allows you to test the null hypothesis that the shape parameter is equal to 0.) When $\kappa > 0$, the GEVD is the same as the Type II extreme value distribution, and when $\kappa < 0$ it is the same as the Type III extreme value distribution.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.
- Jenkinson, A.F. (1955). The Frequency Distribution of the Annual Maximum (or Minimum) of Meteorological Events. *Quarterly Journal of the Royal Meteorological Society*, **81**, 158–171.
- Johnson, N. L., S. Kotz, and N. Balakrishnan. (1995). *Continuous Univariate Distributions, Volume 2*. Second Edition. John Wiley and Sons, New York.

See Also

[egevd](#), [zTestGevdShape](#), [EVD](#), [Probability Distributions and Random Numbers](#).

Examples

```
# Density of a generalized extreme value distribution with
# location=0, scale=1, and shape=0, evaluated at 0.5:

dgevd(.5)
#[1] 0.3307043

#-----

# The cdf of a generalized extreme value distribution with
# location=1, scale=2, and shape=0.25, evaluated at 0.5:

pgevd(.5, 1, 2, 0.25)
#[1] 0.2795905

#-----
```

```
# The 90'th percentile of a generalized extreme value distribution with
# location=-2, scale=0.5, and shape=-0.25:

qgev(.9, -2, 0.5, -0.25)
#[1] -0.4895683

#-----

# Random sample of 4 observations from a generalized extreme value
# distribution with location=5, scale=2, and shape=1.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(20)
rgev(4, 5, 2, 1)
#[1] 6.738692 6.473457 4.446649 5.727085
```

Gibbons.et.al.09.Alkilinity.vec

Alkilinity Data from Gibbons et al. (2009)

Description

Alkilinity concentrations (mg/L) in groundwater.

Usage

```
data(Gibbons.et.al.09.Alkilinity.vec)
```

Format

A numeric vector with 27 elements.

Source

Gibbons, R.D., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*. Second Edition. John Wiley & Sons, Hoboken. Table 5.5, p. 107.

Gibbons.et.al.09.Vinyl.Chloride.vec

Vinyl Chloride Data from Gibbons et al. (2009)

Description

Vinyl chloride concentrations (*mug/L*) in groundwater from upgradient background monitoring wells.

Usage

```
data(Gibbons.et.al.09.Vinyl.Chloride.vec)
```

Format

A numeric vector with 34 elements.

Source

Gibbons, R.D., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*. Second Edition. John Wiley & Sons, Hoboken. Table 4.3, p. 87.

gof.object	S3 Class "gof"
------------	----------------

Description

Objects of S3 class "gof" are returned by the **EnvStats** function `gofTest` when just the x argument is supplied.

Details

Objects of S3 class "gof" are lists that contain information about the assumed distribution, the estimated or user-supplied distribution parameters, and the test statistic and p-value.

Value**Required Components**

The following components must be included in a legitimate list of class "gof".

distribution	a character string indicating the name of the assumed distribution (see Distribution.df).
dist.abb	a character string containing the abbreviated name of the distribution (see Distribution.df).
distribution.parameters	a numeric vector with a names attribute containing the names and values of the estimated or user-supplied distribution parameters associated with the assumed distribution.
n.param.est	a scalar indicating the number of distribution parameters estimated prior to performing the goodness-of-fit test. The value of this component will be 0 if the parameters were supplied by the user.
estimation.method	a character string indicating the method used to compute the estimated parameters. The value of this component will depend on the available estimation methods (see Distribution.df). The value of this component will be NULL if the parameters were supplied by the user.
statistic	a numeric scalar with a names attribute containing the name and value of the goodness-of-fit statistic.
sample.size	a numeric scalar containing the number of non-missing observations in the sample used for the goodness-of-fit test.
parameters	numeric vector with a names attribute containing the name(s) and value(s) of the parameter(s) associated with the test statistic given in the <code>statistic</code> component.

<code>z.value</code>	(except when <code>test="chisq"</code> or <code>test="ks"</code>) numeric scalar containing the z-value associated with the goodness-of-fit statistic.
<code>p.value</code>	numeric scalar containing the p-value associated with the goodness-of-fit statistic.
<code>alternative</code>	character string indicating the alternative hypothesis.
<code>method</code>	character string indicating the name of the goodness-of-fit test (e.g., "Shapiro-Wilk GOF").
<code>data</code>	numeric vector containing the data actually used for the goodness-of-fit test (i.e., the original data without any missing or infinite values).
<code>data.name</code>	character string indicating the name of the data object used for the goodness-of-fit test.

NOTE: when the function `gofTest` is called with both arguments `x` and `y` and `test="ks"`, it returns an object of class `"gofTwoSample"`. No specific parametric distribution is assumed, so the value of the component `distribution` is `"Equal"` and the following components are omitted: `dist.abb`, `distribution.parameters`, `n.param.est`, `estimation.method`, and `z.value`.

Optional Components

The following component is included when the data object contains missing (NA), undefined (NaN) and/or infinite (Inf, -Inf) values.

<code>bad.obs</code>	numeric scalar indicating the number of missing (NA), undefined (NaN) and/or infinite (Inf, -Inf) values that were removed from the data object prior to performing the goodness-of-fit test.
----------------------	---

The following components are included in the result of calling `gofTest` with the argument `test="chisq"` and may be used by the function `plot.gof`:

<code>cut.points</code>	numeric vector containing the cutpoints used to define the cells.
<code>counts</code>	numeric vector containing the observed number of counts for each cell.
<code>expected</code>	numeric vector containing the expected number of counts for each cell.
<code>X2.components</code>	numeric vector containing the contribution of each cell to the chi-square statistic.

Methods

Generic functions that have methods for objects of class `"gof"` include: `print`, `plot`.

Note

Since objects of class `"gof"` are lists, you may extract their components with the `$` and `[[` operators.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

See Also

`gofTest`, `print.gof`, `plot.gof`, `Goodness-of-Fit Tests`, `Distribution.df`, `gofCensored.object`.

Examples

[illegible]

```

# Extract the p-value
#-----

gof.obj$p.value
#[1] 0.6279872

#=====

# Plot the results of the test
#-----

windows()
plot(gof.obj)

#=====

# Clean up
#-----
rm(gof.obj)
graphics.off()

```

gofGroup.object	<i>S3 Class "gofGroup"</i>
-----------------	----------------------------

Description

Objects of S3 class "gofGroup" are returned by the **EnvStats** function [gofGroupTest](#).

Details

Objects of S3 class "gofGroup" are lists that contain information about the assumed distribution, the estimated or user-supplied distribution parameters, and the test statistic and p-value.

Value

Required Components

The following components must be included in a legitimate list of class "gofGroup".

distribution	a character string indicating the name of the assumed distribution (see Distribution.df).
dist.abb	a character string containing the abbreviated name of the distribution (see Distribution.df).
statistic	a numeric scalar with a names attribute containing the name and value of the goodness-of-fit statistic.
sample.size	a numeric scalar containing the number of non-missing observations in the sample used for the goodness-of-fit test.
parameters	numeric vector with a names attribute containing the name(s) and value(s) of the parameter(s) associated with the test statistic given in the statistic component.
p.value	numeric scalar containing the p-value associated with the goodness-of-fit statistic.

alternative	character string indicating the alternative hypothesis.
method	character string indicating the name of the goodness-of-fit test (e.g., "Wilk-Shapiro GOF (Normal Scores)").
data.name	character string indicating the name of the data object used for the goodness-of-fit test.
grouping.variable	character string indicating the name of the variable defining the groups.
bad.obs	numeric vector indicating the number of missing (NA), undefined (NaN) and/or infinite (Inf, -Inf) values that were removed from each group and the grouping variable prior to performing the goodness-of-fit test.
n.groups	numeric scalar containing the number of groups.
group.names	character vector containing the levels of the grouping variable, i.e., the names of each of the groups.
group.scores	numeric vector containing the individual statistics for each group.

Optional Component

The following component is included when `gofGroupTest` is called with a formula for the first argument and a data argument.

`parent.of.data` character string indicating the name of the object supplied in the data argument.

Methods

Generic functions that have methods for objects of class "gofGroup" include: `print`, `plot`.

Note

Since objects of class "gofGroup" are lists, you may extract their components with the `$` and `[[` operators.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

See Also

`gofGroupTest`, `print.gofGroup`, `plot.gofGroup`, [Goodness-of-Fit Tests](#), [Distribution.df](#).

Examples

```
# Create an object of class "gofGroup", then print it out.

# Example 10-4 of USEPA (2009, page 10-20) gives an example of
# simultaneously testing the assumption of normality for nickel
# concentrations (ppb) in groundwater collected at 4 monitoring
# wells over 5 months. The data for this example are stored in
# EPA.09.Ex.10.1.nickel.df.

gofGroup.obj <- gofGroupTest(Nickel.ppb ~ Well,
  data = EPA.09.Ex.10.1.nickel.df)

mode(gofGroup.obj)
```

```

#[1] "list"

class(gofGroup.obj)
#[1] "gofGroup"

names(gofGroup.obj)
# [1] "distribution"      "dist.abb"      "statistic"
# [4] "sample.size"      "parameters"    "p.value"
# [7] "alternative"      "method"        "data.name"
#[10] "grouping.variable" "parent.of.data" "bad.obs"
#[13] "n.groups"        "group.names"   "group.scores"

gofGroup.obj
#Results of Group Goodness-of-Fit Test
#-----
#
#Test Method:                Wilk-Shapiro GOF (Normal Scores)
#
#Hypothesized Distribution:   Normal
#
#Data:                        Nickel.ppb
#
#Grouping Variable:          Well
#
#Data Source:                 EPA.09.Ex.10.1.nickel.df
#
#Number of Groups:           4
#
#Sample Sizes:                Well.1 = 5
#                             Well.2 = 5
#                             Well.3 = 5
#                             Well.4 = 5
#
#Test Statistic:              z (G) = -3.658696
#
#P-values for
#Individual Tests:            Well.1 = 0.03510747
#                             Well.2 = 0.02385344
#                             Well.3 = 0.01120775
#                             Well.4 = 0.10681461
#
#P-value for
#Group Test:                  0.0001267509
#
#Alternative Hypothesis:      At least one group
#                             does not come from a
#                             Normal Distribution.

#=====

# Extract the p-values
#-----

gofGroup.obj$p.value
#      Well.1      Well.2      Well.3      Well.4      z (G)
#0.0351074733 0.0238534406 0.0112077511 0.1068146088 0.0001267509

```

```

#####

# Plot the results of the test
#-----

windows()
plot(gofGroup.obj)

#####

# Clean up
#-----
rm(gofGroup.obj)
graphics.off()

```

gofGroupTest	<i>Goodness-of-Fit Test for a Specified Probability Distribution for Groups</i>
--------------	---

Description

Perform a goodness-of-fit test to determine whether data in a set of groups appear to all come from the same probability distribution (with possibly different parameters for each group).

Usage

```

gofGroupTest(object, ...)

## S3 method for class 'formula'
gofGroupTest(object, data = NULL, subset,
  na.action = na.pass, ...)

## Default S3 method:
gofGroupTest(object, group, test = "sw",
  distribution = "norm", est.arg.list = NULL, n.classes = NULL,
  cut.points = NULL, param.list = NULL,
  estimate.params = ifelse(is.null(param.list), TRUE, FALSE),
  n.param.est = NULL, correct = NULL, digits = .Options$digits,
  exact = NULL, ws.method = "normal scores",
  data.name = NULL, group.name = NULL, parent.of.data = NULL,
  subset.expression = NULL, ...)

## S3 method for class 'data.frame'
gofGroupTest(object, ...)

## S3 method for class 'matrix'
gofGroupTest(object, ...)

## S3 method for class 'list'
gofGroupTest(object, ...)

```

Arguments

object	an object containing data for 2 or more groups to be compared to the hypothesized distribution specified by <code>distribution</code> . In the default method, the argument object must be a numeric vector. When object is a data frame, all columns must be numeric. When object is a matrix, it must be a numeric matrix. When object is a list, all components must be numeric vectors. In the formula method, a symbolic specification of the form $y \sim g$ can be given, indicating the observations in the vector y are to be grouped according to the levels of the factor g . Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
data	when object is a formula, data specifies an optional data frame, list or environment (or object coercible by <code>as.data.frame</code> to a data frame) containing the variables in the model. If not found in data, the variables are taken from <code>environment(formula)</code> , typically the environment from which <code>summaryStats</code> is called.
subset	when object is a formula, subset specifies an optional vector specifying a subset of observations to be used.
na.action	when object is a formula, na.action specifies a function which indicates what should happen when the data contain NAs. The default is na.pass .
group	when object is a numeric vector, group is a factor or character vector indicating which group each observation belongs to. When object is a matrix or data frame this argument is ignored and the columns define the groups. When object is a list this argument is ignored and the components define the groups. When object is a formula, this argument is ignored and the right-hand side of the formula specifies the grouping variable.
test	character string defining which goodness-of-fit test to perform on each group. Possible values are: "sw" (Shapiro-Wilk; the default), "sf" (Shapiro-Francia), "ppcc" (Probability Plot Correlation Coefficient), "skew" (Zero-skew), "chisq" (Chi-squared), "ks" (Kolmogorov-Smirnov), and "ws" (Wilk-Shapiro test for Uniform [0, 1] distribution).
distribution	<p>a character string denoting the distribution abbreviation. See the help file for Distribution.df for a list of distributions and their abbreviations. The default value is <code>distribution="norm"</code> (Normal distribution).</p> <p>When <code>test="sw"</code>, <code>test="sf"</code>, or <code>test="ppcc"</code>, any continuous distribution is allowed (e.g., "norm" (normal), "lnorm" (lognormal), "gamma" (gamma), etc.), as well as mixed distributions involving the normal distribution (i.e., "zmnorm" (zero-modified normal), "zmlnorm" (zero-modified lognormal (delta)), and "zmlnorm.alt" (zero-modified lognormal with alternative parameterization)).</p> <p>When <code>test="skew"</code>, only the values "norm" (normal), "lnorm" (lognormal), "lnorm.alt" (lognormal with alternative parameterization), "zmnorm" (zero-modified normal), "zmlnorm" (zero-modified lognormal (delta)), and "zmlnorm.alt" (zero-modified lognormal with alternative parameterization) are allowed.</p> <p>When <code>test="ks"</code>, any continuous distribution is allowed.</p> <p>When <code>test="chisq"</code>, any distribution is allowed.</p> <p>When <code>test="ws"</code>, this argument is ignored.</p>
est.arg.list	a list of arguments to be passed to the function estimating the distribution parameters for each group of observations. For example, if <code>test="sw"</code> and <code>distribution="gamma"</code> , setting <code>est.arg.list=list(method="bcmle")</code> indicates using the bias-corrected maximum-likelihood estimators of shape and

scale (see the help file for [egamma](#). See the help file [Estimating Distribution Parameters](#) for a list of estimating functions. The default value is `est.arg.list=NULL` so that all default values for the estimating function are used. This argument is ignored if `estimate.params=FALSE`.

When `test="sw"`, `test="sf"`, `test="ppcc"`, or `test="skew"`, and you are testing for some form of normality (i.e., [Normal](#), [Lognormal](#), [Three-Parameter Lognormal](#), [Zero-Modified Normal](#), or [Zero-Modified Lognormal \(Delta\)](#)), the estimated parameters are provided in the output merely for information, and the choice of the method of estimation has no effect on the goodness-of-fit test statistics or p-values.

When `test="ks"`, and `estimate.params=TRUE`, the estimated parameters are used to specify the null hypothesis of which distribution the data are assumed to come from.

When `test="chisq"` and `estimate.params=TRUE`, the estimated parameters are used to specify the null hypothesis of which distribution the data are assumed to come from.

When `test="ws"`, this argument is ignored.

<code>n.classes</code>	for the case when <code>test="chisq"</code> , the number of cells into which the observations within each group are to be allocated. If the argument <code>cut.points</code> is supplied, then <code>n.classes</code> is set to <code>length(cut.points)-1</code> . The default value is <code>ceiling(2* (length(x)^(2/5)))</code> and is recommended by Moore (1986).
<code>cut.points</code>	for the case when <code>test="chisq"</code> , a vector of cutpoints that defines the cells for each group of observations. The element <code>x[i]</code> is allocated to cell <code>j</code> if <code>cut.points[j] < x[i] ≤ cut.points[j+1]</code> . If <code>x[i]</code> is less than or equal to the first cutpoint or greater than the last cutpoint, then <code>x[i]</code> is treated as missing. If the hypothesized distribution is discrete, <code>cut.points</code> must be supplied. The default value is <code>cut.points=NULL</code> , in which case the cutpoints are determined by <code>n.classes</code> equi-probable intervals.
<code>param.list</code>	for the case when <code>test="ks"</code> or <code>test="chisq"</code> , a list with values for the parameters of the specified distribution. See the help file for Distribution.df for the names and possible values of the parameters associated with each distribution. The default value is <code>NULL</code> , which forces estimation of the distribution parameters. This argument is ignored if <code>estimate.params=TRUE</code> .
<code>estimate.params</code>	for the case when <code>test="ks"</code> or <code>test="chisq"</code> , a logical scalar indicating whether to perform the goodness-of-fit test based on estimating the distribution parameters (<code>estimate.params=TRUE</code>) or using the user-supplied distribution parameters specified by <code>param.list</code> (<code>estimate.params=FALSE</code>). The default value of <code>estimate.params</code> is <code>TRUE</code> if <code>param.list=NULL</code> , otherwise it is <code>FALSE</code> .
<code>n.param.est</code>	for the case when <code>test="ks"</code> or <code>test="chisq"</code> , an integer indicating the number of parameters estimated from the data. If <code>estimate.params=TRUE</code> , the default value is the number of parameters associated with the distribution specified by <code>distribution</code> (e.g., 2 for a normal distribution). If <code>estimate.params=FALSE</code> , the default value is <code>n.param.est=0</code> .
<code>correct</code>	for the case when <code>test="chisq"</code> , a logical scalar indicating whether to use the continuity correction. The default value is <code>correct=FALSE</code> unless <code>n.classes=2</code> .

<code>digits</code>	a scalar indicating how many significant digits to print out for the parameters associated with the hypothesized distribution. The default value is <code>.Options\$digits</code> .
<code>exact</code>	for the case when <code>test="ks"</code> , <code>exact=NULL</code> by default, but can be set to a logical scalar indicating whether an exact p-value should be computed. See the help file for ks.test for more information.
<code>ws.method</code>	character string indicating which method to use when performing the Wilk-Shapiro test for a Uniform [0,1] distribution on the p-values from the goodness-of-fit tests on each group. Possible values are <code>ws.method="normal scores"</code> (the default) or <code>ws.method="chi-square scores"</code> . See the subsection <i>Wilk-Shapiro goodness-of-fit test for Uniform [0, 1] distribution</i> under the DETAILS section of the help file for gofTest for more information. NOTE: In the case where you are testing whether each group comes from a Uniform [0,1] distribution (i.e., when you set <code>test="ws"</code>), the argument <code>ws.method</code> determines which score types are used for each individual test of the groups as well.
<code>data.name</code>	character string indicating the name of the data used for the goodness-of-fit tests. The default value is <code>data.name=deparse(substitute(object))</code> .
<code>group.name</code>	character string indicating the name of the data used to create the groups. The default value is <code>group.name=deparse(substitute(group))</code> .
<code>parent.of.data</code>	character string indicating the source of the data used for the goodness-of-fit tests.
<code>subset.expression</code>	character string indicating the expression used to subset the data.
<code>...</code>	additional arguments affecting the goodness-of-fit test.

Details

The function `gofGroupTest` performs a goodness-of-fit test for each group of data by calling the function [gofTest](#). Using the p-values from these goodness-of-fit tests, it then calls the function [gofTest](#) with the argument `test="ws"` to test whether the p-values appear to come from a [Uniform \[0,1\] distribution](#).

Value

a list of class "gofGroup" containing the results of the group goodness-of-fit test. Objects of class "gofGroup" have special printing and plotting methods. See the help file for [gofGroup.object](#) for details.

Note

The Wilk-Shapiro (1968) tests for a Uniform [0, 1] distribution were introduced in the context of testing whether several independent samples all come from normal distributions, with possibly different means and variances. The function `gofGroupTest` extends this idea to allow you to test whether several independent samples come from the same distribution (e.g., gamma, extreme value, etc.), with possibly different parameters.

Examples of simultaneously assessing whether several groups come from the same distribution are given in USEPA (2009) and Gibbons et al. (2009).

In practice, almost any goodness-of-fit test will *not* reject the null hypothesis if the number of observations is relatively small. Conversely, almost any goodness-of-fit test *will* reject the null hypothesis if the number of observations is very large, since "real" data are never distributed according to any

theoretical distribution (Conover, 1980, p.367). For most cases, however, the distribution of “real” data is close enough to some theoretical distribution that fairly accurate results may be provided by assuming that particular theoretical distribution. One way to assess the goodness of the fit is to use goodness-of-fit tests. Another way is to look at quantile-quantile (Q-Q) plots (see [qqPlot](#)).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Gibbons, R.D., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*, Second Edition. John Wiley & Sons, Hoboken.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.17-17.
- USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.
- Wilk, M.B., and S.S. Shapiro. (1968). The Joint Assessment of Normality of Several Independent Samples. *Technometrics*, **10**(4), 825-839.

See Also

[gofTest](#), [gofGroup.object](#), [print.gofGroup](#), [plot.gofGroup](#), [qqPlot](#).

Examples

```
# Example 10-4 of USEPA (2009, page 10-20) gives an example of
# simultaneously testing the assumption of normality for nickel
# concentrations (ppb) in groundwater collected at 4 monitoring
# wells over 5 months. The data for this example are stored in
# EPA.09.Ex.10.1.nickel.df.
```

```
EPA.09.Ex.10.1.nickel.df
#   Month  Well Nickel.ppb
#1      1 Well.1      58.8
#2      3 Well.1       1.0
#3      6 Well.1     262.0
#4      8 Well.1      56.0
#5     10 Well.1       8.7
#6      1 Well.2      19.0
#7      3 Well.2     81.5
#8      6 Well.2    331.0
#9      8 Well.2      14.0
#10     10 Well.2     64.4
#11     1 Well.3      39.0
#12     3 Well.3     151.0
#13     6 Well.3      27.0
#14     8 Well.3      21.4
#15     10 Well.3    578.0
#16     1 Well.4       3.1
#17     3 Well.4     942.0
#18     6 Well.4     85.6
```

```

#19      8 Well.4      10.0
#20     10 Well.4     637.0

# Test for a normal distribution at each well:
#-----

gofGroup.list <- gofGroupTest(Nickel.ppb ~ Well,
  data = EPA.09.Ex.10.1.nickel.df)

gofGroup.list

#Results of Group Goodness-of-Fit Test
#-----
#
#Test Method:                Wilk-Shapiro GOF (Normal Scores)
#
#Hypothesized Distribution:   Normal
#
#Data:                       Nickel.ppb
#
#Grouping Variable:          Well
#
#Data Source:                EPA.09.Ex.10.1.nickel.df
#
#Number of Groups:           4
#
#Sample Sizes:               Well.1 = 5
#                             Well.2 = 5
#                             Well.3 = 5
#                             Well.4 = 5
#
#Test Statistic:             z (G) = -3.658696
#
#P-values for
#Individual Tests:           Well.1 = 0.03510747
#                             Well.2 = 0.02385344
#                             Well.3 = 0.01120775
#                             Well.4 = 0.10681461
#
#P-value for
#Group Test:                 0.0001267509
#
#Alternative Hypothesis:     At least one group
#                             does not come from a
#                             Normal Distribution.

windows()
plot(gofGroup.list)

#-----

# Test for a lognormal distribution at each well:
#-----

gofGroupTest(Nickel.ppb ~ Well, data = EPA.09.Ex.10.1.nickel.df,
  dist = "lnorm")

```



```

#Results of Group Goodness-of-Fit Test
#-----
#
#Test Method:                Wilk-Shapiro GOF (Normal Scores)
#
#Hypothesized Distribution:   Lognormal
#
#Data:                       Nickel.ppb
#
#Grouping Variable:          Well
#
#Data Source:                 EPA.09.Ex.10.1.nickel.df
#
#Number of Groups:           4
#
#Sample Sizes:               Well.1 = 5
#                             Well.2 = 5
#                             Well.3 = 5
#                             Well.4 = 5
#
#Test Statistic:              z (G) = 0.2401720
#
#P-values for
#Individual Tests:           Well.1 = 0.6898164
#                             Well.2 = 0.6700394
#                             Well.3 = 0.3208299
#                             Well.4 = 0.5041375
#
#P-value for
#Group Test:                  0.5949015
#
#Alternative Hypothesis:      At least one group
#                             does not come from a
#                             Lognormal Distribution.

#-----
# Clean up
rm(gofGroup.list)
graphics.off()

```

gofTest

Goodness-of-Fit Test

Description

Perform a goodness-of-fit test to determine whether a data set appears to come from a specified probability distribution or if two data sets appear to come from the same distribution.

Usage

```
gofTest(y, ...)
```

```
## S3 method for class 'formula'
```

```

gofTest(y, data = NULL, subset,
        na.action = na.pass, ...)

## Default S3 method:
gofTest(y, x = NULL,
        test = ifelse(is.null(x), "sw", "ks"),
        distribution = "norm", est.arg.list = NULL,
        alternative = "two.sided", n.classes = NULL,
        cut.points = NULL, param.list = NULL,
        estimate.params = ifelse(is.null(param.list), TRUE, FALSE),
        n.param.est = NULL, correct = NULL, digits = .Options$digits,
        exact = NULL, ws.method = "normal scores", warn = TRUE,
        data.name = NULL, data.name.x = NULL, parent.of.data = NULL,
        subset.expression = NULL, ...)

```

Arguments

<code>y</code>	an object containing data for the goodness-of-fit test. In the default method, the argument <code>y</code> must be numeric vector of observations. In the formula method, <code>y</code> must be a formula of the form <code>y ~ 1</code> or <code>y ~ x</code> . The form <code>y ~ 1</code> indicates use the observations in the vector <code>y</code> for a one-sample goodness-of-fit test. The form <code>y ~ x</code> is only relevant to the case of the two-sample Kolmogorov-Smirnov test (<code>test="ks"</code>) and indicates use the observations in the vector <code>y</code> as the second sample and use the observations in the vector <code>x</code> as the first sample. Note that for the formula method, <code>x</code> and <code>y</code> must be the same length but this is not a requirement of the test and you can use vectors of different lengths via the default method. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
<code>data</code>	specifies an optional data frame, list or environment (or object coercible by <code>as.data.frame</code> to a data frame) containing the variables in the model. If not found in data, the variables are taken from <code>environment(formula)</code> , typically the environment from which <code>gofTest</code> is called.
<code>subset</code>	specifies an optional vector specifying a subset of observations to be used.
<code>na.action</code>	specifies a function which indicates what should happen when the data contain NAs. The default is na.pass .
<code>x</code>	numeric vector of values for the first sample in the case of a two-sample Kolmogorov-Smirnov goodness-of-fit test (<code>test="ks"</code>). Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
<code>test</code>	character string defining which goodness-of-fit test to perform. Possible values are: "sw" (Shapiro-Wilk; the default when <code>x</code> is NOT supplied), "sf" (Shapiro-Francia), "ppcc" (Probability Plot Correlation Coefficient), "skew" (Zero-skew), "chisq" (Chi-squared), "ks" (Kolmogorov-Smirnov; the default when <code>x</code> IS supplied), and "ws" (Wilk-Shapiro test for Uniform [0, 1] distribution). When the argument <code>x</code> is supplied, you must set <code>test="ks"</code> , which is what <code>gofTest</code> does by default.
<code>distribution</code>	a character string denoting the distribution abbreviation. See the help file for Distribution.df for a list of distributions and their abbreviations. The default value is <code>distribution="norm"</code> (Normal distribution). When <code>test="sw"</code> , <code>test="sf"</code> , or <code>test="ppcc"</code> , any continuous distribuiton is allowed (e.g., "norm" (normal), "lnorm" (lognormal), "gamma" (gamma), etc.), as well as mixed distributions involving the normal distribution (i.e., "zmnorm"

	<p>(zero-modified normal), "zmlnorm" (zero-modified lognormal (delta)), and "zmlnorm.alt" (zero-modified lognormal with alternative parameterization)).</p> <p>When test="skew", only the values "norm" (normal), "lnorm" (lognormal), "lnorm.alt" (lognormal with alternative parameterization), "zmnorm" (zero-modified normal), "zmlnorm" (zero-modified lognormal (delta)), and "zmlnorm.alt" (zero-modified lognormal with alternative parameterization) are allowed.</p> <p>When test="ks", any continuous distribution is allowed.</p> <p>When test="chisq", any distribution is allowed.</p> <p>When test="ws", this argument is ignored.</p>
est.arg.list	<p>a list of arguments to be passed to the function estimating the distribution parameters. For example, if test="sw" and distribution="gamma", setting est.arg.list=list(method="bcmle") indicates using the bias-corrected maximum-likelihood estimators of shape and scale (see the help file for egamma. See the help file Estimating Distribution Parameters for a list of estimating functions. The default value is est.arg.list=NULL so that all default values for the estimating function are used. This argument is ignored if estimate.params=FALSE.</p> <p>When test="sw", test="sf", test="ppcc", or test="skew", and you are testing for some form of normality (i.e., Normal, Lognormal, Three-Parameter Lognormal, Zero-Modified Normal, or Zero-Modified Lognormal (Delta)), the estimated parameters are provided in the output merely for information, and the choice of the method of estimation has no effect on the goodness-of-fit test statistic or p-value.</p> <p>When test="ks", x is not supplied, and estimate.params=TRUE, the estimated parameters are used to specify the null hypothesis of which distribution the data are assumed to come from.</p> <p>When test="chisq" and estimate.params=TRUE, the estimated parameters are used to specify the null hypothesis of which distribution the data are assumed to come from.</p> <p>When test="ws", this argument is ignored.</p>
alternative	<p>for the case when test="ks", test="skew", or test="ws", character string specifying the alternative hypothesis. When test="ks" or test="skew", the possible values are "two-sided" (the default), "greater", or "less". When test="ws", the possible values are "greater" (the default), or "less". See the DETAILS section of the help file for ks.test for more explanation of the meaning of this argument.</p>
n.classes	<p>for the case when test="chisq", the number of cells into which the observations are to be allocated. If the argument cut.points is supplied, then n.classes is set to length(cut.points)-1. The default value is $\text{ceiling}(2 * (\text{length}(x)^{(2/5)}))$ and is recommended by Moore (1986).</p>
cut.points	<p>for the case when test="chisq", a vector of cutpoints that defines the cells. The element x[i] is allocated to cell j if $\text{cut.points}[j] < x[i] \leq \text{cut.points}[j+1]$. If x[i] is less than or equal to the first cutpoint or greater than the last cutpoint, then x[i] is treated as missing. If the hypothesized distribution is discrete, cut.points must be supplied. The default value is cut.points=NULL, in which case the cutpoints are determined by n.classes equi-probable intervals.</p>
param.list	<p>for the case when test="ks" and x is not supplied, or when test="chisq", a list with values for the parameters of the specified distribution. See the help file for Distribution.df for the names and possible values of the parameters</p>

associated with each distribution. The default value is `param.list=NULL`, which forces estimation of the distribution parameters. This argument is ignored if `estimate.params=TRUE`.

<code>estimate.params</code>	for the case when <code>test="ks"</code> and <code>x</code> is not supplied, or when <code>test="chisq"</code> , a logical scalar indicating whether to perform the goodness-of-fit test based on estimating the distribution parameters (<code>estimate.params=TRUE</code>) or using the user-supplied distribution parameters specified by <code>param.list</code> (<code>estimate.params=FALSE</code>). The default value of <code>estimate.params</code> is <code>TRUE</code> if <code>param.list=NULL</code> , otherwise it is <code>FALSE</code> .
<code>n.param.est</code>	for the case when <code>test="ks"</code> and <code>x</code> is not supplied, or when <code>test="chisq"</code> , an integer indicating the number of parameters estimated from the data. If <code>estimate.params=TRUE</code> , the default value is the number of parameters associated with the distribution specified by <code>distribution</code> (e.g., 2 for a normal distribution). If <code>estimate.params=FALSE</code> , the default value is <code>n.param.est=0</code> .
<code>correct</code>	for the case when <code>test="chisq"</code> , a logical scalar indicating whether to use the continuity correction. The default value is <code>correct=FALSE</code> unless <code>n.classes=2</code> .
<code>digits</code>	a scalar indicating how many significant digits to print out for the parameters associated with the hypothesized distribution. The default value is <code>.Options\$digits</code> .
<code>exact</code>	for the case when <code>test="ks"</code> , <code>exact=NULL</code> by default, but can be set to a logical scalar indicating whether an exact p-value should be computed. See the help file for ks.test for more information.
<code>ws.method</code>	for the case when <code>test="ws"</code> , this argument specifies whether to perform the test based on normal scores (<code>ws.method="normal scores"</code> , the default) or chi-square scores (<code>ws.method="chi-square scores"</code>). See the DETAILS section for more information.
<code>warn</code>	logical scalar indicating whether to print a warning message when observations with NAs, NaNs, or Infs in <code>y</code> or <code>x</code> are removed. The default value is <code>TRUE</code> .
<code>data.name</code>	character string indicating the name of the data used for argument <code>y</code> .
<code>data.name.x</code>	character string indicating the name of the data used for argument <code>x</code> .
<code>parent.of.data</code>	character string indicating the source of the data used for the goodness-of-fit test.
<code>subset.expression</code>	character string indicating the expression used to subset the data.
<code>...</code>	additional arguments affecting the goodness-of-fit test.

Details

- Shapiro-Wilk goodness-of-fit test (`test="sw"`).

The Shapiro-Wilk goodness-of-fit test (Shapiro and Wilk, 1965; Royston, 1992a) is one of the most commonly used goodness-of-fit tests for normality. You can use it to test the following hypothesized distributions: [Normal](#), [Lognormal](#), [Three-Parameter Lognormal](#), [Zero-Modified Normal](#), or [Zero-Modified Lognormal \(Delta\)](#). **In addition, you can also use it to test the null hypothesis of any continuous distribution that is available** (see the help file for [Distribution.df](#)). This is done by using the idea of Chen and Balakrishnan (1995), who proposed a general purpose approximate goodness-of-fit test based on the Cramer-von Mises or Anderson-Darling goodness-of-fit tests for normality. The function `gofTest` modifies the approach of Chen and Balakrishnan (1995) by using the same first 2 steps, and then applying the Shapiro-Wilk test:

1. Let x denote the vector of n observations. Compute cumulative probabilities for each value of x based on the cumulative distribution function for the hypothesized distribution. That is, compute $p_i = F(x_i, \hat{\theta})$ where $F(x, \theta)$ denotes the hypothesized cumulative distribution function with parameter(s) θ , x_i denotes the i 'th ordered observation, and $\hat{\theta}$ denotes the estimated parameter(s).
 2. Compute standard normal deviates based on the computed cumulative probabilities:

$$y_i = \Phi^{-1}(p_i)$$
 3. Perform the Shapiro-Wilk goodness-of-fit test on the y_i 's.
- Shapiro-Francia goodness-of-fit test (`test="sf"`).
 The Shapiro-Francia goodness-of-fit test (Shapiro and Francia, 1972; Weisberg and Bingham, 1975; Royston, 1992c) is also one of the most commonly used goodness-of-fit tests for normality. You can use it to test the following hypothesized distributions: [Normal](#), [Lognormal](#), [Zero-Modified Normal](#), or [Zero-Modified Lognormal \(Delta\)](#). In addition, you can also use it to test the null hypothesis of any continuous distribution that is available (see the help file for [Distribution.df](#)). See the explanation for the Shapiro-Wilk test above for how this is done.
 - Probability Plot Correlation Coefficient (PPCC) goodness-of-fit test (`test="ppcc"`).
 The PPCC goodness-of-fit test (Filliben, 1975; Looney and Gudge, 1985) can be used to test the following hypothesized distributions: [Normal](#), [Lognormal](#), [Zero-Modified Normal](#), or [Zero-Modified Lognormal \(Delta\)](#). In addition, you can also use it to test the null hypothesis of any continuous distribution that is available (see the help file for [Distribution.df](#)). The function `gofTest` computes the PPCC test statistic using Blom plotting positions. It can be shown that the square of this statistic is equivalent to the Weisberg-Bingham Approximation to the Shapiro-Francia W' -Test (Weisberg and Bingham, 1975; Royston, 1993). Thus the PPCC goodness-of-fit test is equivalent to the Shapiro-Francia goodness-of-fit test.
 - Zero-skew goodness-of-fit test (`test="skew"`).
 The Zero-skew goodness-of-fit test (D'Agostino, 1970) can be used to test the following hypothesized distributions: [Normal](#), [Lognormal](#), [Zero-Modified Normal](#), or [Zero-Modified Lognormal \(Delta\)](#).
 - Kolmogorov-Smirnov goodness-of-fit test (`test="ks"`).
 When `test="ks"`, the function `gofTest` calls the R function `ks.test` to compute the test statistic and p-value. Note that for the one-sample case, the distribution parameters should be pre-specified and not estimated from the data, and if the distribution parameters are estimated from the data you will receive a warning that this test is very conservative (Type I error smaller than assumed; high Type II error) in this case.
 - Chi-squared goodness-of-fit test (`test="chisq"`).
 The method used by `gofTest` is a modification of what is used for `chisq.test`. If the hypothesized distribution function is completely specified, the degrees of freedom are $m - 1$ where m denotes the number of classes. If any parameters are estimated, the degrees of freedom depend on the method of estimation. The function `gofTest` follows the convention of computing degrees of freedom as $m - 1 - k$, where k is the number of parameters estimated. It can be shown that if the parameters are estimated by maximum likelihood, the degrees of freedom are bounded between $m - 1$ and $m - 1 - k$. Therefore, especially when the sample size is small, it is important to compare the test statistic to the chi-squared distribution with both $m - 1$ and $m - 1 - k$ degrees of freedom. See Kendall and Stuart (1991, Chapter 30) for a more complete discussion.
 The distribution theory of chi-square statistics is a large sample theory. The expected cell counts are assumed to be at least moderately large. As a rule of thumb, each should be at least 5. Although authors have found this rule to be conservative (especially when the class probabilities are not too different from each other), the user should regard p-values with caution when expected cell counts are small.

- Wilk-Shapiro goodness-of-fit test for Uniform [0, 1] distribution (test="ws").

Wilk and Shapiro (1968) suggested this test in the context of jointly testing several independent samples for normality simultaneously. If p_1, p_2, \dots, p_n denote the p-values associated with the test for normality of n independent samples, then under the null hypothesis that all n samples come from a normal distribution, the p-values are a random sample of n observations from a Uniform [0,1] distribution, that is a Uniform distribution with minimum 0 and maximum 1. Wilk and Shapiro (1968) suggested two different methods for testing whether the p-values come from a Uniform [0, 1] distribution:

- *Test Based on Normal Scores.* Under the null hypothesis, the normal scores

$$\Phi^{-1}(p_1), \Phi^{-1}(p_2), \dots, \Phi^{-1}(p_n)$$

are a random sample of n observations from a standard normal distribution. Wilk and Shapiro (1968) denote the i 'th normal score by

$$G_i = \Phi^{-1}(p_i)$$

and note that under the null hypothesis, the quantity G defined as

$$G = \frac{1}{\sqrt{n}} \sum_{i=1}^n G_i$$

has a standard normal distribution. Wilk and Shapiro (1968) were interested in the alternative hypothesis that some of the n independent samples did not come from a normal distribution and hence would be associated with smaller p-values than expected under the null hypothesis, which translates to the alternative that the cdf for the distribution of the p-values is greater than the cdf of a Uniform [0, 1] distribution (alternative="greater"). In terms of the test statistic G , this alternative hypothesis would tend to make G smaller than expected, so the p-value is given by $\Phi(G)$. For the one-sided lower alternative that the cdf for the distribution of p-values is less than the cdf for a Uniform [0, 1] distribution, the p-value is given by $1 - \Phi(G)$.

- *Test Based on Chi-Square Scores.* Under the null hypothesis, the chi-square scores

$$-2 \log(p_1), -2 \log(p_2), \dots, -2 \log(p_n)$$

are a random sample of n observations from a chi-square distribution with 2 degrees of freedom (Fisher, 1950). Wilk and Shapiro (1968) denote the i 'th chi-square score by

$$C_i = -2 \log(p_i)$$

and note that under the null hypothesis, the quantity C defined as

$$C = \sum_{i=1}^n C_i$$

has a chi-square distribution with $2n$ degrees of freedom. Wilk and Shapiro (1968) were interested in the alternative hypothesis that some of the n independent samples did not come from a normal distribution and hence would be associated with smaller p-values than expected under the null hypothesis, which translates to the alternative that the cdf for the distribution of the p-values is greater than the cdf of a Uniform [0, 1] distribution (alternative="greater"). In terms of the test statistic C , this alternative hypothesis would tend to make C larger than expected, so the p-value is given by $1 - F_{2n}(C)$, where F_{2n} denotes the cumulative distribution function of the chi-square distribution with $2n$ degrees of freedom. For the one-sided lower alternative that the cdf for the distribution of p-values is less than the cdf for a Uniform [0, 1] distribution, the p-value is given by $F_{2n}(C)$.

Value

a list of class "gof" containing the results of the goodness-of-fit test, unless the two-sample Kolmogorov-Smirnov test is used, in which case the value is a list of class "gofTwoSample". Objects of class "gof" and "gofTwoSample" have special printing and plotting methods. See the help files for [gof.object](#) and [gofTwoSample.object](#) for details.

Note

The Shapiro-Wilk test (Shapiro and Wilk, 1965) and the Shapiro-Francia test (Shapiro and Francia, 1972) are probably the two most commonly used hypothesis tests to test departures from normality. The Shapiro-Wilk test is most powerful at detecting short-tailed (platykurtic) and skewed distributions, and least powerful against symmetric, moderately long-tailed (leptokurtic) distributions. Conversely, the Shapiro-Francia test is more powerful against symmetric long-tailed distributions and less powerful against short-tailed distributions (Royston, 1992b; 1993).

Kolmogorov (1933) introduced a goodness-of-fit test to test the hypothesis that a random sample of n observations \mathbf{x} comes from a specific hypothesized distribution with cumulative distribution function H . This test is now usually called the one-sample Kolmogorov-Smirnov goodness-of-fit test. Smirnov (1939) introduced a goodness-of-fit test to test the hypothesis that a random sample of n observations \mathbf{x} comes from the same distribution as a random sample of m observations \mathbf{y} . This test is now usually called the two-sample Kolmogorov-Smirnov goodness-of-fit test. Both tests are based on the maximum vertical distance between two cumulative distribution functions. For the one-sample problem with a small sample size, the Kolmogorov-Smirnov test may be preferred over the chi-squared goodness-of-fit test since the KS-test is exact, while the chi-squared test is based on an asymptotic approximation.

The chi-squared test, introduced by Pearson in 1900, is the oldest and best known goodness-of-fit test. The idea is to reduce the goodness-of-fit problem to a multinomial setting by comparing the observed cell counts with their expected values under the null hypothesis. Grouping the data sacrifices information, especially if the hypothesized distribution is continuous. On the other hand, chi-squared tests can be applied to any type of variable: continuous, discrete, or a combination of these.

The Wilk-Shapiro (1968) tests for a Uniform $[0, 1]$ distribution were introduced in the context of testing whether several independent samples all come from normal distributions, with possibly different means and variances. The function [gofGroupTest](#) extends this idea to allow you to test whether several independent samples come from the same distribution (e.g., gamma, extreme value, etc.), with possibly different parameters.

In practice, almost any goodness-of-fit test will *not* reject the null hypothesis if the number of observations is relatively small. Conversely, almost any goodness-of-fit test *will* reject the null hypothesis if the number of observations is very large, since "real" data are never distributed according to any theoretical distribution (Conover, 1980, p.367). For most cases, however, the distribution of "real" data is close enough to some theoretical distribution that fairly accurate results may be provided by assuming that particular theoretical distribution. One way to assess the goodness of the fit is to use goodness-of-fit tests. Another way is to look at quantile-quantile (Q-Q) plots (see [qqPlot](#)).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Birnbaum, Z.W., and F.H. Tingey. (1951). One-Sided Confidence Contours for Probability Distribution Functions. *Annals of Mathematical Statistics* **22**, 592-596.

- Blom, G. (1958). *Statistical Estimates and Transformed Beta Variables*. John Wiley and Sons, New York.
- Conover, W.J. (1980). *Practical Nonparametric Statistics*. Second Edition. John Wiley and Sons, New York.
- Dallal, G.E., and L. Wilkinson. (1986). An Analytic Approximation to the Distribution of Lilliefors's Test for Normality. *The American Statistician* **40**, 294-296.
- D'Agostino, R.B. (1970). Transformation to Normality of the Null Distribution of g_1 . *Biometrika* **57**, 679-681.
- D'Agostino, R.B. (1971). An Omnibus Test of Normality for Moderate and Large Size Samples. *Biometrika* **58**, 341-348.
- D'Agostino, R.B. (1986b). Tests for the Normal Distribution. In: D'Agostino, R.B., and M.A. Stephens, eds. *Goodness-of Fit Techniques*. Marcel Dekker, New York.
- D'Agostino, R.B., and E.S. Pearson (1973). Tests for Departures from Normality. Empirical Results for the Distributions of b_2 and $\sqrt{b_1}$. *Biometrika* **60**(3), 613-622.
- D'Agostino, R.B., and G.L. Tietjen (1973). Approaches to the Null Distribution of $\sqrt{b_1}$. *Biometrika* **60**(1), 169-173.
- Fisher, R.A. (1950). *Statistical Methods for Research Workers*. 11'th Edition. Hafner Publishing Company, New York, pp.99-100.
- Gibbons, R.D., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*, Second Edition. John Wiley & Sons, Hoboken.
- Kendall, M.G., and A. Stuart. (1991). *The Advanced Theory of Statistics, Volume 2: Inference and Relationship*. Fifth Edition. Oxford University Press, New York.
- Kim, P.J., and R.I. Jennrich. (1973). Tables of the Exact Sampling Distribution of the Two Sample Kolmogorov-Smirnov Criterion. In Harter, H.L., and D.B. Owen, eds. *Selected Tables in Mathematical Statistics, Vol. 1*. American Mathematical Society, Providence, Rhode Island, pp.79-170.
- Kolmogorov, A.N. (1933). Sulla determinazione empirica di una legge di distribuzione. *Giornale dell' Istituto Italiano degli Attuari* **4**, 83-91.
- Marsaglia, G., W.W. Tsang, and J. Wang. (2003). Evaluating Kolmogorov's distribution. *Journal of Statistical Software*, **8**(18). <http://www.jstatsoft.org/v08/i18/>.
- Moore, D.S. (1986). Tests of Chi-Squared Type. In D'Agostino, R.B., and M.A. Stephens, eds. *Goodness-of Fit Techniques*. Marcel Dekker, New York, pp.63-95.
- Pomeranz, J. (1973). Exact Cumulative Distribution of the Kolmogorov-Smirnov Statistic for Small Samples (Algorithm 487). *Collected Algorithms from ACM* ??, ??-??.
- Royston, J.P. (1992a). Approximating the Shapiro-Wilk W-Test for Non-Normality. *Statistics and Computing* **2**, 117-119.
- Royston, J.P. (1992b). Estimation, Reference Ranges and Goodness of Fit for the Three-Parameter Log-Normal Distribution. *Statistics in Medicine* **11**, 897-912.
- Royston, J.P. (1992c). A Pocket-Calculator Algorithm for the Shapiro-Francia Test of Non-Normality: An Application to Medicine. *Statistics in Medicine* **12**, 181-184.
- Royston, P. (1993). A Toolkit for Testing for Non-Normality in Complete and Censored Samples. *The Statistician* **42**, 37-43.
- Ryan, T., and B. Joiner. (1973). *Normal Probability Plots and Tests for Normality*. Technical Report, Pennsylvania State University, Department of Statistics.
- Shapiro, S.S., and R.S. Francia. (1972). An Approximate Analysis of Variance Test for Normality. *Journal of the American Statistical Association* **67**(337), 215-219.

- Shapiro, S.S., and M.B. Wilk. (1965). An Analysis of Variance Test for Normality (Complete Samples). *Biometrika* **52**, 591-611.
- Smirnov, N.V. (1939). Estimate of Deviation Between Empirical Distribution Functions in Two Independent Samples. *Bulletin Moscow University* **2**(2), 3-16.
- Smirnov, N.V. (1948). Table for Estimating the Goodness of Fit of Empirical Distributions. *Annals of Mathematical Statistics* **19**, 279-281.
- Stephens, M.A. (1970). Use of the Kolmogorov-Smirnov, Cramer-von Mises and Related Statistics Without Extensive Tables. *Journal of the Royal Statistical Society, Series B*, **32**, 115-122.
- Stephens, M.A. (1986a). Tests Based on EDF Statistics. In D'Agostino, R. B., and M.A. Stevens, eds. *Goodness-of-Fit Techniques*. Marcel Dekker, New York.
- Verrill, S., and R.A. Johnson. (1987). The Asymptotic Equivalence of Some Modified Shapiro-Wilk Statistics – Complete and Censored Sample Cases. *The Annals of Statistics* **15**(1), 413-419.
- Verrill, S., and R.A. Johnson. (1988). Tables and Large-Sample Distribution Theory for Censored-Data Correlation Statistics for Testing Normality. *Journal of the American Statistical Association* **83**, 1192-1197.
- Weisberg, S., and C. Bingham. (1975). An Approximate Analysis of Variance Test for Non-Normality Suitable for Machine Calculation. *Technometrics* **17**, 133-134.
- Wilk, M.B., and S.S. Shapiro. (1968). The Joint Assessment of Normality of Several Independent Samples. *Technometrics*, **10**(4), 825-839.
- Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ.

See Also

[gof.object](#), [print.gof](#), [plot.gof](#), [shapiro.test](#), [ks.test](#), [chisq.test](#), [Normal](#), [Lognormal](#), [Lognormal3](#), [Zero-Modified Normal](#), [Zero-Modified Lognormal \(Delta\)](#), [enorm](#), [elnorm](#), [elnormAlt](#), [elnorm3](#), [ezmnorm](#), [ezmlnorm](#), [ezmlnormAlt](#), [qqPlot](#).

Examples

```
# Generate 20 observations from a gamma distribution with
# parameters shape = 2 and scale = 3 then run various
# goodness-of-fit tests.
# (Note: the call to set.seed lets you reproduce this example.)

set.seed(47)
dat <- rgamma(20, shape = 2, scale = 3)

# Shapiro-Wilk generalized goodness-of-fit test
#-----
gof.list <- gofTest(dat, distribution = "gamma")
gof.list

#Results of Goodness-of-Fit Test
#-----
#
#Test Method:                Shapiro-Wilk GOF Based on
#                           Chen & Balakrisnan (1995)
#
#Hypothesized Distribution:   Gamma
#
#Estimated Parameter(s):     shape = 1.909462
#                           scale = 4.056819
```

```

#
#Estimation Method:      mle
#
#Data:                   dat
#
#Sample Size:            20
#
#Test Statistic:         W = 0.9834958
#
#Test Statistic Parameter:  n = 20
#
#P-value:                 0.970903
#
#Alternative Hypothesis:  True cdf does not equal the
#                          Gamma Distribution.

windows()
plot(gof.list)

#-----

# Redo the example above, but use the bias-corrected mle

gofTest(dat, distribution = "gamma",
        est.arg.list = list(method = "bcmle"))

#Results of Goodness-of-Fit Test
#-----
#
#Test Method:            Shapiro-Wilk GOF Based on
#                          Chen & Balakrisnan (1995)
#
#Hypothesized Distribution:  Gamma
#
#Estimated Parameter(s):   shape = 1.656376
#                          scale = 4.676680
#
#Estimation Method:       bcmle
#
#Data:                    dat
#
#Sample Size:             20
#
#Test Statistic:          W = 0.9834346
#
#Test Statistic Parameter:  n = 20
#
#P-value:                  0.9704046
#
#Alternative Hypothesis:   True cdf does not equal the
#                          Gamma Distribution.

#-----

# Komogorov-Smirnov goodness-of-fit test (pre-specified parameters)
#-----

```



```

# Clean up

rm(dat, gof.list)
graphics.off()

#-----

# Example 10-2 of USEPA (2009, page 10-14) gives an example of
# using the Shapiro-Wilk test to test the assumption of normality
# for nickel concentrations (ppb) in groundwater collected over
# 4 years. The data for this example are stored in
# EPA.09.Ex.10.1.nickel.df.

EPA.09.Ex.10.1.nickel.df
#   Month   Well Nickel.ppb
#1      1 Well.1      58.8
#2      3 Well.1       1.0
#3      6 Well.1     262.0
#4      8 Well.1      56.0
#5     10 Well.1       8.7
#6      1 Well.2      19.0
#7      3 Well.2     81.5
#8      6 Well.2    331.0
#9      8 Well.2      14.0
#10     10 Well.2     64.4
#11     1 Well.3     39.0
#12     3 Well.3    151.0
#13     6 Well.3     27.0
#14     8 Well.3     21.4
#15     10 Well.3    578.0
#16     1 Well.4       3.1
#17     3 Well.4    942.0
#18     6 Well.4     85.6
#19     8 Well.4     10.0
#20     10 Well.4    637.0

# Test for a normal distribution:
#-----

gof.list <- gofTest(Nickel.ppb ~ 1,
  data = EPA.09.Ex.10.1.nickel.df)
gof.list

#Results of Goodness-of-Fit Test
#-----
#
#Test Method:                Shapiro-Wilk GOF
#
#Hypothesized Distribution:   Normal
#
#Estimated Parameter(s):     mean = 169.5250
#                             sd   = 259.7175
#
#Estimation Method:          mvue
#
#Data:                        Nickel.ppb
#

```

```

#Data Source:          EPA.09.Ex.10.1.nickel.df
#
#Sample Size:          20
#
#Test Statistic:        W = 0.6788888
#
#Test Statistic Parameter:  n = 20
#
#P-value:              2.17927e-05
#
#Alternative Hypothesis: True cdf does not equal the
#                        Normal Distribution.

```

```

windows()
plot(gof.list)

```

```

#-----

```

```

# Test for a lognormal distribution:
#-----

```

```

gofTest(Nickel.ppb ~ 1,
  data = EPA.09.Ex.10.1.nickel.df,
  dist = "lnorm")

```

```

#Results of Goodness-of-Fit Test
#-----

```

```

#
#Test Method:          Shapiro-Wilk GOF
#
#Hypothesized Distribution: Lognormal
#
#Estimated Parameter(s): meanlog = 3.918529
#                        sdlog  = 1.801404
#
#Estimation Method:    mvue
#
#Data:                 Nickel.ppb
#
#Data Source:          EPA.09.Ex.10.1.nickel.df
#
#Sample Size:          20
#
#Test Statistic:        W = 0.978946
#
#Test Statistic Parameter:  n = 20
#
#P-value:              0.9197735
#
#Alternative Hypothesis: True cdf does not equal the
#                        Lognormal Distribution.

```

```

#-----

```

```

# Test for a lognormal distribution, but use the
# Mean and CV parameterization:
#-----

```

```

gofTest(Nickel.ppb ~ 1,
  data = EPA.09.Ex.10.1.nickel.df,
  dist = "lnormAlt")

#Results of Goodness-of-Fit Test
#-----
#
#Test Method:                Shapiro-Wilk GOF
#
#Hypothesized Distribution:   Lognormal
#
#Estimated Parameter(s):     mean = 213.415628
#                             cv   = 2.809377
#
#Estimation Method:          mvue
#
#Data:                        Nickel.ppb
#
#Data Source:                 EPA.09.Ex.10.1.nickel.df
#
#Sample Size:                 20
#
#Test Statistic:              W = 0.978946
#
#Test Statistic Parameter:    n = 20
#
#P-value:                     0.9197735
#
#Alternative Hypothesis:      True cdf does not equal the
#                             Lognormal Distribution.

#-----
# Clean up

rm(gof.list)
graphics.off()

#-----

# Generate 20 observations from a normal distribution with mean=3 and sd=2, and
# generate 10 observaions from a normal distribution with mean=2 and sd=2 then
# test whether these sets of observations come from the same distribution.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(300)
dat1 <- rnorm(20, mean = 3, sd = 2)
dat2 <- rnorm(10, mean = 1, sd = 2)
gofTest(x = dat1, y = dat2, test = "ks")

#Results of Goodness-of-Fit Test
#-----
#
#Test Method:                2-Sample K-S GOF
#
#Hypothesized Distribution:   Equal
#

```

```

#Data:                x = dat1
#                    y = dat2
#
#Sample Sizes:        n.x = 20
#                    n.y = 10
#
#Test Statistic:       ks = 0.7
#
#Test Statistic Parameters:  n = 20
#                          m = 10
#
#P-value:              0.001669561
#
#Alternative Hypothesis: The cdf of 'dat1' does not equal
#                      the cdf of 'dat2'.

#-----
# Clean up

rm(dat1, dat2)

```

gofTwoSample.object *S3 Class "gofTwoSample"*

Description

Objects of S3 class "gofTwoSample" are returned by the **EnvStats** function [gofTest](#) when both the `x` and `y` arguments are supplied.

Details

Objects of S3 class "gofTwoSample" are lists that contain information about the assumed distribution, the estimated or user-supplied distribution parameters, and the test statistic and p-value.

Value

Required Components

The following components must be included in a legitimate list of class "gofTwoSample".

<code>distribution</code>	a character string with the value "Equal".
<code>statistic</code>	a numeric scalar with a <code>names</code> attribute containing the name and value of the goodness-of-fit statistic.
<code>sample.size</code>	a numeric scalar containing the number of non-missing observations in the sample used for the goodness-of-fit test.
<code>parameters</code>	numeric vector with a <code>names</code> attribute containing the name(s) and value(s) of the parameter(s) associated with the test statistic given in the <code>statistic</code> component.
<code>p.value</code>	numeric scalar containing the p-value associated with the goodness-of-fit statistic.
<code>alternative</code>	character string indicating the alternative hypothesis.
<code>method</code>	character string indicating the name of the goodness-of-fit test.

data	a list of length 2 containing the numeric vectors actually used for the goodness-of-fit test (i.e., the original data but with any missing or infinite values removed).
data.name	a character vector of length 2 indicating the name of the data object used for the x argument and the name of the data object used for the y argument in the goodness-of-fit test.

Optional Component

The following component is included when the arguments x and/or y contain missing (NA), undefined (NaN) and/or infinite (Inf, -Inf) values.

bad.obs	numeric vector of length 2 indicating the number of missing (NA), undefined (NaN) and/or infinite (Inf, -Inf) values that were removed from the data in the x and y arguments prior to performing the goodness-of-fit test.
---------	---

Methods

Generic functions that have methods for objects of class "gofTwoSample" include: [print](#), [plot](#).

Note

Since objects of class "gofTwoSample" are lists, you may extract their components with the \$ and [[operators.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

See Also

[print.gofTwoSample](#), [plot.gofTwoSample](#), [Goodness-of-Fit Tests](#).

Examples

```
# Create an object of class "gofTwoSample", then print it out.

# Generate 20 observations from a normal distribution with mean=3 and sd=2, and
# generate 10 observations from a normal distribution with mean=2 and sd=2 then
# test whether these sets of observations come from the same distribution.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(300)
dat1 <- rnorm(20, mean = 3, sd = 2)
dat2 <- rnorm(10, mean = 1, sd = 2)
gofTest(x = dat1, y = dat2, test = "ks")

#Results of Goodness-of-Fit Test
#-----
#
#Test Method:                2-Sample K-S GOF
#
#Hypothesized Distribution:   Equal
#
#Data:                      x = dat1
#                           y = dat2
```



```

#
#Sample Sizes:                n.x = 20
#                             n.y = 10
#
#Test Statistic:              ks = 0.7
#
#Test Statistic Parameters:    n = 20
#                             m = 10
#
#P-value:                     0.001669561
#
#Alternative Hypothesis:       The cdf of 'dat1' does not equal
#                             the cdf of 'dat2'.

#-----
# Clean up

rm(dat1, dat2)

```

gpqCiNormCensored	<i>Generalized Pivotal Quantity for Confidence Interval for the Mean of a Normal Distribution Based on Censored Data</i>
-------------------	--

Description

Generate a generalized pivotal quantity (GPQ) for a confidence interval for the mean of a [Normal distribution](#) based on singly or multiply censored data.

Usage

```

gpqCiNormSinglyCensored(n, n.cen, probs, nmc, method = "mle",
  censoring.side = "left", seed = NULL, names = TRUE)

gpqCiNormMultiplyCensored(n, cen.index, probs, nmc, method = "mle",
  censoring.side = "left", seed = NULL, names = TRUE)

```

Arguments

n	positive integer ≥ 3 indicating the sample size.
n.cen	for the case of singly censored data, a positive integer indicating the number of censored observations. The value of n.cen must be between 1 and n-2, inclusive.
cen.index	for the case of multiply censored data, a sorted vector of unique integers indicating the indices of the censored observations when the observations are “ordered”. The length of cen.index must be between 1 and n-2, inclusive, and the values of cen.index must be between 1 and n.
probs	numeric vector of values between 0 and 1 indicating the confidence level(s) associated with the GPQ(s).
nmc	positive integer ≥ 10 indicating the number of Monte Carlo trials to run in order to compute the GPQ(s).

method	character string indicating the method to use for parameter estimation. For singly censored data, possible values are "mle" (the default), "bcmle", "qq.reg", "qq.reg.w.cen.level", "impute.w.qq.reg", "impute.w.qq.reg.w.cen.level", "impute.w.mle", "iterative.impute.w.qq.reg", "m.est", and "half.cen.level". See the help file for enormCensored for details. For multiply censored data, possible values are "mle" (the default), "qq.reg", "impute.w.qq.reg", and "half.cen.level". See the help file for enormCensored for details.
censoring.side	character string indicating on which side the censoring occurs. The possible values are "left" (the default) and "right".
seed	positive integer to pass to the function set.seed . This argument is ignored if seed=NULL (the default). Using the seed argument lets you reproduce the exact same result if all other arguments stay the same.
names	a logical scalar passed to quantile indicating whether to add a names attribute to the resulting GPQ(s). The default value is names=TRUE.

Details

The functions `gpqCiNormSinglyCensored` and `gpqCiNormMultiplyCensored` are called by [enormCensored](#) when `ci.method="gpq"`. They are used to construct generalized pivotal quantities to create confidence intervals for the mean μ of an assumed normal distribution.

This idea was introduced by Schmee et al. (1985) in the context of Type II singly censored data. The function `gpqCiNormSinglyCensored` generates GPQs using a modification of Algorithm 12.1 of Krishnamoorthy and Mathew (2009, p. 329). Algorithm 12.1 is used to generate GPQs for a tolerance interval. The modified algorithm for generating GPQs for confidence intervals for the mean μ is as follows:

1. Generate a random sample of n observations from a standard normal (i.e., $N(0,1)$) distribution and let $z_{(1)}, z_{(2)}, \dots, z_{(n)}$ denote the ordered (sorted) observations.
2. Set the smallest $n.cen$ observations as censored.
3. Compute the estimates of μ and σ by calling [enormCensored](#) using the method specified by the `method` argument, and denote these estimates as $\hat{\mu}^*, \hat{\sigma}^*$.
4. Compute the t-like pivotal quantity $\hat{t} = \hat{\mu}^* / \hat{\sigma}^*$.
5. Repeat steps 1-4 nmc times to produce an empirical distribution of the t-like pivotal quantity.

A two-sided $(1 - \alpha)100\%$ confidence interval for μ is then computed as:

$$[\hat{\mu} - \hat{t}_{1-(\alpha/2)}\hat{\sigma}, \hat{\mu} - \hat{t}_{\alpha/2}\hat{\sigma}]$$

where \hat{t}_p denotes the p 'th empirical quantile of the nmc generated \hat{t} values.

Schmee et al. (1985) derived this method in the context of Type II singly censored data (for which these limits are exact within Monte Carlo error), but state that according to Regal (1982) this method produces confidence intervals that are close approximations to the correct limits for Type I censored data.

The function `gpqCiNormMultiplyCensored` is an extension of this idea to multiply censored data. The algorithm is the same as for singly censored data, except Step 2 changes to:

2. Set observations as censored for elements of the argument `cen.index` that have the value TRUE.

The functions `gpqCiNormSinglyCensored` and `gpqCiNormMultiplyCensored` are computationally intensive and provided to the user to allow you to create your own tables.

Value

a numeric vector containing the GPQ(s).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Krishnamoorthy K., and T. Mathew. (2009). *Statistical Tolerance Regions: Theory, Applications, and Computation*. John Wiley and Sons, Hoboken.

Regal, R. (1982). Applying Order Statistic Censored Normal Confidence Intervals to Time Censored Data. Unpublished manuscript, University of Minnesota, Duluth, Department of Mathematical Sciences.

Schmee, J., D.Gladstein, and W. Nelson. (1985). Confidence Limits for Parameters of a Normal Distribution from Singly Censored Samples, Using Maximum Likelihood. *Technometrics* **27**(2) 119–128.

See Also

[enormCensored](#), [estimateCensored.object](#).

Examples

```
# Reproduce the entries for n=10 observations with n.cen=6 in Table 4
# of Schmee et al. (1985, p.122).
#
# Notes:
# 1. This table applies to right-censored data, and the
#    quantity "r" in this table refers to the number of
#    uncensored observations.
#
# 2. Passing a value for the argument "seed" simply allows
#    you to reproduce this example.
#
# Here we will specify just 1,000 Monte Carlos, whereas
# Krishnamoorthy and Mathew (2009) suggest 10,000.

# Here are the values given in Schmee et al. (1985):
Schmee.values <- c(-3.59, -2.60, -1.73, -0.24, 0.43, 0.58, 0.73)
probs = c(0.025, 0.05, 0.1, 0.5, 0.9, 0.95, 0.975)
names(Schmee.values) <- paste(probs * 100, "%", sep = "")

Schmee.values
# 2.5%   5%   10%   50%   90%   95% 97.5%
#-3.59 -2.60 -1.73 -0.24  0.43  0.58  0.73

gpqs <- gpqCiNormSinglyCensored(n = 10, n.cen = 6,
  probs = c(0.025, 0.05, 0.1, 0.5, 0.9, 0.95, 0.975),
  nmc = 1000, censoring.side = "right", seed = 529)

round(gpqs, 2)
# 2.5%   5%   10%   50%   90%   95% 97.5%
#-3.50 -2.49 -1.67 -0.25  0.41  0.57  0.71
```

```

# Clean up
#-----
rm(Schmee.values, probs, gpqs)

#-----

# Example of using gpqCiNormMultiplyCensored
#-----

# Consider the following set of multiply left-censored data:
dat <- 12:16
censored <- c(TRUE, FALSE, TRUE, FALSE, FALSE)

# Since the data are "ordered" we can identify the indices of the
# censored observations in the ordered data as follow:

cen.index <- (1:length(dat))[censored]
cen.index
#[1] 1 3

# Now we can generate a GPQ using gpqCiNormMultiplyCensored.
# Here we'll generate a GPQs to use to create a
# 95% confidence interval for left-censored data.

gpqCiNormMultiplyCensored(n = 5, cen.index = cen.index,
  probs = c(0.025, 0.975), nmc = 1000, seed = 237)
#      2.5%      97.5%
#-1.303667  1.936694

# Clean up
#-----
rm(dat, censored, cen.index)

```

gpqTolIntNormCensored *Generalized Pivotal Quantity for Tolerance Interval for a Normal Distribution Based on Censored Data*

Description

Generate a generalized pivotal quantity (GPQ) for a tolerance interval for a Normal distribution based on singly or multiply censored data.

Usage

```

gpqTolIntNormSinglyCensored(n, n.cen, p, probs, nmc, method = "mle",
  censoring.side = "left", seed = NULL, names = TRUE)

gpqTolIntNormMultiplyCensored(n, cen.index, p, probs, nmc, method = "mle",
  censoring.side = "left", seed = NULL, names = TRUE)

```

Arguments

n positive integer ≥ 3 indicating the sample size.

n.cen	for the case of singly censored data, a positive integer indicating the number of censored observations. The value of n.cen must be between 1 and n-2, inclusive.
cen.index	for the case of multiply censored data, a sorted vector of unique integers indicating the indices of the censored observations when the observations are “ordered”. The length of cen.index must be between 1 and n-2, inclusive, and the values of cen.index must be between 1 and n.
p	numeric scalar strictly greater than 0 and strictly less than 1 indicating the quantile for which to generate the GPQ(s) (i.e., the coverage associated with a one-sided tolerance interval).
probs	numeric vector of values between 0 and 1 indicating the confidence level(s) associated with the GPQ(s).
nmc	positive integer ≥ 10 indicating the number of Monte Carlo trials to run in order to compute the GPQ(s).
method	character string indicating the method to use for parameter estimation.

For singly censored data, possible values are "mle" (the default), "bcmle", "qq.reg", "qq.reg.w.cen.level", "impute.w.qq.reg", "impute.w.qq.reg.w.cen.level", "impute.w.mle", "iterative.impute.w.qq.reg", "m.est", and "half.cen.level". See the help file for [enormCensored](#) for details.

For multiply censored data, possible values are "mle" (the default), "qq.reg", "impute.w.qq.reg", and "half.cen.level". See the help file for [enormCensored](#) for details.

censoring.side	character string indicating on which side the censoring occurs. The possible values are "left" (the default) and "right".
seed	positive integer to pass to the function set.seed . This argument is ignored if seed=NULL (the default). Using the seed argument lets you reproduce the exact same result if all other arguments stay the same.
names	a logical scalar passed to quantile indicating whether to add a names attribute to the resulting GPQ(s). The default value is names=TRUE.

Details

The function `gpqTolIntNormSinglyCensored` generates GPQs as described in Algorithm 12.1 of Krishnamoorthy and Mathew (2009, p. 329). The function `gpqTolIntNormMultiplyCensored` is an extension of this idea to multiply censored data. These functions are called by [tolIntNormCensored](#) when `ti.method="gpq"`, and also by [eqnormCensored](#) when `ci=TRUE` and `ci.method="gpq"`. See the help files for these functions for an explanation of GPQs.

Note that technically these are only GPQs if the data are Type II censored. However, Krishnamoorthy and Mathew (2009, p. 328) state that in the case of Type I censored data these quantities should approximate the true GPQs and the results appear to be satisfactory, even for small sample sizes.

The functions `gpqTolIntNormSinglyCensored` and `gpqTolIntNormMultiplyCensored` are computationally intensive and provided to the user to allow you to create your own tables.

Value

a numeric vector containing the GPQ(s).

Note

Tolerance intervals have long been applied to quality control and life testing problems (Hahn, 1970b,c; Hahn and Meeker, 1991; Krishnamoorthy and Mathew, 2009). References that discuss tolerance intervals in the context of environmental monitoring include: Berthouex and Brown (2002, Chapter 21), Gibbons et al. (2009), Millard and Neerchal (2001, Chapter 6), Singh et al. (2010b), and USEPA (2009).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Krishnamoorthy K., and T. Mathew. (2009). *Statistical Tolerance Regions: Theory, Applications, and Computation*. John Wiley and Sons, Hoboken.

See Also

[tolIntNormCensored](#), [eqnormCensored](#), [enormCensored](#), [estimateCensored.object](#).

Examples

```
# Reproduce the entries for n=10 observations with n.cen=1 in Table 12.2
# of Krishnamoorthy and Mathew (2009, p.331).
#
# (Note: passing a value for the argument "seed" simply allows you to
# reproduce this example.)
#
# Here we will specify just 1,000 Monte Carlos, whereas
# Krishnamoorthy and Mathew (2009) suggest 10,000.

gpqTolIntNormSinglyCensored(n = 10, n.cen = 1, p = 0.05, probs = 0.05,
  nmc = 1000, seed = 529)
#      5%
#-3.316002

gpqTolIntNormSinglyCensored(n = 10, n.cen = 1, p = 0.1, probs = 0.05,
  nmc = 1000, seed = 497)
#      5%
#-2.590449

gpqTolIntNormSinglyCensored(n = 10, n.cen = 1, p = 0.9, probs = 0.95,
  nmc = 1000, seed = 623)
#     95%
#2.547389

gpqTolIntNormSinglyCensored(n = 10, n.cen = 1, p = 0.95, probs = 0.95,
  nmc = 1000, seed = 623)
#     95%
#3.152924

#-----

# Example of using gpqTolIntNormMultiplyCensored
#-----
```

```

# Consider the following set of multiply left-censored data:
dat <- 12:16
censored <- c(TRUE, FALSE, TRUE, FALSE, FALSE)

# Since the data are "ordered" we can identify the indices of the
# censored observations in the ordered data as follow:

cen.index <- (1:length(dat))[censored]
cen.index
#[1] 1 3

# Now we can generate a GPQ using gpqTolIntNormMultiplyCensored.
# Here we'll generate a GPQ corresponding to an upper tolerance
# interval with coverage 90% with 95% confidence for
# left-censored data.

gpqTolIntNormMultiplyCensored(n = 5, cen.index = cen.index, p = 0.9,
  probs = 0.95, nmc = 1000, seed = 237)
#      95%
#4.093022

# Clean up
#-----

rm(dat, censored, cen.index)

```

Graham.et.al.75.etu.df

Ethylene Thiourea Dose-Response Data

Description

These data are the results of an experiment in which different groups of rats were exposed to different concentration levels of ethylene thiourea (ETU), which is a decomposition product of a certain class of fungicides that can be found in treated foods (Graham et al., 1975; Rodricks, 1992, p.133). In this experiment, the outcome of concern was the number of rats that developed thyroid tumors.

Usage

Graham.et.al.75.etu.df

Format

A data frame with 6 observations on the following 4 variables.

dose a numeric vector of dose (ppm/day) of ETU.

tumors a numeric vector indicating number of rats that developed thyroid tumors.

n a numeric vector indicating the number of rats in the dose group.

proportion a numeric vector indicating proportion of rats that developed thyroid tumors.

Source

Graham, S.L., K.J. Davis, W.H. Hansen, and C.H. Graham. (1975). Effects of Prolonged Ethylene Thiourea Ingestion on the Thyroid of the Rat. *Food and Cosmetics Toxicology*, **13**(5), 493–499.

References

Rodricks, J.V. (1992). *Calculated Risks: The Toxicity and Human Health Risks of Chemicals in Our Environment*. Cambridge University Press, New York, p.133.

Helsel.Cohn.88.app.b.df

Example of Multiply Left-censored Data from Literature

Description

Made up multiply left-censored data. There are 9 observations out of a total of 18 that are reported as $<DL$, where DL denotes a detection limit. There are 2 distinct detection limits.

Usage

Helsel.Cohn.88.app.b.df

Format

A data frame with 18 observations on the following 3 variables.

Conc.orig a character vector of original observations

Conc a numeric vector of observations with censored values coded to censoring levels

Censored a logical vector indicating which values are censored

Source

Helsel, D.R., and T.A. Cohn. (1988). Estimation of Descriptive Statistics for Multiply Censored Water Quality Data. *Water Resources Research* **24**(12), 1997–2004, Appendix B.

Helsel.Cohn.88.silver.df

Silver Concentrations From An Interlab Comparison

Description

Silver concentrations (mg/L) from an interlab comparison. There are 34 observations out of a total of 56 that are reported as $<DL$, where DL denotes a detection limit. There are 12 distinct detection limits.

Usage

Helsel.Cohn.88.silver.df

Format

A data frame with 56 observations on the following 4 variables.

Ag.orig a character vector of original silver concentrations (mg/L)

Ag a numeric vector with nondetects coded to the detection limit

Censored a logical vector indicating which observations are censored

log.Ag the natural logarithm of Ag

Source

Helsel, D.R., and T.A. Cohn. (1988). Estimation of Descriptive Statistics for Multiply Censored Water Quality Data. *Water Resources Research* **24**(12), 1997–2004.

References

Janzer, V.J. (1986). *Report of the U.S. Geological Survey's Analytical Evaluation Program—Standard Reference Water Samples M6, M94, T95, N16, P8, and SED3*. Technical Report, Branch of Quality Assurance, U.S. Geological Survey, Arvada, CO.

HoskingEtAl1985

Abstract: Hosking et al. (1985)

Description

Detailed abstract of the manuscript:

Hosking, J.R.M., J.R. Wallis, and E.F. Wood. (1985). Estimation of the Generalized Extreme-Value Distribution by the Method of Probability-Weighted Moments. *Technometrics* **27**(3), 251–261.

Details**Abstract**

Hosking et al. (1985) use the method of probability-weighted moments, introduced by Greenwood et al. (1979), to estimate the parameters of the [generalized extreme value distribution](#) (GEVD) with parameters location= η , scale= θ , and shape= κ . Hosking et al. (1985) derive the asymptotic distributions of the probability-weighted moment estimators (PWME), and compare the asymptotic and small-sample statistical properties (via computer simulation) of the PWME with maximum likelihood estimators (MLE) and Jenkinson's (1969) method of sextiles estimators (JSE). They also compare the statistical properties of quantile estimators (which are based on the distribution parameter estimators). Finally, they derive a test of the null hypothesis that the shape parameter is zero, and assess its performance via computer simulation.

Hosking et al. (1985) note that when $\kappa \leq -1$, the moments and probability-weighted moments of the GEVD do not exist. They also note that in practice the shape parameter usually lies between $-1/2$ and $1/2$.

Hosking et al. (1985) found that the asymptotic efficiency of the PWME (the limit as the sample size approaches infinity of the ratio of the variance of the MLE divided by the variance of the PWME) tends to 0 as the shape parameter approaches $1/2$ or $-1/2$. For values of κ within the range $[-0.2, 0.2]$, however, the efficiency of the estimator of location is close to 100 and greater than 70. Hosking et al. (1985) found that the asymptotic efficiency of the PWME is poor for κ outside the range $[-0.2, 0.2]$.

For the small sample results, Hosking et al. (1985) considered several possible forms of the PWME (see equations (8)-(10) below). The best overall results were given by the plotting-position PWME defined by equations (9) and (10) with $a = 0.35$ and $b = 0$.

Small sample results for estimating the parameters show that for $n \geq 50$ all three methods give almost identical results. For $n < 50$ the results for the different estimators are a bit different, but not dramatically so. The MLE tends to be slightly less biased than the other two methods. For estimating the shape parameter, the MLE has a slightly larger standard deviation, and the PWME has consistently the smallest standard deviation.

Small sample results for estimating large quantiles show that for $n \geq 100$ all three methods are comparable. For $n < 100$ the PWME and JSE are comparable and in general have much smaller standard deviations than the MLE. All three methods are very inaccurate for estimating large quantiles in small samples, especially when $\kappa < 0$.

Hosking et al. (1985) derive a test of the null hypothesis $H_0 : \kappa = 0$ based on the PWME of κ . The test is performed by computing the statistic:

$$z = \frac{\hat{\kappa}_{pwme}}{\sqrt{0.5663/n}} \quad (1)$$

and comparing z to a standard normal distribution (see [zTestGevdShape](#)). Based on computer simulations using the plotting-position PWME, they found that a sample size of $n \geq 25$ ensures an adequate normal approximation. They also found this test has power comparable to the modified likelihood-ratio test, which was found by Hosking (1984) to be the best overall test of $H_0 : \kappa = 0$ of the thirteen tests he considered.

More Details

Probability-Weighted Moments and Parameters of the GEVD

The definition of a [probability-weighted moment](#), introduced by Greenwood et al. (1979), is as follows. Let X denote a random variable with cdf F , and let $x(p)$ denote the p 'th quantile of the distribution. Then the ijk 'th probability-weighted moment is given by:

$$M(i, j, k) = E[X^i F^j (1 - F)^k] = \int_0^1 [x(F)]^i F^j (1 - F)^k dF \quad (2)$$

where i, j , and k are real numbers.

Hosking et al. (1985) set

$$\beta_j = M(i, j, 0) \quad (3)$$

and Greenwood et al. (1979) show that

$$\beta_j = \frac{1}{j+1} E[X_{j+1:j+1}] \quad (4)$$

where

$$E[X_{j+1:j+1}]$$

denotes the expected value of the $j+1$ 'th order statistic (i.e., the maximum) in a sample of size $j+1$. Hosking et al. (1985) show that if X has a GEVD with parameters location= η , scale= θ , and shape= κ , where $\kappa \neq 0$, then

$$\beta_j = \frac{1}{j+1} \left\{ \eta + \frac{\theta[1 - (j+1)^{-\kappa} \Gamma(1+\kappa)]}{\kappa} \right\} \quad (5)$$

for $\kappa > -1$, where $\Gamma()$ denotes the [gamma function](#). Thus,

$$\beta_0 = \eta + \frac{\theta[1 - \Gamma(1+\kappa)]}{\kappa} \quad (6)$$

$$2\beta_1 - \beta_0 = \frac{\theta[\Gamma(1 + \kappa)](1 - 2^{-\kappa})}{\kappa} \quad (7)$$

$$\frac{3\beta_2 - \beta_0}{2\beta_1 - \beta_0} = \frac{1 - 3^{-\kappa}}{1 - 2^{-\kappa}} \quad (8)$$

Estimating Distribution Parameters

Using the results of Landwehr et al. (1979), Hosking et al. (1985) show that given a random sample of n values from some arbitrary distribution, an unbiased, distribution-free, and parameter-free estimator of the probability-weighted moment $\beta_j = M(i, j, 0)$ defined above is given by:

$$b_j = \frac{1}{n} \sum_{i=j+1}^n x_{i,n} \frac{\binom{i-1}{j}}{\binom{n-1}{j}} \quad (9)$$

where the quantity $x_{i,n}$ denotes the i 'th order statistic in the random sample of size n . Hosking et al. (1985) note that this estimator is closely related to U-statistics (Hoeffding, 1948; Lehmann, 1975, pp. 362-371).

An alternative "plotting position" estimator is given by:

$$\hat{\beta}_j[p_{i,n}] = \frac{1}{n} \sum_{i=1}^n p_{i,n}^j x_{i,n} \quad (10)$$

where

$$p_{i,n} = \hat{F}(x_{i,n}) \quad (11)$$

denotes the plotting position of the i 'th order statistic in the random sample of size n , that is, a distribution-free estimate of the cdf of X evaluated at the i 'th order statistic. Typically, plotting positions have the form:

$$p_{i,n} = \frac{i - a}{n + b} \quad (12)$$

where $b > -a > -1$. For this form of plotting position, the plotting-position estimators in (10) are asymptotically equivalent to the U-statistic estimators in (9).

Although the unbiased and plotting position estimators are asymptotically equivalent (Hosking, 1990), Hosking and Wallis (1995) recommend using the unbiased estimator for almost all applications because of its superior performance in small and moderate samples.

Using equations (6)-(8) above, i.e., the three equations involving β_0 , β_1 , and β_2 , Hosking et al. (1985) define the probability-weighted moment estimators of η , θ , and κ as the solutions to these three simultaneous equations, with the values of the probability-weighted moments replaced by their estimated values (using either the unbiased or plotting position estimators in (9) and (10) above). Hosking et al. (1985) note that the third equation (equation (8)) must be solved iteratively for the PWME of κ . Using the unbiased estimators of the PWMEs to solve for κ , the PWMEs of η and θ are given by:

$$\hat{\eta}_{pwme} = b_0 + \frac{\hat{\theta}_{pwme}[\Gamma(1 + \hat{\kappa}_{pwme}) - 1]}{\hat{\kappa}_{pwme}} \quad (13)$$

$$\hat{\theta}_{pwme} = \frac{(2b_1 - b_0)\hat{\kappa}_{pwme}}{\Gamma(1 + \hat{\kappa}_{pwme})(1 - 2^{-\hat{\kappa}_{pwme}})} \quad (14)$$

Hosking et al. (1985) show that when the unbiased estimates of the PWMEs are used to estimate the probability-weighted moments, the estimates of θ and κ satisfy the feasibility criteria

$$\hat{\theta}_{pwme} > 0; \hat{\kappa}_{pwme} > -1$$

almost surely.

Hosking et al. (1985) show that the asymptotic distribution of the PWME is multivariate normal with mean equal to (η, θ, κ) , and they derive the formula for the asymptotic variance-covariance matrix as:

$$V_{\hat{\eta}, \hat{\theta}, \hat{\kappa}} = \frac{1}{n} G V_{\hat{\beta}_0, \hat{\beta}_1, \hat{\beta}_2} G^T \quad (15)$$

where

$$V_{\hat{\beta}_0, \hat{\beta}_1, \hat{\beta}_2}$$

denotes the variance-covariance matrix of the estimators of the probability-weighted moments defined in either equation (9) or (10) above (recall that these two estimators are asymptotically equivalent), and the matrix G is defined by:

$$G_{i1} = \frac{\partial \eta}{\partial \beta_{i-1}}, G_{i2} = \frac{\partial \theta}{\partial \beta_{i-1}}, G_{i3} = \frac{\partial \kappa}{\partial \beta_{i-1}} \quad (16)$$

for $i = 1, 2, 3$. Hosking et al. (1985) provide formulas for the matrix

$$V_{\hat{\beta}_0, \hat{\beta}_1, \hat{\beta}_2}$$

in Appendix C of their manuscript. Note that there is a typographical error in equation (C.11) (Jon Hosking, personal communication, 1996). In the second line of this equation, the quantity $-(r+s)^{-k}$ should be replaced with $-(r+s)^{-2k}$.

The matrix G in equation (16) is not easily computed. Its inverse, however, is easy to compute and then can be inverted numerically (Jon Hosking, 1996, personal communication). The inverse of G is given by:

$$G_{i1}^{-1} = \frac{\partial \beta_{i-1} \partial \eta}{\partial \beta_{i-1}}, G_{i2}^{-1} = \frac{\partial \beta_{i-1} \partial \theta}{\partial \beta_{i-1}}, G_{i3}^{-1} = \frac{\partial \beta_{i-1} \partial \kappa}{\partial \beta_{i-1}} \quad (17)$$

and by equation (5) above it can be shown that:

$$\frac{\partial \beta_j}{\partial \eta} = \frac{1}{j+1} \quad (18)$$

$$\frac{\partial \beta_j}{\partial \theta} = \frac{1 - (j+1)^{-\kappa} \Gamma(1+\kappa)}{(j+1)\kappa} \quad (19)$$

$$\frac{\partial \beta_j}{\partial \kappa} = \frac{\theta}{j+1} \left\{ \frac{(j+1)^{-\kappa} [\log(j+1) \Gamma(1+\kappa) - \Gamma'(1+\kappa)]}{\kappa} - \frac{1 - (j+1)^{-\kappa} \Gamma(1+\kappa)}{\kappa^2} \right\} \quad (20)$$

for $i = 1, 2, 3$.

Estimating Distribution Quantiles

If X has a GEVD with parameters location= η , scale= θ , and shape= κ , where $\kappa \neq 0$, then the p 'th quantile of the distribution is given by:

$$x(p) = \eta + \frac{\theta \{1 - [-\log(p)]^\kappa\}}{\kappa} \quad (21)$$

($0 \leq p \leq 1$). Given estimated values of the location, scale, and shape parameters, the p 'th quantile of the distribution is estimated as:

$$\hat{x}(p) = \hat{\eta} + \frac{\hat{\theta} \{1 - [-\log(p)]^{\hat{\kappa}}\}}{\hat{\kappa}} \quad (22)$$

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.
- Greenwood, J.A., J.M. Landwehr, N.C. Matalas, and J.R. Wallis. (1979). Probability Weighted Moments: Definition and Relation to Parameters of Several Distributions Expressible in Inverse Form. *Water Resources Research* **15**(5), 1049–1054.
- Hoeffding, W. (1948). A Class of Statistics with Asymptotically Normal Distribution. *Annals of Mathematical Statistics* **19**, 293–325.
- Hosking, J.R.M. (1985). Algorithm AS 215: Maximum-Likelihood Estimation of the Parameters of the Generalized Extreme-Value Distribution. *Applied Statistics* **34**(3), 301–310.
- Hosking, J.R.M. (1990). *L*-Moments: Analysis and Estimation of Distributions Using Linear Combinations of Order Statistics. *Journal of the Royal Statistical Society, Series B* **52**(1), 105–124.
- Hosking, J.R.M., and J.R. Wallis (1995). A Comparison of Unbiased and Plotting-Position Estimators of *L* Moments. *Water Resources Research* **31**(8), 2019–2025.
- Jenkinson, A.F. (1969). *Statistics of Extremes*. *Technical Note 98*, World Meteorological Office, Geneva.
- Johnson, N. L., S. Kotz, and A.W. Kemp. (1992). *Univariate Discrete Distributions*. Second Edition. John Wiley and Sons, New York, pp.4-8.
- Johnson, N. L., S. Kotz, and N. Balakrishnan. (1995). *Continuous Univariate Distributions, Volume 2*. Second Edition. John Wiley and Sons, New York.
- Lehmann, E.L. (1975). *Nonparametrics: Statistical Methods Based on Ranks*. Holden-Day, Oakland, CA, 457pp.

See Also

[Generalized Extreme Value Distribution](#), `egev`.

inversePredictCalibrate

Predict Concentration Using Calibration

Description

Predict concentration using a calibration line (or curve) and inverse regression.

Usage

```
inversePredictCalibrate(object, obs.y = NULL,
  n.points = ifelse(is.null(obs.y), 100, length(obs.y)),
  intervals = FALSE, coverage = 0.99, simultaneous = FALSE,
  individual = FALSE, trace = FALSE)
```

Arguments

<code>object</code>	an object that is the result of calling the function <code>calibrate</code> .
<code>obs.y</code>	optional numeric vector of observed values for the machine signal. The default value is <code>obs.y=NULL</code> , in which case <code>obs.y</code> is set equal to a vector of values (of length <code>n.points</code>) ranging from the minimum to the maximum of the fitted values from the <code>calibrate</code> object.
<code>n.points</code>	optional integer indicating the number of points at which to predict concentrations (i.e., perform inverse regression). The default value is <code>n.points=100</code> . This argument is ignored when <code>obs.y</code> is supplied.
<code>intervals</code>	optional logical scalar indicating whether to compute confidence intervals for the predicted concentrations. The default value is <code>intervals=FALSE</code> .
<code>coverage</code>	optional numeric scalar between 0 and 1 indicating the confidence level associated with the confidence intervals for the predicted concentrations. The default value is <code>coverage=0.99</code> .
<code>simultaneous</code>	optional logical scalar indicating whether to base the confidence intervals for the predicted values on simultaneous or non-simultaneous prediction limits. The default value is <code>simultaneous=FALSE</code> .
<code>individual</code>	optional logical scalar indicating whether to base the confidence intervals for the predicted values on prediction limits for the mean (<code>individual=FALSE</code>) or prediction limits for an individual observation (<code>individual=TRUE</code>). The default value is <code>individual=FALSE</code> .
<code>trace</code>	optional logical scalar indicating whether to print out (trace) the progress of the inverse prediction for each of the specified values of <code>obs.y</code> . The default value is <code>trace=FALSE</code> .

Details

A simple and frequently used calibration model is a straight line where the response variable S denotes the signal of the machine and the predictor variable C denotes the true concentration in the physical sample. The error term is assumed to follow a normal distribution with mean 0. Note that the average value of the signal for a blank ($C = 0$) is the intercept. Other possible calibration models include higher order polynomial models such as a quadratic or cubic model.

In a typical setup, a small number of samples (e.g., $n = 6$) with known concentrations are measured and the signal is recorded. A sample with no chemical in it, called a blank, is also measured. (You have to be careful to define exactly what you mean by a “blank.” A blank could mean a container from the lab that has nothing in it but is prepared in a similar fashion to containers with actual samples in them. Or it could mean a field blank: the container was taken out to the field and subjected to the same process that all other containers were subjected to, except a physical sample of soil or water was not placed in the container.) Usually, replicate measures at the same known concentrations are taken. (The term “replicate” must be well defined to distinguish between for example the same physical samples that are measured more than once vs. two different physical samples of the same known concentration.)

The function `calibrate` initially fits a linear calibration line or curve. Once the calibration line is fit, samples with unknown concentrations are measured and their signals are recorded. In order to produce estimated concentrations, you have to use inverse regression to map the signals to the estimated concentrations. We can quantify the uncertainty in the estimated concentration by combining inverse regression with prediction limits for the signal S .

Value

A numeric matrix containing the results of the inverse calibration. The first two columns are labeled `obs.y` and `pred.x` containing the values of the argument `obs.y` and the predicted values of `x` (the concentration), respectively. If `intervals=TRUE`, then the matrix also contains the columns `lpl.x` and `upl.x` corresponding to the lower and upper prediction limits for `x`. Also, if `intervals=TRUE`, then the matrix has the attributes `coverage` (the value of the argument `coverage`) and `simultaneous` (the value of the argument `simultaneous`).

Note

Almost always the process of determining the concentration of a chemical in a soil, water, or air sample involves using some kind of machine that produces a signal, and this signal is related to the concentration of the chemical in the physical sample. The process of relating the machine signal to the concentration of the chemical is called **calibration** (see [calibrate](#)). Once calibration has been performed, estimated concentrations in physical samples with unknown concentrations are computed using inverse regression. The uncertainty in the process used to estimate the concentration may be quantified with decision, detection, and quantitation limits.

In practice, only the point estimate of concentration is reported (along with a possible qualifier), without confidence bounds for the true concentration C . This is most unfortunate because it gives the impression that there is no error associated with the reported concentration. Indeed, both the International Organization for Standardization (ISO) and the International Union of Pure and Applied Chemistry (IUPAC) recommend always reporting both the estimated concentration and the uncertainty associated with this estimate (Currie, 1997).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Currie, L.A. (1997). Detection: International Update, and Some Emerging Di-Lemmas Involving Calibration, the Blank, and Multiple Detection Decisions. *Chemometrics and Intelligent Laboratory Systems* **37**, 151–181.
- Draper, N., and H. Smith. (1998). *Applied Regression Analysis*. Third Edition. John Wiley and Sons, New York, Chapter 3 and p.335.
- Hubaux, A., and G. Vos. (1970). Decision and Detection Limits for Linear Calibration Curves. *Annals of Chemistry* **42**, 849–855.
- Millard, S.P., and N.K. Neerchal. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, FL, pp.562–575.

See Also

[pointwise](#), [calibrate](#), [detectionLimitCalibrate](#), [lm](#)

Examples

```
# The data frame EPA.97.cadmium.111.df contains calibration data
# for cadmium at mass 111 (ng/L) that appeared in
# Gibbons et al. (1997b) and were provided to them by the U.S. EPA.
# Here we
# 1. Display a plot of these data along with the fitted calibration
#    line and 99% non-simultaneous prediction limits.
```

```

# 2. Then based on an observed signal of 60 from a sample with
#   unknown concentration, we use the calibration line to estimate
#   the true concentration and use the prediction limits to compute
#   confidence bounds for the true concentration.
# An observed signal of 60 results in an estimated value of cadmium
# of 59.97 ng/L and a confidence interval of [53.83, 66.15].
# See Millard and Neerchal (2001, pp.566-569) for more details on
# this example.

Cadmium <- EPA.97.cadmium.111.df$Cadmium

Spike <- EPA.97.cadmium.111.df$Spike

calibrate.list <- calibrate(Cadmium ~ Spike,
  data=EPA.97.cadmium.111.df)

newdata <- data.frame(Spike = seq(min(Spike), max(Spike), len=100))

pred.list <- predict(calibrate.list, newdata=newdata, se.fit=TRUE)

pointwise.list <- pointwise(pred.list, coverage=0.99,
  individual=TRUE)

plot(Spike, Cadmium, ylim=c(min(pointwise.list$lower),
  max(pointwise.list$upper)), xlab="True Concentration (ng/L)",
  ylab="Observed Concentration (ng/L)")

abline(calibrate.list, lwd=2)

lines(newdata$Spike, pointwise.list$lower, lty=8, lwd=2)

lines(newdata$Spike, pointwise.list$upper, lty=8, lwd=2)

title(paste("Calibration Line and 99% Prediction Limits",
  "for US EPA Cadmium 111 Data", sep="\n"))

# Now estimate the true concentration based on
# an observed signal of 60 ng/L.

inversePredictCalibrate(calibrate.list, obs.y=60,
  intervals=TRUE, coverage=0.99, individual=TRUE)

#   obs.y   pred.x   lpl.x   upl.x
#[1,]    60 59.97301 53.8301 66.15422
#attr(, "coverage"):
#[1] 0.99
#attr(, "simultaneous"):
#[1] FALSE

rm(Cadmium, Spike, calibrate.list, newdata, pred.list, pointwise.list)

```


Description

Compute the interquartile range for a set of data.

Usage

```
iqr(x, na.rm = FALSE)
```

Arguments

<code>x</code>	numeric vector of observations.
<code>na.rm</code>	logical scalar indicating whether to remove missing values from <code>x</code> . If <code>na.rm=FALSE</code> (the default) and <code>x</code> contains missing values, then a missing value (NA) is returned. If <code>na.rm=TRUE</code> , missing values are removed from <code>x</code> prior to computing the coefficient of variation.

Details

Let \underline{x} denote a random sample of n observations from some distribution associated with a random variable X . The sample interquartile range is defined as:

$$IQR = \hat{X}_{0.75} - \hat{X}_{0.25} \quad (1)$$

where X_p denotes the p 'th quantile of the distribution and \hat{X}_p denotes the estimate of this quantile (i.e., the sample p 'th quantile).

See the R help file for [quantile](#) for information on how sample quantiles are computed.

Value

A numeric scalar – the interquartile range.

Note

The interquartile range is a robust estimate of the spread of the distribution. It is the distance between the two ends of a boxplot (see the R help file for [boxplot](#)). For a normal distribution with standard deviation σ it can be shown that:

$$IQR = 1.34898\sigma \quad (2)$$

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Chambers, J.M., W.S. Cleveland, B. Kleiner, and P.A. Tukey. (1983). *Graphical Methods for Data Analysis*. Duxbury Press, Boston, MA.
- Cleveland, W.S. (1993). *Visualizing Data*. Hobart Press, Summit, New Jersey.
- Helsel, D.R., and R.M. Hirsch. (1992). *Statistical Methods in Water Resources Research*. Elsevier, New York, NY.
- Hirsch, R.M., D.R. Helsel, T.A. Cohn, and E.J. Gilroy. (1993). Statistical Analysis of Hydrologic Data. In: Maidment, D.R., ed. *Handbook of Hydrology*. McGraw-Hill, New York, Chapter 17, pp.5–7.
- Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ.

See Also

[Summary Statistics](#), [summaryFull](#), [var](#), [sd](#).

Examples

```
# Generate 20 observations from a normal distribution with parameters
# mean=10 and sd=2, and compute the standard deviation and
# interquartile range.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
dat <- rnorm(20, mean=10, sd=2)

sd(dat)
#[1] 1.180226

iqr(dat)
#[1] 1.489932

#-----

# Repeat the last example, but add a couple of large "outliers" to the
# data. Note that the estimated standard deviation is greatly affected
# by the outliers, while the interquartile range is not.

summaryStats(dat, quartiles = TRUE)
#      N   Mean      SD Median   Min     Max 1st Qu. 3rd Qu.
#dat 20  9.8612  1.1802  9.6978  7.6042 11.8756  9.1618 10.6517

new.dat <- c(dat, 20, 50)

sd(dat)
#[1] 1.180226

sd(new.dat)
#[1] 8.79796

iqr(dat)
#[1] 1.489932

iqr(new.dat)
#[1] 1.851472

#-----
# Clean up
rm(dat, new.dat)
```

kendallSeasonalTrendTest

*Nonparametric Test for Monotonic Trend Within Each Season Based
on Kendall's Tau Statistic*

Description

Perform a nonparametric test for a monotonic trend within each season based on Kendall's tau statistic, and optionally compute a confidence interval for the slope across all seasons.

Usage

```
kendallSeasonalTrendTest(y, ...)

## S3 method for class 'formula'
kendallSeasonalTrendTest(y, data = NULL, subset,
  na.action = na.pass, ...)

## Default S3 method:
kendallSeasonalTrendTest(y, season, year,
  alternative = "two.sided", correct = TRUE, ci.slope = TRUE, conf.level = 0.95,
  independent.obs = TRUE, data.name = NULL, season.name = NULL, year.name = NULL,
  parent.of.data = NULL, subset.expression = NULL, ...)

## S3 method for class 'data.frame'
kendallSeasonalTrendTest(y, ...)

## S3 method for class 'matrix'
kendallSeasonalTrendTest(y, ...)
```

Arguments

<code>y</code>	an object containing data for the trend test. In the default method, the argument <code>y</code> must be numeric vector of observations. When <code>y</code> is a data frame, all columns must be numeric. When <code>y</code> is a matrix, it must be a numeric matrix. In the formula method, <code>y</code> must be a formula of the form <code>y ~ season + year</code> , where <code>y</code> , <code>season</code> , and <code>year</code> specify what variables to use for these arguments in the call to <code>kendallSeasonalTrendTest.default</code> . Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
<code>data</code>	specifies an optional data frame, list or environment (or object coercible by <code>as.data.frame</code> to a data frame) containing the variables in the model. If not found in <code>data</code> , the variables are taken from <code>environment(formula)</code> , typically the environment from which <code>kendallTrendTest</code> is called.
<code>subset</code>	specifies an optional vector specifying a subset of observations to be used.
<code>na.action</code>	specifies a function which indicates what should happen when the data contain NAs. The default is na.pass .
<code>season</code>	numeric or character vector or a factor indicating the seasons in which the observations in <code>y</code> were taken. The length of <code>season</code> must equal the length of <code>y</code> .
<code>year</code>	numeric vector indicating the years in which the observations in <code>y</code> were taken. The length of <code>year</code> must equal the length of <code>y</code> .
<code>alternative</code>	character string indicating the kind of alternative hypothesis. The possible values are "two.sided" (tau not equal to 0; the default), "less" (tau less than 0), and "greater" (tau greater than 0).
<code>correct</code>	logical scalar indicating whether to use the correction for continuity in computing the z -statistic that is based on the test statistic S' . The default value is TRUE.

<code>ci.slope</code>	logical scalar indicating whether to compute a confidence interval for the slope. The default value is TRUE.
<code>conf.level</code>	numeric scalar between 0 and 1 indicating the confidence level associated with the confidence interval for the slope. The default value is 0.95.
<code>independent.obs</code>	logical scalar indicating whether to assume the observations in <code>y</code> are serially independent. The default value is TRUE.
<code>data.name</code>	character string indicating the name of the data used for the trend test. The default value is <code>deparse(substitute(y))</code> .
<code>season.name</code>	character string indicating the name of the data used for the season. The default value is <code>deparse(substitute(season))</code> .
<code>year.name</code>	character string indicating the name of the data used for the year. The default value is <code>deparse(substitute(year))</code> .
<code>parent.of.data</code>	character string indicating the source of the data used for the trend test.
<code>subset.expression</code>	character string indicating the expression used to subset the data.
<code>...</code>	additional arguments affecting the test for trend.

Details

Hirsch et al. (1982) introduced a modification of Kendall's test for trend (see [kendallTrendTest](#)) that allows for seasonality in observations collected over time. They call this test the *seasonal Kendall test*. Their test is appropriate for testing for trend in each season when the trend is always in the same direction across all seasons. van Belle and Hughes (1984) introduced a heterogeneity test for trend which is appropriate for testing for trend in any direction in any season. Hirsch and Slack (1984) proposed an extension to the seasonal Kendall test that allows for serial dependence in the observations. The function [kendallSeasonalTrendTest](#) includes all of these tests, as well as an extension of the van Belle-Hughes heterogeneity test to the case of serial dependence.

Value

A list of class "htest" containing the results of the hypothesis test. See the help file for [htest.object](#) for details. In addition, the following components are part of the list returned by `kendallSeasonalTrendTest`:

<code>seasonal.S</code>	numeric vector. The value of the Kendall S-statistic for each season.
<code>var.seasonal.S</code>	numeric vector. The variance of the Kendall S-statistic for each season. This component only appears when <code>independent.obs=TRUE</code> .
<code>var.cov.seasonal.S</code>	numeric matrix. The estimated variance-covariance matrix of the Kendall S-statistics for each season. This component only appears when <code>independent.obs=FALSE</code> .
<code>seasonal.estimates</code>	numeric matrix. The estimated Kendall's tau, slope, and intercept for each season.

Note

Kendall's test for independence or trend is a nonparametric test. No assumptions are made about the distribution of the X and Y variables. Hirsch et al. (1982) introduced the seasonal Kendall test

to test for trend within each season. They note that Kendall's test for trend is easy to compute, even in the presence of missing values, and can also be used with censored values.

van Belle and Hughes (1984) note that the seasonal Kendall test introduced by Hirsch et al. (1982) is similar to a multivariate extension of the sign test proposed by Jonckheere (1954). Jonckheere's test statistic is based on the unweighted sum of the seasonal tau statistics, while Hirsch et al.'s test is based on the weighted sum (weighted by number of observations within a season) of the seasonal tau statistics.

van Belle and Hughes (1984) also note that Kendall's test for trend is slightly less powerful than the test based on Spearman's rho, but it converges to normality faster. Also, Bradley (1968, p.288) shows that for the case of a linear model with normal (Gaussian) errors, the asymptotic relative efficiency of Kendall's test for trend versus the parametric test for a zero slope is 0.98.

Based on the work of Dietz and Killeen (1981), Hirsch and Slack (1984) describe a modified version of the seasonal Kendall test that allows for serial dependence in the observations. They performed a Monte Carlo study to determine the empirical significance level and power of this modified test vs. the test that assumes independent observations and found a trade-off between power and the correct significance level. For $p = 12$ seasons, they found the modified test gave correct significance levels for $n \geq 10$ as long as the lag-one autocorrelation was 0.6 or less, while the original test that assumes independent observations yielded highly inflated significance levels. On the other hand, if in fact the observations are serially independent, the original test is more powerful than the modified test.

Hirsch and Slack (1984) also looked at the performance of the test for trend introduced by Dietz and Killeen (1981), which is a weighted sums of squares of the seasonal Kendall S-statistics, where the matrix of weights is the inverse of the covariance matrix. The Dietz-Killeen test statistic, unlike the one proposed by Hirsch and Slack (1984), tests for trend in either direction in any season, and is asymptotically distributed as a chi-square random variable with p (number of seasons) degrees of freedom. Hirsch and Slack (1984), however, found that the test based on this statistic is quite conservative (i.e., the significance level is much smaller than the assumed significance level) and has poor power even for moderate sample sizes. The chi-square approximation becomes reasonably close only when $n > 40$ if $p = 12$, $n > 30$ if $p = 4$, and $n > 20$ if $p = 2$.

Lettenmaier (1988) notes the poor power of the test proposed by Dietz and Killeen (1981) and states the poor power apparently results from an upward bias in the estimated variance of the statistic, which can be traced to the inversion of the estimated covariance matrix. He suggests an alternative test statistic (to test trend in either direction in any season) that is the sum of the squares of the scaled seasonal Kendall S-statistics (scaled by their standard deviations). Note that this test statistic ignores information about the covariance between the seasonal Kendall S-statistics, although its distribution depends on these covariances. In the case of no serial dependence, Lettenmaier's test statistic is exactly the same as the Dietz-Killeen test statistic. In the case of serial dependence, Lettenmaier (1988) notes his test statistic is a quadratic form of a multivariate normal random variable and therefore all the moments of this random variable are easily computed. Lettenmaier (1988) approximates the distribution of his test statistic as a scaled non-central chi-square distribution (with fractional degrees of freedom). Based on extensive Monte Carlo studies, Lettenmaier (1988) shows that for the case when the trend is the same in all seasons, the seasonal Kendall's test of Hirsch and Slack (1984) is superior to his test and far superior to the Dietz-Killeen test. The power of Lettenmaier's test approached that of the seasonal Kendall test for large trend magnitudes.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Bradley, J.V. (1968). *Distribution-Free Statistical Tests*. Prentice-Hall, Englewood Cliffs, NJ.

- Conover, W.J. (1980). *Practical Nonparametric Statistics*. Second Edition. John Wiley and Sons, New York, pp.256-272.
- Gibbons, R.D., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*, Second Edition. John Wiley & Sons, Hoboken.
- Gilbert, R.O. (1987). *Statistical Methods for Environmental Pollution Monitoring*. Van Nostrand Reinhold, New York, NY, Chapter 16.
- Helsel, D.R. and R.M. Hirsch. (1988). Discussion of Applicability of the t-test for Detecting Trends in Water Quality Variables. *Water Resources Bulletin* **24**(1), 201-204.
- Helsel, D.R., and R.M. Hirsch. (1992). *Statistical Methods in Water Resources Research*. Elsevier, NY.
- Helsel, D.R., and R. M. Hirsch. (2002). *Statistical Methods in Water Resources*. Techniques of Water Resources Investigations, Book 4, chapter A3. U.S. Geological Survey. Available on-line at <http://www.practicalstats.com/aes/aes/AESbook.html>.
- Hirsch, R.M., J.R. Slack, and R.A. Smith. (1982). Techniques of Trend Analysis for Monthly Water Quality Data. *Water Resources Research* **18**(1), 107-121.
- Hirsch, R.M. and J.R. Slack. (1984). A Nonparametric Trend Test for Seasonal Data with Serial Dependence. *Water Resources Research* **20**(6), 727-732.
- Hirsch, R.M., R.B. Alexander, and R.A. Smith. (1991). Selection of Methods for the Detection and Estimation of Trends in Water Quality. *Water Resources Research* **27**(5), 803-813.
- Hollander, M., and D.A. Wolfe. (1999). *Nonparametric Statistical Methods, Second Edition*. John Wiley and Sons, New York.
- Kendall, M.G. (1938). A New Measure of Rank Correlation. *Biometrika* **30**, 81-93.
- Kendall, M.G. (1975). *Rank Correlation Methods*. Charles Griffin, London.
- Mann, H.B. (1945). Nonparametric Tests Against Trend. *Econometrica* **13**, 245-259.
- Millard, S.P., and Neerchal, N.K. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, Florida.
- Sen, P.K. (1968). Estimates of the Regression Coefficient Based on Kendall's Tau. *Journal of the American Statistical Association* **63**, 1379-1389.
- Theil, H. (1950). A Rank-Invariant Method of Linear and Polynomial Regression Analysis, I-III. *Proc. Kon. Ned. Akad. v. Wetensch. A.* **53**, 386-392, 521-525, 1397-1412.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.
- USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.
- van Belle, G., and J.P. Hughes. (1984). Nonparametric Tests for Trend in Water Quality. *Water Resources Research* **20**(1), 127-136.

See Also

[kendallTrendTest](#), [htest.object](#), [cor.test](#).

Examples

```

# Reproduce Example 14-10 on page 14-38 of USEPA (2009). This example
# tests for trend in analyte concentrations (ppm) collected monthly
# between 1983 and 1985.

head(EPA.09.Ex.14.8.df)
#      Month Year Unadj.Conc Adj.Conc
#1  January 1983      1.99    2.11
#2  February 1983      2.10    2.14
#3   March 1983      2.12    2.10
#4   April 1983      2.12    2.13
#5    May 1983      2.11    2.12
#6   June 1983      2.15    2.12

tail(EPA.09.Ex.14.8.df)
#      Month Year Unadj.Conc Adj.Conc
#31   July 1985      2.31    2.23
#32  August 1985      2.32    2.24
#33 September 1985      2.28    2.23
#34  October 1985      2.22    2.24
#35 November 1985      2.19    2.25
#36 December 1985      2.22    2.23

# Plot the data
#-----
Unadj.Conc <- EPA.09.Ex.14.8.df$Unadj.Conc
Adj.Conc    <- EPA.09.Ex.14.8.df$Adj.Conc
Month       <- EPA.09.Ex.14.8.df$Month
Year        <- EPA.09.Ex.14.8.df$Year
Time        <- paste(substring(Month, 1, 3), Year - 1900, sep = "-")
n           <- length(Unadj.Conc)
Three.Yr.Mean <- mean(Unadj.Conc)

windows()
par(mar = c(7, 4, 3, 1) + 0.1, cex.lab = 1.25)
plot(1:n, Unadj.Conc, type = "n", xaxt = "n",
     xlab = "Time (Month)",
     ylab = "ANALYTE CONCENTRATION (mg/L)",
     main = "Figure 14-15. Seasonal Time Series Over a Three Year Period",
     cex.main = 1.1)
axis(1, at = 1:n, labels = rep("", n))
at <- rep(c(1, 5, 9), 3) + rep(c(0, 12, 24), each = 3)
axis(1, at = at, labels = Time[at])
points(1:n, Unadj.Conc, pch = 0, type = "o", lwd = 2)
points(1:n, Adj.Conc, pch = 3, type = "o", col = 8, lwd = 2)
abline(h = Three.Yr.Mean, lwd = 2)
legend("topleft", c("Unadjusted", "Adjusted", "3-Year Mean"), bty = "n",
      pch = c(0, 3, -1), lty = c(1, 1, 1), lwd = 2, col = c(1, 8, 1),
      inset = c(0.05, 0.01))

# Clean up
rm(Unadj.Conc, Adj.Conc, Month, Year, Time, n, Three.Yr.Mean, at)

# Perform the seasonal Kendall trend test

```

```

#-----

kendallSeasonalTrendTest(Unadj.Conc ~ Month + Year,
  data = EPA.09.Ex.14.8.df)

#Results of Hypothesis Test
#-----
#
#Null Hypothesis:          All 12 values of tau = 0
#
#Alternative Hypothesis:   The seasonal taus are not all equal
#                           (Chi-Square Heterogeneity Test)
#                           At least one seasonal tau != 0
#                           and all non-zero tau's have the
#                           same sign (z Trend Test)
#
#Test Name:                Seasonal Kendall Test for Trend
#                           (with continuity correction)
#
#Estimated Parameter(s):   tau      = 0.9722222
#                           slope    = 0.0600000
#                           intercept = -131.7350000
#
#Estimation Method:        tau:      Weighted Average of
#                           Seasonal Estimates
#                           slope:    Hirsch et al.'s
#                           Modification of
#                           Thiel/Sen Estimator
#                           intercept: Median of
#                           Seasonal Estimates
#
#Data:                     y      = Unadj.Conc
#                           season = Month
#                           year   = Year
#
#Data Source:              EPA.09.Ex.14.8.df
#
#Sample Sizes:             January  = 3
#                           February = 3
#                           March    = 3
#                           April    = 3
#                           May      = 3
#                           June     = 3
#                           July     = 3
#                           August   = 3
#                           September = 3
#                           October  = 3
#                           November = 3
#                           December = 3
#                           Total    = 36
#
#Test Statistics:          Chi-Square (Het) = 0.1071882
#                           z (Trend)      = 5.1849514
#
#Test Statistic Parameter: df = 11
#
#P-values:                 Chi-Square (Het) = 1.000000e+00

```



```

#                               z (Trend)          = 2.160712e-07
#
#Confidence Interval for:      slope
#
#Confidence Interval Method:    Gilbert's Modification of
#                               Theil/Sen Method
#
#Confidence Interval Type:      two-sided
#
#Confidence Level:              95%
#
#Confidence Interval:           LCL = 0.05786914
#                               UCL = 0.07213086

```

kendallTrendTest	<i>Kendall's Nonparametric Test for Monotonic Trend</i>
------------------	---

Description

Perform a nonparametric test for a monotonic trend based on Kendall's tau statistic, and optionally compute a confidence interval for the slope.

Usage

```

kendallTrendTest(y, ...)

## S3 method for class 'formula'
kendallTrendTest(y, data = NULL, subset,
  na.action = na.pass, ...)

## Default S3 method:
kendallTrendTest(y, x = seq(along = y),
  alternative = "two.sided", correct = TRUE, ci.slope = TRUE,
  conf.level = 0.95, warn = TRUE, data.name = NULL, data.name.x = NULL,
  parent.of.data = NULL, subset.expression = NULL, ...)

```

Arguments

<code>y</code>	an object containing data for the trend test. In the default method, the argument <code>y</code> must be numeric vector of observations. In the formula method, <code>y</code> must be a formula of the form <code>y ~ 1</code> or <code>y ~ x</code> . The form <code>y ~ 1</code> indicates use the observations in the vector <code>y</code> for the test for trend and use the default value of the argument <code>x</code> in the call to <code>kendallTrendTest.default</code> . The form <code>y ~ x</code> indicates use the observations in the vector <code>y</code> for the test for trend and use the specified value of the argument <code>x</code> in the call to <code>kendallTrendTest.default</code> . Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
<code>data</code>	specifies an optional data frame, list or environment (or object coercible by <code>as.data.frame</code> to a data frame) containing the variables in the model. If not found in <code>data</code> , the variables are taken from <code>environment(formula)</code> , typically the environment from which <code>kendallTrendTest</code> is called.
<code>subset</code>	specifies an optional vector specifying a subset of observations to be used.

<code>na.action</code>	specifies a function which indicates what should happen when the data contain NAs. The default is na.pass .
<code>x</code>	numeric vector of "predictor" values. The length of <code>x</code> must equal the length of <code>y</code> . Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed. The default value of <code>x</code> is the vector of numbers $1, 2, \dots, n$ where n is the number of elements in <code>y</code> .
<code>alternative</code>	character string indicating the kind of alternative hypothesis. The possible values are "two.sided" (tau not equal to 0; the default), "less" (tau less than 0), and "greater" (tau greater than 0).
<code>correct</code>	logical scalar indicating whether to use the correction for continuity in computing the z -statistic that is based on the test statistic S . The default value is TRUE.
<code>ci.slope</code>	logical scalar indicating whether to compute a confidence interval for the slope. The default value is TRUE.
<code>conf.level</code>	numeric scalar between 0 and 1 indicating the confidence level associated with the confidence interval for the slope. The default value is 0.95.
<code>warn</code>	logical scalar indicating whether to print a warning message when <code>y</code> does not contain at least two non-missing values, or when <code>x</code> does not contain at least two unique non-missing values. The default value is TRUE.
<code>data.name</code>	character string indicating the name of the data used for the trend test. The default value is <code>deparse(substitute(y))</code> .
<code>data.name.x</code>	character string indicating the name of the data used for the predictor variable <code>x</code> . If <code>x</code> is not supplied this argument is ignored. When <code>x</code> is supplied, the default value is <code>deparse(substitute(x))</code> .
<code>parent.of.data</code>	character string indicating the source of the data used for the trend test.
<code>subset.expression</code>	character string indicating the expression used to subset the data.
<code>...</code>	additional arguments affecting the test for trend.

Details

`kendallTrendTest` performs Kendall's nonparametric test for a monotonic trend, which is a special case of the test for independence based on Kendall's tau statistic (see [cor.test](#)). The slope is estimated using the method of Theil (1950) and Sen (1968). When `ci.slope=TRUE`, the confidence interval for the slope is computed using Gilbert's (1987) Modification of the Theil/Sen Method.

Value

A list of class "htest" containing the results of the hypothesis test. See the help file for [htest.object](#) for details. In addition, the following components are part of the list returned by `kendallTrendTest`:

<code>S</code>	The value of the Kendall S-statistic.
<code>var.S</code>	The variance of the Kendall S-statistic.
<code>slopes</code>	A numeric vector of all possible two-point slope estimates. This component is used by the function kendallSeasonalTrendTest .

Note

Kendall's test for independence or trend is a nonparametric test. No assumptions are made about the distribution of the X and Y variables. Hirsch et al. (1982) introduced the "seasonal Kendall test" to test for trend within each season. They note that Kendall's test for trend is easy to compute, even in the presence of missing values, and can also be used with censored values.

van Belle and Hughes (1984) note that Kendall's test for trend is slightly less powerful than the test based on Spearman's rho, but it converges to normality faster. Also, Bradley (1968, p.288) shows that for the case of a linear model with normal (Gaussian) errors, the asymptotic relative efficiency of Kendall's test for trend versus the parametric test for a zero slope is 0.98.

The results of the function `kendallTrendTest` are similar to the results of the built-in R function `cor.test` with the argument `method="kendall"` except that `cor.test` 1) computes exact p-values when the number of pairs is less than 50 and there are no ties, and 2) does not return a confidence interval for the slope.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Bradley, J.V. (1968). *Distribution-Free Statistical Tests*. Prentice-Hall, Englewood Cliffs, NJ.
- Conover, W.J. (1980). *Practical Nonparametric Statistics*. Second Edition. John Wiley and Sons, New York, pp.256-272.
- Gibbons, R.D., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*, Second Edition. John Wiley & Sons, Hoboken.
- Gilbert, R.O. (1987). *Statistical Methods for Environmental Pollution Monitoring*. Van Nostrand Reinhold, New York, NY, Chapter 16.
- Helsel, D.R. and R.M. Hirsch. (1988). Discussion of Applicability of the t-test for Detecting Trends in Water Quality Variables. *Water Resources Bulletin* **24**(1), 201-204.
- Helsel, D.R., and R.M. Hirsch. (1992). *Statistical Methods in Water Resources Research*. Elsevier, NY.
- Helsel, D.R., and R. M. Hirsch. (2002). *Statistical Methods in Water Resources*. Techniques of Water Resources Investigations, Book 4, chapter A3. U.S. Geological Survey. Available on-line at <http://www.practicalstats.com/aes/aes/AESbook.html>.
- Hirsch, R.M., J.R. Slack, and R.A. Smith. (1982). Techniques of Trend Analysis for Monthly Water Quality Data. *Water Resources Research* **18**(1), 107-121.
- Hirsch, R.M. and J.R. Slack. (1984). A Nonparametric Trend Test for Seasonal Data with Serial Dependence. *Water Resources Research* **20**(6), 727-732.
- Hirsch, R.M., R.B. Alexander, and R.A. Smith. (1991). Selection of Methods for the Detection and Estimation of Trends in Water Quality. *Water Resources Research* **27**(5), 803-813.
- Hollander, M., and D.A. Wolfe. (1999). *Nonparametric Statistical Methods, Second Edition*. John Wiley and Sons, New York.
- Kendall, M.G. (1938). A New Measure of Rank Correlation. *Biometrika* **30**, 81-93.
- Kendall, M.G. (1975). *Rank Correlation Methods*. Charles Griffin, London.
- Mann, H.B. (1945). Nonparametric Tests Against Trend. *Econometrica* **13**, 245-259.
- Millard, S.P., and Neerchal, N.K. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, Florida.

Sen, P.K. (1968). Estimates of the Regression Coefficient Based on Kendall's Tau. *Journal of the American Statistical Association* **63**, 1379-1389.

Theil, H. (1950). A Rank-Invariant Method of Linear and Polynomial Regression Analysis, I-III. *Proc. Kon. Ned. Akad. v. Wetensch. A.* **53**, 386-392, 521-525, 1397-1412.

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

van Belle, G., and J.P. Hughes. (1984). Nonparametric Tests for Trend in Water Quality. *Water Resources Research* **20**(1), 127-136.

See Also

[cor.test](#), [kendallSeasonalTrendTest](#), [htest.object](#).

Examples

```
# Reproduce Example 17-6 on page 17-33 of USEPA (2009). This example
# tests for trend in sulfate concentrations (ppm) collected at various
# months between 1989 and 1996.

head(EPA.09.Ex.17.6.sulfate.df)
# Sample.No Year Month Sampling.Date Date Sulfate.ppm
#1          1   89    6          89.6 1989-06-01      480
#2          2   89    8          89.8 1989-08-01      450
#3          3   90    1          90.1 1990-01-01      490
#4          4   90    3          90.3 1990-03-01      520
#5          5   90    6          90.6 1990-06-01      485
#6          6   90    8          90.8 1990-08-01      510

# Plot the data
#-----
with(EPA.09.Ex.17.6.sulfate.df,
     plot(Sampling.Date, Sulfate.ppm, pch = 15, ylim = c(400, 900),
          xlab = "Sampling Date", ylab = "Sulfate Conc (ppm)",
          main = "Figure 17-6. Time Series Plot of \nSulfate Concentrations (ppm)"))
)
Sulfate.fit <- lm(Sulfate.ppm ~ Sampling.Date,
                 data = EPA.09.Ex.17.6.sulfate.df)
abline(Sulfate.fit, lty = 2)

# Perform the Kendall test for trend
#-----
kendallTrendTest(Sulfate.ppm ~ Sampling.Date,
                 data = EPA.09.Ex.17.6.sulfate.df)

#Results of Hypothesis Test
#-----
#
#Null Hypothesis:                tau = 0
```

```

#
#Alternative Hypothesis:      True tau is not equal to 0
#
#Test Name:                  Kendall's Test for Trend
#                             (with continuity correction)
#
#Estimated Parameter(s):    tau      =      0.7667984
#                             slope    =      26.6666667
#                             intercept = -1909.3333333
#
#Estimation Method:         slope:      Theil/Sen Estimator
#                             intercept: Conover's Estimator
#
#Data:                       y = Sulfate.ppm
#                             x = Sampling.Date
#
#Data Source:                EPA.09.Ex.17.6.sulfate.df
#
#Sample Size:                23
#
#Test Statistic:             z = 5.107322
#
#P-value:                    3.267574e-07
#
#Confidence Interval for:    slope
#
#Confidence Interval Method: Gilbert's Modification
#                             of Theil/Sen Method
#
#Confidence Interval Type:    two-sided
#
#Confidence Level:            95%
#
#Confidence Interval:         LCL = 20.00000
#                             UCL = 35.71182

# Clean up
#-----
rm(Sulfate.fit)

```

kurtosis

Coefficient of (Excess) Kurtosis

Description

Compute the sample coefficient of kurtosis or excess kurtosis.

Usage

```

kurtosis(x, na.rm = FALSE, method = "fisher", l.moment.method = "unbiased",
  plot.pos.cons = c(a = 0.35, b = 0), excess = TRUE)

```

Arguments

<code>x</code>	numeric vector of observations.
<code>na.rm</code>	logical scalar indicating whether to remove missing values from <code>x</code> . If <code>na.rm=FALSE</code> (the default) and <code>x</code> contains missing values, then a missing value (NA) is returned. If <code>na.rm=TRUE</code> , missing values are removed from <code>x</code> prior to computing the coefficient of variation.
<code>method</code>	character string specifying what method to use to compute the sample coefficient of kurtosis. The possible values are "fisher" (ratio of unbiased moment estimators; the default), "moments" (ratio of product moment estimators), or "l.moments" (ratio of <i>L</i> -moment estimators).
<code>l.moment.method</code>	character string specifying what method to use to compute the <i>L</i> -moments when <code>method="l.moments"</code> . The possible values are "ubiased" (method based on the <i>U</i> -statistic; the default), or "plotting.position" (method based on the plotting position formula).
<code>plot.pos.cons</code>	numeric vector of length 2 specifying the constants used in the formula for the plotting positions when <code>method="l.moments"</code> and <code>l.moment.method="plotting.position"</code> . The default value is <code>plot.pos.cons=c(a=0.35, b=0)</code> . If this vector has a names attribute with the value <code>c("a", "b")</code> or <code>c("b", "a")</code> , then the elements will be matched by name in the formula for computing the plotting positions. Otherwise, the first element is mapped to the name "a" and the second element to the name "b".
<code>excess</code>	logical scalar indicating whether to compute the kurtosis (<code>excess=FALSE</code>) or excess kurtosis (<code>excess=TRUE</code> ; the default).

Details

Let \underline{x} denote a random sample of n observations from some distribution with mean μ and standard deviation σ .

Product Moment Coefficient of Kurtosis

(`method="moment"` or `method="fisher"`)

The **coefficient of kurtosis** of a distribution is the fourth standardized moment about the mean:

$$\eta_4 = \beta_2 = \frac{\mu_4}{\sigma^4} \quad (1)$$

where

$$\eta_r = E\left[\left(\frac{X - \mu}{\sigma}\right)^r\right] = \frac{1}{\sigma^r} E[(X - \mu)^r] = \frac{\mu_r}{\sigma^r} \quad (2)$$

and

$$\mu_r = E[(X - \mu)^r] \quad (3)$$

denotes the r 'th moment about the mean (central moment).

The **coefficient of excess kurtosis** is defined as:

$$\beta_2 - 3 \quad (4)$$

For a normal distribution, the coefficient of kurtosis is 3 and the coefficient of excess kurtosis is 0. Distributions with kurtosis less than 3 (excess kurtosis less than 0) are called **platykurtic**: they have shorter tails than a normal distribution. Distributions with kurtosis greater than 3 (excess kurtosis greater than 0) are called **leptokurtic**: they have heavier tails than a normal distribution.

When method="moment", the coefficient of kurtosis is estimated using the method of moments estimator for the fourth central moment and the method of moments estimator for the variance:

$$\hat{\eta}_4 = \frac{\hat{\mu}_4}{\sigma^4} = \frac{\frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^4}{\left[\frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2 \right]^2} \quad (5)$$

where

$$\hat{\sigma}_m^2 = s_m^2 = \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2 \quad (6)$$

This form of estimation should be used when resampling (bootstrap or jackknife).

When method="fisher", the coefficient of kurtosis is estimated using the unbiased estimator for the fourth central moment (Serfling, 1980, p.73) and the unbiased estimator for the variance.

$$\hat{\sigma}^2 = s^2 = \frac{1}{n-1} \sum_{i=1}^n (x_i - \bar{x})^2 \quad (7)$$

L-Moment Coefficient of Kurtosis (method="l.moments")

Hosking (1990) defines the *L*-moment analog of the coefficient of kurtosis as:

$$\tau_4 = \frac{\lambda_4}{\lambda_2} \quad (8)$$

that is, the fourth *L*-moment divided by the second *L*-moment. He shows that this quantity lies in the interval (-1, 1).

When l.moment.method="unbiased", the *L*-kurtosis is estimated by:

$$t_4 = \frac{l_4}{l_2} \quad (9)$$

that is, the unbiased estimator of the fourth *L*-moment divided by the unbiased estimator of the second *L*-moment.

When l.moment.method="plotting.position", the *L*-kurtosis is estimated by:

$$\tilde{\tau}_4 = \frac{\tilde{\lambda}_4}{\tilde{\lambda}_2} \quad (10)$$

that is, the plotting-position estimator of the fourth *L*-moment divided by the plotting-position estimator of the second *L*-moment.

See the help file for [lMoment](#) for more information on estimating *L*-moments.

Value

A numeric scalar – the sample coefficient of kurtosis or excess kurtosis.

Note

Traditionally, the coefficient of kurtosis has been estimated using product moment estimators. Sometimes an estimate of kurtosis is used in a goodness-of-fit test for normality (D'Agostino and Stephens, 1986). Hosking (1990) introduced the idea of *L*-moments and *L*-kurtosis.

Vogel and Fennessey (1993) argue that *L*-moment ratios should replace product moment ratios because of their superior performance (they are nearly unbiased and better for discriminating between distributions). They compare product moment diagrams with *L*-moment diagrams.

Hosking and Wallis (1995) recommend using unbiased estimators of *L*-moments (vs. plotting-position estimators) for almost all applications.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Berthouex, P.M., and L.C. Brown. (2002). *Statistics for Environmental Engineers, Second Edition*. Lewis Publishers, Boca Raton, FL.
- Ott, W.R. (1995). *Environmental Statistics and Data Analysis*. Lewis Publishers, Boca Raton, FL.
- Taylor, J.K. (1990). *Statistical Techniques for Data Analysis*. Lewis Publishers, Boca Raton, FL.
- Vogel, R.M., and N.M. Fennessey. (1993). *L Moment Diagrams Should Replace Product Moment Diagrams*. *Water Resources Research* **29**(6), 1745–1752.
- Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ.

See Also

[var](#), [sd](#), [cv](#), [skewness](#), [summaryFull](#), [Summary Statistics](#).

Examples

```
# Generate 20 observations from a lognormal distribution with parameters
# mean=10 and cv=1, and estimate the coefficient of kurtosis and
# coefficient of excess kurtosis.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)

dat <- rlnormAlt(20, mean = 10, cv = 1)

# Compute standard kurtosis first
#-----
kurtosis(dat, excess = FALSE)
#[1] 2.964612

kurtosis(dat, method = "moment", excess = FALSE)
#[1] 2.687146

kurtosis(dat, method = "l.moment", excess = FALSE)
#[1] 0.1444779

# Now compute excess kurtosis
#-----
kurtosis(dat)
#[1] -0.0353876

kurtosis(dat, method = "moment")
#[1] -0.3128536

kurtosis(dat, method = "l.moment")
#[1] -2.855522

#-----
# Clean up
rm(dat)
```


Lin.Evans.80.df

*Fecal Coliform Data from the Illinois River***Description**

Lin and Evans (1980) reported fecal coliform measures (organisms per 100 ml) from the Illinois River taken between 1971 and 1976. The object Lin.Evans.80.df is a small subset of these data that were reported by Helsel and Hirsch (1992, p.162).

Usage

Lin.Evans.80.df

Format

A data frame with 24 observations on the following 2 variables.

Fecal.Coliform a numeric vector of fecal coliform measure (organisms per 100 ml).

Season an ordered factor indicating the season of collection

Source

Helsel, D.R., and R.M. Hirsch. (1992). *Statistical Methods in Water Resources Research*. Elsevier, New York, NY, p.162.

References

Lin, S.D., and R.L. Evans. (1980). *Coliforms and fecal streptococcus in the Illinois River at Peoria, 1971-1976*. Illinois State Water Survey Report of Investigations No. 93. Urbana, IL, 28pp.

lMoment

*Estimate L-Moments***Description**

Estimate the r 'th L -moment from a random sample.

Usage

```
lMoment(x, r = 1, method = "unbiased",
        plot.pos.cons = c(a = 0.35, b = 0), na.rm = FALSE)
```

Arguments

x	numeric vector of observations.
r	positive integer specifying the order of the moment.
method	character string specifying what method to use to compute the L -moment. The possible values are "unbiased" (method based on the U-statistic; the default), or "plotting.position" (method based on the plotting position formula). See the DETAILS section for more information.
plot.pos.cons	numeric vector of length 2 specifying the constants used in the formula for the plotting positions when method="plotting.position". The default value is plot.pos.cons=c(a=0.35, b=0). If this vector has a names attribute with the value c("a", "b") or c("b", "a"), then the elements will be matched by name in the formula for computing the plotting positions. Otherwise, the first element is mapped to the name "a" and the second element to the name "b". See the DETAILS section for more information. This argument is ignored if method="unbiased".
na.rm	logical scalar indicating whether to remove missing values from x. If na.rm=FALSE (the default) and x contains missing values, then a missing value (NA) is returned. If na.rm=TRUE, missing values are removed from x prior to computing the L -moment.

Details**Definitions: L -Moments and L -Moment Ratios**

The definition of an L -moment given by Hosking (1990) is as follows. Let X denote a random variable with cdf F , and let $x(p)$ denote the p 'th quantile of the distribution. Furthermore, let

$$x_{1:n} \leq x_{2:n} \leq \dots \leq x_{n:n}$$

denote the order statistics of a random sample of size n drawn from the distribution of X . Then the r 'th L -moment is given by:

$$\lambda_r = \frac{1}{r} \sum_{k=0}^{r-1} (-1)^k \binom{r-1}{k} E[X_{r-k:r}]$$

for $r = 1, 2, \dots$

Hosking (1990) shows that the above equation can be rewritten as:

$$\lambda_r = \int_0^1 x(u) P_{r-1}^*(u) du$$

where

$$P_r^*(u) = \sum_{k=0}^r p_{r,k}^* u^k$$

$$p_{r,k}^* = (-1)^{r-k} \binom{r}{k} \binom{r+k}{k} = \frac{(-1)^{r-k} (r+k)!}{(k!)^2 (r-k)!}$$

The first four L -moments are given by:

$$\lambda_1 = E[X]$$

$$\lambda_2 = \frac{1}{2} E[X_{2:2} - X_{1:2}]$$

$$\lambda_3 = \frac{1}{3}E[X_{3:3} - 2X_{2:3} + X_{1:3}]$$

$$\lambda_4 = \frac{1}{4}E[X_{4:4} - 3X_{3:4} + 3X_{2:4} - X_{1:4}]$$

Thus, the first L -moment is a measure of location, and the second L -moment is a measure of scale. Hosking (1990) defines the L -moment ratios of X to be:

$$\tau_r = \frac{\lambda_r}{\lambda_2}$$

for $r = 2, 3, \dots$. He shows that for a non-degenerate random variable with a finite mean, these quantities lie in the interval $(-1, 1)$. The quantity

$$\tau_3 = \frac{\lambda_3}{\lambda_2}$$

is the L -moment analog of the coefficient of skewness, and the quantity

$$\tau_4 = \frac{\lambda_4}{\lambda_2}$$

is the L -moment analog of the coefficient of kurtosis. Hosking (1990) also defines an L -moment analog of the coefficient of variation (denoted the L -CV) as:

$$\lambda = \frac{\lambda_2}{\lambda_1}$$

He shows that for a positive-valued random variable, the L -CV lies in the interval $(0, 1)$.

Relationship Between L -Moments and Probability-Weighted Moments

Hosking (1990) and Hosking and Wallis (1995) show that L -moments can be written as linear combinations of probability-weighted moments:

$$\lambda_r = (-1)^{r-1} \sum_{k=0}^{r-1} p_{r-1,k}^* \alpha_k = \sum_{j=0}^{r-1} p_{r-1,j}^* \beta_j$$

where

$$\alpha_k = M(1, 0, k) = \frac{1}{k+1} E[X_{1:k+1}]$$

$$\beta_j = M(1, j, 0) = \frac{1}{j+1} E[X_{j+1:j+1}]$$

See the help file for [pwMoment](#) for more information on probability-weighted moments.

Estimating L -Moments

The two commonly used methods for estimating L -moments are the “unbiased” method based on U-statistics (Hoeffding, 1948; Lehmann, 1975, pp. 362-371), and the “plotting-position” method. Hosking and Wallis (1995) recommend using the unbiased method for almost all applications.

Unbiased Estimators (method="unbiased")

Using the relationship between L -moments and probability-weighted moments explained above, the unbiased estimator of the r 'th L -moment is based on unbiased estimators of probability-weighted moments and is given by:

$$l_r = (-1)^{r-1} \sum_{k=0}^{r-1} p_{r-1,k}^* a_k = \sum_{j=0}^{r-1} p_{r-1,j}^* b_j$$

where

$$a_k = \frac{1}{n} \sum_{i=1}^{n-k} x_{i:n} \frac{\binom{n-i}{k}}{\binom{n-1}{k}}$$

$$b_j = \frac{1}{n} \sum_{i=j+1}^n x_{i:n} \frac{\binom{i-1}{j}}{\binom{n-1}{j}}$$

Plotting-Position Estimators (method="plotting.position")

Using the relationship between L -moments and probability-weighted moments explained above, the plotting-position estimator of the r 'th L -moment is based on the plotting-position estimators of probability-weighted moments and is given by:

$$\tilde{\lambda}_r = (-1)^{r-1} \sum_{k=0}^{r-1} p_{r-1,k}^* \tilde{\alpha}_k = \sum_{j=0}^{r-1} p_{r-1,j}^* \tilde{\beta}_j$$

where

$$\tilde{\alpha}_k = \frac{1}{n} \sum_{i=1}^n (1 - p_{i:n})^k x_{i:n}$$

$$\tilde{\beta}_j = \frac{1}{n} \sum_{i=1}^n p_{i:n}^j x_{i:n}$$

and

$$p_{i:n} = \hat{F}(x_{i:n})$$

denotes the plotting position of the i 'th order statistic in the random sample of size n , that is, a distribution-free estimate of the cdf of X evaluated at the i 'th order statistic. Typically, plotting positions have the form:

$$p_{i:n} = \frac{i - a}{n + b}$$

where $b > -a > -1$. For this form of plotting position, the plotting-position estimators are asymptotically equivalent to their unbiased estimator counterparts.

Estimating L -Moment Ratios

L -moment ratios are estimated by simply replacing the population L -moments with the estimated L -moments. The estimated ratios based on the unbiased estimators are given by:

$$t_r = \frac{l_r}{l_2}$$

and the estimated ratios based on the plotting-position estimators are given by:

$$\tilde{\tau}_r = \frac{\tilde{\lambda}_r}{\tilde{\lambda}_2}$$

In particular, the L -moment skew is estimated by:

$$t_3 = \frac{l_3}{l_2}$$

or

$$\tilde{\tau}_3 = \frac{\tilde{\lambda}_3}{\tilde{\lambda}_2}$$

and the L -moment kurtosis is estimated by:

$$t_4 = \frac{l_4}{l_2}$$

or

$$\tilde{\tau}_4 = \frac{\tilde{\lambda}_4}{\tilde{\lambda}_2}$$

Similarly, the L -moment coefficient of variation can be estimated using the unbiased L -moment estimators:

$$l = \frac{l_2}{l_1}$$

or using the plotting-position L -moment estimators:

$$\tilde{\lambda} = \frac{\tilde{\lambda}_2}{\tilde{\lambda}_1}$$

Value

A numeric scalar—the value of the r 'th L -moment as defined by Hosking (1990).

Note

Hosking (1990) introduced the idea of L -moments, which are expectations of certain linear combinations of order statistics, as the basis of a general theory of describing theoretical probability distributions, computing summary statistics from observed data, estimating distribution parameters and quantiles, and performing hypothesis tests. The theory of L -moments parallels the theory of conventional moments. L -moments have several advantages over conventional moments, including:

- L -moments can characterize a wider range of distributions because they always exist as long as the distribution has a finite mean.
- L -moments are estimated by linear combinations of order statistics, so estimators based on L -moments are more robust to the presence of outliers than estimators based on conventional moments.
- Based on the author's and others' experience, L -moment estimators are less biased and approximate their asymptotic distribution more closely in finite samples than estimators based on conventional moments.
- L -moment estimators are sometimes more efficient (smaller RMSE) than even maximum likelihood estimators for small samples.

Hosking (1990) presents a table with formulas for the L -moments of common probability distributions. Articles that illustrate the use of L -moments include Fill and Stedinger (1995), Hosking and Wallis (1995), and Vogel and Fennessey (1993).

Hosking (1990) and Hosking and Wallis (1995) show the relationship between probability-weighted moments and L -moments.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Fill, H.D., and J.R. Stedinger. (1995). L Moment and Probability Plot Correlation Coefficient Goodness-of-Fit Tests for the Gumbel Distribution and Impact of Autocorrelation. *Water Resources Research* **31**(1), 225–229.
- Hosking, J.R.M. (1990). L -Moments: Analysis and Estimation of Distributions Using Linear Combinations of Order Statistics. *Journal of the Royal Statistical Society, Series B* **52**(1), 105–124.

Hosking, J.R.M., and J.R. Wallis (1995). A Comparison of Unbiased and Plotting-Position Estimators of L Moments. *Water Resources Research* **31**(8), 2019–2025.

Vogel, R.M., and N.M. Fennessey. (1993). L Moment Diagrams Should Replace Product Moment Diagrams. *Water Resources Research* **29**(6), 1745–1752.

See Also

[cv](#), [skewness](#), [kurtosis](#), [pwMoment](#).

Examples

```
# Generate 20 observations from a generalized extreme value distribution
# with parameters location=10, scale=2, and shape=.25, then compute the
# first four L-moments.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
dat <- rgevd(20, location = 10, scale = 2, shape = 0.25)

lMoment(dat)
#[1] 10.59556

lMoment(dat, 2)
#[1] 1.0014

lMoment(dat, 3)
#[1] 0.1681165

lMoment(dat, 4)
#[1] 0.08732692

#-----

# Now compute some L-moments based on the plotting-position estimators:

lMoment(dat, method = "plotting.position")
#[1] 10.59556

lMoment(dat, 2, method = "plotting.position")
#[1] 1.110264

lMoment(dat, 3, method="plotting.position", plot.pos.cons = c(.325,1))
#[1] -0.4430792

#-----

# Clean up
#-----
rm(dat)
```

Description

Density, distribution function, quantile function, and random generation for the three-parameter lognormal distribution with parameters meanlog, sdlog, and threshold.

Usage

```
dlnorm3(x, meanlog = 0, sdlog = 1, threshold = 0)
plnorm3(q, meanlog = 0, sdlog = 1, threshold = 0)
qlnorm3(p, meanlog = 0, sdlog = 1, threshold = 0)
rlnorm3(n, meanlog = 0, sdlog = 1, threshold = 0)
```

Arguments

x	vector of quantiles.
q	vector of quantiles.
p	vector of probabilities between 0 and 1.
n	sample size. If length(n) is larger than 1, then length(n) random values are returned.
meanlog	vector of means of the distribution of the random variable on the log scale. The default is meanlog=0.
sdlog	vector of (positive) standard deviations of the random variable on the log scale. The default is sdlog=1.
threshold	vector of thresholds of the random variable on the log scale. The default is threshold=0.

Details

The three-parameter lognormal distribution is simply the usual [two-parameter lognormal distribution](#) with a location shift.

Let X be a random variable with a three-parameter lognormal distribution with parameters meanlog= μ , sdlog= σ , and threshold= γ . Then the random variable $Y = X - \gamma$ has a [lognormal distribution](#) with parameters meanlog= μ and sdlog= σ . Thus,

- dlnorm3 calls [dlnorm](#) using the arguments $x = x - \text{threshold}$, meanlog = meanlog, sdlog = sdlog
- plnorm3 calls [plnorm](#) using the arguments $q = q - \text{threshold}$, meanlog = meanlog, sdlog = sdlog
- qlnorm3 calls [qlnorm](#) using the arguments $q = q$, meanlog = meanlog, sdlog = sdlog and then adds the argument threshold to the result.
- rlnorm3 calls [rlnorm](#) using the arguments $n = n$, meanlog = meanlog, sdlog = sdlog and then adds the argument threshold to the result.

The threshold parameter γ affects only the location of the three-parameter lognormal distribution; it has no effect on the variance or the shape of the distribution.

Denote the mean, variance, and coefficient of variation of $Y = X - \gamma$ by:

$$E(Y) = \theta$$

$$Var(Y) = \eta^2$$

$$CV(Y) = \tau = \eta/\theta$$

Then the mean, variance, and coefficient of variation of X are given by:

$$E(X) = \theta + \eta$$

$$Var(X) = \eta^2$$

$$CV(X) = \frac{\eta}{\theta + \gamma} = \frac{\tau\theta}{\theta + \gamma}$$

The relationships between the parameters μ , σ , θ , η , and τ are as follows:

$$\theta = \beta\sqrt{\omega}$$

$$\eta = \beta\sqrt{\omega(\omega - 1)}$$

$$\tau = \sqrt{\omega - 1}$$

$$\mu = \log\left(\frac{\theta}{\sqrt{\tau^2 + 1}}\right)$$

$$\sigma = \sqrt{\log(\tau^2 + 1)}$$

where

$$\beta = e^\mu, \omega = \exp(\sigma^2)$$

Since quantiles of a distribution are preserved under monotonic transformations, the median of X is:

$$Median(X) = \gamma + \beta$$

Value

`dlnorm3` gives the density, `plnorm3` gives the distribution function, `qlnorm3` gives the quantile function, and `rlnorm3` generates random deviates.

Note

The [two-parameter lognormal distribution](#) is the distribution of a random variable whose logarithm is normally distributed. The two major characteristics of the two-parameter lognormal distribution are that it is bounded below at 0, and it is skewed to the right. The three-parameter lognormal distribution is a generalization of the two-parameter lognormal distribution in which the distribution is shifted so that the threshold parameter is some arbitrary number, not necessarily 0.

The three-parameter lognormal distribution was introduced by Wicksell (1917) in a study of the distribution of ages at first marriage. Both the two- and three-parameter lognormal distributions have been used in a variety of fields, including economics and business, industry, biology, ecology, atmospheric science, and geology (Crow and Shimizu, 1988). Royston (1992) has discussed the application of the three-parameter lognormal distribution in the field of medicine.

The two-parameter lognormal distribution is often used to characterize chemical concentrations in the environment. Ott (1990) has shown mathematically how a series of successive random dilutions gives rise to a distribution that can be approximated by a two-parameter lognormal distribution.

The three-parameter lognormal distribution starts to resemble a normal distribution as the parameter σ (the standard deviation of $\log(X - \gamma)$) tends to 0.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Aitchison, J., and J.A.C. Brown (1957). *The Lognormal Distribution (with special references to its uses in economics)*. Cambridge University Press, London, 176pp.
- Crow, E.L., and K. Shimizu. (1988). *Lognormal Distributions: Theory and Applications*. Marcel Dekker, New York, 387pp.
- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.
- Johnson, N. L., S. Kotz, and N. Balakrishnan. (1994). *Continuous Univariate Distributions, Volume 1*. Second Edition. John Wiley and Sons, New York.
- Ott, W.R. (1990). A Physical Explanation of the Lognormality of Pollutant Concentrations. *Journal of the Air and Waste Management Association* **40**, 1378–1383.
- Ott, W.R. (1995). *Environmental Statistics and Data Analysis*. Lewis Publishers, Boca Raton, FL, Chapter 9.
- Royston, J.P. (1992b). Estimation, Reference Ranges and Goodness of Fit for the Three-Parameter Log-Normal Distribution. *Statistics in Medicine* **11**, 897–912.
- Wicksell, S.D. (1917). On Logarithmic Correlation with an Application to the Distribution of Ages at First Marriage. *Medd. Lunds. Astr. Obs.* **84**, 1–21.

See Also

[Lognormal](#), [elnorm3](#), [Probability Distributions and Random Numbers](#).

Examples

```
# Density of the three-parameter lognormal distribution with
# parameters meanlog=1, sdlog=2, and threshold=10, evaluated at 10.5:

dlnorm3(10.5, 1, 2, 10)
#[1] 0.278794

#-----

# The cdf of the three-parameter lognormal distribution with
# parameters meanlog=2, sdlog=3, and threshold=5, evaluated at 9:

plnorm3(9, 2, 3, 5)
#[1] 0.4189546

#-----

# The median of the three-parameter lognormal distribution with
# parameters meanlog=2, sdlog=3, and threshold=20:

qlnorm3(0.5, 2, 3, 20)
#[1] 27.38906

#-----

# Random sample of 3 observations from the three-parameter lognormal
# distribution with parameters meanlog=2, sdlog=1, and threshold=-5.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(20)
```

```
rlnorm3(3, 2, 1, -5)
#[1] 18.6339749 -0.8873173 39.0561521
```

LognormalAlt

The Lognormal Distribution (Alternative Parameterization)

Description

Density, distribution function, quantile function, and random generation for the lognormal distribution with parameters mean and cv.

Usage

```
dlnormAlt(x, mean = exp(1/2), cv = sqrt(exp(1) - 1), log = FALSE)
plnormAlt(q, mean = exp(1/2), cv = sqrt(exp(1) - 1),
  lower.tail = TRUE, log.p = FALSE)
qlnormAlt(p, mean = exp(1/2), cv = sqrt(exp(1) - 1),
  lower.tail = TRUE, log.p = FALSE)
rlnormAlt(n, mean = exp(1/2), cv = sqrt(exp(1) - 1))
```

Arguments

x	vector of quantiles.
q	vector of quantiles.
p	vector of probabilities between 0 and 1.
n	sample size. If length(n) is larger than 1, then length(n) random values are returned.
mean	vector of (positive) means of the distribution of the random variable.
cv	vector of (positive) coefficients of variation of the random variable.
log, log.p	logical; if TRUE, probabilities/densities p are returned as $\log(p)$.
lower.tail	logical; if TRUE (default), probabilities are $P[X \leq x]$, otherwise, $P[X > x]$.

Details

Let X be a random variable with a [lognormal distribution](#) with parameters $\text{meanlog}=\mu$ and $\text{sdlog}=\sigma$. That is, μ and σ denote the mean and standard deviation of the random variable on the log scale. The relationship between these parameters and the mean ($\text{mean}=\theta$) and coefficient of variation ($\text{cv}=\tau$) of the distribution on the original scale is given by:

$$\mu = \log\left(\frac{\theta}{\sqrt{\tau^2 + 1}}\right) \quad (1)$$

$$\sigma = [\log(\tau^2 + 1)]^{1/2} \quad (2)$$

$$\theta = \exp[\mu + (\sigma^2/2)] \quad (3)$$

$$\tau = [\exp(\sigma^2) - 1]^{1/2} \quad (4)$$

Thus, the functions `dlnormAlt`, `plnormAlt`, `qlnormAlt`, and `rlnormAlt` call the R functions [dlnorm](#), [plnorm](#), [qlnorm](#), and [rlnorm](#), respectively using the following values for the `meanlog` and `sdlog` parameters:

```
sdlog <- sqrt(log(1 + cv^2)),
meanlog <- log(mean) - (sdlog^2)/2
```

Value

`dlnormAlt` gives the density, `plnormAlt` gives the distribution function, `qlnormAlt` gives the quantile function, and `rlnormAlt` generates random deviates.

Note

The two-parameter [lognormal distribution](#) is the distribution of a random variable whose logarithm is normally distributed. The two major characteristics of the lognormal distribution are that it is bounded below at 0, and it is skewed to the right.

Because the empirical distribution of many variables is inherently positive and skewed to the right (e.g., size of organisms, amount of rainfall, size of income, etc.), the lognormal distribution has been widely applied in several fields, including economics, business, industry, biology, ecology, atmospheric science, and geology (Aitchison and Brown, 1957; Crow and Shimizu, 1988).

Gibrat (1930) derived the lognormal distribution from theoretical assumptions, calling it the "law of proportionate effect", but Kapteyn (1903) had described a machine that was the mechanical equivalent. The basic idea is that the Central Limit Theorem states that the distribution of the sum of several independent random variables tends to look like a normal distribution, no matter what the underlying distribution(s) of the original random variables, hence the product of several independent random variables tends to look like a lognormal distribution.

The lognormal distribution is often used to characterize chemical concentrations in the environment. Ott (1990) has shown mathematically how a series of successive random dilutions gives rise to a distribution that can be approximated by a lognormal distribution.

A lognormal distribution starts to resemble a normal distribution as the parameter σ (the standard deviation of the log of the distribution) tends to 0.

Some EPA guidance documents (e.g., Singh et al., 2002; Singh et al., 2010a,b) discourage using the assumption of a lognormal distribution for some types of environmental data and recommend instead assessing whether the data appear to fit a gamma distribution.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.
- Johnson, N. L., S. Kotz, and N. Balakrishnan. (1994). *Continuous Univariate Distributions, Volume 1*. Second Edition. John Wiley and Sons, New York.
- Limpert, E., W.A. Stahel, and M. Abbt. (2001). Log-Normal Distributions Across the Sciences: Keys and Clues. *BioScience* **51**, 341–352.
- Ott, W.R. (1995). *Environmental Statistics and Data Analysis*. Lewis Publishers, Boca Raton, FL.
- Singh, A., R. Maichle, and N. Armbya. (2010a). *ProUCL Version 4.1.00 User Guide (Draft)*. EPA/600/R-07/041, May 2010. Office of Research and Development, U.S. Environmental Protection Agency, Washington, D.C.
- Singh, A., N. Armbya, and A. Singh. (2010b). *ProUCL Version 4.1.00 Technical Guide (Draft)*. EPA/600/R-07/041, May 2010. Office of Research and Development, U.S. Environmental Protection Agency, Washington, D.C.

See Also

[Lognormal](#), [elnormAlt](#), [Probability Distributions and Random Numbers](#).

Examples

```
# Density of the lognormal distribution with parameters
# mean=10 and cv=1, evaluated at 5:

dlnormAlt(5, mean = 10, cv = 1)
#[1] 0.08788173

#-----

# The cdf of the lognormal distribution with parameters mean=2 and cv=3,
# evaluated at 4:

plnormAlt(4, 2, 3)
#[1] 0.8879132

#-----

# The median of the lognormal distribution with parameters
# mean=10 and cv=1:

qlnormAlt(0.5, mean = 10, cv = 1)
#[1] 7.071068

#-----

# Random sample of 3 observations from a lognormal distribution with
# parameters mean=10 and cv=1.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(20)
rlnormAlt(3, mean = 10, cv = 1)
#[1] 18.615797  4.341402 31.265293
```

LognormalMix

Mixture of Two Lognormal Distributions

Description

Density, distribution function, quantile function, and random generation for a mixture of two log-normal distribution with parameters meanlog1, sdlog1, meanlog2, sdlog2, and p.mix.

Usage

```
dlnormMix(x, meanlog1 = 0, sdlog1 = 1, meanlog2 = 0, sdlog2 = 1, p.mix = 0.5)
plnormMix(q, meanlog1 = 0, sdlog1 = 1, meanlog2 = 0, sdlog2 = 1, p.mix = 0.5)
qlnormMix(p, meanlog1 = 0, sdlog1 = 1, meanlog2 = 0, sdlog2 = 1, p.mix = 0.5)
rlnormMix(n, meanlog1 = 0, sdlog1 = 1, meanlog2 = 0, sdlog2 = 1, p.mix = 0.5)
```

Arguments

x	vector of quantiles.
q	vector of quantiles.
p	vector of probabilities between 0 and 1.
n	sample size. If length(n) is larger than 1, then length(n) random values are returned.
meanlog1	vector of means of the first lognormal random variable on the log scale. The default is meanlog1=0.
sdlog1	vector of standard deviations of the first lognormal random variable on the log scale. The default is sdlog1=1.
meanlog2	vector of means of the second lognormal random variable on the log scale. The default is meanlog2=0.
sdlog2	vector of standard deviations of the second lognormal random variable on the log scale. The default is sdlog2=1.
p.mix	vector of probabilities between 0 and 1 indicating the mixing proportion. For rlnormMix this must be a single, non-missing number.

Details

Let $f(x; \mu, \sigma)$ denote the density of a [lognormal random variable](#) with parameters meanlog= μ and sdlog= σ . The density, g , of a lognormal mixture random variable with parameters meanlog1= μ_1 , sdlog1= σ_1 , meanlog2= μ_2 , sdlog2= σ_2 , and p.mix= p is given by:

$$g(x; \mu_1, \sigma_1, \mu_2, \sigma_2, p) = (1 - p)f(x; \mu_1, \sigma_1) + pf(x; \mu_2, \sigma_2)$$

Value

dlnormMix gives the density, plnormMix gives the distribution function, qlnormMix gives the quantile function, and rlnormMix generates random deviates.

Note

A lognormal mixture distribution is often used to model positive-valued data that appear to be “contaminated”; that is, most of the values appear to come from a single lognormal distribution, but a few “outliers” are apparent. In this case, the value of meanlog2 would be larger than the value of meanlog1, and the mixing proportion p.mix would be fairly close to 0 (e.g., p.mix=0.1). The value of the second standard deviation (sdlog2) may or may not be the same as the value for the first (sdlog1).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Gilliom, R.J., and D.R. Helsel. (1986). Estimation of Distributional Parameters for Censored Trace Level Water Quality Data: 1. Estimation Techniques. *Water Resources Research* **22**, 135-146.
- Johnson, N. L., S. Kotz, and A.W. Kemp. (1992). *Univariate Discrete Distributions*. Second Edition. John Wiley and Sons, New York, pp.53-54, and Chapter 8.
- Johnson, N. L., S. Kotz, and N. Balakrishnan. (1994). *Continuous Univariate Distributions, Volume 1*. Second Edition. John Wiley and Sons, New York.

See Also

[Lognormal](#), [NormalMix](#), [Probability Distributions and Random Numbers](#).

Examples

```
# Density of a lognormal mixture with parameters meanlog1=0, sdlog1=1,
# meanlog2=2, sdlog2=3, p.mix=0.5, evaluated at 1.5:

dlnormMix(1.5, meanlog1 = 0, sdlog1 = 1, meanlog2 = 2, sdlog2 = 3, p.mix = 0.5)
#[1] 0.1609746

#-----

# The cdf of a lognormal mixture with parameters meanlog1=0, sdlog1=1,
# meanlog2=2, sdlog2=3, p.mix=0.2, evaluated at 4:

plnormMix(4, 0, 1, 2, 3, 0.2)
#[1] 0.8175281

#-----

# The median of a lognormal mixture with parameters meanlog1=0, sdlog1=1,
# meanlog2=2, sdlog2=3, p.mix=0.2:

qlnormMix(0.5, 0, 1, 2, 3, 0.2)
#[1] 1.156891

#-----

# Random sample of 3 observations from a lognormal mixture with
# parameters meanlog1=0, sdlog1=1, meanlog2=3, sdlog2=4, p.mix=0.2.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(20)
rlnormMix(3, 0, 1, 2, 3, 0.2)
#[1] 0.08975283 1.07591103 7.85482514
```

LognormalMixAlt

Mixture of Two Lognormal Distributions (Alternative Parameterization)

Description

Density, distribution function, quantile function, and random generation for a mixture of two log-normal distribution with parameters mean1, cv1, mean2, cv2, and p.mix.

Usage

```
dlnormMixAlt(x, mean1 = exp(1/2), cv1 = sqrt(exp(1) - 1),
             mean2 = exp(1/2), cv2 = sqrt(exp(1) - 1), p.mix = 0.5)
plnormMixAlt(q, mean1 = exp(1/2), cv1 = sqrt(exp(1) - 1),
             mean2 = exp(1/2), cv2 = sqrt(exp(1) - 1), p.mix = 0.5)
qlnormMixAlt(p, mean1 = exp(1/2), cv1 = sqrt(exp(1) - 1),
```

```

mean2 = exp(1/2), cv2 = sqrt(exp(1) - 1), p.mix = 0.5)
rlnormMixAlt(n, mean1 = exp(1/2), cv1 = sqrt(exp(1) - 1),
mean2 = exp(1/2), cv2 = sqrt(exp(1) - 1), p.mix = 0.5)

```

Arguments

x	vector of quantiles.
q	vector of quantiles.
p	vector of probabilities between 0 and 1.
n	sample size. If length(n) is larger than 1, then length(n) random values are returned.
mean1	vector of means of the first lognormal random variable. The default is meanlog1=sqrt(exp(1) - 1).
cv1	vector of coefficient of variations of the first lognormal random variable. The default is sdlog1=sqrt(exp(1) - 1).
mean2	vector of means of the second lognormal random variable. The default is mean2=sqrt(exp(1) - 1).
cv2	vector of coefficient of variations of the second lognormal random variable. The default is sdlog2=sqrt(exp(1) - 1).
p.mix	vector of probabilities between 0 and 1 indicating the mixing proportion. For rlnormMixAlt this must be a single, non-missing number.

Details

Let $f(x; \eta, \theta)$ denote the density of a [lognormal random variable](#) with parameters mean= η and cv= θ . The density, g , of a lognormal mixture random variable with parameters mean1= η_1 , cv1= θ_1 , mean2= η_2 , cv2= θ_2 , and p.mix= p is given by:

$$g(x; \eta_1, \theta_1, \eta_2, \theta_2, p) = (1 - p)f(x; \eta_1, \theta_1) + pf(x; \eta_2, \theta_2)$$

The default values for mean1 and cv1 correspond to a [lognormal distribution](#) with parameters meanlog=0 and sdlog=1. Similarly for the default values of mean2 and cv2.

Value

dlnormMixAlt gives the density, plnormMixAlt gives the distribution function, qlnormMixAlt gives the quantile function, and rlnormMixAlt generates random deviates.

Note

A lognormal mixture distribution is often used to model positive-valued data that appear to be “contaminated”; that is, most of the values appear to come from a single lognormal distribution, but a few “outliers” are apparent. In this case, the value of mean2 would be larger than the value of mean1, and the mixing proportion p.mix would be fairly close to 0 (e.g., p.mix=0.1).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Gilliom, R.J., and D.R. Helsel. (1986). Estimation of Distributional Parameters for Censored Trace Level Water Quality Data: 1. Estimation Techniques. *Water Resources Research* **22**, 135-146.
- Johnson, N. L., S. Kotz, and A.W. Kemp. (1992). *Univariate Discrete Distributions*. Second Edition. John Wiley and Sons, New York, pp.53-54, and Chapter 8.
- Johnson, N. L., S. Kotz, and N. Balakrishnan. (1994). *Continuous Univariate Distributions, Volume 1*. Second Edition. John Wiley and Sons, New York.

See Also

[LognormalAlt](#), [LognormalMix](#), [Lognormal](#), [NormalMix](#), [Probability Distributions and Random Numbers](#).

Examples

```
# Density of a lognormal mixture with parameters mean=2, cv1=3,
# mean2=4, cv2=5, p.mix=0.5, evaluated at 1.5:

dlnormMixAlt(1.5, mean1 = 2, cv1 = 3, mean2 = 4, cv2 = 5, p.mix = 0.5)
#[1] 0.1436045

#-----

# The cdf of a lognormal mixture with parameters mean=2, cv1=3,
# mean2=4, cv2=5, p.mix=0.5, evaluated at 1.5:

plnormMixAlt(1.5, mean1 = 2, cv1 = 3, mean2 = 4, cv2 = 5, p.mix = 0.5)
#[1] 0.6778064

#-----

# The median of a lognormal mixture with parameters mean=2, cv1=3,
# mean2=4, cv2=5, p.mix=0.5:

qlnormMixAlt(0.5, 2, 3, 4, 5, 0.5)
#[1] 0.6978355

#-----

# Random sample of 3 observations from a lognormal mixture with
# parameters mean1=2, cv1=3, mean2=4, cv2=5, p.mix=0.5.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(20)
rlnormMixAlt(3, 2, 3, 4, 5, 0.5)
#[1] 0.70672151 14.43226313 0.05521329
```

LognormalTrunc

The Truncated Lognormal Distribution

Description

Density, distribution function, quantile function, and random generation for the truncated lognormal distribution with parameters meanlog, sdlog, min, and max.

Usage

```

dlnormTrunc(x, meanlog = 0, sdlog = 1, min = 0, max = Inf)
plnormTrunc(q, meanlog = 0, sdlog = 1, min = 0, max = Inf)
qlnormTrunc(p, meanlog = 0, sdlog = 1, min = 0, max = Inf)
rlnormTrunc(n, meanlog = 0, sdlog = 1, min = 0, max = Inf)

```

Arguments

x	vector of quantiles.
q	vector of quantiles.
p	vector of probabilities between 0 and 1.
n	sample size. If length(n) is larger than 1, then length(n) random values are returned.
meanlog	vector of means of the distribution of the non-truncated random variable on the log scale. The default is meanlog=0.
sdlog	vector of (positive) standard deviations of the non-truncated random variable on the log scale. The default is sdlog=1.
min	vector of minimum values for truncation on the left. The default value is min=0.
max	vector of maximum values for truncation on the right. The default value is max=Inf.

Details

See the help file for [the lognormal distribution](#) for information about the density and cdf of a log-normal distribution.

Probability Density and Cumulative Distribution Function

Let X denote a random variable with density function $f(x)$ and cumulative distribution function $F(x)$, and let Y denote the truncated version of X where Y is truncated below at $\text{min}=A$ and above at $\text{max}=B$. Then the density function of Y , denoted $g(y)$, is given by:

$$g(y) = \frac{f(y)}{F(B) - F(A)}, A \leq y \leq B$$

and the cdf of Y , denoted $G(y)$, is given by:

$$G(y) = \begin{cases} 0 & \text{for } y < A \\ \frac{F(y) - F(A)}{F(B) - F(A)} & \text{for } A \leq y \leq B \\ 1 & \text{for } y > B \end{cases}$$

Quantiles

The p^{th} quantile y_p of Y is given by:

$$y_p = \begin{cases} A & \text{for } p = 0 \\ F^{-1}\{p[F(B) - F(A)] + F(A)\} & \text{for } 0 < p < 1 \\ B & \text{for } p = 1 \end{cases}$$

Random Numbers

Random numbers are generated using the inverse transformation method:

$$y = G^{-1}(u)$$

where u is a random deviate from a uniform $[0, 1]$ distribution.

Value

`dlnormTrunc` gives the density, `plnormTrunc` gives the distribution function, `qlnormTrunc` gives the quantile function, and `rlnormTrunc` generates random deviates.

Note

A truncated lognormal distribution is sometimes used as an input distribution for probabilistic risk assessment.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.
- Johnson, N. L., S. Kotz, and N. Balakrishnan. (1994). *Continuous Univariate Distributions, Volume 1*. Second Edition. John Wiley and Sons, New York.
- Schneider, H. (1986). *Truncated and Censored Samples from Normal Populations*. Marcel Dekker, New York, Chapter 2.

See Also

[Lognormal, Probability Distributions and Random Numbers.](#)

Examples

```
# Density of a truncated lognormal distribution with parameters
# meanlog=1, sdlog=0.75, min=0, max=10, evaluated at 2 and 4:

dlnormTrunc(c(2, 4), 1, 0.75, 0, 10)
#[1] 0.2551219 0.1214676

#-----

# The cdf of a truncated lognormal distribution with parameters
# meanlog=1, sdlog=0.75, min=0, max=10, evaluated at 2 and 4:

plnormTrunc(c(2, 4), 1, 0.75, 0, 10)
#[1] 0.3558867 0.7266934

#-----

# The median of a truncated lognormal distribution with parameters
# meanlog=1, sdlog=0.75, min=0, max=10:

qlnormTrunc(.5, 1, 0.75, 0, 10)
```

```

#[1] 2.614945

#-----

# A random sample of 3 observations from a truncated lognormal distribution
# with parameters meanlog=1, sdlog=0.75, min=0, max=10.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(20)
rlnormTrunc(3, 1, 0.75, 0, 10)
#[1] 5.754805 4.372218 1.706815

```

LognormalTruncAlt	<i>The Truncated Lognormal Distribution (Alternative Parameterization)</i>
-------------------	--

Description

Density, distribution function, quantile function, and random generation for the truncated lognormal distribution with parameters mean, cv, min, and max.

Usage

```

dlnormTruncAlt(x, mean = exp(1/2), cv = sqrt(exp(1) - 1), min = 0, max = Inf)
plnormTruncAlt(q, mean = exp(1/2), cv = sqrt(exp(1) - 1), min = 0, max = Inf)
qlnormTruncAlt(p, mean = exp(1/2), cv = sqrt(exp(1) - 1), min = 0, max = Inf)
rlnormTruncAlt(n, mean = exp(1/2), cv = sqrt(exp(1) - 1), min = 0, max = Inf)

```

Arguments

x	vector of quantiles.
q	vector of quantiles.
p	vector of probabilities between 0 and 1.
n	sample size. If length(n) is larger than 1, then length(n) random values are returned.
mean	vector of means of the distribution of the non-truncated random variable. The default is mean=exp(1/2).
cv	vector of (positive) coefficient of variations of the non-truncated random variable. The default is cv=sqrt(exp(1)-1).
min	vector of minimum values for truncation on the left. The default value is min=0.
max	vector of maximum values for truncation on the right. The default value is max=Inf.

Details

See the help file for [LognormalAlt](#) for information about the density and cdf of a lognormal distribution with this alternative parameterization.

Let X denote a random variable with density function $f(x)$ and cumulative distribution function $F(x)$, and let Y denote the truncated version of X where Y is truncated below at $\min=A$ and above at $\max=B$. Then the density function of Y , denoted $g(y)$, is given by:

$$g(y) = \frac{f(y)F(B) - F(A)}{F(B) - F(A)}, A \leq y \leq B$$

and the cdf of Y , denoted $G(y)$, is given by:

$$G(y) = \begin{cases} 0 & \text{for } y < A \\ \frac{F(y)-F(A)}{F(B)-F(A)} & \text{for } A \leq y \leq B \\ 1 & \text{for } y > B \end{cases}$$

The p^{th} quantile y_p of Y is given by:

$$y_p = \begin{cases} A & \text{for } p = 0 \\ F^{-1}\{p[F(B) - F(A)] + F(A)\} & \text{for } 0 < p < 1 \\ B & \text{for } p = 1 \end{cases}$$

Random numbers are generated using the inverse transformation method:

$$y = G^{-1}(u)$$

where u is a random deviate from a uniform $[0, 1]$ distribution.

Value

dlnormTruncAlt gives the density, plnormTruncAlt gives the distribution function, qlnormTruncAlt gives the quantile function, and rlnormTruncAlt generates random deviates.

Note

A truncated lognormal distribution is sometimes used as an input distribution for probabilistic risk assessment.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.
- Johnson, N. L., S. Kotz, and N. Balakrishnan. (1994). *Continuous Univariate Distributions, Volume 1*. Second Edition. John Wiley and Sons, New York.
- Schneider, H. (1986). *Truncated and Censored Samples from Normal Populations*. Marcel Dekker, New York, Chapter 2.

See Also

[LognormalAlt, Probability Distributions and Random Numbers.](#)

Examples

```
# Density of a truncated lognormal distribution with parameters
# mean=10, cv=1, min=0, max=20, evaluated at 2 and 12:

dlnormTruncAlt(c(2, 12), 10, 1, 0, 20)
#[1] 0.08480874 0.03649884
```

```
#-----

# The cdf of a truncated lognormal distribution with parameters
# mean=10, cv=1, min=0, max=20, evaluated at 2 and 12:

plnormTruncAlt(c(2, 4), 10, 1, 0, 20)
#[1] 0.07230627 0.82467603

#-----

# The median of a truncated lognormal distribution with parameters
# mean=10, cv=1, min=0, max=20:

qlnormTruncAlt(.5, 10, 1, 0, 20)
#[1] 6.329505

#-----

# A random sample of 3 observations from a truncated lognormal distribution
# with parameters mean=10, cv=1, min=0, max=20.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(20)
rlnormTruncAlt(3, 10, 1, 0, 20)
#[1] 6.685391 17.445387 18.543553
```

longToWide

Convert a Long Format Data Set into a Wide Format

Description

Given a data frame or matrix in long format, convert it to wide format based on the levels of two variables in the data frame. This is a simplified version of the R function [reshape](#) with the argument `direction="wide"`.

Usage

```
longToWide(x, data.var, row.var, col.var,
  row.labels = levels(factor(x[, row.var])),
  col.labels = levels(factor(x[, col.var])),
  paste.row.name = FALSE, paste.col.name = FALSE, sep = ".",
  check.names = FALSE, ...)
```

Arguments

x	data frame or matrix to convert to wide format. Must have at least 3 columns corresponding to the data variable, row variable, and column variable, respectively.
data.var	character string or numeric scalar indicating column variable name in x for data values.
row.var	character string or numeric scalar indicating column variable name in x for defining rows of output. The indicated column in x cannot have missing values.

<code>col.var</code>	character string or numeric scalar indicating column variable name in <code>x</code> for defining columns of output. The indicated column in <code>x</code> cannot have missing values.
<code>row.labels</code>	optional character vector indicating labels to use for rows. The default value is the levels of the variable indicated by <code>row.var</code> when coerced to a factor.
<code>col.labels</code>	optional character vector indicating labels to use for columns. The default value is the levels of the variable indicated by <code>col.var</code> when coerced to a factor.
<code>paste.row.name</code>	logical scalar indicating whether to paste the name of the variable used to define the row names (i.e., the value of <code>row.var</code>) in front of the values defining the row names. The default value is <code>paste.row.name=FALSE</code> .
<code>paste.col.name</code>	logical scalar indicating whether to paste the name of the variable used to define the column names (i.e., the value of <code>col.var</code>) in front of the values defining the column names. The default value is <code>paste.col.name=FALSE</code> .
<code>sep</code>	character string separator used when <code>paste.row.name=TRUE</code> and/or <code>paste.col.name=TRUE</code> . The default value is <code>sep="."</code> .
<code>check.names</code>	argument to data.frame . Used to convert the return value to a data frame when the argument <code>x</code> is a data frame. This argument is ignored if <code>x</code> is a matrix.
<code>...</code>	other arguments to data.frame . This argument is ignored if <code>x</code> is a matrix.

Details

The combination of values in `x[, row.var]` and `x[, col.var]` must yield n unique values, where n is the number of rows in `x`.

Value

`longToWide` returns a matrix when `x` is a matrix and a data frame when `x` is a data frame. The number of rows is equal to the number of unique values in `x[, row.var]` and the number of columns is equal to the number of unique values in `x[, col.var]`.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>), based on a template from Phil Dixon.

See Also

[reshape](#), [data.frame](#), [matrix](#).

Examples

```
EPA.09.Ex.10.1.nickel.df
#  Month  Well Nickel.ppb
#1     1 Well.1      58.8
#2     3 Well.1       1.0
#3     6 Well.1    262.0
#4     8 Well.1     56.0
#5    10 Well.1      8.7
#6     1 Well.2     19.0
#7     3 Well.2     81.5
#8     6 Well.2    331.0
#9     8 Well.2     14.0
#10    10 Well.2     64.4
```

```
#11    1 Well.3    39.0
#12    3 Well.3   151.0
#13    6 Well.3    27.0
#14    8 Well.3    21.4
#15   10 Well.3   578.0
#16    1 Well.4     3.1
#17    3 Well.4   942.0
#18    6 Well.4    85.6
#19    8 Well.4    10.0
#20   10 Well.4   637.0
```

```
longToWide(EPA.09.Ex.10.1.nickel.df,
  "Nickel.ppb", "Month", "Well", paste.row.name = TRUE)
#      Well.1 Well.2 Well.3 Well.4
#Month.1    58.8   19.0   39.0    3.1
#Month.3     1.0   81.5  151.0  942.0
#Month.6   262.0  331.0   27.0   85.6
#Month.8    56.0   14.0   21.4   10.0
#Month.10    8.7   64.4  578.0  637.0
```

Millard.Deverel.88.df *Copper and Zinc Concentrations in Shallow Ground Water*

Description

Copper and zinc concentrations (mg/L) in shallow ground water from two geological zones (Alluvial Fan and Basin-Trough) in the San Joaquin Valley, CA. There are 68 samples from the Alluvial Fan zone and 50 from the Basin-Trough zone. Some observations are reported as $<DL$, where DL denotes a detection limit. There are multiple detection limits for both the copper and zinc data in each of the geological zones.

Usage

Millard.Deverel.88.df

Format

A data frame with 118 observations on the following 8 variables.

`Cu.orig` a character vector of original copper concentrations (mg/L)

`Cu` a numeric vector of copper concentrations with nondetects coded to their detection limit

`Cu.censored` a logical vector indicating which copper concentrations are censored

`Zn.orig` a character vector of original zinc concentrations (mg/L)

`Zn` a numeric vector of zinc concentrations with nondetects coded to their detection limit

`Zn.censored` a logical vector indicating which zinc concentrations are censored

`Zone` a factor indicating the zone (alluvial fan vs. basin trough)

`Location` a numeric vector indicating the sampling location

Source

Millard, S.P., and S.J. Deverel. (1988). Nonparametric Statistical Methods for Comparing Two Sites Based on Data With Multiple Nondetect Limits. *Water Resources Research*, **24**(12), 2087-2098.

References

Deverel, S.J., R.J. Gilliom, R. Fujii, J.A. Izbicki, and J.C. Fields. (1984). *Areal Distribution of Selenium and Other Inorganic Constituents in Shallow Ground Water of the San Luis Drain Service Area, San Joaquin, California: A Preliminary Study*. U.S. Geological Survey Water Resources Investigative Report 84-4319.

Modified.TcCB.df

Modified 1,2,3,4-Tetrachlorobenzene Data with Censored Values

Description

Artificial 1,2,3,4-Tetrachlorobenzene (TcCB) concentrations with censored values; based on the reference area data stored in [EPA.94b.tccb.df](#). The data frame [EPA.94b.tccb.df](#) contains TcCB concentrations (ppb) in soil samples at a reference area and a cleanup area. The data frame [Modified.TcCB.df](#) contains a modified version of the data from the reference area. For this data set, the concentrations of TcCB less than 0.5 ppb have been recoded as <0.5.

Usage

[Modified.TcCB.df](#)

Format

A data frame with 47 observations on the following 3 variables.

[TcCB.orig](#) a character vector of original TcCB concentrations (ppb)

[TcCB](#) a numeric vector with censored observations set to their detection level

[Censored](#) a logical vector indicating which observations are censored

Source

Millard, S.P., and N.K. Neerchal. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, FL, p.595.

References

USEPA. (1994b). *Statistical Methods for Evaluating the Attainment of Cleanup Standards, Volume 3: Reference-Based Standards for Soils and Solid Media*. EPA/230-R-94-004. Office of Policy, Planning, and Evaluation, U.S. Environmental Protection Agency, Washington, D.C.

See Also

[EPA.94b.tccb.df](#).

NIOSH.89.air.lead.vec *NIOSH Air Lead Levels Data*

Description

Air lead levels collected by the National Institute for Occupational Safety and Health (NIOSH) at 15 different areas within the Alma American Labs, Fairplay, CO, for health hazard evaluation (HETA 89-052) on February 23, 1989.

Usage

NIOSH.89.air.lead.vec

Format

A numeric vector with 15 elements containing air lead concentrations ($\mu\text{g}/\text{m}^3$).

Source

Krishnamoorthy, K., T. Matthew, and G. Ramachandran. (2006). Generalized P-Values and Confidence Intervals: A Novel Approach for Analyzing Lognormally Distributed Exposure Data. *Journal of Occupational and Environmental Hygiene*, **3**, 642–650.

References

Zou, G.Y., C.Y. Huo, and J. Taleban. (2009). Simple Confidence Intervals for Lognormal Means and their Differences with Environmental Applications. *Environmetrics*, **20**, 172–180.

NormalMix

Mixture of Two Normal Distributions

Description

Density, distribution function, quantile function, and random generation for a mixture of two normal distribution with parameters mean1, sd1, mean2, sd2, and p.mix.

Usage

```
dnormMix(x, mean1 = 0, sd1 = 1, mean2 = 0, sd2 = 1, p.mix = 0.5)
pnormMix(q, mean1 = 0, sd1 = 1, mean2 = 0, sd2 = 1, p.mix = 0.5)
qnormMix(p, mean1 = 0, sd1 = 1, mean2 = 0, sd2 = 1, p.mix = 0.5)
rnormMix(n, mean1 = 0, sd1 = 1, mean2 = 0, sd2 = 1, p.mix = 0.5)
```

Arguments

x	vector of quantiles.
q	vector of quantiles.
p	vector of probabilities between 0 and 1.
n	sample size. If length(n) is larger than 1, then length(n) random values are returned.
mean1	vector of means of the first normal random variable. The default is mean1=0.
sd1	vector of standard deviations of the first normal random variable. The default is sd1=1.
mean2	vector of means of the second normal random variable. The default is mean2=0.
sd2	vector of standard deviations of the second normal random variable. The default is sd2=1.
p.mix	vector of probabilities between 0 and 1 indicating the mixing proportion. For rnormMix this must be a single, non-missing number.

Details

Let $f(x; \mu, \sigma)$ denote the density of a [normal random variable](#) with parameters mean= μ and sd= σ . The density, g , of a normal mixture random variable with parameters mean1= μ_1 , sd1= σ_1 , mean2= μ_2 , sd2= σ_2 , and p.mix= p is given by:

$$g(x; \mu_1, \sigma_1, \mu_2, \sigma_2, p) = (1 - p)f(x; \mu_1, \sigma_1) + pf(x; \mu_2, \sigma_2)$$

Value

dnormMix gives the density, pnormMix gives the distribution function, qnormMix gives the quantile function, and rnormMix generates random deviates.

Note

A normal mixture distribution is sometimes used to model data that appear to be “contaminated”; that is, most of the values appear to come from a single normal distribution, but a few “outliers” are apparent. In this case, the value of mean2 would be larger than the value of mean1, and the mixing proportion p.mix would be fairly close to 0 (e.g., p.mix=0.1). The value of the second standard deviation (sd2) may or may not be the same as the value for the first (sd1).

Another application of the normal mixture distribution is to bi-modal data; that is, data exhibiting two modes.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Johnson, N. L., S. Kotz, and A.W. Kemp. (1992). *Univariate Discrete Distributions*. Second Edition. John Wiley and Sons, New York, pp.53-54, and Chapter 8.
- Johnson, N. L., S. Kotz, and N. Balakrishnan. (1994). *Continuous Univariate Distributions, Volume 1*. Second Edition. John Wiley and Sons, New York.

See Also

[Normal](#), [LognormalMix](#), [Probability Distributions and Random Numbers](#).

Examples

```
# Density of a normal mixture with parameters mean1=0, sd1=1,
# mean2=4, sd2=2, p.mix=0.5, evaluated at 1.5:

dnormMix(1.5, mean2=4, sd2=2)
#[1] 0.1104211

#-----

# The cdf of a normal mixture with parameters mean1=10, sd1=2,
# mean2=20, sd2=2, p.mix=0.1, evaluated at 15:

pnormMix(15, 10, 2, 20, 2, 0.1)
#[1] 0.8950323

#-----

# The median of a normal mixture with parameters mean1=10, sd1=2,
# mean2=20, sd2=2, p.mix=0.1:

qnormMix(0.5, 10, 2, 20, 2, 0.1)
#[1] 10.27942

#-----

# Random sample of 3 observations from a normal mixture with
# parameters mean1=0, sd1=1, mean2=4, sd2=2, p.mix=0.5.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(20)
rnormMix(3, mean2=4, sd2=2)
#[1] 0.07316778 2.06112801 1.05953620
```

NormalTrunc

The Truncated Normal Distribution

Description

Density, distribution function, quantile function, and random generation for the truncated normal distribution with parameters mean, sd, min, and max.

Usage

```
dnormTrunc(x, mean = 0, sd = 1, min = -Inf, max = Inf)
pnormTrunc(q, mean = 0, sd = 1, min = -Inf, max = Inf)
qnormTrunc(p, mean = 0, sd = 1, min = -Inf, max = Inf)
rnormTrunc(n, mean = 0, sd = 1, min = -Inf, max = Inf)
```

Arguments

x	vector of quantiles.
q	vector of quantiles.
p	vector of probabilities between 0 and 1.
n	sample size. If length(n) is larger than 1, then length(n) random values are returned.
mean	vector of means of the distribution of the non-truncated random variable. The default is mean=0.
sd	vector of (positive) standard deviations of the non-truncated random variable. The default is sd=1.
min	vector of minimum values for truncation on the left. The default value is min=-Inf.
max	vector of maximum values for truncation on the right. The default value is max=Inf.

Details

See the help file for [the normal distribution](#) for information about the density and cdf of a normal distribution.

Probability Density and Cumulative Distribution Function

Let X denote a random variable with density function $f(x)$ and cumulative distribution function $F(x)$, and let Y denote the truncated version of X where Y is truncated below at $\min=A$ and above at $\max=B$. Then the density function of Y , denoted $g(y)$, is given by:

$$g(y) = \frac{f(y)(F(B) - F(A))}{F(B) - F(A)}, A \leq y \leq B$$

and the cdf of Y , denoted $G(y)$, is given by:

$$G(y) = \begin{cases} 0 & \text{for } y < A \\ \frac{F(y) - F(A)}{F(B) - F(A)} & \text{for } A \leq y \leq B \\ 1 & \text{for } y > B \end{cases}$$

Quantiles

The p^{th} quantile y_p of Y is given by:

$$y_p = \begin{cases} A & \text{for } p = 0 \\ F^{-1}\{p[F(B) - F(A)] + F(A)\} & \text{for } 0 < p < 1 \\ B & \text{for } p = 1 \end{cases}$$

Random Numbers

Random numbers are generated using the inverse transformation method:

$$y = G^{-1}(u)$$

where u is a random deviate from a uniform $[0, 1]$ distribution.

Mean and Variance

The expected value of a truncated normal random variable with parameters $\text{mean}=\mu$, $\text{sd}=\sigma$, $\min=A$,

and $\max=B$ is given by:

$$E(Y) = \mu + \sigma^2 \frac{f(A) - f(B)}{F(B) - F(A)}$$

(Johnson et al., 1994, p.156; Schneider, 1986, p.17).

The variance of this random variable is given by:

$$\sigma^2 + \sigma^3 \{z_A f(A) - z_B f(B) - \sigma[f(A) - f(B)]^2\}$$

where

$$z_A = \frac{A - \mu}{\sigma}; z_B = \frac{B - \mu}{\sigma}$$

(Johnson et al., 1994, p.158; Schneider, 1986, p.17).

Value

`dnormTrunc` gives the density, `pnormTrunc` gives the distribution function, `qnormTrunc` gives the quantile function, and `rnormTrunc` generates random deviates.

Note

A truncated normal distribution is sometimes used as an input distribution for probabilistic risk assessment.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.
- Johnson, N. L., S. Kotz, and N. Balakrishnan. (1994). *Continuous Univariate Distributions, Volume 1*. Second Edition. John Wiley and Sons, New York.
- Schneider, H. (1986). *Truncated and Censored Samples from Normal Populations*. Marcel Dekker, New York, Chapter 2.

See Also

[Normal, Probability Distributions and Random Numbers.](#)

Examples

```
# Density of a truncated normal distribution with parameters
# mean=10, sd=2, min=8, max=13, evaluated at 10 and 11.5:

dnormTrunc(c(10, 11.5), 10, 2, 8, 13)
#[1] 0.2575358 0.1943982

#-----

# The cdf of a truncated normal distribution with parameters
# mean=10, sd=2, min=8, max=13, evaluated at 10 and 11.5:

pnormTrunc(c(10, 11.5), 10, 2, 8, 13)
```

```

#[1] 0.4407078 0.7936573

#-----

# The median of a truncated normal distribution with parameters
# mean=10, sd=2, min=8, max=13:

qnormTrunc(.5, 10, 2, 8, 13)
#[1] 10.23074

#-----

# A random sample of 3 observations from a truncated normal distribution
# with parameters mean=10, sd=2, min=8, max=13.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(20)
rnormTrunc(3, 10, 2, 8, 13)
#[1] 11.975223 11.373711 9.361258

```

Olympic.NH4.df

Ammonium Concentration in Precipitation Measured at Olympic National Park Hoh Ranger Station

Description

Ammonium (NH₄) concentration (mg/L) in precipitation measured at Olympic National Park, Hoh Ranger Station (WA14), weekly or every other week from January 6, 2009 through December 20, 2011.

Usage

```
Olympic.NH4.df
```

Format

A data frame with 102 observations on the following 6 variables.

Date.On Start of collection period. Date on which the sample bucket was installed on the collector.

Date.Off End of collection period. Date on which the sample bucket was removed from the collector.

Week a numeric vector indicating the cumulative week number starting from January 1, 2009.

NH4.Orig.mg.per.L a character vector of the original NH₄ concentrations reported either as the observed value or less than some detection limit. For values reported as less than a detection limit, the value reported is the actual limit of detection or, in the case of a diluted sample, the product of the detection limit value and the dilution factor.

NH4.mg.per.L a numeric vector of NH₄ concentrations with non-detects coded to their detection limit.

Censored a logical vector indicating which observations are censored.

Details

- StationOlympic National Park-Hoh Ranger Station (WA14)
- LocationJefferson County, Washington
- Latitude47.8597
- Longitude-123.9325
- Elevation182 meters
- USGS 1:24000 Map NameOwl Mountain
- Operating AgencyOlympic National Park
- Sponsoring AgencyNPS-Air Resources Division

Source

National Atmospheric Deposition Program, National Trends Network (NADP/NTN).
<http://nadp.isws.illinois.edu/sites/siteinfo.asp?id=WA14&net=NTN>
<http://nadp.isws.illinois.edu/NTN/>

Ozone.NE.df

*Ozone Concentrations in the Northeast U.S.***Description**

Ozone concentrations in 41 U.S. cities based on daily maxima collected between June and August 1974.

Usage

Ozone.NE.df

Format

A data frame with 41 observations on the following 5 variables.

Median median of daily maxima ozone concentration (ppb).

Quartile Upper quartile (i.e., 75th percentile) of daily maxima ozone concentration (ppb).

City a factor indicating the city

Longitude negative longitude of the city

Latitude latitude of the city

Source

Cleveland, W.S., Kleiner, B., McRae, J.E., Warner, J.L., and Pasceri, P.E. (1975). *The Analysis of Ground-Level Ozone Data from New Jersey, New York, Connecticut, and Massachusetts: Data Quality Assessment and Temporal and Geographical Properties*. Bell Laboratories Memorandum.

The original data were collected by the New Jersey Department of Environmental Protection, the New York State Department of Environmental Protection, the Boyce Thompson Institute (Yonkers, for New York data), the Connecticut Department of Environmental Protection, and the Massachusetts Department of Public Health.

Examples

```
summary(Ozone.NE.df)
#      Median      Quartile      City      Longitude
# Min.   : 34.00   Min.   : 48.00   Asbury Park: 1   Min.   : -74.71
# 1st Qu.: 58.00   1st Qu.: 79.75   Babylon   : 1   1st Qu.: -73.74
# Median : 65.00   Median : 90.00   Bayonne   : 1   Median : -73.17
# Mean   : 68.15   Mean   : 95.10   Boston    : 1   Mean   : -72.94
# 3rd Qu.: 80.00   3rd Qu.:112.25   Bridgeport: 1   3rd Qu.: -72.08
# Max.   :100.00   Max.   :145.00   Cambridge : 1   Max.   : -71.05
#
#      NA's      : 1.00   (Other)   :35
#      Latitude
# Min.   :40.22
# 1st Qu.:40.97
# Median :41.56
# Mean   :41.60
# 3rd Qu.:42.25
# Max.   :43.32
```

Pareto

The Pareto Distribution

Description

Density, distribution function, quantile function, and random generation for the Pareto distribution with parameters location and shape.

Usage

```
dpareto(x, location, shape = 1)
ppareto(q, location, shape = 1)
qpareto(p, location, shape = 1)
rpareto(n, location, shape = 1)
```

Arguments

x	vector of quantiles.
q	vector of quantiles.
p	vector of probabilities between 0 and 1.
n	sample size. If length(n) is larger than 1, then length(n) random values are returned.
location	vector of (positive) location parameters.
shape	vector of (positive) shape parameters. The default is shape=1.

Details

Let X be a Pareto random variable with parameters location= η and shape= θ . The density function of X is given by:

$$f(x; \eta, \theta) = \frac{\theta \eta^\theta}{x^{\theta+1}}, \quad \eta > 0, \theta > 0, x \geq \eta$$

The cumulative distribution function of X is given by:

$$F(x; \eta, \theta) = 1 - \left(\frac{\eta}{x}\right)^\theta$$

and the p 'th quantile of X is given by:

$$x_p = \eta(1 - p)^{-1/\theta}, \quad 0 \leq p \leq 1$$

The mode, mean, median, variance, and coefficient of variation of X are given by:

$$\text{Mode}(X) = \eta$$

$$E(X) = \frac{\theta\eta}{\theta - 1}, \quad \theta > 1$$

$$\text{Median}(X) = x_{0.5} = 2^{1/\theta}\eta$$

$$\text{Var}(X) = \frac{\theta\eta^2}{(\theta - 1)^2(\theta - 1)}, \quad \theta > 2$$

$$\text{CV}(X) = [\theta(\theta - 2)]^{-1/2}, \quad \theta > 2$$

Value

dpareto gives the density, ppareto gives the distribution function, qpareto gives the quantile function, and rpareto generates random deviates.

Note

The Pareto distribution is named after Vilfredo Pareto (1848-1923), a professor of economics. It is derived from Pareto's law, which states that the number of persons N having income $\geq x$ is given by:

$$N = Ax^{-\theta}$$

where θ denotes Pareto's constant and is the shape parameter for the probability distribution.

The Pareto distribution takes values on the positive real line. All values must be larger than the "location" parameter η , which is really a threshold parameter. There are three kinds of Pareto distributions. The one described here is the Pareto distribution of the first kind. Stable Pareto distributions have $0 < \theta < 2$. Note that the r 'th moment only exists if $r < \theta$.

The Pareto distribution is related to the [exponential distribution](#) and [logistic distribution](#) as follows. Let X denote a Pareto random variable with location= η and shape= θ . Then $\log(X/\eta)$ has an exponential distribution with parameter rate= θ , and $-\log\{[(X/\eta)^\theta] - 1\}$ has a logistic distribution with parameters location=0 and scale=1.

The Pareto distribution has a very long right-hand tail. It is often applied in the study of socioeconomic data, including the distribution of income, firm size, population, and stock price fluctuations.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.
- Johnson, N. L., S. Kotz, and N. Balakrishnan. (1994). *Continuous Univariate Distributions, Volume 1*. Second Edition. John Wiley and Sons, New York.

See Also

[epareto](#), [eqpareto](#), [Exponential](#), [Probability Distributions and Random Numbers](#).

Examples

```
# Density of a Pareto distribution with parameters location=1 and shape=1,
# evaluated at 2, 3 and 4:

dpareto(2:4, 1, 1)
#[1] 0.2500000 0.1111111 0.0625000

#-----

# The cdf of a Pareto distribution with parameters location=2 and shape=1,
# evaluated at 3, 4, and 5:

ppareto(3:5, 2, 1)
#[1] 0.3333333 0.5000000 0.6000000

#-----

# The 25'th percentile of a Pareto distribution with parameters
# location=1 and shape=1:

qpareto(0.25, 1, 1)
#[1] 1.333333

#-----

# A random sample of 4 numbers from a Pareto distribution with parameters
# location=3 and shape=2.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(10)
rpareto(4, 3, 2)
#[1] 4.274728 3.603148 3.962862 5.415322
```

pdfPlot

Plot Probability Density Function

Description

Produce a probability density function (pdf) plot for a user-specified distribution.

Usage

```
pdfPlot(distribution = "norm", param.list = list(mean = 0, sd = 1),
  left.tail.cutoff = ifelse(is.finite(supp.min), 0, 0.001),
  right.tail.cutoff = ifelse(is.finite(supp.max), 0, 0.001),
  plot.it = TRUE, add = FALSE, n.points = 1000, pdf.col = "black",
  pdf.lwd = 3 * par("cex"), pdf.lty = 1, curve.fill = !add,
  curve.fill.col = "cyan", x.ticks.at.all.x.max = 15,
  hist.col = ifelse(add, "black", "cyan"), density = 5,
```

```
digits = .Options$digits, ..., type = "l", main = NULL, xlab = NULL,
ylab = NULL, xlim = NULL, ylim = NULL)
```

Arguments

- | | |
|----------------------|--|
| distribution | a character string denoting the distribution abbreviation. The default value is <code>distribution="norm"</code> . See the help file for Distribution.df for a list of possible distribution abbreviations. |
| param.list | a list with values for the parameters of the distribution. The default value is <code>param.list=list(mean=0, sd=1)</code> . See the help file for Distribution.df for the names and possible values of the parameters associated with each distribution. |
| left.tail.cutoff | a numeric scalar indicating what proportion of the left-tail of the probability distribution to omit from the plot. For densities with a finite support minimum (e.g., Lognormal) the default value is 0; for all other densities the default value is 0.001. |
| right.tail.cutoff | a scalar indicating what proportion of the right-tail of the probability distribution to omit from the plot. For densities with a finite support maximum (e.g., Binomial) the default value is 0; for all other densities the default value is 0.001. |
| plot.it | a logical scalar indicating whether to create a plot or add to the existing plot (see <code>add</code>) on the current graphics device. If <code>plot.it=FALSE</code> , no plot is produced, but a list of (x, y) values is returned (see the section <code>VALUE</code> below). The default value is <code>plot.it=TRUE</code> . |
| add | a logical scalar indicating whether to add the probability density curve to the existing plot (<code>add=TRUE</code>), or to create a new plot (<code>add=FALSE</code> ; the default). This argument is ignored if <code>plot.it=FALSE</code> . |
| n.points | a numeric scalar specifying at how many evenly-spaced points the probability density function will be evaluated. The default value is <code>n.points=1000</code> . |
| pdf.col | for continuous distributions, a numeric scalar or character string determining the color of the pdf line in the plot. The default value is <code>pdf.col="black"</code> . See the entry for <code>col</code> in the help file for par for more information. |
| pdf.lwd | for continuous distributions, a numeric scalar determining the width of the pdf line in the plot. The default value is <code>pdf.lwd=3*par("cex")</code> . See the entry for <code>lwd</code> in the help file for par for more information. |
| pdf.lty | for continuous distributions, a numeric scalar determining the line type of the pdf line in the plot. The default value is <code>pdf.lty=1</code> . See the entry for <code>lty</code> in the help file for par for more information. |
| curve.fill | for continuous distributions, a logical value indicating whether to fill in the area below the probability density curve with the color specified by <code>curve.fill.col</code> . The default value is <code>TRUE</code> unless <code>add=TRUE</code> . |
| curve.fill.col | for continuous distributions, when <code>curve.fill=TRUE</code> , a numeric scalar or character string indicating what color to use to fill in the area below the probability density curve. The default value is <code>curve.fill.col="cyan"</code> . See the entry for <code>col</code> in the help file for par for more information. |
| x.ticks.at.all.x.max | a numeric scalar indicating the maximum number of ticks marks on the x -axis. The default value is <code>x.ticks.at.all.x.max=15</code> . |

hist.col	for discrete distributions, a numeric scalar or character string indicating what color to use to fill in the histogram if add=FALSE, or the color of the shading lines if add=TRUE. The default is "cyan" if add=FALSE and "black" if add=TRUE. See the entry for col in the help file for par for more information.
density	for discrete distributions, a scalar indicating the density of line shading for the histogram when add=TRUE. This argument is ignored if add=FALSE.
digits	a scalar indicating how many significant digits to print for the distribution parameters. The default value is digits=.Options\$digits.
type, main, xlab, ylab, xlim, ylim, ...	additional graphical parameters. See plot.default and par).

Details

The **probability density function (pdf)** of a random variable X , usually denoted f , is defined as:

$$f(x) = \frac{dF(x)}{dx} \quad (1)$$

where F is the cumulative distribution function (cdf) of X . That is, $f(x)$ is the derivative of the cdf F with respect to x (where this derivative exists).

For discrete distributions, the probability density function is simply:

$$f(x) = Pr(X = x) \quad (2)$$

In this case, f is sometimes called the **probability function** or **probability mass function**.

The probability that the random variable X takes on a value in the interval $[a, b]$ is simply the (Lebesgue) integral of the pdf evaluated between a and b . That is,

$$Pr(a \leq X \leq b) = \int_a^b f(x)dx \quad (3)$$

For discrete distributions, Equation (3) translates to summing up the probabilities of all values in this interval:

$$Pr(a \leq X \leq b) = \sum_{x \in [a, b]} f(x) = \sum_{x \in [a, b]} Pr(X = x) \quad (4)$$

A **probability density function (pdf) plot** plots the values of the pdf against quantiles of the specified distribution. Theoretical pdf plots are sometimes plotted along with [empirical pdf plots](#) (density plots), histograms or bar graphs to visually assess whether data have a particular distribution.

Value

pdfPlot invisibly returns a list giving coordinates of the points that have been or would have been plotted:

Quantiles	The quantiles used for the plot.
Probability.Densities	The values of the pdf associated with the quantiles.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.
- Johnson, N. L., S. Kotz, and A.W. Kemp. (1992). *Univariate Discrete Distributions, Second Edition*. John Wiley and Sons, New York.
- Johnson, N. L., S. Kotz, and N. Balakrishnan. (1994). *Continuous Univariate Distributions, Volume 1*. Second Edition. John Wiley and Sons, New York.
- Johnson, N. L., S. Kotz, and N. Balakrishnan. (1995). *Continuous Univariate Distributions, Volume 2*. Second Edition. John Wiley and Sons, New York.

See Also

[Distribution.df](#), [epdfPlot](#), [cdfPlot](#).

Examples

```
# Plot the pdf of the standard normal distribution
#-----
windows()
pdfPlot()

#=====

# Plot the pdf of the standard normal distribution
# and a N(2, 2) distribution on the sample plot.
#-----
windows()
pdfPlot(param.list = list(mean=2, sd=2),
        curve.fill = FALSE, ylim = c(0, dnorm(0)), main = "")

pdfPlot(add = TRUE, pdf.col = "red")

legend("topright", legend = c("N(2,2)", "N(0,1)"),
      col = c("black", "red"), lwd = 3 * par("cex"))

title("PDF Plots for Two Normal Distributions")

#=====

# Clean up
#-----
graphics.off()
```

plot.boxcox

Plot Results of Box-Cox Transformations

Description

Plot the results of calling the function [boxcox](#), which returns an object of class "boxcox". Three different kinds of plots are available.

The function `plot.boxcox` is automatically called by [plot](#) when given an object of class "boxcox". The names of other functions associated with Box-Cox transformations are listed under [Data Transformations](#).

Usage

```
## S3 method for class 'boxcox'
plot(x, plot.type = "Objective vs. lambda", same.window = TRUE,
      ask = same.window & plot.type != "Objective vs. lambda",
      plot.pos.con = 0.375, estimate.params = FALSE,
      equal.axes = qq.line.type == "0-1" || estimate.params, add.line = TRUE,
      qq.line.type = "least squares", duplicate.points.method = "standard",
      points.col = 1, line.col = 1, line.lwd = par("cex"), line.lty = 1,
      digits = .Options$digits, cex.main = 1.4 * par("cex"), cex.sub = par("cex"),
      main = NULL, sub = NULL, xlab = NULL, ylab = NULL, xlim = NULL,
      ylim = NULL, ...)
```

Arguments

<code>x</code>	an object of class "boxcox". See boxcox.object for details.
<code>plot.type</code>	character string indicating what kind of plot to create. Only one particular plot type will be created, unless <code>plot.type="All"</code> , in which case all plots will be created sequentially. The possible values of <code>plot.type</code> are: "Objective vs. lambda" (the default), "Q-Q Plots", "Tukey M-D Q-Q Plots", and "All".
<code>same.window</code>	logical scalar indicating whether to produce all plots in the same graphics window (<code>same.window=TRUE</code> ; the default), or to create a new graphics window for each separate plot (<code>same.window=FALSE</code>). The argument is relevant only when <code>plot.type</code> produces more than one plot (i.e., when <code>plot.type</code> is not equal to "Objective vs. lambda").
<code>ask</code>	logical scalar supplied to the function devAskNewPage , indicating whether to prompt the user before creating a new plot within a single graphics window. This argument is ignored when <code>plot.type="Objective vs. lambda"</code> (since only one plot is produced) or when <code>same.window=FALSE</code> , otherwise the default value is <code>ask=TRUE</code> .
<code>points.col</code>	numeric scalar determining the color of the points in the plot. The default value is <code>points.col=1</code> . See the entry for <code>col</code> in the R help file for par for more information. Arguments associated with <code>plot.type="Q-Q Plots"</code> , <code>plot.type="Tukey M-D Q-Q Plots"</code> , or <code>plot.type="All"</code> (supplied to qqPlot):
<code>plot.pos.con</code>	numeric scalar between 0 and 1 containing the value of the plotting position constant used to construct the Q-Q plots and/or Tukey Mean-Difference Q-Q plots. The default value is <code>plot.pos.con=0.375</code> . See the help files for qqPlot for more information and the motivation for this choice.
<code>estimate.params</code>	logical scalar indicating whether to compute quantiles based on estimating the distribution parameters (<code>estimate.params=TRUE</code>) or using the distribution parameters for a standard normal distribution (i.e, <code>mean=0</code> , <code>sd=1</code>). The default value is <code>estimate.params=FALSE</code> because a standard normal Q-Q plot will yield roughly a straight line if the observations are from <i>any</i> normal distribution. If you specify <code>plot.type="Tukey M-D Q-Q Plots"</code> , then you need to set <code>estimate.params=TRUE</code> unless you want to assume the transformed data come from a standard normal distribution.
<code>equal.axes</code>	logical scalar indicating whether to use the same range on the <i>x</i> - and <i>y</i> -axes when <code>plot.type="Q-Q Plots"</code> . The default value is <code>TRUE</code> if <code>qq.line.type="0-1"</code> or <code>estimate.params=TRUE</code> , otherwise it is <code>FALSE</code> .

<code>add.line</code>	logical scalar indicating whether to add a line to the plot. If <code>add.line=TRUE</code> and <code>plot.type="Q-Q Plots"</code> , a line determined by the value of <code>qq.line.type</code> is added to the plot. If <code>add.line=TRUE</code> and <code>plot.type="Tukey M-D Q-Q Plots"</code> , a horizontal line at $y = 0$ is added to the plot. The default value is <code>add.line=TRUE</code> .
<code>qq.line.type</code>	character string determining what kind of line to add to the plot when <code>plot.type="Q-Q Plots"</code> . Possible values are "least squares" (a least squares line; the default), "0-1" (a line with intercept 0 and slope 1), and "robust" (a line is fit through the first and third quartiles of the x and y data). This argument is ignored if <code>add.line=FALSE</code> .
<code>duplicate.points.method</code>	a character string denoting how to plot points with duplicate (x, y) values. Possible values are "standard" (a single plotting symbol is plotted; the default), "jitter" (a separate plotting symbol is plotted for each duplicate point, where the plotting symbols cluster around the true value of x and y), and "number" (a single number is plotted at (x, y) that represents how many duplicate points are at that (x, y) coordinate).
<code>line.col</code>	numeric scalar determining the color of the line in the plot. The default value is <code>line.col=1</code> . See the entry for <code>col</code> in the R help file for par for more information. This argument is ignored if <code>add.line=FALSE</code> .
<code>line.lwd</code>	numeric scalar determining the width of the line in the plot. The default value is <code>line.lwd=par("cex")</code> . See the entry for <code>lwd</code> in the R help file for par for more information. This argument is ignored if <code>add.line=FALSE</code> .
<code>line.lty</code>	numeric scalar determining the line type (style) of the line in the plot. The default value is <code>line.lty=1</code> . See the entry for <code>lty</code> in the R help file for par for more information. This argument is ignored if <code>add.line=FALSE</code> .
<code>digits</code>	scalar indicating how many significant digits to print for the distribution parameters and the value of the objective in the sub-title. The default value is the current setting of <code>options("digits")</code> .
<code>cex.main, cex.sub, main, sub, xlab, ylab, xlim, ylim, ...</code>	Graphics parameters: graphics parameters; see par for more information. The default value of <code>cex.main</code> is <code>cex.main=1.4 * par("cex")</code> . The default value of <code>cex.sub</code> is <code>cex.sub=par("cex")</code> .

Details

The function `plot.bboxcox` is a method for the generic function `plot` for the class "bboxcox" (see `boxcox.object`). It can be invoked by calling `plot` and giving it an object of class "bboxcox" as the first argument, or by calling `plot.bboxcox` directly, regardless of the class of the object given as the first argument to `plot.bboxcox`.

Plots associated with Box-Cox transformations are produced on the current graphics device. These can be one or all of the following:

- Objective vs. λ .
- Observed Quantiles vs. Normal Quantiles (Q-Q Plot) for the transformed observations for each of the values of λ .
- Tukey Mean-Difference Q-Q Plots for the transformed observations for each of the values of λ .

See the help files for `boxcox` and `qqPlot` for more information.

Value

`plot.bboxcox` invisibly returns the first argument, x .

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Chambers, J. M. and Hastie, T. J. (1992). *Statistical Models in S*. Wadsworth & Brooks/Cole.

See Also

[qqPlot](#), [boxcox](#), [boxcox.object](#), [print.boxcox](#), [Data Transformations](#), [plot](#).

Examples

```
# Generate 30 observations from a lognormal distribution with
# mean=10 and cv=2, call the function boxcox, and then plot
# the results.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
x <- rlnormAlt(30, mean = 10, cv = 2)

# Plot the results based on the PPCC objective
#-----
boxcox.list <- boxcox(x)
windows()
plot(boxcox.list)

# Look at Q-Q Plots for the candidate values of lambda
#-----
plot(boxcox.list, plot.type = "Q-Q Plots", same.window = FALSE)

# Look at Tukey Mean-Difference Q-Q Plots
# for the candidate values of lambda
#-----
plot(boxcox.list, plot.type = "Tukey M-D Q-Q Plots", same.window = FALSE)

#=====

# Clean up
#-----
rm(x, boxcox.list)
graphics.off()
```

plot.boxcoxCensored	<i>Plot Results of Box-Cox Transformations Based on Type I Censored Data</i>
---------------------	--

Description

Plot the results of calling the function [boxcoxCensored](#), which returns an object of class "boxcoxCensored". Three different kinds of plots are available.

The function `plot.boxcoxCensored` is automatically called by [plot](#) when given an object of class "boxcoxCensored".

Usage

```
## S3 method for class 'boxcoxCensored'
plot(x, plot.type = "Objective vs. lambda", same.window = TRUE,
     ask = same.window & plot.type != "Objective vs. lambda",
     prob.method = "michael-schucany", plot.pos.con = 0.375, estimate.params = FALSE,
     equal.axes = qq.line.type == "0-1" || estimate.params, add.line = TRUE,
     qq.line.type = "least squares", duplicate.points.method = "standard",
     points.col = 1, line.col = 1, line.lwd = par("cex"), line.lty = 1,
     digits = .Options$digits, cex.main = 1.4 * par("cex"), cex.sub = par("cex"),
     main = NULL, sub = NULL, xlab = NULL, ylab = NULL, xlim = NULL,
     ylim = NULL, ...)
```

Arguments

x	an object of class "boxcoxCensored". See boxcoxCensored.object for details.
plot.type	character string indicating what kind of plot to create. Only one particular plot type will be created, unless plot.type="All", in which case all plots will be created sequentially. The possible values of plot.type are: "Objective vs. lambda" (the default), "Q-Q Plots", "Tukey M-D Q-Q Plots", and "All".
same.window	logical scalar indicating whether to produce all plots in the same graphics window (same.window=TRUE; the default), or to create a new graphics window for each separate plot (same.window=FALSE). The argument is relevant only when plot.type produces more than one plot (i.e., when plot.type is not equal to "Objective vs. lambda").
ask	logical scalar supplied to the function devAskNewPage , indicating whether to prompt the user before creating a new plot within a single graphics window. This argument is ignored when plot.type="Objective vs. lambda" (since only one plot is produced) or when same.window=FALSE, otherwise the default value is ask=TRUE.
points.col	numeric scalar determining the color of the points in the plot. The default value is points.col=1. See the entry for col in the R help file for par for more information. Arguments associated with plot.type="Q-Q Plots", plot.type="Tukey M-D Q-Q Plots", or plot.type="All" (supplied to qqPlot):
prob.method	character string indicating what method to use to compute the plotting positions for Q-Q plots or Tukey Mean-Difference Q-Q plots. Possible values are "kaplan-meier" (product-limit method of Kaplan and Meier (1958)), "nelson" (hazard plotting method of Nelson (1972)), "michael-schucany" (generalization of the product-limit method due to Michael and Schucany (1986)), and "hirsch-stedinger" (generalization of the product-limit method due to Hirsch and Stedinger (1987)). The default value is prob.method="michael-schucany". The "nelson" method is only available for objects that are the result of calling boxcoxCensored with the argument censoring.side="right". See the help file for qqPlotCensored for more information. This argument is ignored if plot.type="Objective vs. lambda".
plot.pos.con	numeric scalar between 0 and 1 containing the value of the plotting position constant used to construct the Q-Q plots and/or Tukey Mean-Difference Q-Q plots. The default value is plot.pos.con=0.375. See the help file for qqPlotCensored for more information.

estimate.params	logical scalar indicating whether to compute quantiles based on estimating the distribution parameters (estimate.params=TRUE) or using the distribution parameters for a standard normal distribution (i.e, mean=0, sd=1). The default value is estimate.params=FALSE because a standard normal Q-Q plot will yield roughly a straight line if the observations are from <i>any</i> normal distribution. If you specify plot.type="Tukey M-D Q-Q Plots", then you need to set estimate.params=TRUE unless you want to assume the transformed data come from a standard normal distribution.
equal.axes	logical scalar indicating whether to use the same range on the <i>x</i> - and <i>y</i> -axes when plot.type="Q-Q Plots". The default value is TRUE if qq.line.type="0-1" or estimate.params=TRUE, otherwise it is FALSE.
add.line	logical scalar indicating whether to add a line to the plot. If add.line=TRUE and plot.type="Q-Q Plots", a line determined by the value of qq.line.type is added to the plot. If add.line=TRUE and plot.type="Tukey M-D Q-Q Plots", a horizontal line at $y = 0$ is added to the plot. The default value is add.line=TRUE.
qq.line.type	character string determining what kind of line to add to the plot when plot.type="Q-Q Plots". Possible values are "least squares" (a least squares line; the default), "0-1" (a line with intercept 0 and slope 1), and "robust" (a line is fit through the first and third quartiles of the <i>x</i> and <i>y</i> data). This argument is ignored if add.line=FALSE.
duplicate.points.method	a character string denoting how to plot points with duplicate (<i>x</i> , <i>y</i>) values. Possible values are "standard" (a single plotting symbol is plotted; the default), "jitter" (a separate plotting symbol is plotted for each duplicate point, where the plotting symbols cluster around the true value of <i>x</i> and <i>y</i>), and "number" (a single number is plotted at (<i>x</i> , <i>y</i>) that represents how many duplicate points are at that (<i>x</i> , <i>y</i>) coordinate).
line.col	numeric scalar determining the color of the line in the plot. The default value is line.col=1. See the entry for col in the R help file for par for more information. This argument is ignored if add.line=FALSE.
line.lwd	numeric scalar determining the width of the line in the plot. The default value is line.lwd=par("cex"). See the entry for lwd in the R help file for par for more information. This argument is ignored if add.line=FALSE.
line.lty	numeric scalar determining the line type (style) of the line in the plot. The default value is line.lty=1. See the entry for lty in the R help file for par for more information. This argument is ignored if add.line=FALSE.
digits	scalar indicating how many significant digits to print for the distribution parameters and the value of the objective in the sub-title. The default value is the current setting of options("digits"). Graphics parameters:
cex.main, cex.sub, main, sub, xlab, ylab, xlim, ylim, ...	graphics parameters; see par for more information. The default value of cex.main is cex.main=1.4 * par("cex"). The default value of cex.sub is cex.sub=par("cex").

Details

The function `plot.boxcoxCensored` is a method for the generic function `plot` for the class "`boxcoxCensored`" (see `boxcoxCensored.object`). It can be invoked by calling `plot` and giving it an object of class "`boxcoxCensored`" as the first argument, or by calling `plot.boxcoxCensored` directly, regardless of the class of the object given as the first argument to `plot.boxcoxCensored`.

Plots associated with Box-Cox transformations are produced on the current graphics device. These can be one or all of the following:

- Objective vs. λ .
- Observed Quantiles vs. Normal Quantiles (Q-Q Plot) for the transformed observations for each of the values of λ .
- Tukey Mean-Difference Q-Q Plots for the transformed observations for each of the values of λ .

See the help files for [boxcoxCensored](#) and [qqPlotCensored](#) for more information.

Value

`plot.boxcoxCensored` invisibly returns the first argument, `x`.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Chambers, J. M. and Hastie, T. J. (1992). *Statistical Models in S*. Wadsworth & Brooks/Cole.

See Also

[qqPlotCensored](#), [boxcoxCensored](#), [boxcoxCensored.object](#), [print.boxcoxCensored](#), [Data Transformations](#), [plot](#).

Examples

```
# Generate 15 observations from a lognormal distribution with
# mean=10 and cv=2 and censor the observations less than 2.
# Then generate 15 more observations from this distribution and
# censor the observations less than 4.
# Then call the function boxcoxCensored, and then plot the results.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)

x.1 <- rlnormAlt(15, mean = 10, cv = 2)
censored.1 <- x.1 < 2
x.1[censored.1] <- 2

x.2 <- rlnormAlt(15, mean = 10, cv = 2)
censored.2 <- x.2 < 4
x.2[censored.2] <- 4

x <- c(x.1, x.2)
censored <- c(censored.1, censored.2)

# Plot the results based on the PPCC objective
#-----
boxcox.list <- boxcoxCensored(x, censored)
windows()
plot(boxcox.list)
```

```

# Look at Q-Q Plots for the candidate values of lambda
#-----
plot(boxcox.list, plot.type = "Q-Q Plots", same.window = FALSE)

# Look at Tukey Mean-Difference Q-Q Plots
# for the candidate values of lambda
#-----
plot(boxcox.list, plot.type = "Tukey M-D Q-Q Plots", same.window = FALSE)

#=====

# Clean up
#-----
rm(x.1, censored.1, x.2, censored.2, x, censored, boxcox.list)
graphics.off()

```

plot.boxcoxLm

Plot Results of Box-Cox Transformations for a Linear Model

Description

Plot the results of calling the function `boxcox` when the argument `x` supplied to `boxcox` is an object of class `"lm"`. Three different kinds of plots are available.

The function `plot.boxcoxLm` is automatically called by `plot` when given an object of class `"boxcoxLm"`. The names of other functions associated with Box-Cox transformations are listed under [Data Transformations](#).

Usage

```

## S3 method for class 'boxcoxLm'
plot(x, plot.type = "Objective vs. lambda", same.window = TRUE,
     ask = same.window & plot.type != "Objective vs. lambda",
     plot.pos.con = 0.375, estimate.params = FALSE,
     equal.axes = qq.line.type == "0-1" || estimate.params, add.line = TRUE,
     qq.line.type = "least squares", duplicate.points.method = "standard",
     points.col = 1, line.col = 1, line.lwd = par("cex"), line.lty = 1,
     digits = .Options$digits, cex.main = 1.4 * par("cex"), cex.sub = par("cex"),
     main = NULL, sub = NULL, xlab = NULL, ylab = NULL, xlim = NULL,
     ylim = NULL, ...)

```

Arguments

<code>x</code>	an object of class <code>"boxcoxLm"</code> . See boxcoxLm.object for details.
<code>plot.type</code>	character string indicating what kind of plot to create. Only one particular plot type will be created, unless <code>plot.type="All"</code> , in which case all plots will be created sequentially. The possible values of <code>plot.type</code> are: <code>"Objective vs. lambda"</code> (the default), <code>"Q-Q Plots"</code> , <code>"Tukey M-D Q-Q Plots"</code> , and <code>"All"</code> .
<code>same.window</code>	logical scalar indicating whether to produce all plots in the same graphics window (<code>same.window=TRUE</code> ; the default), or to create a new graphics window for each separate plot (<code>same.window=FALSE</code>). The argument is relevant only when <code>plot.type</code> produces more than one plot (i.e., when <code>plot.type</code> is not equal to <code>"Objective vs. lambda"</code>).

ask	logical scalar supplied to the function devAskNewPage , indicating whether to prompt the user before creating a new plot within a single graphics window. This argument is ignored when <code>plot.type="Objective vs. lambda"</code> (since only one plot is produced) or when <code>same.window=FALSE</code> , otherwise the default value is <code>ask=TRUE</code> .
points.col	numeric scalar determining the color of the points in the plot. The default value is <code>points.col=1</code> . See the entry for <code>col</code> in the R help file for par for more information. Arguments associated with <code>plot.type="Q-Q Plots"</code> , <code>plot.type="Tukey M-D Q-Q Plots"</code> , or <code>plot.type="All"</code> (supplied to qqPlot):
plot.pos.con	numeric scalar between 0 and 1 containing the value of the plotting position constant used to construct the Q-Q plots and/or Tukey Mean-Difference Q-Q plots. The default value is <code>plot.pos.con=0.375</code> . See the help files for qqPlot for more information and the motivation for this choice.
estimate.params	logical scalar indicating whether to compute quantiles based on estimating the distribution parameters (<code>estimate.params=TRUE</code>) or using the distribution parameters for a standard normal distribution (i.e., <code>mean=0</code> , <code>sd=1</code>). The default value is <code>estimate.params=FALSE</code> because a standard normal Q-Q plot will yield roughly a straight line if the observations are from <i>any</i> normal distribution. If you specify <code>plot.type="Tukey M-D Q-Q Plots"</code> , then you need to set <code>estimate.params=TRUE</code> unless you want to assume the transformed data come from a standard normal distribution.
equal.axes	logical scalar indicating whether to use the same range on the <i>x</i> - and <i>y</i> -axes when <code>plot.type="Q-Q Plots"</code> . The default value is <code>TRUE</code> if <code>qq.line.type="0-1"</code> or <code>estimate.params=TRUE</code> , otherwise it is <code>FALSE</code> .
add.line	logical scalar indicating whether to add a line to the plot. If <code>add.line=TRUE</code> and <code>plot.type="Q-Q Plots"</code> , a line determined by the value of <code>qq.line.type</code> is added to the plot. If <code>add.line=TRUE</code> and <code>plot.type="Tukey M-D Q-Q Plots"</code> , a horizontal line at $y = 0$ is added to the plot. The default value is <code>add.line=TRUE</code> .
qq.line.type	character string determining what kind of line to add to the plot when <code>plot.type="Q-Q Plots"</code> . Possible values are "least squares" (a least squares line; the default), "0-1" (a line with intercept 0 and slope 1), and "robust" (a line is fit through the first and third quartiles of the <i>x</i> and <i>y</i> data). This argument is ignored if <code>add.line=FALSE</code> .
duplicate.points.method	a character string denoting how to plot points with duplicate (<i>x</i> , <i>y</i>) values. Possible values are "standard" (a single plotting symbol is plotted; the default), "jitter" (a separate plotting symbol is plotted for each duplicate point, where the plotting symbols cluster around the true value of <i>x</i> and <i>y</i>), and "number" (a single number is plotted at (<i>x</i> , <i>y</i>) that represents how many duplicate points are at that (<i>x</i> , <i>y</i>) coordinate).
line.col	numeric scalar determining the color of the line in the plot. The default value is <code>line.col=1</code> . See the entry for <code>col</code> in the R help file for par for more information. This argument is ignored if <code>add.line=FALSE</code> .
line.lwd	numeric scalar determining the width of the line in the plot. The default value is <code>line.lwd=par("cex")</code> . See the entry for <code>lwd</code> in the R help file for par for more information. This argument is ignored if <code>add.line=FALSE</code> .
line.lty	numeric scalar determining the line type (style) of the line in the plot. The default value is <code>line.lty=1</code> . See the entry for <code>lty</code> in the R help file for par for more information. This argument is ignored if <code>add.line=FALSE</code> .

`digits` scalar indicating how many significant digits to print for the distribution parameters and the value of the objective in the sub-title. The default value is the current setting of `options("digits")`.

Graphics parameters:

`cex.main`, `cex.sub`, `main`, `sub`, `xlab`, `ylab`, `xlim`, `ylim`, ...
 graphics parameters; see [par](#) for more information. The default value of `cex.main` is `cex.main=1.4 * par("cex")`. The default value of `cex.sub` is `cex.sub=par("cex")`.

Details

The function `plot.boxcoxLm` is a method for the generic function [plot](#) for the class `"boxcoxLm"` (see [boxcoxLm.object](#)). It can be invoked by calling [plot](#) and giving it an object of class `"boxcoxLm"` as the first argument, or by calling `plot.boxcoxLm` directly, regardless of the class of the object given as the first argument to `plot.boxcoxLm`.

Plots associated with Box-Cox transformations are produced on the current graphics device. These can be one or all of the following:

- Objective vs. λ .
- Observed Quantiles vs. Normal Quantiles (Q-Q Plot) for the residuals of the linear model based on transformed values of the response variable for each of the values of λ .
- Tukey Mean-Difference Q-Q Plots for the residuals of the linear model based on transformed values of the response variable for each of the values of λ .

See the help files for [boxcox](#) and [qqPlot](#) for more information.

Value

`plot.boxcoxLm` invisibly returns the first argument, `x`.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Chambers, J. M. and Hastie, T. J. (1992). *Statistical Models in S*. Wadsworth & Brooks/Cole.

See Also

[qqPlot](#), [boxcox](#), [boxcoxLm.object](#), [print.boxcoxLm](#), [Data Transformations](#), [plot](#).

Examples

```
# Create an object of class "boxcoxLm", then plot the results.

# The data frame Environmental.df contains daily measurements of
# ozone concentration, wind speed, temperature, and solar radiation
# in New York City for 153 consecutive days between May 1 and
# September 30, 1973. In this example, we'll model ozone as a
# function of temperature.

# Fit the model with the raw Ozone data
#-----
ozone.fit <- lm(ozone ~ temperature, data = Environmental.df)
```

```

boxcox.list <- boxcox(ozone.fit)

# Plot PPCC vs. lambda based on Q-Q plots of residuals
#-----
windows()
plot(boxcox.list)

# Look at Q-Q plots of residuals for the various transformation
#-----
plot(boxcox.list, plot.type = "Q-Q Plots", same.window = FALSE)

# Look at Tukey Mean-Difference Q-Q plots of residuals
# for the various transformation
#-----
plot(boxcox.list, plot.type = "Tukey M-D Q-Q Plots", same.window = FALSE)

#=====

# Clean up
#-----
rm(ozone.fit, boxcox.list)
graphics.off()

```

plot.gof

Plot Results of Goodness-of-Fit Test

Description

Plot the results of calling the function `gofTest`, which returns an object of class "gof" when testing the goodness-of-fit of a set of data to a distribution (i.e., when supplied with the `y` argument but not the `x` argument). Five different kinds of plots are available.

The function `plot.gof` is automatically called by `plot` when given an object of class "gof". The names of other functions associated with goodness-of-fit test are listed under [Goodness-of-Fit Tests](#).

Usage

```

## S3 method for class 'gof'
plot(x, plot.type = "Summary",
     captions = list(PDFs = NULL, CDFs = NULL, QQ = NULL, MDQQ = NULL, Results = NULL),
     x.labels = list(PDFs = NULL, CDFs = NULL, QQ = NULL, MDQQ = NULL),
     y.labels = list(PDFs = NULL, CDFs = NULL, QQ = NULL, MDQQ = NULL),
     same.window = FALSE, ask = same.window & plot.type == "All", hist.col = "cyan",
     fitted.pdf.col = "black", fitted.pdf.lwd = 3 * par("cex"), fitted.pdf.lty = 1,
     plot.pos.con = switch(dist.abb, norm = , lnorm = , lnormAlt = , lnorm3 = 0.375,
     evd = 0.44, 0.4), ecdf.col = "cyan", fitted.cdf.col = "black",
     ecdf.lwd = 3 * par("cex"), fitted.cdf.lwd = 3 * par("cex"), ecdf.lty = 1,
     fitted.cdf.lty = 2, add.line = TRUE,
     digits = ifelse(plot.type == "Summary", 2, .Options$digits), test.result.font = 1,
     test.result.cex = ifelse(plot.type == "Summary", 0.9, 1) * par("cex"),
     test.result.mar = c(0, 0, 3, 0) + 0.1,
     cex.main = ifelse(plot.type == "Summary", 1.2, 1.5) * par("cex"),

```



```

cex.axis = ifelse(plot.type == "Summary", 0.9, 1) * par("cex"),
cex.lab = ifelse(plot.type == "Summary", 0.9, 1) * par("cex"),
main = NULL, xlab = NULL, ylab = NULL, xlim = NULL, ylim = NULL,
add.om.title = TRUE,
oma = if (plot.type == "Summary" & add.om.title) c(0, 0, 2.5, 0) else c(0, 0, 0, 0),
om.title = NULL, om.font = 2, om.cex.main = 1.75 * par("cex"), om.line = 0.5, ...)

```

Arguments

<code>x</code>	an object of class "gof". See gof.object for details.
<code>plot.type</code>	character string indicating what kind of plot to create. Only one particular plot type will be created, unless <code>plot.type="All"</code> , in which case all plots will be created sequentially. The possible values of <code>plot.type</code> are: "Summary" (the default), "PDFs: Observed and Fitted", "CDFs: Observed and Fitted", "Q-Q Plot", "Tukey M-D Q-Q Plot", "Test Results", and "All". See the DETAILS section for more information.
<code>captions</code>	a list with 1 to 5 components with the names "PDFs", "CDFs", "QQ", "MDQQ", and/or "Results". Each component either has the value NULL or else it is a character string containing the title for that particular kind of plot. When the component has the value NULL (the default), a default title is used. This argument is useful when you are creating more than one kind of plot with a single call to <code>plot.gof</code> (i.e., when <code>plot.type="Summary"</code> or <code>plot.type="All"</code>) and you want to specify titles different from the default ones. If you are creating only one kind of plot, then you can just use the <code>main</code> argument to specify a title different from the default one.
<code>x.labels</code>	a list of 1 to 4 components with the names "PDFs", "CDFs", "QQ", and/or "MDQQ". Each component either has the value NULL or else it is a character string containing the label for the <i>x</i> -axis for that particular kind of plot. When the component has the value NULL (the default), a default <i>x</i> -axis label is used. This argument is useful when you are creating more than one kind of plot with a single call to <code>plot.gof</code> (i.e., when <code>plot.type="Summary"</code> or <code>plot.type="All"</code>) and you want to specify <i>x</i> -axis labels different from the default ones. If you are creating only one plot, then you can just use the <code>xlab</code> argument to specify an <i>x</i> -axis label different from the default one.
<code>y.labels</code>	a list of 1 to 4 components with the names "PDFs", "CDFs", "QQ", and/or "MDQQ". Each component either has the value NULL or else it is a character string containing the label for the <i>y</i> -axis for that particular kind of plot. When the component has the value NULL (the default), a default <i>y</i> -axis label is used. This argument is useful when you are creating more than one kind of plot with a single call to <code>plot.gof</code> (i.e., when <code>plot.type="Summary"</code> or <code>plot.type="All"</code>) and you want to specify <i>y</i> -axis labels different from the default ones. If you are creating only one plot, then you can just use the <code>ylab</code> argument to specify a <i>y</i> -axis label different from the default one.
<code>same.window</code>	logical scalar indicating whether to produce all plots in the same graphics window (<code>same.window=TRUE</code>), or to create a new graphics window for each separate plot (<code>same.window=FALSE</code> ; the default). The argument is relevant only when <code>plot.type="All"</code> .
<code>ask</code>	logical scalar supplied to the function devAskNewPage , indicating whether to prompt the user before creating a new plot within a single graphics window. The default value is FALSE unless <code>same.window=TRUE</code> and <code>plot.type == "All"</code> .

digits	<p>scalar indicating how many significant digits to print for the distribution parameters. If <code>plot.type == "Summary"</code>, the default value is <code>digits=2</code>, otherwise it is <code>.Options\$digits</code> (i.e., the current setting of <code>options("digits")</code>). This argument is ignored when <code>plot.type="PDFs: Observed and Fitted"</code>.</p> <p>Arguments associated with <code>plot.type="PDFs: Observed and Fitted"</code>:</p>
hist.col	<p>a character string or numeric scalar determining the color of the histogram used to display the distribution of the observed values. The default value is <code>hist.col="cyan"</code>. See the entry for <code>col</code> in the R help file for par for more information.</p>
fitted.pdf.col	<p>a character string or numeric scalar determining the color of the fitted PDF (which is displayed as a line for continuous distributions and a histogram for discrete distributions). The default value is <code>fitted.pdf.col="black"</code>. See the entry for <code>col</code> in the R help file for par for more information.</p>
fitted.pdf.lwd	<p>numeric scalar determining the width of the line used to display the fitted PDF. The default value is <code>fitted.pdf.lwd=3*par("cex")</code>. See the entry for <code>lwd</code> in the R help file for par for more information.</p>
fitted.pdf.lty	<p>numeric scalar determining the line type used to display the fitted PDF. The default value is <code>fitted.pdf.lty=1</code>. See the entry for <code>lty</code> in the R help file for par for more information.</p> <p>Arguments associated with <code>plot.type="CDFs: Observed and Fitted"</code>:</p>
plot.pos.con	<p>numeric scalar between 0 and 1 containing the value of the plotting position constant used to construct the observed (empirical) CDF. The default value of <code>plot.pos.con</code> depends on the value of <code>gof.obj\$distribution</code> (i.e., the distribution assumed for the goodness-of-fit test). For the normal, lognormal, and three-parameter lognormal distributions, the default value is <code>plot.pos.con=0.375</code>. For the Type I extreme value (Gumbel) distribution, the default value is <code>plot.pos.con=0.44</code>. For all other distributions, the default value is <code>plot.pos.con=0.4</code>. See the help files for ecdfPlot and qqPlot for more information and the motivation for these choices of values.</p> <p>NOTE: This argument is also used to determine the value of the plotting position constant for the Q-Q plot (<code>plot.type="Q-Q Plot"</code>), or the Tukey Mean-Difference Q-Q plot (<code>plot.type="Tukey M-D Q-Q Plot"</code>).</p>
ecdf.col	<p>a character string or numeric scalar determining the color of the line used to display the empirical CDF. The default value is <code>ecdf.col="cyan"</code>. See the entry for <code>col</code> in the R help file for par for more information.</p>
fitted.cdf.col	<p>a character string or numeric scalar determining the color of the line used to display the fitted CDF. The default value is <code>fitted.cdf.col="black"</code>. See the entry for <code>col</code> in the R help file for par for more information.</p>
ecdf.lwd	<p>numeric scalar determining the width of the line used to display the empirical CDF. The default value is <code>ecdf.lwd=3*par("cex")</code>. See the entry for <code>lwd</code> in the R help file for par for more information.</p>
fitted.cdf.lwd	<p>numeric scalar determining the width of the line used to display the fitted CDF. The default value is <code>fitted.cdf.lwd=3*par("cex")</code>. See the entry for <code>lwd</code> in the R help file for par for more information.</p>
ecdf.lty	<p>numeric scalar determining the line type used to display the empirical CDF. The default value is <code>ecdf.lty=1</code>. See the entry for <code>lty</code> in the R help file for par for more information.</p>

- `fitted.cdf.lty` numeric scalar determining the line type used to display the fitted CDF. The default value is `fitted.cdf.lty=2`. See the entry for `lty` in the R help file for [par](#) for more information.
Arguments associated with `plot.type="Q-Q Plot"` or `plot.type="Tukey M-D Q-Q Plot"`:
As explained above, `plot.pos.con` is used for these plot types. Also:
- `add.line` logical scalar indicating whether to add a line to the plot. If `add.line=TRUE` and `plot.type="Q-Q Plot"`, a 0-1 line is added to the plot. If `add.line=TRUE` and `plot.type="Tukey M-D Q-Q Plot"`, a horizontal line at $y = 0$ is added to the plot. The default value is `add.line=TRUE`.
Arguments associated with `plot.type="Test Results"`
- `test.result.font` numeric scalar indicating which font to use to print out the test results. The default value is `test.result.font=1`. See the description of the `font` argument in the help file for [par](#) for more information. You may get better results if you use a font number that corresponds to a fixed font (e.g., `courier`).
- `test.result.cex` numeric scalar indicating the value of `cex` to use to print out the test results. The default value is `0.9*par("cex")` when `plot.type="Summary"`, otherwise it is `par("cex")`. See the description of the `cex` argument in the help file for [par](#) for more information.
- `test.result.mar` numeric vector indicating the value of `mar` to use to print out the test results. The default value is `test.result.mar=c(0, 0, 3, 0)+0.1`. See the description of the `mar` argument in the help file for [par](#) for more information.
Arguments associated with `plot.type="Summary"`
- `add.om.title` logical scalar indicating whether to add a title in the outer margin when `plot.type="Summary"`. The default value is `add.om.title=TRUE`.
- `om.title` character string containing the outer margin title. The default value is `om.title=NULL`, which will result in a default title.
- `om.font` numeric scalar indicating the font to use for the outer margin. The default value is `om.font=2`.
- `om.cex.main` numeric scalar indicating the value of `cex` for the outer margin title. The default value is `1.75 * par("cex")`.
- `om.line` numeric scalar indicating the line to place the outer margin title on. The default value is `om.line=0.5`.
Graphics parameters:
- `cex.main`, `cex.axis`, `cex.lab`, `main`, `xlab`, `ylab`, `xlim`, `ylim`, `oma`, ...
additional graphics parameters. See the help file for [par](#).

Details

The function `plot.gof` is a method for the generic function [plot](#) for the class "gof" (see [gof.object](#)). It can be invoked by calling [plot](#) and giving it an object of class "gof" as the first argument, or by calling `plot.gof` directly, regardless of the class of the object given as the first argument to `plot.gof`.

Plots associated with the goodness-of-fit test are produced on the current graphics device. These can be one or all of the following:

- Observed distribution overlaid with fitted distribution (`plot.type="PDFs: Observed and Fitted"`). See the help files for [hist](#) and [pdfPlot](#).
- Observed empirical distribution overlaid with fitted cumulative distribution (`plot.type="CDFs: Observed and Fitted"`). See the help file for [cdfCompare](#).
- Observed quantiles vs. fitted quantiles (Q-Q Plot) (`plot.type="Q-Q Plot"`). See the help file for [qqPlot](#).
- Tukey mean-difference Q-Q plot (`plot.type="Tukey M-D Q-Q Plot"`). See the help file for [qqPlot](#).
- Results of the goodness-of-fit test (`plot.type="Test Results"`). See the help file for [print.gof](#).

See the help file for [gofTest](#) for more information.

Value

`plot.gof` invisibly returns the first argument, `x`.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Chambers, J. M. and Hastie, T. J. (1992). *Statistical Models in S*. Wadsworth & Brooks/Cole.

See Also

[gofTest](#), [gof.object](#), [print.gof](#), [Goodness-of-Fit Tests](#), [plot](#).

Examples

```
# Create an object of class "gof" then plot the results.
# (Note: the call to set.seed simply allows you to reproduce
# this example.)

set.seed(250)
dat <- rnorm(20, mean = 3, sd = 2)
gof.obj <- gofTest(dat)

# Summary plot (the default)
#-----
windows()
plot(gof.obj)

# Make your own titles for the summary plot
#-----
windows()
plot(gof.obj, captions = list(PDFs = "Compare PDFs",
  CDFs = "Compare CDFs", QQ = "Q-Q Plot", Results = "Results"),
  om.title = "Summary")

# Just the Q-Q Plot
#-----
windows()
```

```

plot(gof.obj, plot.type="Q-Q")

# Make your own title for the Q-Q Plot
#-----
windows()
plot(gof.obj, plot.type="Q-Q", main = "Q-Q Plot")

#=====

# Clean up
#-----
rm(dat, gof.obj)
graphics.off()

```

plot.gofGroup

Plot Results of Group Goodness-of-Fit Test

Description

Plot the results of calling the function [gofGroupTest](#), which returns an object of class "gofGroup" when performing a goodness-of-fit test to determine whether data in a set of groups appear to all come from the same probability distribution (with possibly different parameters for each group). Five different kinds of plots are available.

The function `plot.gofGroup` is automatically called by [plot](#) when given an object of class "gofGroup". The names of other functions associated with goodness-of-fit test are listed under [Goodness-of-Fit Tests](#).

Usage

```

## S3 method for class 'gofGroup'
plot(x, plot.type = "Summary",
      captions = list(QQ = NULL, MDQQ = NULL, ScoresQQ = NULL, ScoresMDQQ = NULL, Results = NULL),
      x.labels = list(QQ = NULL, MDQQ = NULL, ScoresQQ = NULL, ScoresMDQQ = NULL),
      y.labels = list(QQ = NULL, MDQQ = NULL, ScoresQQ = NULL, ScoresMDQQ = NULL),
      same.window = FALSE, ask = same.window & plot.type == "All", add.line = TRUE,
      digits = ifelse(plot.type == "Summary", 2, .Options$digits), test.result.font = 1,
      test.result.cex = ifelse(plot.type == "Summary", 0.9, 1) * par("cex"),
      test.result.mar = c(0, 0, 3, 0) + 0.1, individual.p.values = FALSE,
      cex.main = ifelse(plot.type == "Summary", 1.2, 1.5) * par("cex"),
      cex.axis = ifelse(plot.type == "Summary", 0.9, 1) * par("cex"),
      cex.lab = ifelse(plot.type == "Summary", 0.9, 1) * par("cex"),
      main = NULL, xlab = NULL, ylab = NULL, xlim = NULL, ylim = NULL, add.om.title = TRUE,
      oma = if (plot.type == "Summary" & add.om.title) c(0, 0, 5, 0) else c(0, 0, 0, 0),
      om.title = NULL, om.font = 2, om.cex.main = 1.5 * par("cex"), om.line = 1, ...)

```

Arguments

<code>x</code>	an object of class "gofGroup". See gofGroup.object for details.
<code>plot.type</code>	character string indicating what kind of plot to create. Only one particular plot type will be created, unless <code>plot.type="All"</code> , in which case all plots will be

created sequentially. The possible values of `plot.type` are: "Summary" (the default), "Q-Q Plot", "Tukey M-D Q-Q Plot", "Scores Q-Q Plot", "Scores Tukey M-D Q-Q Plot", "Test Results", and "All". See the DETAILS section for more information.

<code>captions</code>	a list with 1 to 5 components with the names "QQ", "MDQQ", "ScoresQQ", "ScoresMDQQ", and/or "Results". Each component either has the value NULL or else it is a character string containing the title for that particular kind of plot. When the component has the value NULL (the default), a default title is used. This argument is useful when you are creating more than one kind of plot with a single call to <code>plot.gofGroup</code> (i.e., when <code>plot.type="Summary"</code> or <code>plot.type="All"</code>) and you want to specify titles different from the default ones. If you are creating only one kind of plot, then you can just use the <code>main</code> argument to specify a title different from the default one.
<code>x.labels</code>	a list of 1 to 4 components with the names "QQ", "MDQQ", "ScoresQQ", and/or "ScoresMDQQ". Each component either has the value NULL or else it is a character string containing the label for the <i>x</i> -axis for that particular kind of plot. When the component has the value NULL (the default), a default <i>x</i> -axis label is used. This argument is useful when you are creating more than one kind of plot with a single call to <code>plot.gofGroup</code> (i.e., when <code>plot.type="Summary"</code> or <code>plot.type="All"</code>) and you want to specify <i>x</i> -axis labels different from the default ones. If you are creating only one plot, then you can just use the <code>xlab</code> argument to specify an <i>x</i> -axis label different from the default one.
<code>y.labels</code>	a list of 1 to 4 components with the names "QQ", "MDQQ", "ScoresQQ", and/or "ScoresMDQQ". Each component either has the value NULL or else it is a character string containing the label for the <i>y</i> -axis for that particular kind of plot. When the component has the value NULL (the default), a default <i>y</i> -axis label is used. This argument is useful when you are creating more than one kind of plot with a single call to <code>plot.gofGroup</code> (i.e., when <code>plot.type="Summary"</code> or <code>plot.type="All"</code>) and you want to specify <i>y</i> -axis labels different from the default ones. If you are creating only one plot, then you can just use the <code>ylab</code> argument to specify a <i>y</i> -axis label different from the default one.
<code>same.window</code>	logical scalar indicating whether to produce all plots in the same graphics window (<code>same.window=TRUE</code>), or to create a new graphics window for each separate plot (<code>same.window=FALSE</code> ; the default). The argument is relevant only when <code>plot.type="All"</code> .
<code>ask</code>	logical scalar supplied to the function devAskNewPage , indicating whether to prompt the user before creating a new plot within a single graphics window. The default value is FALSE unless <code>same.window=TRUE</code> and <code>plot.type == "All"</code> .
<code>add.line</code>	logical scalar indicating whether to add a line to the plot. If <code>add.line=TRUE</code> and <code>plot.type="Q-Q Plot"</code> or <code>plot.type="Scores Q-Q Plot"</code> , a 0-1 line is added to the plot. If <code>add.line=TRUE</code> and <code>plot.type="Tukey M-D Q-Q Plot"</code> or <code>plot.type="Scores Tukey M-D Q-Q Plot"</code> , a horizontal line at $y = 0$ is added to the plot. The default value is <code>add.line=TRUE</code> . This argument is ignored if <code>plot.type="Test Results"</code> . Arguments associated with <code>plot.type="Test Results"</code>
<code>digits</code>	scalar indicating how many significant digits to print for the test results when <code>plot.type="Summary"</code> or <code>plot.type="Test Results"</code> . If <code>plot.type == "Summary"</code> , the default value is <code>digits=2</code> , otherwise it is <code>.Options\$digits</code> (i.e., the current setting of <code>options("digits")</code>).

<code>individual.p.values</code>	logical scalar indicating whether to display the p-values associated with each individual group. The default value is <code>individual.p.values=FALSE</code> .
<code>test.result.font</code>	numeric scalar indicating which font to use to print out the test results. The default value is <code>test.result.font=1</code> . See the description of the <code>font</code> argument in the help file for par for more information. You may get better results if you use a font number that corresponds to a fixed font (e.g., <code>courier</code>).
<code>test.result.cex</code>	numeric scalar indicating the value of <code>cex</code> to use to print out the test results. The default value is <code>0.9*par("cex")</code> when <code>plot.type="Summary"</code> , otherwise it is <code>par("cex")</code> . See the description of the <code>cex</code> argument in the help file for par for more information.
<code>test.result.mar</code>	numeric vector indicating the value of <code>mar</code> to use to print out the test results. The default value is <code>test.result.mar=c(0, 0, 3, 0)+0.1</code> . See the description of the <code>mar</code> argument in the help file for par for more information. Arguments associated with <code>plot.type="Summary"</code>
<code>add.om.title</code>	logical scalar indicating whether to add a title in the outer margin when <code>plot.type="Summary"</code> . The default value is <code>add.om.title=TRUE</code> .
<code>om.title</code>	character string containing the outer margin title. The default value is <code>om.title=NULL</code> , which will result in a default title.
<code>om.font</code>	numeric scalar indicating the font to use for the outer margin. The default value is <code>om.font=2</code> .
<code>om.cex.main</code>	numeric scalar indicating the value of <code>cex</code> for the outer margin title. The default value is <code>1.5 * par("cex")</code> .
<code>om.line</code>	numeric scalar indicating the line to place the outer margin title on. The default value is <code>om.line=1</code> .
	Graphics parameters:
<code>cex.main, cex.axis, cex.lab, main, xlab, ylab, xlim, ylim, oma, ...</code>	additional graphics parameters. See the help file for par .

Details

The function `plot.gofGroup` is a method for the generic function [plot](#) for the class "gofGroup" (see [gofGroup.object](#)). It can be invoked by calling [plot](#) and giving it an object of class "gofGroup" as the first argument, or by calling `plot.gofGroup` directly, regardless of the class of the object given as the first argument to `plot.gofGroup`.

Plots associated with the goodness-of-fit test are produced on the current graphics device. These can be one or all of the following:

- `plot.type="Q-Q Plot"`. Q-Q Plot of observed p-values vs. quantiles from a [Uniform \[0,1\] distribution](#). See the help file for [qqPlot](#).
- `plot.type="Tukey M-D Q-Q Plot"`. Tukey mean-difference Q-Q plot for observed p-values and quantiles from a [Uniform \[0,1\] distribution](#). See the help file for [qqPlot](#).
- `plot.type="Scores Q-Q Plot"`. Q-Q Plot of Normal scores vs. quantiles from a [Normal\(0,1\) distribution](#) or Q-Q Plot of Chisquare scores vs. quantiles from a [Chisquare distribution](#) with 2 degrees of freedom. See the help file for [qqPlot](#).

- `plot.type="Scores Tukey M-D Q-Q Plot"`. Tukey mean-difference Q-Q plot based on Normal scores or Chisquare scores. See the help file for [qqPlot](#).
- Results of the goodness-of-fit test (`plot.type="Test Results"`). See the help file for [print.gofGroup](#).

See the help file for [gofGroupTest](#) for more information.

Value

`plot.gofGroup` invisibly returns the first argument, `x`.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Chambers, J. M. and Hastie, T. J. (1992). *Statistical Models in S*. Wadsworth & Brooks/Cole.

See Also

[gofGroupTest](#), [gofGroup.object](#), [print.gofGroup](#), [Goodness-of-Fit Tests](#), [plot](#).

Examples

```
# Create an object of class "gofGroup" then plot it.

# Example 10-4 of USEPA (2009, page 10-20) gives an example of
# simultaneously testing the assumption of normality for nickel
# concentrations (ppb) in groundwater collected at 4 monitoring
# wells over 5 months. The data for this example are stored in
# EPA.09.Ex.10.1.nickel.df.

EPA.09.Ex.10.1.nickel.df
#   Month  Well Nickel.ppb
#1      1 Well.1      58.8
#2      3 Well.1       1.0
#3      6 Well.1    262.0
#...
#18     6 Well.4     85.6
#19     8 Well.4     10.0
#20    10 Well.4    637.0

# Test for a normal distribution at each well:
#-----

gofGroup.obj <- gofGroupTest(Nickel.ppb ~ Well,
  data = EPA.09.Ex.10.1.nickel.df)

windows()
plot(gofGroup.obj)

# Make your own titles for the summary plot
#-----
windows()
plot(gofGroup.obj, captions = list(QQ = "Q-Q Plot",
```



```

ScoresQQ = "Scores Q-Q Plot", Results = "Results"),
om.title = "Summary Plot")

# Just the Q-Q Plot
#-----
windows()
plot(gofGroup.obj, plot.type="Q-Q")

# Make your own title for the Q-Q Plot
#-----
windows()
plot(gofGroup.obj, plot.type="Q-Q", main = "Q-Q Plot")

#=====

# Clean up
#-----
rm(gofGroup.obj)
graphics.off()

```

plot.gofTwoSample

Plot Results of Goodness-of-Fit Test to Compare Two Samples

Description

Plot the results of calling the function `gofTest` to compare two samples. `gofTest` returns an object of class "gofTwoSample" when supplied with both the arguments `y` and `x`. `plot.gofTwoSample` provides five different kinds of plots.

The function `plot.gofTwoSample` is automatically called by `plot` when given an object of class "gofTwoSample". The names of other functions associated with goodness-of-fit test are listed under [Goodness-of-Fit Tests](#).

Usage

```

## S3 method for class 'gofTwoSample'
plot(x, plot.type = "Summary",
      captions = list(PDFs = NULL, CDFs = NULL, QQ = NULL, MDQQ = NULL, Results = NULL),
      x.labels = list(PDFs = NULL, CDFs = NULL, QQ = NULL, MDQQ = NULL),
      y.labels = list(PDFs = NULL, CDFs = NULL, QQ = NULL, MDQQ = NULL),
      same.window = FALSE, ask = same.window & plot.type == "All", x.points.col = "blue",
      y.points.col = "black", points.pch = 1, jitter.points = TRUE, discrete = FALSE,
      plot.pos.con = 0.375, x.ecdf.col = "blue", y.ecdf.col = "black",
      x.ecdf.lwd = 3 * par("cex"), y.ecdf.lwd = 3 * par("cex"), x.ecdf.lty = 1,
      y.ecdf.lty = 4, add.line = TRUE,
      digits = ifelse(plot.type == "Summary", 2, .Options$digits), test.result.font = 1,
      test.result.cex = ifelse(plot.type == "Summary", 0.9, 1) * par("cex"),
      test.result.mar = c(0, 0, 3, 0) + 0.1,
      cex.main = ifelse(plot.type == "Summary", 1.2, 1.5) * par("cex"),
      cex.axis = ifelse(plot.type == "Summary", 0.9, 1) * par("cex"),
      cex.lab = ifelse(plot.type == "Summary", 0.9, 1) * par("cex"),
      main = NULL, xlab = NULL, ylab = NULL, xlim = NULL, ylim = NULL,

```

```
add.om.title = TRUE,
oma = if (plot.type == "Summary" & add.om.title) c(0, 0, 4, 0) else c(0, 0, 0, 0),
om.title = NULL, om.font = 2, om.cex.main = 1.5 * par("cex"), om.line = 0, ...)
```

Arguments

<code>x</code>	an object of class "gof". See gof.object for details.
<code>plot.type</code>	character string indicating what kind of plot to create. Only one particular plot type will be created, unless <code>plot.type="All"</code> , in which case all plots will be created sequentially. The possible values of <code>plot.type</code> are: "Summary" (the default), "PDFs: Observed", "CDFs: Observed", "Q-Q Plot", "Tukey M-D Q-Q Plot", "Test Results", and "All". See the DETAILS section for more information.
<code>captions</code>	a list with 1 to 5 components with the names "PDFs", "CDFs", "QQ", "MDQQ", and/or "Results". Each component either has the value NULL or else it is a character string containing the title for that particular kind of plot. When the component has the value NULL (the default), a default title is used. This argument is useful when you are creating more than one kind of plot with a single call to <code>plot.gofTwoSample</code> (i.e., when <code>plot.type="Summary"</code> or <code>plot.type="All"</code>) and you want to specify titles different from the default ones. If you are creating only one kind of plot, then you can just use the <code>main</code> argument to specify a title different from the default one.
<code>x.labels</code>	a list of 1 to 4 components with the names "PDFs", "CDFs", "QQ", and/or "MDQQ". Each component either has the value NULL or else it is a character string containing the label for the <i>x</i> -axis for that particular kind of plot. When the component has the value NULL (the default), a default <i>x</i> -axis label is used. This argument is useful when you are creating more than one kind of plot with a single call to <code>plot.gofTwoSample</code> (i.e., when <code>plot.type="Summary"</code> or <code>plot.type="All"</code>) and you want to specify <i>x</i> -axis labels different from the default ones. If you are creating only one plot, then you can just use the <code>xlab</code> argument to specify an <i>x</i> -axis label different from the default one.
<code>y.labels</code>	a list of 1 to 4 components with the names "PDFs", "CDFs", "QQ", and/or "MDQQ". Each component either has the value NULL or else it is a character string containing the label for the <i>y</i> -axis for that particular kind of plot. When the component has the value NULL (the default), a default <i>y</i> -axis label is used. This argument is useful when you are creating more than one kind of plot with a single call to <code>plot.gofTwoSample</code> (i.e., when <code>plot.type="Summary"</code> or <code>plot.type="All"</code>) and you want to specify <i>y</i> -axis labels different from the default ones. If you are creating only one plot, then you can just use the <code>ylab</code> argument to specify a <i>y</i> -axis label different from the default one.
<code>same.window</code>	logical scalar indicating whether to produce all plots in the same graphics window (<code>same.window=TRUE</code>), or to create a new graphics window for each separate plot (<code>same.window=FALSE</code> ; the default). The argument is relevant only when <code>plot.type="All"</code> .
<code>ask</code>	logical scalar supplied to the function devAskNewPage , indicating whether to prompt the user before creating a new plot within a single graphics window. The default value is FALSE unless <code>same.window=TRUE</code> and <code>plot.type == "All"</code> . Arguments associated with <code>plot.type="PDFs: Observed"</code> :
<code>x.points.col</code>	a character string or numeric scalar determining the color of the plotting symbol used to display the distribution of the observed <i>x</i> values that were supplied to

	gofTest . The default value is <code>x.points.col="blue"</code> . See the entry for <code>col</code> in the R help file for par for more information.
<code>y.points.col</code>	a character string or numeric scalar determining the color of the plotting symbol used to display the distribution of the observed <code>y</code> values that were supplied to gofTest . The default value is <code>y.points.col="black"</code> . See the entry for <code>col</code> in the R help file for par for more information.
<code>points.pch</code>	a character string or numeric scalar determining the plotting symbol used to display the distribution of the observed <code>x</code> and <code>y</code> values that were supplied to gofTest . The default value is <code>points.pch=1</code> . See the entry for <code>pch</code> in the R help file for par for more information.
<code>jitter.points</code>	logical scalar indicating whether to jitter the points in the strip chart. The default value is <code>jitter.points=TRUE</code> . Arguments associated with <code>plot.type="CDFs: Observed"</code> :
<code>discrete</code>	logical scalar indicating whether the two distributions are considered to be discrete (<code>discrete=TRUE</code>) or not (<code>discrete=FALSE</code> ; the default). When <code>discrete=TRUE</code> , the empirical CDFs are plotted as step functions.
<code>plot.pos.con</code>	numeric scalar between 0 and 1 containing the value of the plotting position constant used to construct the observed (empirical) CDFs. The default value is <code>plot.pos.con=0.375</code> . See the help files for ecdfPlot and qqPlot for more information and the motivation for this choice of values. NOTE: This argument is also used to determine the value of the plotting position constant for the Q-Q plot (<code>plot.type="Q-Q Plot"</code>), or the Tukey Mean-Difference Q-Q plot (<code>plot.type="Tukey M-D Q-Q Plot"</code>).
<code>x.ecdf.col</code>	a character string or numeric scalar determining the color of the line used to display the empirical CDF for the <code>x</code> values that were supplied to gofTest . The default value is <code>x.ecdf.col="blue"</code> . See the entry for <code>col</code> in the R help file for par for more information.
<code>y.ecdf.col</code>	a character string or numeric scalar determining the color of the line used to display the empirical CDF for the <code>y</code> values that were supplied to gofTest . The default value is <code>y.ecdf.col="black"</code> . See the entry for <code>col</code> in the R help file for par for more information.
<code>x.ecdf.lwd</code>	numeric scalar determining the width of the line used to display the empirical CDF for the <code>x</code> values that were supplied to gofTest . The default value is <code>x.ecdf.lwd=3*par("cex")</code> . See the entry for <code>lwd</code> in the R help file for par for more information.
<code>y.ecdf.lwd</code>	numeric scalar determining the width of the line used to display the empirical CDF for the <code>y</code> values that were supplied to gofTest . The default value is <code>y.ecdf.lwd=3*par("cex")</code> . See the entry for <code>lwd</code> in the R help file for par for more information.
<code>x.ecdf.lty</code>	numeric scalar determining the line type used to display the empirical CDF for the <code>x</code> values that were supplied to gofTest . The default value is <code>x.ecdf.lty=1</code> . See the entry for <code>lty</code> in the R help file for par for more information.
<code>y.ecdf.lty</code>	numeric scalar determining the line type used to display the empirical CDF for the <code>y</code> values that were supplied to gofTest . The default value is <code>y.ecdf.lty=4</code> . See the entry for <code>lty</code> in the R help file for par for more information. Arguments associated with <code>plot.type="Q-Q Plot"</code> or <code>plot.type="Tukey M-D Q-Q Plot"</code> : As explained above, <code>plot.pos.con</code> is used for these plot types. Also:

<code>add.line</code>	logical scalar indicating whether to add a line to the plot. If <code>add.line=TRUE</code> and <code>plot.type="Q-Q Plot"</code> , a 0-1 line is added to the plot. If <code>add.line=TRUE</code> and <code>plot.type="Tukey M-D Q-Q Plot"</code> , a horizontal line at $y = 0$ is added to the plot. The default value is <code>add.line=TRUE</code> . Arguments associated with <code>plot.type="Test Results"</code>
<code>digits</code>	scalar indicating how many significant digits to print for the test results when <code>plot.type="Summary"</code> or <code>plot.type="Test Results"</code> . If <code>plot.type == "Summary"</code> , the default value is <code>digits=2</code> , otherwise it is <code>.Options\$digits</code> (i.e., the current setting of <code>options("digits")</code>).
<code>test.result.font</code>	numeric scalar indicating which font to use to print out the test results. The default value is <code>test.result.font=1</code> . See the description of the <code>font</code> argument in the help file for par for more information. You may get better results if you use a font number that corresponds to a fixed font (e.g., <code>courier</code>).
<code>test.result.cex</code>	numeric scalar indicating the value of <code>cex</code> to use to print out the test results. The default value is <code>0.9*par("cex")</code> when <code>plot.type="Summary"</code> , otherwise it is <code>par("cex")</code> . See the description of the <code>cex</code> argument in the help file for par for more information.
<code>test.result.mar</code>	numeric vector indicating the value of <code>mar</code> to use to print out the test results. The default value is <code>test.result.mar=c(0, 0, 3, 0)+0.1</code> . See the description of the <code>mar</code> argument in the help file for par for more information. Arguments associated with <code>plot.type="Summary"</code>
<code>add.om.title</code>	logical scalar indicating whether to add a title in the outer margin when <code>plot.type="Summary"</code> . The default value is <code>add.om.title=TRUE</code> .
<code>om.title</code>	character string containing the outer margin title. The default value is <code>om.title=NULL</code> , which will result in a default title.
<code>om.font</code>	numeric scalar indicating the font to use for the outer margin. The default value is <code>om.font=2</code> .
<code>om.cex.main</code>	numeric scalar indicating the value of <code>cex</code> for the outer margin title. The default value is <code>1.75 * par("cex")</code> .
<code>om.line</code>	numeric scalar indicating the line to place the outer margin title on. The default value is <code>om.line=0.5</code> . Graphics parameters:
<code>cex.main, cex.axis, cex.lab, main, xlab, ylab, xlim, ylim, oma, ...</code>	additional graphics parameters. See the help file for par .

Details

The function `plot.gofTwoSample` is a method for the generic function [plot](#) for the class `"gofTwoSample"` (see [gofTwoSample.object](#)). It can be invoked by calling [plot](#) and giving it an object of class `"gofTwoSample"` as the first argument, or by calling `plot.gofTwoSample` directly, regardless of the class of the object given as the first argument to `plot.gofTwoSample`.

Plots associated with the goodness-of-fit test are produced on the current graphics device. These can be one or all of the following:

- Observed distributions (plot.type="PDFs: Observed").
- Observed CDFs (plot.type="CDFs: Observed"). See the help file for [cdfCompare](#).
- Q-Q Plot (plot.type="Q-Q Plot"). See the help file for [qqPlot](#).
- Tukey mean-difference Q-Q plot (plot.type="Tukey M-D Q-Q Plot"). See the help file for [qqPlot](#).
- Results of the goodness-of-fit test (plot.type="Test Results"). See the help file for [print.gofTwoSample](#).

See the help file for [gofTest](#) for more information.

Value

plot.gofTwoSample invisibly returns the first argument, x.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Chambers, J. M. and Hastie, T. J. (1992). *Statistical Models in S*. Wadsworth & Brooks/Cole.

See Also

[gofTest](#), [gofTwoSample.object](#), [print.gofTwoSample](#), [Goodness-of-Fit Tests](#), [plot](#).

Examples

```
# Create an object of class "gofTwoSample" then plot the results.
# (Note: the call to set.seed simply allows you to reproduce
# this example.)

set.seed(300)
dat1 <- rnorm(20, mean = 3, sd = 2)
dat2 <- rnorm(10, mean = 1, sd = 2)
gof.obj <- gofTest(x = dat1, y = dat2)

# Summary plot (the default)
#-----
windows()
plot(gof.obj)

# Make your own titles for the summary plot
#-----
windows()
plot(gof.obj, captions = list(PDFs = "Compare PDFs",
  CDFs = "Compare CDFs", QQ = "Q-Q Plot", Results = "Results"),
  om.title = "Summary Plot")

# Just the Q-Q Plot
#-----
windows()
plot(gof.obj, plot.type="Q-Q")
```

```
# Make your own title for the Q-Q Plot
#-----
windows()
plot(gof.obj, plot.type="Q-Q", main = "Q-Q Plot")

#=====

# Clean up
#-----
rm(dat1, dat2, gof.obj)
graphics.off()
```

plotAovDesign

Create Plots for a Sampling Design Based on a One-Way Fixed-Effects Analysis of Variance

Description

Create plots involving sample size, power, scaled difference, and significance level for a one-way fixed-effects analysis of variance.

Usage

```
plotAovDesign(x.var = "n", y.var = "power",
  range.x.var = c(2, 50), n.vec = c(25, 25), mu.vec = c(0, 1),
  sigma = 1, alpha = 0.05, power = 0.95, round.up = FALSE,
  tol = 1e-07, plot.it = TRUE, add = FALSE, n.points = 50,
  plot.col = 1, plot.lwd = 3 * par("cex"), plot.lty = 1,
  digits = .Options$digits, main = NULL, xlab = NULL,
  ylab = NULL, type = "l", ...)
```

Arguments

- | | |
|-------------|--|
| x.var | character string indicating what variable to use for the x-axis. Possible values are "n" (sample size; the default), "power" (power of the test), and "alpha" (significance level of the test). |
| y.var | character string indicating what variable to use for the y-axis. Possible values are "power" (power of the test; the default) and "n" (sample size). |
| range.x.var | numeric vector of length 2 indicating the range of the x-variable to use for the plot. The default value depends on the value of x.var. When x.var="n" the default value is c(2,50). When x.var="power" the default value is c(alpha+.Machine\$double.eps, 0.95). When x.var="alpha", the default value is c(0.01, 0.2). |
| n.vec | numeric vector indicating the sample size for each group. The default value is n.vec=c(25, 25). Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed. This argument must be the same length as mu.vec. This argument is ignored if either x.var="n" or y.var="n". |
| mu.vec | numeric vector indicating the population mean for each group. The default value is mu.vec=c(0, 1). Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed. This argument must be the same length as n.vec. |

<code>sigma</code>	numeric scalar indicating the population standard deviation for all groups. The default value is <code>sigma=1</code> . Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed.
<code>alpha</code>	numeric scalar between 0 and 1 indicating the Type I error level associated with the hypothesis test. The default value is <code>alpha=0.05</code> . This argument is ignored when <code>x.var="alpha"</code> .
<code>power</code>	numeric scalar between 0 and 1 indicating the power associated with the hypothesis test. The default value is <code>power=0.95</code> . This argument is ignored when <code>x.var="power"</code> or <code>y.var="power"</code> .
<code>round.up</code>	logical scalar indicating whether to round up the values of the computed sample size(s) to the next smallest integer. The default value is FALSE. This argument is ignored unless <code>y.var="n"</code> .
<code>tol</code>	optional numeric scalar indicating the tolerance to use in the search algorithm for the sample size when <code>y.var="n"</code> . The default value is <code>tol=1e-7</code> .
<code>plot.it</code>	a logical scalar indicating whether to create a plot or add to the existing plot (see <code>add</code>) on the current graphics device. If <code>plot.it=FALSE</code> , no plot is produced, but a list of (x,y) values is returned (see <code>VALUE</code>). The default value is <code>plot.it=TRUE</code> .
<code>add</code>	a logical scalar indicating whether to add the design plot to the existing plot (<code>add=TRUE</code>), or to create a new plot (<code>add=FALSE</code>). The default value is <code>add=FALSE</code> . This argument is ignored if <code>plot.it=FALSE</code> .
<code>n.points</code>	a numeric scalar specifying how many (x,y) pairs to use to produce the plot. There are <code>n.points</code> x-values evenly spaced between <code>range.x.var[1]</code> and <code>range.x.var[2]</code> . The default value is <code>n.points=50</code> .
<code>plot.col</code>	a numeric scalar or character string determining the color of the plotted line or points. The default value is <code>plot.col=1</code> . See the entry for <code>col</code> in the help file for par for more information.
<code>plot.lwd</code>	a numeric scalar determining the width of the plotted line. The default value is <code>3*par("cex")</code> . See the entry for <code>lwd</code> in the help file for par for more information.
<code>plot.lty</code>	a numeric scalar determining the line type of the plotted line. The default value is <code>plot.lty=1</code> . See the entry for <code>lty</code> in the help file for par for more information.
<code>digits</code>	a scalar indicating how many significant digits to print out on the plot. The default value is the current setting of <code>options("digits")</code> .
<code>main, xlab, ylab, type, ...</code>	additional graphical parameters (see par).

Details

See the help files for [aovPower](#) and [aovN](#) for information on how to compute the power and sample size for a one-way fixed-effects analysis of variance.

Value

`plotAovDesign` invisibly returns a list with components:

<code>x.var</code>	x-coordinates of the points that have been or would have been plotted
<code>y.var</code>	y-coordinates of the points that have been or would have been plotted

Note

The normal and lognormal distribution are probably the two most frequently used distributions to model environmental data. Sometimes it is necessary to compare several means to determine whether any are significantly different from each other (e.g., USEPA, 2009, p.6-38). In this case, assuming normally distributed data, you perform a one-way parametric analysis of variance.

In the course of designing a sampling program, an environmental scientist may wish to determine the relationship between sample size, Type I error level, power, and differences in means if one of the objectives of the sampling program is to determine whether a particular mean differs from a group of means. The functions [aovPower](#), [aovN](#), and [plotAovDesign](#) can be used to investigate these relationships for the case of normally-distributed observations.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Berthouex, P.M., and L.C. Brown. (1994). *Statistics for Environmental Engineers*. Lewis Publishers, Boca Raton, FL, Chapter 17.
- Helsel, D.R., and R.M. Hirsch. (1992). *Statistical Methods in Water Resources Research*. Elsevier, New York, NY, Chapter 7.
- Johnson, N. L., S. Kotz, and N. Balakrishnan. (1995). *Continuous Univariate Distributions, Volume 2*. Second Edition. John Wiley and Sons, New York, Chapters 27, 29, 30.
- Scheffe, H. (1959). *The Analysis of Variance*. John Wiley and Sons, New York, 477pp.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.
- Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ, Chapter 10.

See Also

[aovPower](#), [aovN](#), [Normal](#), [aov](#).

Examples

```
# Look at the relationship between power and sample size
# for a one-way ANOVA, assuming k=2 groups, group means of
# 0 and 1, a population standard deviation of 1, and a
# 5% significance level:

windows()
plotAovDesign()

#-----

# Plot power vs. sample size for various levels of significance:

windows()
plotAovDesign(mu.vec = c(0, 0.5, 1), ylim=c(0, 1), main="")
```



```

plotAovDesign(mu.vec = c(0, 0.5, 1), alpha=0.1, add=TRUE, plot.col=2)

plotAovDesign(mu.vec = c(0, 0.5, 1), alpha=0.2, add=TRUE, plot.col=3)

legend(35, 0.6, c("20%", "10%", " 5%"), lty=1, lwd = 3, col=3:1,
      bty = "n")

mtext("Power vs. Sample Size for One-Way ANOVA", line = 3, cex = 1.25)
mtext(expression(paste("with ", mu, "=(0, 0.5, 1), ", sigma,
      "=1, and Various Significance Levels", sep="")),
      line = 1.5, cex = 1.25)

#-----

# The example on pages 5-11 to 5-14 of USEPA (1989b) shows
# log-transformed concentrations of lead (mg/L) at two
# background wells and four compliance wells, where
# observations were taken once per month over four months
# (the data are stored in EPA.89b.loglead.df).
# Assume the true mean levels at each well are
# 3.9, 3.9, 4.5, 4.5, 4.5, and 5, respectively. Plot the
# power vs. sample size of a one-way ANOVA to test for mean
# differences between wells. Use alpha=0.05, and assume the
# true standard deviation is equal to the one estimated
# from the data in this example.

names(EPA.89b.loglead.df)
#[1] "LogLead" "Month" "Well" "Well.type"

# Perform the ANOVA and get the estimated sd
aov.list <- aov(LogLead ~ Well, data=EPA.89b.loglead.df)

summary(aov.list)
#           Df Sum Sq Mean Sq F value Pr(>F)
#Well          5  5.7447  1.14895   3.3469 0.02599 *
#Residuals    18  6.1791  0.34328
#---
#Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

# Now create the plot
windows()
plotAovDesign(range.x.var = c(2, 20),
  mu.vec = c(3.9,3.9,4.5,4.5,4.5,5),
  sigma=sqrt(0.34),
  ylim = c(0, 1), digits=2)

# Clean up
#-----
rm(aov.list)
graphics.off()

```

Description

Create plots for a sampling design based on a confidence interval for a binomial proportion or the difference between two proportions.

Usage

```
plotCiBinomDesign(x.var = "n", y.var = "half.width",
  range.x.var = c(10, 50), n.or.n1 = 25, p.hat.or.p1.hat = 0.5,
  n2 = n.or.n1, p2.hat = 0.4, ratio = 1, half.width = 0.05,
  conf = 0.95, sample.type = "one.sample", ci.method = "score",
  correct = TRUE, warn = TRUE, n.or.n1.min = 2,
  n.or.n1.max = 10000, tol.half.width = 0.005, tol.p.hat = 0.005,
  maxiter = 10000, plot.it = TRUE, add = FALSE, n.points = 100,
  plot.col = 1, plot.lwd = 3 * par("cex"), plot.lty = 1,
  digits = .Options$digits,
  main = NULL, xlab = NULL, ylab = NULL, type = "l", ...)
```

Arguments

- | | |
|-----------------|---|
| x.var | character string indicating what variable to use for the x-axis. Possible values are "n" (sample size; the default), "half.width" (the half-width of the confidence interval), "p.hat" (the estimated probability of dQuotesuccess), and "conf" (the confidence level). |
| y.var | character string indicating what variable to use for the y-axis. Possible values are "half.width" (the half-width of the confidence interval; the default), and "n" (sample size). |
| range.x.var | numeric vector of length 2 indicating the range of the x-variable to use for the plot. The default value depends on the value of x.var.
When x.var="n" the default value is c(10, 50).
When x.var="half.width", the default value is c(0.03, 0.1).
When x.var="p.hat", the default value is c(0.5, 0.9).
When x.var="conf", the default value is c(0.8, 0.99). |
| n.or.n1 | numeric scalar indicating the sample size. The default value is n.or.n1=25. When sample.type="one.sample", n.or.n1 denotes the number of observations in the single sample. When sample.type="two.sample", n.or.n1 denotes the number of observations from group 1. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed. This argument is ignored if either x.var="n" or y.var="n". |
| p.hat.or.p1.hat | numeric scalar indicating an estimated proportion.
When sample.type="one.sample", p.hat.or.p1.hat denotes the estimated value of p , the probability of dQuotesuccess.
When sample.type="two.sample", p.hat.or.p1.hat denotes the estimated value of p_1 , the probability of dQuotesuccess in group 1.
Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed. This argument is ignored if x.var="p.hat". |
| n2 | numeric scalar indicating the sample size for group 2. The default value is the value of n.or.n1. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed. This argument is ignored when sample.type="one.sample". |

p2.hat	numeric scalar indicating the estimated proportion for group 2. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed. This argument is ignored if sample.type="one.sample".
ratio	numeric vector indicating the ratio of sample size in group 2 to sample size in group 1 (n_2/n_1). The default value is ratio=1. All values of ratio must be greater than or equal to 1. This argument is only used when sample.type="two.sample" and either x.var="n" or y.var="n".
half.width	positive numeric scalar indicating the half-width of the confidence interval. The default value is half.width=0.05. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed. This argument is ignored if either x.var="half.width" or y.var="half.width".
conf	a numeric scalar between 0 and 1 indicating the confidence level associated with the confidence intervals. The default value is conf=0.95. This argument is ignored when x.var="conf".
sample.type	character string indicating whether this is a one-sample or two-sample confidence interval. When sample.type="one.sample", the computations for the plot are based on a confidence interval for a single proportion. When sample.type="two.sample", the computations for the plot are based on a confidence interval for the difference between two proportions. The default value is sample.type="one.sample" unless the arguments n2, p2.hat, and/or ratio are supplied.
ci.method	character string indicating which method to use to construct the confidence interval. Possible values are "score" (the default), "exact", "adjusted Wald", and "Wald" (the "Wald" method is never recommended but is included for historical purposes). The exact method is only available for the one-sample case, i.e., when sample.type="one.sample".
correct	logical scalar indicating whether to use the continuity correction when ci.method="score" or ci.method="Wald". The default value is correct=TRUE. This argument is ignored if ci.method="exact" or ci.method="adjusted Wald".
warn	logical scalar indicating whether to issue a warning when ci.method="Wald" for cases when the normal approximation to the binomial distribution probably is not accurate. The default value is warn=TRUE.
n.or.n1.min	for the case when y.var="n", integer indicating the minimum allowed value for n (sample.type="one.sample") or n_1 (sample.type="two.sample"). The default value is n.or.n1.min=2.
n.or.n1.max	for the case when y.var="n", integer indicating the maximum allowed value for n (sample.type="one.sample") or n_1 (sample.type="two.sample"). The default value is n.or.n1.max=10000.
tol.half.width	for the case when y.var="n", numeric scalar indicating the tolerance to use for the half width for the search algorithm. The sample sizes are computed so that the actual half width is less than or equal to half.width + tol.half.width. The default value is tol.half.width=0.005.
tol.p.hat	for the case when y.var="n", numeric scalar indicating the tolerance to use for the estimated proportion(s) for the search algorithm. For the one-sample case, the sample sizes are computed so that the absolute value of the difference

between the user supplied value of \hat{p} or \hat{p}_1 and the actual estimated proportion is less than or equal to tol.p.hat . For the two-sample case, the sample sizes are computed so that the absolute value of the difference between the user supplied value of \hat{p} or \hat{p}_1 and the actual estimated proportion for group 1 is less than or equal to tol.p.hat , and the absolute value of the difference between the user supplied value of \hat{p}_2 and the actual estimated proportion for group 2 is less than or equal to tol.p.hat . The default value is $\text{tol.p.hat}=0.005$.

<code>maxiter</code>	for the case when <code>y.var="n"</code> , integer indicating the maximum number of iterations to use for the search algorithm. The default value is <code>maxiter=1000</code> .
<code>plot.it</code>	a logical scalar indicating whether to create a plot or add to the existing plot (see description of the argument <code>add</code> below) on the current graphics device. If <code>plot.it=FALSE</code> , no plot is produced, but a list of (x,y) values is returned (see the VALUE section below). The default value is <code>plot.it=TRUE</code> .
<code>add</code>	a logical scalar indicating whether to add the design plot to the existing plot (<code>add=TRUE</code>), or to create a plot from scratch (<code>add=FALSE</code>). The default value is <code>add=FALSE</code> . This argument is ignored if <code>plot.it=FALSE</code> .
<code>n.points</code>	a numeric scalar specifying how many (x,y) pairs to use to produce the plot. There are <code>n.points</code> x-values evenly spaced between <code>range.x.var[1]</code> and <code>range.x.var[2]</code> . The default value is <code>n.points=100</code> . This argument is ignored when <code>x.var="n"</code> , in which case the x-values are all the integers between <code>range.x.var[1]</code> and <code>range.x.var[2]</code> .
<code>plot.col</code>	a numeric scalar or character string determining the color of the plotted line or points. The default value is <code>plot.col=1</code> . See the entry for <code>col</code> in the help file for par for more information.
<code>plot.lwd</code>	a numeric scalar determining the width of the plotted line. The default value is <code>3*par("cex")</code> . See the entry for <code>lwd</code> in the help file for par for more information.
<code>plot.lty</code>	a numeric scalar determining the line type of the plotted line. The default value is <code>plot.lty=1</code> . See the entry for <code>lty</code> in the help file for par for more information.
<code>digits</code>	a scalar indicating how many significant digits to print out on the plot. The default value is the current setting of <code>options("digits")</code> .
<code>main, xlab, ylab, type, ...</code>	additional graphical parameters (see par).

Details

See the help files for [ciBinomHalfWidth](#) and [ciBinomN](#) for information on how to compute a one-sample confidence interval for a single binomial proportion or a two-sample confidence interval for the difference between two proportions, how the half-width is computed when other quantities are fixed, and how the sample size is computed when other quantities are fixed.

Value

`plotCiBinomDesign` invisibly returns a list with components:

<code>x.var</code>	x-coordinates of the points that have been or would have been plotted
<code>y.var</code>	y-coordinates of the points that have been or would have been plotted

Note

The binomial distribution is used to model processes with binary (Yes-No, Success-Failure, Heads-Tails, etc.) outcomes. It is assumed that the outcome of any one trial is independent of any other trial, and that the probability of “success”, p , is the same on each trial. A binomial discrete random variable X is the number of “successes” in n independent trials. A special case of the binomial distribution occurs when $n = 1$, in which case X is also called a Bernoulli random variable.

In the context of environmental statistics, the binomial distribution is sometimes used to model the proportion of times a chemical concentration exceeds a set standard in a given period of time (e.g., Gilbert, 1987, p.143), or to compare the proportion of detects in a compliance well vs. a background well (e.g., USEPA, 1989b, Chapter 8, p.3-7). (However, USEPA 2009, p.8-27 recommends using the Wilcoxon rank sum test ([wilcox.test](#)) instead of comparing proportions.)

In the course of designing a sampling program, an environmental scientist may wish to determine the relationship between sample size, confidence level, and half-width if one of the objectives of the sampling program is to produce confidence intervals. The functions `ciBinomHalfWidth`, `ciBinomN`, and `plotCiBinomDesign` can be used to investigate these relationships for the case of binomial proportions.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Agresti, A., and B.A. Coull. (1998). Approximate is Better than "Exact" for Interval Estimation of Binomial Proportions. *The American Statistician*, **52**(2), 119–126.
- Agresti, A., and B. Caffo. (2000). Simple and Effective Confidence Intervals for Proportions and Differences of Proportions Result from Adding Two Successes and Two Failures. *The American Statistician*, **54**(4), 280–288.
- Berthouex, P.M., and L.C. Brown. (1994). *Statistics for Environmental Engineers*. Lewis Publishers, Boca Raton, FL, Chapters 2 and 15.
- Cochran, W.G. (1977). *Sampling Techniques*. John Wiley and Sons, New York, Chapter 3.
- Fisher, R.A., and F. Yates. (1963). *Statistical Tables for Biological, Agricultural, and Medical Research*. 6th edition. Hafner, New York, 146pp.
- Fleiss, J. L. (1981). *Statistical Methods for Rates and Proportions*. Second Edition. John Wiley and Sons, New York, Chapters 1-2.
- Gilbert, R.O. (1987). *Statistical Methods for Environmental Pollution Monitoring*. Van Nostrand Reinhold, New York, NY, Chapter 11.
- Newcombe, R.G. (1998a). Two-Sided Confidence Intervals for the Single Proportion: Comparison of Seven Methods. *Statistics in Medicine*, **17**, 857–872.
- Newcombe, R.G. (1998b). Interval Estimation for the Difference Between Independent Proportions: Comparison of Eleven Methods. *Statistics in Medicine*, **17**, 873–890.
- Ott, W.R. (1995). *Environmental Statistics and Data Analysis*. Lewis Publishers, Boca Raton, FL, Chapter 4.
- USEPA. (1989b). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities, Interim Final Guidance*. EPA/530-SW-89-026. Office of Solid Waste, U.S. Environmental Protection Agency, Washington, D.C.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery

Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.

Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ, Chapter 24.

See Also

[ciBinomHalfWidth](#), [ciBinomN](#), [ebinom](#), [binom.test](#), [prop.test](#), [par](#).

Examples

```
# Look at the relationship between half-width and sample size
# for a one-sample confidence interval for a binomial proportion,
# assuming an estimated proportion of 0.5 and a confidence level of
# 95%. The jigsaw appearance of the plot is the result of using the
# score method:

windows()
plotCiBinomDesign()

#-----

# Redo the example above, but use the traditional (and inaccurate)
# Wald method.

windows()
plotCiBinomDesign(ci.method = "Wald")

#-----

# Plot sample size vs. the estimated proportion for various half-widths,
# using a 95% confidence level and the adjusted Wald method:

windows()
plotCiBinomDesign(x.var = "p.hat", y.var = "n",
  half.width = 0.04, ylim = c(0, 600), main = "",
  xlab = expression(hat(p)))

plotCiBinomDesign(x.var = "p.hat", y.var = "n",
  half.width = 0.05, add = TRUE, plot.col = 2)

plotCiBinomDesign(x.var = "p.hat", y.var = "n",
  half.width = 0.06, add = TRUE, plot.col = 3)

legend(0.5, 150, paste("Half-Width =", c(0.04, 0.05, 0.06)),
  lty = rep(1, 3), lwd = rep(2, 3), col=1:3, bty = "n")

mtext(expression(paste("Sample Size vs. ", hat(p),
  " for Confidence Interval for p")), line = 2.5, cex = 1.25)
mtext("with Confidence=95% and Various Values of Half-Width",
  line = 1.5, cex = 1.25)
mtext(paste("CI Method = Score Normal Approximation",
  "with Continuity Correction"), line = 0.5)

#-----
```

```

# Modifying the example on pages 8-5 to 8-7 of USEPA (1989b),
# look at the relationship between half-width and sample size
# for a 95% confidence interval for the difference between the
# proportion of detects at the background and compliance wells.
# Use the estimated proportion of detects from the original data.
# (The data are stored in EPA.89b.cadmium.df.)
# Assume equal sample sizes at each well.

EPA.89b.cadmium.df
#   Cadmium.orig Cadmium Censored Well.type
#1          0.1    0.100     FALSE Background
#2          0.12   0.120     FALSE Background
#3          BDL    0.000      TRUE Background
# .....
#86          BDL    0.000      TRUE Compliance
#87          BDL    0.000      TRUE Compliance
#88          BDL    0.000      TRUE Compliance

p.hat.back <- with(EPA.89b.cadmium.df,
  mean(!Censored[Well.type=="Background"]))

p.hat.back
#[1] 0.3333333

p.hat.comp <- with(EPA.89b.cadmium.df,
  mean(!Censored[Well.type=="Compliance"]))

p.hat.comp
#[1] 0.375

windows()
plotCiBinomDesign(p.hat.or.p1.hat = p.hat.back,
  p2.hat = p.hat.comp, digits=3)

#=====

# Clean up
#-----
rm(p.hat.back, p.hat.comp)
graphics.off()

```

plotCiNormDesign

Create Plots for Sampling Design Based on Confidence Interval for Mean of a Normal Distribution or Difference Between Two Means

Description

Create plots involving sample size, half-width, estimated standard deviation, and confidence level for a confidence interval for the mean of a normal distribution or the difference between two means.

Usage

```

plotCiNormDesign(x.var = "n", y.var = "half.width",
  range.x.var = c(2, 50), n.or.n1 = 25, n2 = n.or.n1,

```

```

half.width = sigma.hat/2, sigma.hat = 1, conf = 0.95,
sample.type = ifelse(missing(n2), "one.sample", "two.sample"),
round.up = FALSE, plot.it = TRUE, add = FALSE, n.points = 100,
plot.col = 1, plot.lwd = 3 * par("cex"), plot.lty = 1,
digits = .Options$digits,
main = NULL, xlab = NULL, ylab = NULL, type = "l", ...)

```

Arguments

<code>x.var</code>	character string indicating what variable to use for the x-axis. Possible values are "n" (sample size; the default), "half.width" (the half-width of the confidence interval), "sigma.hat" (the estimated standard deviation), and "conf" (the confidence level).
<code>y.var</code>	character string indicating what variable to use for the y-axis. Possible values are "half.width" (the half-width of the confidence interval; the default), and "n" (sample size).
<code>range.x.var</code>	numeric vector of length 2 indicating the range of the x-variable to use for the plot. The default value depends on the value of <code>x.var</code> . When <code>x.var="n"</code> the default value is <code>c(2,50)</code> . When <code>x.var="half.width"</code> the default value is <code>c(0.1/sigma.hat, 2/sigma.hat)</code> . When <code>x.var="sigma.hat"</code> , the default value is <code>c(0.1, 2)</code> . When <code>x.var="conf"</code> , the default value is <code>c(0.5, 0.99)</code> .
<code>n.or.n1</code>	numeric scalar indicating the sample size. The default value is <code>n.or.n1=25</code> . When <code>sample.type="one.sample"</code> , this argument denotes the number of observations in the single sample. When <code>sample.type="two.sample"</code> , this argument denotes the number of observations from group 1. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed. This argument is ignored if either <code>x.var="n"</code> or <code>y.var="n"</code> .
<code>n2</code>	numeric scalar indicating the sample size for group 2. The default value is the value of <code>n.or.n1</code> . Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed. This argument is ignored when <code>sample.type="one.sample"</code> .
<code>half.width</code>	positive numeric scalar indicating the half-width of the confidence interval. The default value is <code>sigma.hat/2</code> . Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed. This argument is ignored if either <code>x.var="half.width"</code> or <code>y.var="half.width"</code> .
<code>sigma.hat</code>	positive numeric scalar specifying the estimated standard deviation. The default value is <code>sigma.hat=1</code> . This argument is ignored if <code>x.var="sigma.hat"</code> .
<code>conf</code>	a scalar between 0 and 1 indicating the confidence level associated with the confidence interval. The default value is <code>conf=0.95</code> . This argument is ignored if <code>x.var="conf"</code> .
<code>sample.type</code>	character string indicating whether this is a one-sample or two-sample confidence interval. When <code>sample.type="one.sample"</code> , the computations for the plot are based on a confidence interval for a single mean. When <code>sample.type="two.sample"</code> , the computations for the plot are based on a confidence interval for the difference between two means. The default value is <code>sample.type="one.sample"</code> unless the argument <code>n2</code> is supplied.
<code>round.up</code>	logical scalar indicating whether to round up the computed sample sizes to the next smallest integer. The default value is <code>round.up=FALSE</code> . This argument is ignored unless <code>y.var="n"</code> .

<code>plot.it</code>	a logical scalar indicating whether to create a plot or add to the existing plot (see explanation of the argument <code>add</code> below) on the current graphics device. If <code>plot.it=FALSE</code> , no plot is produced, but a list of (x,y) values is returned (see the section <code>VALUE</code>). The default value is <code>plot.it=TRUE</code> .
<code>add</code>	a logical scalar indicating whether to add the design plot to the existing plot (<code>add=TRUE</code>), or to create a plot from scratch (<code>add=FALSE</code>). The default value is <code>add=FALSE</code> . This argument is ignored if <code>plot.it=FALSE</code> .
<code>n.points</code>	a numeric scalar specifying how many (x,y) pairs to use to produce the plot. There are <code>n.points</code> x-values evenly spaced between <code>range.x.var[1]</code> and <code>range.x.var[2]</code> . The default value is <code>n.points=100</code> .
<code>plot.col</code>	a numeric scalar or character string determining the color of the plotted line or points. The default value is <code>plot.col=1</code> . See the entry for <code>col</code> in the help file for par for more information.
<code>plot.lwd</code>	a numeric scalar determining the width of the plotted line. The default value is <code>3*par("cex")</code> . See the entry for <code>lwd</code> in the help file for par for more information.
<code>plot.lty</code>	a numeric scalar determining the line type of the plotted line. The default value is <code>plot.lty=1</code> . See the entry for <code>lty</code> in the help file for par for more information.
<code>digits</code>	a scalar indicating how many significant digits to print out on the plot. The default value is the current setting of <code>options("digits")</code> .
<code>main, xlab, ylab, type, ...</code>	additional graphical parameters (see par).

Details

See the help files for [ciNormHalfWidth](#) and [ciNormN](#) for information on how to compute a one-sample confidence interval for the mean of a normal distribution or a two-sample confidence interval for the difference between two means, how the half-width is computed when other quantities are fixed, and how the sample size is computed when other quantities are fixed.

Value

`plotCiNormDesign` invisibly returns a list with components:

<code>x.var</code>	x-coordinates of points that have been or would have been plotted.
<code>y.var</code>	y-coordinates of points that have been or would have been plotted.

Note

The normal distribution and lognormal distribution are probably the two most frequently used distributions to model environmental data. In order to make any kind of probability statement about a normally-distributed population (of chemical concentrations for example), you have to first estimate the mean and standard deviation (the population parameters) of the distribution. Once you estimate these parameters, it is often useful to characterize the uncertainty in the estimate of the mean. This is done with confidence intervals.

In the course of designing a sampling program, an environmental scientist may wish to determine the relationship between sample size, confidence level, and half-width if one of the objectives of the sampling program is to produce confidence intervals. The functions [ciNormHalfWidth](#), [ciNormN](#), and `plotCiNormDesign` can be used to investigate these relationships for the case of normally-distributed observations.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Berthouex, P.M., and L.C. Brown. (2002). *Statistics for Environmental Engineers*. Second Edition. Lewis Publishers, Boca Raton, FL.
- Gilbert, R.O. (1987). *Statistical Methods for Environmental Pollution Monitoring*. Van Nostrand Reinhold, New York, NY.
- Helsel, D.R., and R.M. Hirsch. (1992). *Statistical Methods in Water Resources Research*. Elsevier, New York, NY, Chapter 7.
- Millard, S.P., and N. Neerchal. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, FL.
- Ott, W.R. (1995). *Environmental Statistics and Data Analysis*. Lewis Publishers, Boca Raton, FL.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. p.21-3.
- Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ, Chapters 7 and 8.

See Also

[plotCiNormHalfWidth](#), [ciNormN](#), [Normal](#), [enorm](#), [t.test](#),
[Estimating Distribution Parameters](#).

Examples

```
# Look at the relationship between half-width and sample size
# for a one-sample confidence interval for the mean, assuming
# an estimated standard deviation of 1 and a confidence level of 95%.

windows()
plotCiNormDesign()

#-----

# Plot sample size vs. the estimated standard deviation for
# various levels of confidence, using a half-width of 0.5.

windows()
plotCiNormDesign(x.var = "sigma.hat", y.var = "n", main = "")

plotCiNormDesign(x.var = "sigma.hat", y.var = "n", conf = 0.9,
  add = TRUE, plot.col = 2)

plotCiNormDesign(x.var = "sigma.hat", y.var = "n", conf = 0.8,
  add = TRUE, plot.col = 3)

legend(0.25, 60, c("95%", "90%", "80%"), lty = 1, lwd = 3, col = 1:3)

mtext("Sample Size vs. Estimated SD for Confidence Interval for Mean",
  font = 2, cex = 1.25, line = 2.75)
```

```

mtext("with Half-Width=0.5 and Various Confidence Levels", font = 2,
      cex = 1.25, line = 1.25)

#-----

# Modifying the example on pages 21-4 to 21-5 of USEPA (2009),
# look at the relationship between half-width and sample size for a
# 95% confidence interval for the mean level of Aldicarb at the
# first compliance well. Use the estimated standard deviation from
# the first four months of data.
# (The data are stored in EPA.09.Ex.21.1.aldicarb.df.)

EPA.09.Ex.21.1.aldicarb.df
#   Month   Well Aldicarb.ppb
#1      1 Well.1      19.9
#2      2 Well.1      29.6
#3      3 Well.1      18.7
#4      4 Well.1      24.2
#...

mu.hat <- with(EPA.09.Ex.21.1.aldicarb.df,
               mean(Aldicarb.ppb[Well=="Well.1"]))

mu.hat
#[1] 23.1

sigma.hat <- with(EPA.09.Ex.21.1.aldicarb.df,
                  sd(Aldicarb.ppb[Well=="Well.1"]))

sigma.hat
#[1] 4.93491

windows()
plotCiNormDesign(sigma.hat = sigma.hat, digits = 2,
                 range.x.var = c(2, 25))

#=====

# Clean up
#-----
rm(mu.hat, sigma.hat)
graphics.off()

```

pointwise

Pointwise Confidence Limits for Predictions

Description

Computes pointwise confidence limits for predictions computed by the function [predict](#).

Usage

```

pointwise(results.predict, coverage = 0.99,
          simultaneous = FALSE, individual = FALSE)

```

Arguments

<code>results.predict</code>	output from a call to <code>predict</code> with <code>se.fit=TRUE</code> .
<code>coverage</code>	optional numeric scalar between 0 and 1 indicating the confidence level associated with the confidence limits. The default value is <code>coverage=0.99</code> .
<code>simultaneous</code>	optional logical scalar indicating whether to base the confidence limits for the predicted values on simultaneous or non-simultaneous prediction limits. The default value is <code>simultaneous=FALSE</code> .
<code>individual</code>	optional logical scalar indicating whether to base the confidence intervals for the predicted values on prediction limits for the mean (<code>individual=FALSE</code>) or prediction limits for an individual observation (<code>individual=TRUE</code>). The default value is <code>individual=FALSE</code> .

Details

This function computes pointwise confidence limits for predictions computed by the function `predict`. The limits are computed at those points specified by the argument `newdata` of `predict`.

The `predict` function is a generic function with methods for several different classes. The function `pointwise` was part of the S language. The modifications to `pointwise` in the package `EnvStats` involve confidence limits for predictions for a linear model (i.e., an object of class `lm`).

Confidence Limits for a Predicted Mean Value (`individual=FALSE`). Consider a standard linear model with p predictor variables. Often, one of the major goals of regression analysis is to predict a future value of the response variable given known values of the predictor variables. The equations for the predicted mean value of the response given fixed values of the predictor variables as well as the equation for a two-sided $(1-\alpha)100\%$ confidence interval for the mean value of the response can be found in Draper and Smith (1998, p.80) and Millard and Neerchal (2001, p.547).

Technically, this formula is a confidence interval for the mean of the response for one set of fixed values of the predictor variables and corresponds to the case when `simultaneous=FALSE`. To create simultaneous confidence intervals over the range of the predictor variables, the critical t-value in the equation has to be replaced with a critical F-value and the modified formula is given in Draper and Smith (1998, p. 83), Miller (1981a, p. 111), and Millard and Neerchal (2001, p. 547). This formula is used in the case when `simultaneous=TRUE`.

Confidence Limits for a Predicted Individual Value (`individual=TRUE`). In the above section we discussed how to create a confidence interval for the mean of the response given fixed values for the predictor variables. If instead we want to create a prediction interval for a single future observation of the response variable, the formula is given in Miller (1981a, p. 115) and Millard and Neerchal (2001, p. 551).

Technically, this formula is a prediction interval for a single future observation for one set of fixed values of the predictor variables and corresponds to the case when `simultaneous=FALSE`. Miller (1981a, p. 115) gives a formula for simultaneous prediction intervals for k future observations. If we are interested in creating an interval that will encompass *all* possible future observations over the range of the predictor variables with some specified probability however, we need to create simultaneous tolerance intervals. A formula for such an interval was developed by Lieberman and Miller (1963) and is given in Miller (1981a, p. 124). This formula is used in the case when `simultaneous=TRUE`.

Value

a list with the following components:

<code>upper</code>	upper limits of pointwise confidence intervals.
--------------------	---

fit	surface values. This is the same as the component fit of the argument <code>results.predict</code> .
lower	lower limits of pointwise confidence intervals.

Note

The function `pointwise` is called by the functions `detectionLimitCalibrate` and `inversePredictCalibrate`, which are used in **calibration**.

Almost always the process of determining the concentration of a chemical in a soil, water, or air sample involves using some kind of machine that produces a signal, and this signal is related to the concentration of the chemical in the physical sample. The process of relating the machine signal to the concentration of the chemical is called **calibration** (see `calibrate`). Once calibration has been performed, estimated concentrations in physical samples with unknown concentrations are computed using inverse regression. The uncertainty in the process used to estimate the concentration may be quantified with decision, detection, and quantitation limits.

In practice, only the point estimate of concentration is reported (along with a possible qualifier), without confidence bounds for the true concentration C . This is most unfortunate because it gives the impression that there is no error associated with the reported concentration. Indeed, both the International Organization for Standardization (ISO) and the International Union of Pure and Applied Chemistry (IUPAC) recommend always reporting both the estimated concentration and the uncertainty associated with this estimate (Currie, 1997).

Author(s)

Authors of S (for code for `pointwise` in S).

Steven P. Millard (for modification to allow the arguments simultaneous and individual);
<EnvStats@ProbStatInfo.com>

References

- Chambers, J.M., and Hastie, T.J., eds. (1992). *Statistical Models in S*. Chapman and Hall/CRC, Boca Raton, FL.
- Draper, N., and H. Smith. (1998). *Applied Regression Analysis*. Third Edition. John Wiley and Sons, New York, Chapter 3.
- Millard, S.P., and N.K. Neerchal. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, FL, pp.546-553.
- Miller, R.G. (1981a). *Simultaneous Statistical Inference*. Springer-Verlag, New York, pp.111, 124.

See Also

`predict`, `lm`, `calibrate`, `inversePredictCalibrate`, `detectionLimitCalibrate`.

Examples

```
# Using the data in the built-in data frame Air.df,
# fit the cube root of ozone as a function of temperature.
# Then compute predicted values for ozone at 70 and 90
# degrees F, and compute 95% confidence intervals for the
# mean value of ozone at these temperatures.

# First create the lm object
#-----
```

```

ozone.fit <- lm(ozone ~ temperature, data = Air.df)

# Now get predicted values and CIs at 70 and 90 degrees
#-----

predict.list <- predict(ozone.fit,
  newdata = data.frame(temperature = c(70, 90)), se.fit = TRUE)

pointwise(predict.list, coverage = 0.95)
# $upper
#      1      2
# 2.839145 4.278533

# $fit
#      1      2
# 2.697810 4.101808

# $lower
#      1      2
# 2.556475 3.925082

#-----

# Continuing with the above example, create a scatterplot of ozone
# vs. temperature, and add the fitted line along with simultaneous
# 95% confidence bands.

x <- Air.df$temperature

y <- Air.df$ozone

windows()
plot(x, y, xlab="Temperature (degrees F)",
  ylab = expression(sqrt("Ozone (ppb)", 3)))

abline(ozone.fit, lwd = 2)

new.x <- seq(min(x), max(x), length=100)

predict.ozone <- predict(ozone.fit,
  newdata = data.frame(temperature = new.x), se.fit = TRUE)

ci.ozone <- pointwise(predict.ozone, coverage=0.95,
  simultaneous=TRUE)

lines(new.x, ci.ozone$lower, lty=2, lwd = 2, col = 2)

lines(new.x, ci.ozone$upper, lty=2, lwd = 2, col = 2)

title(main=paste("Cube Root Ozone vs. Temperature with Fitted Line",
  "and Simultaneous 95% Confidence Bands",
  sep="\n"))

#-----

```

```

# Redo the last example by creating non-simultaneous
# confidence bounds and prediction bounds as well.

windows()
plot(x, y, xlab = "Temperature (degrees F)",
     ylab = expression(sqrt("Ozone (ppb)", 3)))

abline(ozone.fit, lwd = 2)

new.x <- seq(min(x), max(x), length=100)

predict.ozone <- predict(ozone.fit,
  newdata = data.frame(temperature = new.x), se.fit = TRUE)

ci.ozone <- pointwise(predict.ozone, coverage=0.95)

lines(new.x, ci.ozone$lower, lty=2, col = 2, lwd = 2)

lines(new.x, ci.ozone$upper, lty=2, col = 2, lwd = 2)

pi.ozone <- pointwise(predict.ozone, coverage = 0.95,
  individual = TRUE)

lines(new.x, pi.ozone$lower, lty=4, col = 4, lwd = 2)

lines(new.x, pi.ozone$upper, lty=4, col = 4, lwd = 2)

title(main=paste("Cube Root Ozone vs. Temperature with Fitted Line",
  "and 95% Confidence and Prediction Bands",
  sep="\n"))

#-----

# Clean up
rm(predict.list, ozone.fit, x, y, new.x, predict.ozone, ci.ozone,
  pi.ozone)

```

ppointsCensored

Plotting Positions for Type I Censored Data

Description

Returns a list of “ordered” observations and associated plotting positions based on Type I left-censored or right-censored data. These plotting positions may be used to construct empirical cumulative distribution plots or quantile-quantile plots, or to estimate distribution parameters.

Usage

```

ppointsCensored(x, censored, censoring.side = "left",
  prob.method = "michael-schucany", plot.pos.con = 0.375)

```

Arguments

<code>x</code>	numeric vector of observations. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
<code>censored</code>	numeric or logical vector indicating which values of <code>x</code> are censored. This must be the same length as <code>x</code> . If the mode of <code>censored</code> is "logical", TRUE values correspond to elements of <code>x</code> that are censored, and FALSE values correspond to elements of <code>x</code> that are not censored. If the mode of <code>censored</code> is "numeric", it must contain only 1's and 0's; 1 corresponds to TRUE and 0 corresponds to FALSE. Missing (NA) values are allowed but will be removed.
<code>censoring.side</code>	character string indicating on which side the censoring occurs. The possible values are "left" (the default) and "right".
<code>prob.method</code>	character string indicating what method to use to compute the plotting positions (empirical probabilities). Possible values are "kaplan-meier" (product-limit method of Kaplan and Meier (1958)), "nelson" (hazard plotting method of Nelson (1972)), "michael-schucany" (generalization of the product-limit method due to Michael and Schucany (1986)), and "hirsch-stedinger" (generalization of the product-limit method due to Hirsch and Stedinger (1987)). The default value is <code>prob.method="michael-schucany"</code> . The "nelson" method is only available for <code>censoring.side="right"</code> . See the DETAILS section for more explanation.
<code>plot.pos.con</code>	numeric scalar between 0 and 1 containing the value of the plotting position constant. The default value is <code>plot.pos.con=0.375</code> . See the DETAILS section for more information. This argument is used only if <code>prob.method</code> is equal to "michael-schucany" or "hirsch-stedinger".

Details

Methods for computing plotting positions for complete data sets (no censored observations) are discussed in D'Agostino, R.B. (1986a) and Cleveland (1993). For data sets with censored observations, these methods must be modified. The function `ppointsCensored` allows you to compute plotting positions based on the product-limit method of Kaplan and Meier (1958) (`prob.method="kaplan-meier"`), the hazard plotting method of Nelson (1972) (`prob.method="nelson"`), the generalization of the product-limit method due to Michael and Schucany (1986) (`prob.method="michael-schucany"`; the default), and the generalization of the product-limit method due to Hirsch and Stedinger (1987) (`prob.method="hirsch-stedinger"`).

Let \underline{x} denote a random sample of N observations from some distribution. Assume n ($0 < n < N$) of these observations are known and c ($c = N - n$) of these observations are all censored below (left-censored) or all censored above (right-censored) at k fixed censoring levels

$$T_1, T_2, \dots, T_K; K \geq 1 \quad (1)$$

For the case when $K \geq 2$, the data are said to be Type I **multiply censored**. For the case when $K = 1$, set $T = T_1$. If the data are left-censored and all n known observations are greater than or equal to T , or if the data are right-censored and all n known observations are less than or equal to T , then the data are said to be Type I **singly censored** (Nelson, 1982, p.7), otherwise they are considered to be Type I multiply censored.

Let c_j denote the number of observations censored below or above censoring level T_j for $j = 1, 2, \dots, K$, so that

$$\sum_{i=1}^K c_j = c \quad (2)$$

Let $x_{(1)}, x_{(2)}, \dots, x_{(N)}$ denote the “ordered” observations, where now “observation” means either the actual observation (for uncensored observations) or the censoring level (for censored observations). For right-censored data, if a censored observation has the same value as an uncensored one, the uncensored observation should be placed first. For left-censored data, if a censored observation has the same value as an uncensored one, the censored observation should be placed first.

Note that in this case the quantity $x_{(i)}$ does not necessarily represent the i ’th “largest” observation from the (unknown) complete sample.

Finally, let Ω (omega) denote the set of n subscripts in the “ordered” sample that correspond to uncensored observations.

Product-Limit Method of Kaplan and Meier (prob.method="kaplan-meier")

For complete data sets (no censored observations), the **empirical probabilities** estimator of the cumulative distribution function evaluated at the i ’th ordered observation is given by (D’Agostino, 1986a, p.8):

$$\hat{F}[x_{(i)}] = \hat{p}_i = \frac{\#[x_j \leq x_{(i)}]}{n} \quad (3)$$

where $\#[x_j \leq x_{(i)}]$ denotes the number of observations less than or equal to $x_{(i)}$ (see the help file for `ecdfPlot`). Kaplan and Meier (1958) extended this method of computing the empirical cdf to the case of right-censored data.

Right-Censored Data (censoring.side="right")

Let $S(t)$ denote the survival function evaluated at t , that is:

$$S(t) = 1 - F(t) = Pr(X > t) \quad (4)$$

Kaplan and Meier (1958) show that a nonparametric estimate of the survival function at the i ’th ordered observation that is not censored (i.e., $i \in \Omega$), is given by:

$$\begin{aligned} \hat{S}[x_{(i)}] &= \widehat{Pr}[X > x_{(i)}] \\ &= \widehat{Pr}[X > x_{(1)}] \\ &\quad \widehat{Pr}[X > x_{(2)} | X > x_{(1)}] \cdots \\ &\quad \widehat{Pr}[X > x_{(i)} | X > x_{(i-1)}] \\ &= \prod_{j \in \Omega, j \leq i} \frac{n_j - d_j}{n_j}, \quad i \in \Omega \end{aligned} \quad (5)$$

where n_j is the number of observations (uncensored or censored) with values greater than or equal to $x_{(j)}$, and d_j denotes the number of uncensored observations exactly equal to $x_{(j)}$ (if there are no tied uncensored observations then d_j will equal 1 for all values of j). (See also Lee and Wang, 2003, pp. 64–69; Michael and Schucany, 1986). By convention, the estimate of the survival function at a censored observation is set equal to the estimated value of the survival function at the largest uncensored observation less than or equal to that censoring level. If there are no uncensored observations less than or equal to a particular censoring level, the estimate of the survival function is set to 1 for that censoring level.

Thus the Kaplan-Meier plotting position at the i ’th ordered observation that is not censored (i.e., $i \in \Omega$), is given by:

$$\hat{p}_i = \hat{F}[x_{(i)}] = 1 - \prod_{j \in \Omega, j \leq i} \frac{n_j - d_j}{n_j} \quad (6)$$

The plotting position for a censored observation is set equal to the plotting position associated with the largest uncensored observation less than or equal to that censoring level. If there are no

uncensored observations less than or equal to a particular censoring level, the plotting position is set to 0 for that censoring level.

As an example, consider the following right-censored data set:

$$3, \geq 4, \geq 4, 5, 5, 6$$

The table below shows how the plotting positions are computed.

i	$x_{(i)}$	n_i	d_i	$\frac{n_i - d_i}{n_i}$	Plotting Position
1	3	6	1	5/6	$1 - (5/6) = 0.167$
2	≥ 4				
3	≥ 4				
4	5	3	2	1/3	$1 - (5/6)(1/3) = 0.722$
5	5				0.722
6	6	1	1	0/1	$1 - (5/6)(1/3)(0/1) = 1$

Note that for complete data sets, Equation (6) reduces to Equation (3).

Left-Censored Data (censoring.side="left")

Gillespie et al. (2010) give formulas for the Kaplan-Meier estimator for the case of left-censoring (censoring.side="left"). In this case, the plotting position for the i 'th ordered observation, assuming it is not censored, is computed as:

$$\hat{p}_i = \hat{F}[x_{(i)}] = \prod_{j \in \Omega, j > i} \frac{n_j - d_j}{n_j} \quad (7)$$

where n_j is the number of observations (uncensored or censored) with values less than or equal to $x_{(j)}$, and d_j denotes the number of uncensored observations exactly equal to $x_{(j)}$ (if there are no tied uncensored observations then d_j will equal 1 for all values of j). The plotting position is equal to 1 for the largest uncensored order statistic.

As an example, consider the following left-censored data set:

$$3, < 4, < 4, 5, 5, 6$$

The table below shows how the plotting positions are computed.

i	$x_{(i)}$	n_i	d_i	$\frac{n_i - d_i}{n_i}$	Plotting Position
1	3	1	1	0/1	$1(5/6)(3/5) = 0.5$
2	< 4				
3	< 4				
4	5	5	2	3/5	0.833
5	5				$1(5/6) = 0.833$
6	6	6	1	5/6	1

Note that for complete data sets, Equation (7) reduces to Equation (3).

Hazard Plotting Method of Nelson (prob.method="nelson")

(Right-Censored Data Only.) For right-censored data, Equation (5) can be re-written as:

$$\hat{S}[x_{(i)}] = \prod_{j \in \Omega, j \leq i} \frac{N - j}{N - j + 1}, \quad i \in \Omega \quad (8)$$

Nelson (1972) proposed the following formula for plotting positions for the uncensored observations in the context of estimating the hazard function (see Michael and Schucany, 1986, p.469):

$$\hat{p}_i = \hat{F}[x_{(i)}] = 1 - \prod_{j \in \Omega, j \leq i} \exp\left(\frac{-1}{N-j+1}\right) \quad (9)$$

See Lee and Wang (2003) for more information about the hazard function.

As for the Kaplan and Meier (1958) method, the plotting position for a censored observation is set equal to the plotting position associated with the largest uncensored observation less than or equal to that censoring level. If there are no uncensored observations less than or equal to a particular censoring level, the plotting position is set to 0 for that censoring level.

Generalization of Product-Limit Method, Michael and Schucany (prob.method="michael-schucany")

For complete data sets, the disadvantage of using Equation (3) above to define plotting positions is that it implies the largest observed value is the maximum possible value of the distribution (the 100'th percentile). This may be satisfactory if the underlying distribution is known to be discrete, but it is usually not satisfactory if the underlying distribution is known to be continuous.

A more frequently used formula for plotting positions for complete data sets is given by:

$$\hat{F}[x_{(i)}] = \hat{p}_i = \frac{i-a}{N-2a+1} \quad (10)$$

where $0 \leq a \leq 1$ (Cleveland, 1993, p. 18; D'Agostino, 1986a, pp. 8,25). The value of a is usually chosen so that the plotting positions are approximately unbiased (i.e., approximate the mean of their distribution) or else approximate the median value of their distribution (see the help file for [ecdfPlot](#)). Michael and Schucany (1986) extended this method for both left- and right-censored data sets.

Right-Censored Data (censoring.side="right")

For right-censored data sets, the plotting positions for the uncensored observations are computed as:

$$\hat{p}_i = 1 - \frac{N-a+1}{N-2a+1} \prod_{j \in \Omega, j \leq i} \frac{N-j-a+1}{N-j-a+2} \quad i \in \Omega \quad (11)$$

Note that the plotting positions proposed by Herd (1960) and Johnson (1964) are a special case of Equation (11) with $a = 0$. Equation (11) reduces to Equation (10) in the case of complete data sets. Note that unlike the Kaplan-Meier method, plotting positions associated with tied uncensored observations are not the same (just as in the case for complete data using Equation (10)).

As for the Kaplan and Meier (1958) method, for right-censored data the plotting position for a censored observation is set equal to the plotting position associated with the largest uncensored observation less than or equal to that censoring level. If there are no uncensored observations less than or equal to a particular censoring level, the plotting position is set to 0 for that censoring level.

Left-Censored Data (censoring.side="left")

For left-censored data sets the plotting positions are computed as:

$$\hat{p}_i = \frac{N-a+1}{N-2a+1} \prod_{j \in \Omega, j \geq i} \frac{j-a}{j-a+1} \quad i \in \Omega \quad (12)$$

Equation (12) reduces to Equation (10) in the case of complete data sets. Note that unlike the Kaplan-Meier method, plotting positions associated with tied uncensored observations are not the same (just as in the case for complete data using Equation (10)).

For left-censored data, the plotting position for a censored observation is set equal to the plotting position associated with the smallest uncensored observation greater than or equal to that censoring level. If there are no uncensored observations greater than or equal to a particular censoring level, the plotting position is set to 1 for that censoring level.

Generalization of Product-Limit Method, Hirsch and Stedinger (prob.method="hirsch-stedinger")

Hirsch and Stedinger (1987) use a slightly different approach than Kaplan and Meier (1958) and Michael and Schucany (1986) to derive a nonparametric estimate of the survival function (probability of exceedance) in the context of left-censored data. First they estimate the value of the survival function at each of the censoring levels. The value of the survival function for an uncensored observation between two adjacent censoring levels is then computed by linear interpolation (in the form of a plotting position). See also Helsel and Cohn (1988).

The discussion below presents an extension of the method of Hirsch and Stedinger (1987) to the case of right-censored data, and then presents the original derivation due to Hirsch and Stedinger (1987) for left-censored data.

Right-Censored Data (censoring.side="right")

For right-censored data, the survival function is estimated as follows. For the j 'th censoring level ($j = 0, 1, \dots, K$), write the value of the survival function as:

$$\begin{aligned} S(T_j) &= Pr[X > T_j] \\ &= Pr[X > T_{j+1}] + Pr[T_j < X \leq T_{j+1}] \\ &= S(T_{j+1}) + Pr[T_j < X \leq T_{j+1} | X > T_j] Pr[X > T_j] \\ &= S(T_{j+1}) + Pr[T_j < X \leq T_{j+1} | X > T_j] S(T_j) \quad (13) \end{aligned}$$

where

$$T_0 = -\infty, \quad (14)$$

$$T_{K+1} = \infty \quad (15)$$

Now set

$$A_j = \# \text{ uncensored observations in } (T_j, T_{j+1}] \quad (16)$$

$$B_j = \# \text{ observations in } (T_{j+1}, \infty) \quad (17)$$

for $j = 0, 1, \dots, K$. Then the method of moments estimator of the conditional probability in Equation (13) is given by:

$$\widehat{Pr}[T_j < X \leq T_{j+1} | X > T_j] = \frac{A_j}{A_j + B_j} \quad (18)$$

Hence, by equations (13) and (18) we have

$$\hat{S}(T_j) = \hat{S}(T_{j+1}) + \left(\frac{A_j}{A_j + B_j}\right) \hat{S}(T_j) \quad (19)$$

which can be rewritten as:

$$\hat{S}(T_{j+1}) = \hat{S}(T_j) \left[1 - \left(\frac{A_j}{A_j + B_j}\right)\right] \quad (20)$$

Equation (20) can be solved iteratively for $j = 1, 2, \dots, K$. Note that

$$\hat{S}(T_0) = \hat{S}(-\infty) = S(-\infty) = 1 \quad (21)$$

$$\hat{S}(T_{K+1}) = \hat{S}(\infty) = S(\infty) = 0 \quad (22)$$

Once the values of the survival function at the censoring levels are computed, the plotting positions for the A_j uncensored observations in the interval $(T_j, T_{j+1}]$ ($j = 0, 1, \dots, K$) are computed as

$$\hat{p}_i = [1 - \hat{S}(T_j)] + [\hat{S}(T_j) - \hat{S}(T_{j+1})] \frac{r - a}{A_j - 2a + 1} \quad (23)$$

where a denotes the plotting position constant, $0 \leq a \leq 1$, and r denotes the rank of the i 'th observation among the A_j uncensored observations in the interval $(T_j, T_{j+1}]$. (Tied observations are given distinct ranks.)

For the c_j observations censored at censoring level T_j ($j = 1, 2, \dots, K$), the plotting positions are computed as:

$$\hat{p}_i = 1 - [\hat{S}(T_j) \frac{r - a}{c_j - 2a + 1}] \quad (24)$$

where r denotes the rank of the i 'th observation among the c_j observations censored at censoring level T_j . Note that all the observations censored at the same censoring level are given distinct ranks, even though there is no way to distinguish between them.

Left-Censored Data (censoring.side="left")

For left-censored data, Hirsch and Stedinger (1987) modify the definition of the survival function as follows:

$$S^*(t) = Pr[X \geq t] \quad (25)$$

For continuous distributions, the functions in Equations (4) and (25) are identical.

Hirsch and Stedinger (1987) show that for the j 'th censoring level ($j = 0, 1, \dots, K$), the value of the survival function can be written as:

$$\begin{aligned} S(T_j) &= Pr[X \geq T_j] \\ &= Pr[X \geq T_{j+1}] + Pr[T_j \leq X < T_{j+1}] \\ &= S^*(T_{j+1}) + Pr[T_j \leq X < T_{j+1} | X < T_{j+1}] Pr[X < T_{j+1}] \\ &= S^*(T_{j+1}) + Pr[T_j \leq X < T_{j+1} | X < T_j] [1 - S^*(T_j)] \end{aligned} \quad (26)$$

where T_0 and T_{K+1} are defined in Equations (14) and (15).

Now set

$$A_j = \# \text{ uncensored observations in } [T_j, T_{j+1}) \quad (27)$$

$$B_j = \# \text{ observations in } (-\infty, T_j) \quad (28)$$

for $j = 0, 1, \dots, K$. Then the method of moments estimator of the conditional probability in Equation (26) is given by:

$$Pr[T_j \leq X < T_{j+1} | X < T_{j+1}] = \frac{A_j}{A_j + B_j} \quad (29)$$

Hence, by Equations (26) and (29) we have

$$\hat{S}(T_j) = \hat{S}(T_{j+1}) + \left(\frac{A_j}{A_j + B_j} \right) \hat{S}(T_j) \quad (30)$$

which can be solved iteratively for $j = 1, 2, \dots, K$. Note that

$$\widehat{S}^*(T_{K+1}) = \widehat{S}^*(\infty) = S^*(\infty) = 0 \quad (31)$$

$$\widehat{S}^*(T_0) = \widehat{S}^*(-\infty) = S^*(-\infty) = 1 \quad (32)$$

Once the values of the survival function at the censoring levels are computed, the plotting positions for the A_j uncensored observations in the interval $[T_j, T_{j+1})$ ($j = 0, 1, \dots, K$) are computed as

$$\hat{p}_i = [1 - \widehat{S}^*(T_j)] + [\widehat{S}^*(T_j) - \widehat{S}^*(T_{j+1})] \frac{r - a}{A_j - 2a + 1} \quad (33)$$

where a denotes the plotting position constant, $0 \leq a \leq 0.5$, and r denotes the rank of the i 'th observation among the A_j uncensored observations in the interval $[T_j, T_{j+1})$. (Tied observations are given distinct ranks.)

For the c_j observations censored at censoring level T_j ($j = 1, 2, \dots, K$), the plotting positions are computed as:

$$\hat{p}_i = [1 - \widehat{S}^*(T_j)] \frac{r - a}{c_j - 2a + 1} \quad (34)$$

where r denotes the rank of the i 'th observation among the c_j observations censored at censoring level T_j . Note that all the observations censored at the same censoring level are given distinct ranks, even though there is no way to distinguish between them.

Value

ppointsCensored returns a list with the following components:

Order.Statistics	numeric vector of the “ordered” observations.
Cumulative.Probabilities	numeric vector of the associated plotting positions.
Censored	logical vector indicating which of the ordered observations are censored.
Censoring.Side	character string indicating whether the data are left- or right-censored. This is same value as the argument censoring.side.
Prob.Method	character string indicating what method was used to compute the plotting positions. This is the same value as the argument prob.method.
Optional Component (only present when prob.method="michael-schucany" or prob.method="hirsch-stedinger"):	
Plot.Pos.Con	numeric scalar containing the value of the plotting position constant that was used. This is the same as the argument plot.pos.con.

Note

For censored data sets, plotting positions may be used to construct empirical cumulative distribution plots (see [ecdfPlotCensored](#)), construct quantile-quantile plots (see [qqPlotCensored](#)), or to estimate distribution parameters (see [FcnsByCatCensoredData](#)).

The function [survfit](#) in the built-in R library **survival** computes the survival function for right-censored, left-censored, or interval-censored data. Calling [survfit](#) with type="kaplan-meier" will produce similar results to calling ppointsCensored with prob.method="kaplan-meier". Also, calling [survfit](#) with type="fh2" will produce similar results to calling ppointsCensored with prob.method="nelson".

Helsel and Cohn (1988, p.2001) found very little effect of changing the value of the plotting position constant when using the method of Hirsch and Stedinger (1987) to compute plotting positions for multiply left-censored data. In general, there will be very little difference between plotting positions computed by the different methods except in the case of very small samples and a large amount of censoring.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Chambers, J.M., W.S. Cleveland, B. Kleiner, and P.A. Tukey. (1983). *Graphical Methods for Data Analysis*. Duxbury Press, Boston, MA, pp.11-16.
- Cleveland, W.S. (1993). *Visualizing Data*. Hobart Press, Summit, New Jersey, 360pp.
- D'Agostino, R.B. (1986a). Graphical Analysis. In: D'Agostino, R.B., and M.A. Stephens, eds. *Goodness-of Fit Techniques*. Marcel Dekker, New York, Chapter 2, pp.7-62.
- Gillespie, B.W., Q. Chen, H. Reichert, A. Franzblau, E. Hedgeman, J. Lepkowski, P. Adriaens, A. Demond, W. Luksemburg, and D.H. Garabrant. (2010). Estimating Population Distributions When Some Data Are Below a Limit of Detection by Using a Reverse Kaplan-Meier Estimator. *Epidemiology* **21**(4), S64–S70.
- Helsel, D.R. (2012). *Statistics for Censored Environmental Data Using Minitab and R, Second Edition*. John Wiley & Sons, Hoboken, New Jersey.
- Helsel, D.R., and T.A. Cohn. (1988). Estimation of Descriptive Statistics for Multiply Censored Water Quality Data. *Water Resources Research* **24**(12), 1997-2004.
- Hirsch, R.M., and J.R. Stedinger. (1987). Plotting Positions for Historical Floods and Their Precision. *Water Resources Research* **23**(4), 715-727.
- Kaplan, E.L., and P. Meier. (1958). Nonparametric Estimation From Incomplete Observations. *Journal of the American Statistical Association* **53**, 457-481.
- Lee, E.T., and J. Wang. (2003). *Statistical Methods for Survival Data Analysis, Third Edition*. John Wiley and Sons, New York.
- Michael, J.R., and W.R. Schucany. (1986). Analysis of Data from Censored Samples. In D'Agostino, R.B., and M.A. Stephens, eds. *Goodness-of Fit Techniques*. Marcel Dekker, New York, 560pp, Chapter 11, 461-496.
- Nelson, W. (1972). Theory and Applications of Hazard Plotting for Censored Failure Data. *Technometrics* **14**, 945-966.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. Chapter 15.
- USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

See Also

[ppoints](#), [ecdfPlot](#), [qqPlot](#), [ecdfPlotCensored](#), [qqPlotCensored](#), [survfit](#).

Examples

```
# Generate 20 observations from a normal distribution with mean=20 and sd=5,
# censor all observations less than 18, then compute plotting positions for
# this data set. Compare the plotting positions to the plotting positions
# for the uncensored data set. Note that the plotting positions for the
# censored data set start at the first ordered uncensored observation and
# that for values of x > 18 the plotting positions for the two data sets are
```

```

# exactly the same. This is because there is only one censoring level and
# no uncensored observations fall below the censored observations.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(333)
x <- rnorm(20, mean=20, sd=5)
censored <- x < 18
censored
# [1] FALSE FALSE TRUE FALSE TRUE FALSE FALSE FALSE TRUE TRUE TRUE
#[13] FALSE TRUE TRUE FALSE FALSE FALSE FALSE FALSE

sum(censored)
#[1] 7

new.x <- x
new.x[censored] <- 18
round(sort(new.x),1)
# [1] 18.0 18.0 18.0 18.0 18.0 18.0 18.0 18.1 18.7 19.6 20.2 20.3 20.6 21.4
#[15] 21.8 21.8 23.2 26.2 26.8 29.7

p.list <- ppointsCensored(new.x, censored)
p.list
#$Order.Statistics
# [1] 18.00000 18.00000 18.00000 18.00000 18.00000 18.00000 18.00000 18.09771
# [9] 18.65418 19.58594 20.21931 20.26851 20.55296 21.38869 21.76359 21.82364
#[17] 23.16804 26.16527 26.84336 29.67340
#
#$Cumulative.Probabilities
# [1] 0.3765432 0.3765432 0.3765432 0.3765432 0.3765432 0.3765432 0.3765432
# [8] 0.3765432 0.4259259 0.4753086 0.5246914 0.5740741 0.6234568 0.6728395
#[15] 0.7222222 0.7716049 0.8209877 0.8703704 0.9197531 0.9691358
#
#$Censored
# [1] TRUE TRUE TRUE TRUE TRUE TRUE TRUE TRUE FALSE FALSE FALSE FALSE
#[13] FALSE FALSE FALSE FALSE FALSE FALSE FALSE FALSE
#
#$Censoring.Side
#[1] "left"
#
#$Prob.Method
#[1] "michael-schucany"
#
#$Plot.Pos.Con
#[1] 0.375

#-----

# Round off plotting positions to two decimal places
# and compare to plotting positions that ignore censoring
#-----

round(p.list$Cum, 2)
# [1] 0.38 0.38 0.38 0.38 0.38 0.38 0.38 0.38 0.43 0.48 0.52 0.57 0.62 0.67
#[15] 0.72 0.77 0.82 0.87 0.92 0.97

round(ppoints(x, a=0.375), 2)
# [1] 0.03 0.08 0.13 0.18 0.23 0.28 0.33 0.38 0.43 0.48 0.52 0.57 0.62 0.67

```



```

#[15] 0.72 0.77 0.82 0.87 0.92 0.97

#-----

# Clean up
#-----
rm(x, censored, new.x, p.list)

#-----

# Reproduce the example in Appendix B of Helsel and Cohn (1988). The data
# are stored in Helsel.Cohn.88.appb.df. This data frame contains 18
# observations, of which 9 are censored below one of 2 distinct censoring
# levels.

Helsel.Cohn.88.app.b.df
#   Conc.orig Conc Censored
#1      <1     1      TRUE
#2      <1     1      TRUE
#...
#17      33    33     FALSE
#18      50    50     FALSE

p.list <- with(Helsel.Cohn.88.app.b.df,
  ppointsCensored(Conc, Censored, prob.method="hirsch-stedinger", plot.pos.con=0))
lapply(p.list[1:2], round, 3)
#$Order.Statistics
# [1] 1 1 1 1 1 1 3 7 9 10 10 10 12 15 20 27 33 50
#
#$Cumulative.Probabilities
# [1] 0.063 0.127 0.190 0.254 0.317 0.381 0.500 0.556 0.611 0.167 0.333 0.500
#[13] 0.714 0.762 0.810 0.857 0.905 0.952

# Clean up
#-----
rm(p.list)

#-----

# Example 15-1 of USEPA (2009, page 15-10) gives an example of
# computing plotting positions based on censored manganese
# concentrations (ppb) in groundwater collected at 5 monitoring
# wells. The data for this example are stored in
# EPA.09.Ex.15.1.manganese.df.

EPA.09.Ex.15.1.manganese.df
#   Sample  Well Manganese.Orig.ppb Manganese.ppb Censored
#1      1 Well.1      <5           5.0      TRUE
#2      2 Well.1     12.1          12.1     FALSE
#3      3 Well.1     16.9          16.9     FALSE
#4      4 Well.1     21.6          21.6     FALSE
#5      5 Well.1      <2           2.0      TRUE
#...
#21     1 Well.5     17.9          17.9     FALSE
#22     2 Well.5     22.7          22.7     FALSE
#23     3 Well.5      3.3           3.3     FALSE
#24     4 Well.5      8.4           8.4     FALSE

```

```

#25      5 Well.5      <2      2.0      TRUE

p.list.EPA <- with(EPA.09.Ex.15.1.manganese.df,
  ppointsCensored(Manganese.ppb, Censored,
    prob.method = "kaplan-meier"))
data.frame(Mn = p.list.EPA$Order.Statistics, Censored = p.list.EPA$Censored,
  CDF = p.list.EPA$Cumulative.Probabilities)
#      Mn Censored  CDF
#1      2.0      TRUE 0.21
#2      2.0      TRUE 0.21
#3      2.0      TRUE 0.21
#4      3.3     FALSE 0.28
#5      5.0      TRUE 0.28
#6      5.0      TRUE 0.28
#7      5.0      TRUE 0.28
#8      5.3     FALSE 0.32
#9      6.3     FALSE 0.36
#10     7.7     FALSE 0.40
#11     8.4     FALSE 0.44
#12     9.5     FALSE 0.48
#13    10.0     FALSE 0.52
#14    11.9     FALSE 0.56
#15    12.1     FALSE 0.60
#16    12.6     FALSE 0.64
#17    16.9     FALSE 0.68
#18    17.9     FALSE 0.72
#19    21.6     FALSE 0.76
#20    22.7     FALSE 0.80
#21    34.5     FALSE 0.84
#22    45.9     FALSE 0.88
#23    53.6     FALSE 0.92
#24    77.2     FALSE 0.96
#25   106.3     FALSE 1.00

#-----

# Clean up
#-----
rm(p.list.EPA)

```

predict.lm

Predict method for Linear Model Fits

Description

Predicted values based on linear model object. This function is a modified version of the built-in R version of `predict.lm`. The **EnvStats** function `predict.lm` returns a component called `n.coefs` when the argument `se.fit=TRUE`. The component `n.coefs` is used by the function [pointwise](#) to create simultaneous confidence or prediction limits.

Usage

```

## S3 method for class 'lm'
predict(object, newdata, se.fit = FALSE, scale = NULL, df = Inf,

```

```
interval = c("none", "confidence", "prediction"),
level = 0.95, type = c("response", "terms"),
terms = NULL, na.action = na.pass,
pred.var = res.var/weights, weights = 1, ...)
```

Arguments

<code>object</code>	Object of class "lm" or a class that inherits from "lm"
<code>newdata</code>	An optional data frame in which to look for variables with which to predict. If omitted, the fitted values are used.
<code>se.fit</code>	A logical scalar indicating whether to compute standard errors. The default value is <code>se.fit=FALSE</code> .
<code>scale</code>	Scale parameter for <code>std.err.</code> calculation.
<code>df</code>	Degrees of freedom for scale.
<code>interval</code>	Type of interval calculation. Possible values are "none" (the default), "confidence", and "prediction".
<code>level</code>	Tolerance/confidence level.
<code>type</code>	Type of prediction (response or model term).
<code>terms</code>	If <code>type="terms"</code> , this argument determines which terms (the default is all terms).
<code>na.action</code>	A function determining what should be done with missing values in <code>newdata</code> . The default is to predict NA.
<code>pred.var</code>	The variance(s) for future observations to be assumed for prediction intervals. See 'Details'.
<code>weights</code>	The variance weights for prediction. This can be a numeric vector or a one-sided model formula. In the latter case, it is interpreted as an expression evaluated in <code>newdata</code> .
<code>...</code>	Further arguments passed to or from other methods.

Details

`predict.lm` produces predicted values, obtained by evaluating the regression function in the frame `newdata` (which defaults to `model.frame(object)`). If the logical `se.fit` is TRUE, standard errors of the predictions are calculated. If the numeric argument `scale` is set (with optional `df`), it is used as the residual standard deviation in the computation of the standard errors, otherwise this is extracted from the model fit. Setting `intervals` specifies computation of confidence or prediction (tolerance) intervals at the specified `level`, sometimes referred to as narrow vs. wide intervals.

If the fit is rank-deficient, some of the columns of the design matrix will have been dropped. Prediction from such a fit only makes sense if `newdata` is contained in the same subspace as the original data. That cannot be checked accurately, so a warning is issued.

If `newdata` is omitted the predictions are based on the data used for the fit. In that case how cases with missing values in the original fit are handled is determined by the `na.action` argument of that fit. If `na.action = na.omit`, omitted cases will not appear in the residuals, whereas if `na.action = na.exclude` they will appear (in predictions, standard errors or interval limits), with residual value NA. See also [napredict](#).

The prediction intervals are for a single observation at each case in `newdata` (or by default, the data used for the fit) with error variance(s) `pred.var`. This can be a multiple of `res.var`, the estimated value of σ^2 : the default is to assume that future observations have the same error variance as those used for fitting. If `weights` is supplied, the inverse of this is used as a scale factor. For a weighted

fit, if the prediction is for the original data frame, `weights` defaults to the weights used for the model fit, with a warning since it might not be the intended result. If the fit was weighted and `newdata` is given, the default is to assume constant prediction variance, with a warning.

Value

`predict.lm` produces a vector of predictions or a matrix of predictions and bounds with column names `fit`, `lwr`, and `upr` if `interval` is set. If `se.fit` is `TRUE`, a list with the following components is returned:

<code>fit</code>	vector or matrix as above
<code>se.fit</code>	standard error of predicted means
<code>residual.scale</code>	residual standard deviations
<code>df</code>	degrees of freedom for residual
<code>n.coefs</code>	numeric scalar denoting the number of predictor variables used in the model

Note

Variables are first looked for in `newdata` and then searched for in the usual way (which will include the environment of the formula used in the fit). A warning will be given if the variables found are not of the same length as those in `newdata` if it was supplied.

Notice that prediction variances and prediction intervals always refer to *future* observations, possibly corresponding to the same predictors as used for the fit. The variance of the *residuals* will be smaller.

Strictly speaking, the formula used for prediction limits assumes that the degrees of freedom for the fit are the same as those for the residual variance. This may not be the case if `res.var` is not obtained from the fit.

Almost always the process of determining the concentration of a chemical in a soil, water, or air sample involves using some kind of machine that produces a signal, and this signal is related to the concentration of the chemical in the physical sample. The process of relating the machine signal to the concentration of the chemical is called *calibration* (see [calibrate](#)). Once calibration has been performed, estimated concentrations in physical samples with unknown concentrations are computed using inverse regression. The uncertainty in the process used to estimate the concentration may be quantified with decision, detection, and quantitation limits.

In practice, only the point estimate of concentration is reported (along with a possible qualifier), without confidence bounds for the true concentration C . This is most unfortunate because it gives the impression that there is no error associated with the reported concentration. Indeed, both the International Organization for Standardization (ISO) and the International Union of Pure and Applied Chemistry (IUPAC) recommend always reporting both the estimated concentration and the uncertainty associated with this estimate (Currie, 1997).

Author(s)

R Development Core Team (for code for R version of `predict.lm`).

Steven P. Millard (for modification to add component `n.coefs`; <EnvStats@ProbStatInfo.com>)

References

Chambers, J.M., and Hastie, T.J., eds. (1992). *Statistical Models in S*. Chapman and Hall/CRC, Boca Raton, FL.

Draper, N., and H. Smith. (1998). *Applied Regression Analysis*. Third Edition. John Wiley and Sons, New York, Chapter 3.

Millard, S.P., and N.K. Neerchal. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, FL, pp.546-553.

Miller, R.G. (1981a). *Simultaneous Statistical Inference*. Springer-Verlag, New York, pp.111, 124.

See Also

[predict.lm](#), [calibrate](#), [calibrate](#), [inversePredictCalibrate](#), [detectionLimitCalibrate](#).

[SafePrediction](#) for prediction from polynomial and spline fits.

Examples

```
# Using the data from the built-in data frame Air.df,
# fit the cube-root of ozone as a function of temperature,
# then compute predicted values for ozone at 70 and 90 degrees F,
# along with the standard errors of these predicted values.

# First look at the data
#-----
attach(Air.df)

plot(temperature, ozone, xlab = "Temperature (degrees F)",
     ylab = "Cube-Root Ozone (ppb)")

# Now create the lm object
#-----
ozone.fit <- lm(ozone ~ temperature, data = Air.df)

# Now get predicted values and CIs at 70 and 90 degrees.
# Note the presence of the last component called n.coefs.
#-----
predict.list <- predict(ozone.fit,
  newdata = data.frame(temperature = c(70, 90)), se.fit = TRUE)

predict.list
#$fit
#      1      2
#2.697810 4.101808
#
#$se.fit
#      1      2
#0.07134554 0.08921071
#
#$df
#[1] 114
#
#$residual.scale
#[1] 0.5903046
#
#$n.coefs
#[1] 2
```

```

#-----

#Continuing with the above example, create a scatterplot of
# cube-root ozone vs. temperature, and add the fitted line
# along with simultaneous 95% confidence bands.

plot(temperature, ozone, xlab = "Temperature (degrees F)",
     ylab = "Cube-Root Ozone (ppb)")

abline(ozone.fit, lwd = 3, col = "blue")

new.temp <- seq(min(temperature), max(temperature), length = 100)

predict.list <- predict(ozone.fit,
  newdata = data.frame(temperature = new.temp),
  se.fit = TRUE)

ci.ozone <- pointwise(predict.list, coverage = 0.95,
  simultaneous = TRUE)

lines(new.temp, ci.ozone$lower, lty = 2, lwd = 3, col = "magenta")
lines(new.temp, ci.ozone$upper, lty = 2, lwd = 3, col = "magenta")

title(main=paste("Scatterplot of Cube-Root Ozone vs. Temperature",
  "with Fitted Line and Simultaneous 95% Confidence Bands",
  sep="\n"))

#-----

# Clean up

rm(ozone.fit, predict.list, new.temp, ci.ozone)
detach("Air.df")

#-----

# Examples from the R help file for predict.lm:

require(graphics)

## Predictions
x <- rnorm(15)
y <- x + rnorm(15)
predict(lm(y ~ x))
new <- data.frame(x = seq(-3, 3, 0.5))
predict(lm(y ~ x), new, se.fit = TRUE)
pred.w.plim <- predict(lm(y ~ x), new, interval="prediction")
pred.w.clim <- predict(lm(y ~ x), new, interval="confidence")
matplot(new$x, cbind(pred.w.clim, pred.w.plim[, -1]),
        lty=c(1,2,2,3,3), type="l", ylab="predicted y")

## Prediction intervals, special cases
## The first three of these throw warnings

```

```

w <- 1 + x^2
fit <- lm(y ~ x)
wfit <- lm(y ~ x, weights = w)
predict(fit, interval = "prediction")
predict(wfit, interval = "prediction")
predict(wfit, new, interval = "prediction")
predict(wfit, new, interval = "prediction", weights = (new$x)^2)
predict(wfit, new, interval = "prediction", weights = ~x^2)

```

predIntGamma

Prediction Interval for Gamma Distribution

Description

Construct a prediction interval for the next k observations or next set of k transformed means for a gamma distribution.

Usage

```

predIntGamma(x, n.transmean = 1, k = 1, method = "Bonferroni",
  pi.type = "two-sided", conf.level = 0.95, est.method = "mle",
  normal.approx.transform = "kulkarni.powar")

predIntGammaAlt(x, n.transmean = 1, k = 1, method = "Bonferroni",
  pi.type = "two-sided", conf.level = 0.95, est.method = "mle",
  normal.approx.transform = "kulkarni.powar")

```

Arguments

- | | |
|-------------|--|
| x | numeric vector of non-negative observations. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed. |
| n.transmean | positive integer specifying the sample size associated with the k future transformed means (see the DETAILS section for an explanation of what the transformation is). The default value is <code>n.transmean=1</code> (i.e., predicting future observations). Note that all future transformed means must be based on the same sample size. |
| k | positive integer specifying the number of future observations or means the prediction interval should contain with confidence level <code>conf.level</code> . The default value is <code>k=1</code> . |
| method | character string specifying the method to use if the number of future observations or averages (k) is greater than 1. The possible values are "Bonferroni" (approximate method based on Bonferonni inequality; the default), and "exact" (exact method due to Dunnett, 1955). See the DETAILS section for more information. This argument is ignored if <code>k=1</code> . |
| pi.type | character string indicating what kind of prediction interval to compute. The possible values are "two-sided" (the default), "lower", and "upper". |
| conf.level | a scalar between 0 and 1 indicating the confidence level associated with the prediction interval. The default value is <code>conf.level=0.95</code> . |

`est.method` character string specifying the method of estimation for the shape and scale distribution parameters. The possible values are "mle" (maximum likelihood; the default), "bcmle" (bias-corrected mle), "mme" (method of moments), and "mmue" (method of moments based on the unbiased estimator of variance). See the DETAILS section of the help file for [egamma](#) for more information.

`normal.approx.transform` character string indicating which power transformation to use. Possible values are "kulkarni.powar" (the default), "cube.root", and "fourth.root". See the DETAILS section for more information.

Details

If `x` contains any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

The function `predIntGamma` returns a prediction interval as well as estimates of the shape and scale parameters. The function `predIntGammaAlt` returns a prediction interval as well as estimates of the mean and coefficient of variation.

Following Krishnamoorthy et al. (2008), the prediction interval is computed by:

1. using a power transformation on the original data to induce approximate normality,
2. calling [predIntNorm](#) with the transformed data to compute the prediction interval, and then
3. back-transforming the interval to create a prediction interval on the original scale.

The argument `normal.approx.transform` determines which transformation is used. The value `normal.approx.transform="cube.root"` uses the cube root transformation suggested by Wil-son and Hilferty (1931) and used by Krishnamoorthy et al. (2008) and Singh et al. (2010b), and the value `normal.approx.transform="fourth.root"` uses the fourth root transformation suggested by Hawkins and Wixley (1986) and used by Singh et al. (2010b). The default value `normal.approx.transform="kulkarni.powar"` uses the "Optimum Power Normal Approximation Method" of Kulkarni and Powar (2010). The "optimum" power p is determined by:

$$\begin{aligned} p &= -0.0705 - 0.178 \text{ shape} + 0.475 \sqrt{\text{shape}} && \text{if } \text{shape} \leq 1.5 \\ p &= 0.246 && \text{if } \text{shape} > 1.5 \end{aligned}$$

where *shape* denotes the estimate of the shape parameter. Although Kulkarni and Powar (2010) use the maximum likelihood estimate of shape to determine the power p , for the functions `predIntGamma` and `predIntGammaAlt` the power p is based on whatever estimate of shape is used (e.g., `est.method="mle"`, `est.method="bcmle"`, etc.).

When the argument `n.transmean` is larger than 1 (i.e., you are constructing a prediction interval for future means, not just single observations), in order to properly compare a future mean with the prediction limits, you must follow these steps:

1. Take the observations that will be used to compute the mean and transform them by raising them to the power given by the value in the component `interval$normal.transform.power` (see the section VALUE below).
2. Compute the mean of the transformed observations.
3. Take the mean computed in step 2 above and raise it to the inverse of the power originally used to transform the observations.

Value

A list of class "estimate" containing the estimated parameters, the prediction interval, and other information. See [estimate.object](#) for details.

In addition to the usual components contained in an object of class "estimate", the returned value also includes two additional components within the "interval" component:

n.transmean	the value of n.transmean supplied in the call to predIntGamma or predIntGammaAlt.
normal.transform.power	the value of the power used to transform the original data to approximate normality.

Warning

It is possible for the lower prediction limit based on the transformed data to be less than 0. In this case, the lower prediction limit on the original scale is set to 0 and a warning is issued stating that the normal approximation is not accurate in this case.

Note

The [gamma distribution](#) takes values on the positive real line. Special cases of the gamma are the [exponential](#) distribution and the [chi-square](#) distributions. Applications of the gamma include life testing, statistical ecology, queuing theory, inventory control, and precipitation processes. A gamma distribution starts to resemble a normal distribution as the shape parameter α tends to infinity.

Some EPA guidance documents (e.g., Singh et al., 2002; Singh et al., 2010a,b) strongly recommend against using a lognormal model for environmental data and recommend trying a gamma distribution instead.

Prediction intervals have long been applied to quality control and life testing problems (Hahn, 1970b,c; Hahn and Nelson, 1973), and are often discussed in the context of linear regression (Draper and Smith, 1998; Zar, 2010). Prediction intervals are useful for analyzing data from groundwater detection monitoring programs at hazardous and solid waste facilities. References that discuss prediction intervals in the context of environmental monitoring include: Berthouex and Brown (2002, Chapter 21), Gibbons et al. (2009), Millard and Neerchal (2001, Chapter 6), Singh et al. (2010b), and USEPA (2009).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Berthouex, P.M., and L.C. Brown. (2002). *Statistics for Environmental Engineers*. Lewis Publishers, Boca Raton.
- Draper, N., and H. Smith. (1998). *Applied Regression Analysis*. Third Edition. John Wiley and Sons, New York.
- Evans, M., N. Hastings, and B. Peacock. (1993). *Statistical Distributions*. Second Edition. John Wiley and Sons, New York, Chapter 18.
- Gibbons, R.D., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*, Second Edition. John Wiley & Sons, Hoboken.
- Hawkins, D. M., and R.A.J. Wixley. (1986). A Note on the Transformation of Chi-Squared Variables to Normality. *The American Statistician*, **40**, 296–298.

- Johnson, N.L., S. Kotz, and N. Balakrishnan. (1994). *Continuous Univariate Distributions, Volume 1*. Second Edition. John Wiley and Sons, New York, Chapter 17.
- Krishnamoorthy K., T. Mathew, and S. Mukherjee. (2008). Normal-Based Methods for a Gamma Distribution: Prediction and Tolerance Intervals and Stress-Strength Reliability. *Technometrics*, **50**(1), 69–78.
- Krishnamoorthy K., and T. Mathew. (2009). *Statistical Tolerance Regions: Theory, Applications, and Computation*. John Wiley and Sons, Hoboken.
- Kulkarni, H.V., and S.K. Powar. (2010). A New Method for Interval Estimation of the Mean of the Gamma Distribution. *Lifetime Data Analysis*, **16**, 431–447.
- Millard, S.P., and N.K. Neerchal. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton.
- Singh, A., A.K. Singh, and R.J. Iaci. (2002). *Estimation of the Exposure Point Concentration Term Using a Gamma Distribution*. EPA/600/R-02/084. October 2002. Technology Support Center for Monitoring and Site Characterization, Office of Research and Development, Office of Solid Waste and Emergency Response, U.S. Environmental Protection Agency, Washington, D.C.
- Singh, A., R. Maichle, and N. Armbya. (2010a). *ProUCL Version 4.1.00 User Guide (Draft)*. EPA/600/R-07/041, May 2010. Office of Research and Development, U.S. Environmental Protection Agency, Washington, D.C.
- Singh, A., N. Armbya, and A. Singh. (2010b). *ProUCL Version 4.1.00 Technical Guide (Draft)*. EPA/600/R-07/041, May 2010. Office of Research and Development, U.S. Environmental Protection Agency, Washington, D.C.
- Wilson, E.B., and M.M. Hilferty. (1931). The Distribution of Chi-Squares. *Proceedings of the National Academy of Sciences*, **17**, 684–688.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.
- USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.
- Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, New Jersey.

See Also

[GammaDist](#), [GammaAlt](#), [estimate.object](#), [egamma](#), [predIntNorm](#), [tolIntGamma](#).

Examples

```
# Generate 20 observations from a gamma distribution with parameters
# shape=3 and scale=2, then create a prediction interval for the
# next observation.
# (Note: the call to set.seed simply allows you to reproduce this
# example.)

set.seed(250)
dat <- rgamma(20, shape = 3, scale = 2)
predIntGamma(dat)

#Results of Distribution Parameter Estimation
```



```

#                               transformation to Normality
#                               based on bcmle of 'shape'
#
#Normal Transform Power:       0.246
#
#Prediction Interval Type:     upper
#
#Confidence Level:            95%
#
#Number of Future
#Transformed Means:           3
#
#Sample Size for
#Transformed Means:           2
#
#Prediction Interval:          LPL = 0.00000
#                               UPL = 12.17404

#-----

# Continuing with the above example, assume the distribution shifts
# in the future to a gamma distribution with shape = 5 and scale = 2.
# Create 6 future observations from this distribution, and create 3
# means by pairing the observations sequentially. Note we must first
# transform these observations using the power 0.246, then compute the
# means based on the transformed data, and then transform the means
# back to the original scale and compare them to the upper prediction
# limit of 12.17

set.seed(427)
new.dat <- rgamma(6, shape = 5, scale = 2)

p <- pred.list$interval$normal.transform.power
p
#[1] 0.246

new.dat.trans <- new.dat^p
means.trans <- c(mean(new.dat.trans[1:2]), mean(new.dat.trans[3:4]),
  mean(new.dat.trans[5:6]))
means <- means.trans^(1/p)
means
#[1] 11.74214 17.05299 11.65272

any(means > pred.list$interval$limits["UPL"])
#[1] TRUE

#-----

# Clean up
rm(dat, pred.list, new.dat, p, new.dat.trans, means.trans, means)

#-----

# Reproduce part of the example on page 73 of
# Krishnamoorthy et al. (2008), which uses alkalinity concentrations
# reported in Gibbons (1994) and Gibbons et al. (2009) to construct a
# one-sided upper 90% prediction limit.

```

```

predIntGamma(Gibbons.et.al.09.Alkilinity.vec, pi.type = "upper",
  conf.level = 0.9, normal.approx.transform = "cube.root")

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Gamma
#
#Estimated Parameter(s):      shape = 9.375013
#                              scale = 6.202461
#
#Estimation Method:           mle
#
#Data:                        Gibbons.et.al.09.Alkilinity.vec
#
#Sample Size:                 27
#
#Prediction Interval Method:   exact using
#                              Wilson & Hilferty (1931) cube-root
#                              transformation to Normality
#
#Normal Transform Power:      0.3333333
#
#Prediction Interval Type:     upper
#
#Confidence Level:            90%
#
#Number of Future Observations: 1
#
#Prediction Interval:          LPL = 0.0000
#                              UPL = 85.3495

```

predIntGammaSimultaneous

Simultaneous Prediction Interval for a Gamma Distribution

Description

Estimate the shape and scale parameters for a [gamma distribution](#), or estimate the mean and coefficient of variation for a [gamma distribution \(alternative parameterization\)](#), and construct a simultaneous prediction interval for the next r sampling occasions, based on one of three possible rules: k-of-m, California, or Modified California.

Usage

```

predIntGammaSimultaneous(x, n.transmean = 1, k = 1, m = 2, r = 1,
  rule = "k.of.m", delta.over.sigma = 0, pi.type = "upper", conf.level = 0.95,
  K.tol = 1e-07, est.method = "mle", normal.approx.transform = "kulkarni.powar")

predIntGammaAltSimultaneous(x, n.transmean = 1, k = 1, m = 2, r = 1,
  rule = "k.of.m", delta.over.sigma = 0, pi.type = "upper", conf.level = 0.95,
  K.tol = 1e-07, est.method = "mle", normal.approx.transform = "kulkarni.powar")

```

Arguments

<code>x</code>	numeric vector of non-negative observations. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
<code>n.transmean</code>	positive integer specifying the sample size associated with future transformed means (see the DETAILS section for an explanation of what the transformation is). The default value is <code>n.transmean=1</code> (i.e., individual observations). Note that all future transformed means must be based on the same sample size.
<code>k</code>	for the <i>k</i> -of- <i>m</i> rule (<code>rule="k.of.m"</code>), a positive integer specifying the minimum number of observations (or transformed means) out of <i>m</i> observations (or transformed means) (all obtained on one future sampling “occasion”) the prediction interval should contain with confidence level <code>conf.level</code> . The default value is <code>k=1</code> . This argument is ignored when the argument <code>rule</code> is not equal to <code>"k.of.m"</code> .
<code>m</code>	positive integer specifying the maximum number of future observations (or transformed means) on one future sampling “occasion”. The default value is <code>m=2</code> , except when <code>rule="Modified.CA"</code> , in which case this argument is ignored and <code>m</code> is automatically set equal to 4.
<code>r</code>	positive integer specifying the number of future sampling “occasions”. The default value is <code>r=1</code> .
<code>rule</code>	character string specifying which rule to use. The possible values are <code>"k.of.m"</code> (<i>k</i> -of- <i>m</i> rule; the default), <code>"CA"</code> (California rule), and <code>"Modified.CA"</code> (modified California rule). See the DETAILS section below for more information.
<code>delta.over.sigma</code>	numeric scalar indicating the ratio Δ/σ . The quantity Δ (delta) denotes the difference between the mean of the population (on the transformed scale) that was sampled to construct the prediction interval, and the mean of the population (on the transformed scale) that will be sampled to produce the future observations. The quantity σ (sigma) denotes the population standard deviation (on the transformed scale) for both populations. See the DETAILS section below for more information. The default value is <code>delta.over.sigma=0</code> .
<code>pi.type</code>	character string indicating what kind of prediction interval to compute. The possible values are <code>pi.type="upper"</code> (the default), and <code>pi.type="lower"</code> .
<code>conf.level</code>	a scalar between 0 and 1 indicating the confidence level of the prediction interval. The default value is <code>conf.level=0.95</code> .
<code>K.tol</code>	numeric scalar indicating the tolerance to use in the nonlinear search algorithm to compute <i>K</i> . The default value is <code>K.tol=.Machine\$double.eps^(1/2)</code> . For many applications, the value of <i>K</i> needs to be known only to the second decimal place, in which case setting <code>K.tol=1e-4</code> will speed up computation a bit.
<code>est.method</code>	character string specifying the method of estimation for the shape and scale distribution parameters. The possible values are <code>"mle"</code> (maximum likelihood; the default), <code>"bcmle"</code> (bias-corrected mle), <code>"mme"</code> (method of moments), and <code>"mmue"</code> (method of moments based on the unbiased estimator of variance). See the DETAILS section of the help file for egamma for more information.
<code>normal.approx.transform</code>	character string indicating which power transformation to use. Possible values are <code>"kulkarni.powar"</code> (the default), <code>"cube.root"</code> , and <code>"fourth.root"</code> . See the DETAILS section for more information.

Details

The function `predIntGammaSimultaneous` returns a simultaneous prediction interval as well as estimates of the shape and scale parameters. The function `predIntGammaAltSimultaneous` returns a simultaneous prediction interval as well as estimates of the mean and coefficient of variation.

Following Krishnamoorthy et al. (2008), the simultaneous prediction interval is computed by:

1. using a power transformation on the original data to induce approximate normality,
2. calling `predIntNormSimultaneous` with the transformed data to compute the simultaneous prediction interval, and then
3. back-transforming the interval to create a simultaneous prediction interval on the original scale.

The argument `normal.approx.transform` determines which transformation is used. The value `normal.approx.transform="cube.root"` uses the cube root transformation suggested by Wilson and Hilferty (1931) and used by Krishnamoorthy et al. (2008) and Singh et al. (2010b), and the value `normal.approx.transform="fourth.root"` uses the fourth root transformation suggested by Hawkins and Wixley (1986) and used by Singh et al. (2010b). The default value `normal.approx.transform="kulkarni.powar"` uses the "Optimum Power Normal Approximation Method" of Kulkarni and Powar (2010). The "optimum" power p is determined by:

$$\begin{aligned} p &= -0.0705 - 0.178 \text{ shape} + 0.475 \sqrt{\text{shape}} && \text{if } \text{shape} \leq 1.5 \\ p &= 0.246 && \text{if } \text{shape} > 1.5 \end{aligned}$$

where *shape* denotes the estimate of the shape parameter. Although Kulkarni and Powar (2010) use the maximum likelihood estimate of shape to determine the power p , for the functions `predIntGammaSimultaneous` and `predIntGammaAltSimultaneous` the power p is based on whatever estimate of shape is used (e.g., `est.method="mle"`, `est.method="bcmle"`, etc.).

When the argument `n.transmean` is larger than 1 (i.e., you are constructing a prediction interval for future means, not just single observations), in order to properly compare a future mean with the prediction limits, you must follow these steps:

1. Take the observations that will be used to compute the mean and transform them by raising them to the power given by the value in the component `interval$normal.transform.power` (see the section **Value** below).
2. Compute the mean of the transformed observations.
3. Take the mean computed in step 2 above and raise it to the inverse of the power originally used to transform the observations.

Value

A list of class "estimate" containing the estimated parameters, the simultaneous prediction interval, and other information. See `estimate.object` for details.

In addition to the usual components contained in an object of class "estimate", the returned value also includes two additional components within the "interval" component:

- | | |
|-------------------------------------|---|
| <code>n.transmean</code> | the value of <code>n.transmean</code> supplied in the call to <code>predIntGammaSimultaneous</code> or <code>predIntGammaAltSimultaneous</code> . |
| <code>normal.transform.power</code> | the value of the power used to transform the original data to approximate normality. |

Warning

It is possible for the lower prediction limit based on the transformed data to be less than 0. In this case, the lower prediction limit on the original scale is set to 0 and a warning is issued stating that the normal approximation is not accurate in this case.

Note

The Gamma Distribution

The [gamma distribution](#) takes values on the positive real line. Special cases of the gamma are the [exponential](#) distribution and the [chi-square](#) distributions. Applications of the gamma include life testing, statistical ecology, queuing theory, inventory control, and precipitation processes. A gamma distribution starts to resemble a normal distribution as the shape parameter a tends to infinity.

Some EPA guidance documents (e.g., Singh et al., 2002; Singh et al., 2010a,b) strongly recommend against using a lognormal model for environmental data and recommend trying a gamma distribution instead.

Motivation

Prediction and tolerance intervals have long been applied to quality control and life testing problems (Hahn, 1970b,c; Hahn and Nelson, 1973). In the context of environmental statistics, prediction intervals are useful for analyzing data from groundwater detection monitoring programs at hazardous and solid waste facilities.

One of the main statistical problems that plague groundwater monitoring programs at hazardous and solid waste facilities is the requirement of testing several wells and several constituents at each well on each sampling occasion. This is an obvious multiple comparisons problem, and the naive approach of using a standard t-test at a conventional α -level (e.g., 0.05 or 0.01) for each test leads to a very high probability of at least one significant result on each sampling occasion, when in fact no contamination has occurred. This problem was pointed out years ago by Millard (1987) and others.

Davis and McNichols (1987) proposed simultaneous prediction intervals as a way of controlling the facility-wide false positive rate (FWFPR) while maintaining adequate power to detect contamination in the groundwater. Because of the ubiquitous presence of spatial variability, it is usually best to use simultaneous prediction intervals at each well (Davis, 1998a). That is, by constructing prediction intervals based on background (pre-landfill) data on each well, and comparing future observations at a well to the prediction interval for that particular well. In each of these cases, the individual α -level at each well is equal to the FWFPR divided by the product of the number of wells and constituents.

Often, observations at downgradient wells are not available prior to the construction and operation of the landfill. In this case, upgradient well data can be combined to create a background prediction interval, and observations at each downgradient well can be compared to this prediction interval. If spatial variability is present and a major source of variation, however, this method is not really valid (Davis, 1994; Davis, 1998a).

Chapter 19 of USEPA (2009) contains an extensive discussion of using the 1-of- m rule and the Modified California rule.

Chapters 1 and 3 of Gibbons et al. (2009) discuss simultaneous prediction intervals for the normal and lognormal distributions, respectively.

The k -of- m Rule

For the k -of- m rule, Davis and McNichols (1987) give tables with “optimal” choices of k (in terms of best power for a given overall confidence level) for selected values of m , r , and n . They found that the optimal ratios of k to m (i.e., k/m) are generally small, in the range of 15-50%.

The California Rule

The California rule was mandated in that state for groundwater monitoring at waste disposal facilities when resampling verification is part of the statistical program (Barclay's Code of California Regulations, 1991). The California code mandates a "California" rule with $m \geq 3$. The motivation for this rule may have been a desire to have a majority of the observations in bounds (Davis, 1998a). For example, for a k -of- m rule with $k = 1$ and $m = 3$, a monitoring location will pass if the first observation is out of bounds, the second resample is out of bounds, but the last resample is in bounds, so that 2 out of 3 observations are out of bounds. For the California rule with $m = 3$, either the first observation must be in bounds, or the next 2 observations must be in bounds in order for the monitoring location to pass.

Davis (1998a) states that if the FWFPR is kept constant, then the California rule offers little increased power compared to the k -of- m rule, and can actually decrease the power of detecting contamination.

The Modified California Rule

The Modified California Rule was proposed as a compromise between a 1-of- m rule and the California rule. For a given FWFPR, the Modified California rule achieves better power than the California rule, and still requires at least as many observations in bounds as out of bounds, unlike a 1-of- m rule.

Different Notations Between Different References

For the k -of- m rule described in this help file, both Davis and McNichols (1987) and USEPA (2009, Chapter 19) use the variable p instead of k to represent the minimum number of future observations the interval should contain on each of the r sampling occasions.

Gibbons et al. (2009, Chapter 1) presents extensive lists of the value of K for both k -of- m rules and California rules. Gibbons et al.'s notation reverses the meaning of k and r compared to the notation used in this help file. That is, in Gibbons et al.'s notation, k represents the number of future sampling occasions or monitoring wells, and r represents the minimum number of observations the interval should contain on each sampling occasion.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Barclay's California Code of Regulations.** (1991). Title 22, Section 66264.97 [concerning hazardous waste facilities] and Title 23, Section 2550.7(e)(8) [concerning solid waste facilities]. Barclay's Law Publishers, San Francisco, CA.
- Davis, C.B. (1998a). *Ground-Water Statistics & Regulations: Principles, Progress and Problems*. Second Edition. Environmetrics & Statistics Limited, Henderson, NV.
- Davis, C.B. (1998b). Personal Communication, September 3, 1998.
- Davis, C.B., and R.J. McNichols. (1987). One-sided Intervals for at Least p of m Observations from a Lognormal Population on Each of r Future Occasions. *Technometrics* **29**, 359–370.
- Evans, M., N. Hastings, and B. Peacock. (1993). *Statistical Distributions*. Second Edition. John Wiley and Sons, New York, Chapter 18.
- Gibbons, R.D., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*, Second Edition. John Wiley & Sons, Hoboken.
- Fertig, K.W., and N.R. Mann. (1977). One-Sided Prediction Intervals for at Least p Out of m Future Observations From a Lognormal Population. *Technometrics* **19**, 167–177.

- Hahn, G.J. (1969). Factors for Calculating Two-Sided Prediction Intervals for Samples from a Lognormal Distribution. *Journal of the American Statistical Association* **64**(327), 878-898.
- Hahn, G.J. (1970a). Additional Factors for Calculating Prediction Intervals for Samples from a Lognormal Distribution. *Journal of the American Statistical Association* **65**(332), 1668-1676.
- Hahn, G.J. (1970b). Statistical Intervals for a Lognormal Population, Part I: Tables, Examples and Applications. *Journal of Quality Technology* **2**(3), 115-125.
- Hahn, G.J. (1970c). Statistical Intervals for a Lognormal Population, Part II: Formulas, Assumptions, Some Derivations. *Journal of Quality Technology* **2**(4), 195-206.
- Hahn, G.J., and W.Q. Meeker. (1991). *Statistical Intervals: A Guide for Practitioners*. John Wiley and Sons, New York.
- Hahn, G., and W. Nelson. (1973). A Survey of Prediction Intervals and Their Applications. *Journal of Quality Technology* **5**, 178-188.
- Hall, I.J., and R.R. Prairie. (1973). One-Sided Prediction Intervals to Contain at Least m Out of k Future Observations. *Technometrics* **15**, 897-914.
- Hawkins, D. M., and R.A.J. Wixley. (1986). A Note on the Transformation of Chi-Squared Variables to Normality. *The American Statistician*, **40**, 296-298.
- Johnson, N.L., S. Kotz, and N. Balakrishnan. (1994). *Continuous Univariate Distributions, Volume 1*. Second Edition. John Wiley and Sons, New York, Chapter 17.
- Krishnamoorthy K., T. Mathew, and S. Mukherjee. (2008). Normal-Based Methods for a Gamma Distribution: Prediction and Tolerance Intervals and Stress-Strength Reliability. *Technometrics*, **50**(1), 69-78.
- Krishnamoorthy K., and T. Mathew. (2009). *Statistical Tolerance Regions: Theory, Applications, and Computation*. John Wiley and Sons, Hoboken.
- Kulkarni, H.V., and S.K. Powar. (2010). A New Method for Interval Estimation of the Mean of the Gamma Distribution. *Lifetime Data Analysis*, **16**, 431-447.
- Millard, S.P. (1987). Environmental Monitoring, Statistics, and the Law: Room for Improvement (with Comment). *The American Statistician* **41**(4), 249-259.
- Millard, S.P., and N.K. Neerchal. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton.
- Singh, A., A.K. Singh, and R.J. Iaci. (2002). *Estimation of the Exposure Point Concentration Term Using a Gamma Distribution*. EPA/600/R-02/084. October 2002. Technology Support Center for Monitoring and Site Characterization, Office of Research and Development, Office of Solid Waste and Emergency Response, U.S. Environmental Protection Agency, Washington, D.C.
- Singh, A., R. Maichle, and N. Armbya. (2010a). *ProUCL Version 4.1.00 User Guide (Draft)*. EPA/600/R-07/041, May 2010. Office of Research and Development, U.S. Environmental Protection Agency, Washington, D.C.
- Singh, A., N. Armbya, and A. Singh. (2010b). *ProUCL Version 4.1.00 Technical Guide (Draft)*. EPA/600/R-07/041, May 2010. Office of Research and Development, U.S. Environmental Protection Agency, Washington, D.C.
- Wilson, E.B., and M.M. Hilferty. (1931). The Distribution of Chi-Squares. *Proceedings of the National Academy of Sciences*, **17**, 684-688.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

See Also

[GammaDist](#), [GammaAlt](#), [predIntNorm](#), [predIntNormSimultaneous](#), [predIntNormSimultaneousTestPower](#), [toIntGamma](#), [egamma](#), [egammaAlt](#), [estimate.object](#).

Examples

```
# Generate 8 observations from a gamma distribution with parameters
# mean=10 and cv=1, then use predIntGammaAltSimultaneous to estimate the
# mean and coefficient of variation of the true distribution and construct an
# upper 95% prediction interval to contain at least 1 out of the next
# 3 observations.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(479)
dat <- rgammaAlt(8, mean = 10, cv = 1)

predIntGammaAltSimultaneous(dat, k = 1, m = 3)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:      Gamma
#
#Estimated Parameter(s):   mean = 13.875825
#                           cv   =  1.049504
#
#Estimation Method:        MLE
#
#Data:                     dat
#
#Sample Size:              8
#
#Prediction Interval Method: exact using
#                           Kulkarni & Powar (2010)
#                           transformation to Normality
#                           based on MLE of 'shape'
#
#Normal Transform Power:   0.2204908
#
#Prediction Interval Type:  upper
#
#Confidence Level:         95%
#
#Minimum Number of
#Future Observations
#Interval Should Contain:   1
#
#Total Number of
#Future Observations:      3
#
#Prediction Interval:       LPL =  0.00000
#                           UPL = 15.87101
```

```

#-----

# Compare the 95% 1-of-3 upper prediction limit to the California and
# Modified California upper prediction limits. Note that the upper
# prediction limit for the Modified California rule is between the limit
# for the 1-of-3 rule and the limit for the California rule.

predIntGammaAltSimultaneous(dat, k = 1, m = 3)$interval$limits["UPL"]
#      UPL
#15.87101

predIntGammaAltSimultaneous(dat, m = 3, rule = "CA")$interval$limits["UPL"]
#      UPL
#34.11499

predIntGammaAltSimultaneous(dat, rule = "Modified.CA")$interval$limits["UPL"]
#      UPL
#22.58809

#-----

# Show how the upper 95% simultaneous prediction limit increases
# as the number of future sampling occasions r increases.
# Here, we'll use the 1-of-3 rule.

predIntGammaAltSimultaneous(dat, k = 1, m = 3)$interval$limits["UPL"]
#      UPL
#15.87101

predIntGammaAltSimultaneous(dat, k = 1, m = 3, r = 10)$interval$limits["UPL"]
#      UPL
#37.86825

#-----

# Compare the upper simultaneous prediction limit for the 1-of-3 rule
# based on individual observations versus based on transformed means of
# order 4.

predIntGammaAltSimultaneous(dat, k = 1, m = 3)$interval$limits["UPL"]
#      UPL
#15.87101

predIntGammaAltSimultaneous(dat, n.transmean = 4, k = 1,
  m = 3)$interval$limits["UPL"]
#      UPL
#14.76528

#=====

# Example 19-1 of USEPA (2009, p. 19-17) shows how to compute an
# upper simultaneous prediction limit for the 1-of-3 rule for
# r = 2 future sampling occasions. The data for this example are
# stored in EPA.09.Ex.19.1.sulfate.df.

# We will pool data from 4 background wells that were sampled on

```

```

# a number of different occasions, giving us a sample size of
# n = 25 to use to construct the prediction limit.

# There are 50 compliance wells and we will monitor 10 different
# constituents at each well at each of the r=2 future sampling
# occasions. To determine the confidence level we require for
# the simultaneous prediction interval, USEPA (2009) recommends
# setting the individual Type I Error level at each well to

#  $1 - (1 - \text{SWFPR})^{(1 / (\text{Number of Constituents} * \text{Number of Wells}))}$ 

# which translates to setting the confidence limit to

#  $(1 - \text{SWFPR})^{(1 / (\text{Number of Constituents} * \text{Number of Wells}))}$ 

# where SWFPR = site-wide false positive rate. For this example, we
# will set SWFPR = 0.1. Thus, the confidence level is given by:

nc <- 10
nw <- 50
SWFPR <- 0.1
conf.level <- (1 - SWFPR)^(1 / (nc * nw))

conf.level
#[1] 0.9997893

#-----

# Look at the data:

names(EPA.09.Ex.19.1.sulfate.df)
#[1] "Well"           "Month"           "Day"
#[4] "Year"           "Date"            "Sulfate.mg.per.l"
#[7] "log.Sulfate.mg.per.l"

EPA.09.Ex.19.1.sulfate.df[,
  c("Well", "Date", "Sulfate.mg.per.l", "log.Sulfate.mg.per.l")]

#   Well      Date Sulfate.mg.per.l log.Sulfate.mg.per.l
#1  GW-01 1999-07-08          63.0          4.143135
#2  GW-01 1999-09-12          51.0          3.931826
#3  GW-01 1999-10-16          60.0          4.094345
#4  GW-01 1999-11-02          86.0          4.454347
#5  GW-04 1999-07-09         104.0          4.644391
#6  GW-04 1999-09-14         102.0          4.624973
#7  GW-04 1999-10-12          84.0          4.430817
#8  GW-04 1999-11-15          72.0          4.276666
#9  GW-08 1997-10-12          31.0          3.433987
#10 GW-08 1997-11-16          84.0          4.430817
#11 GW-08 1998-01-28          65.0          4.174387
#12 GW-08 1999-04-20          41.0          3.713572
#13 GW-08 2002-06-04          51.8          3.947390
#14 GW-08 2002-09-16          57.5          4.051785
#15 GW-08 2002-12-02          66.8          4.201703
#16 GW-08 2003-03-24          87.1          4.467057
#17 GW-09 1997-10-16          59.0          4.077537
#18 GW-09 1998-01-28          85.0          4.442651

```

```

#19 GW-09 1998-04-12      75.0      4.317488
#20 GW-09 1998-07-12      99.0      4.595120
#21 GW-09 2000-01-30      75.8      4.328098
#22 GW-09 2000-04-24      82.5      4.412798
#23 GW-09 2000-10-24      85.5      4.448516
#24 GW-09 2002-12-01     188.0      5.236442
#25 GW-09 2003-03-24     150.0      5.010635

# The EPA guidance document constructs the upper simultaneous
# prediction limit for the 1-of-3 plan assuming a lognormal
# distribution for the sulfate data. Here we will compare
# the value of the limit based on assuming a lognormal distribution
# versus assuming a gamma distribution.

Sulfate <- EPA.09.Ex.19.1.sulfate.df$Sulfate.mg.per.l

pred.int.list.lnorm <-
  predIntLnormSimultaneous(x = Sulfate, k = 1, m = 3, r = 2,
    rule = "k.of.m", pi.type = "upper", conf.level = conf.level)

pred.int.list.gamma <-
  predIntGammaSimultaneous(x = Sulfate, k = 1, m = 3, r = 2,
    rule = "k.of.m", pi.type = "upper", conf.level = conf.level)

pred.int.list.lnorm$interval$limits["UPL"]
#      UPL
#159.5497

pred.int.list.gamma$interval$limits["UPL"]
#      UPL
#153.3232

#=====

# Cleanup
#-----
rm(dat, nc, nw, SWFPR, conf.level, Sulfate, pred.int.list.lnorm,
  pred.int.list.gamma)

```

predIntLnorm

Prediction Interval for a Lognormal Distribution

Description

Estimate the mean and standard deviation on the log-scale for a [lognormal distribution](#), or estimate the mean and coefficient of variation for a [lognormal distribution \(alternative parameterization\)](#), and construct a prediction interval for the next k observations or next set of k geometric means.

Usage

```

predIntLnorm(x, n.geomean = 1, k = 1, method = "Bonferroni",
  pi.type = "two-sided", conf.level = 0.95)

```

```
predIntLnormAlt(x, n.geomean = 1, k = 1, method = "Bonferroni",
  pi.type = "two-sided", conf.level = 0.95, est.arg.list = NULL)
```

Arguments

<code>x</code>	For <code>predIntLnorm</code> , <code>x</code> can be a numeric vector of positive observations, or an object resulting from a call to an estimating function that assumes a lognormal distribution (i.e., elnorm or elnormCensored). You <i>cannot</i> supply objects resulting from a call to estimating functions that use the alternative parameterization such as elnormAlt or elnormAltCensored . For <code>predIntLnormAlt</code> , a numeric vector of positive observations. If <code>x</code> is a numeric vector, missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
<code>n.geomean</code>	positive integer specifying the sample size associated with the k future geometric means. The default value is <code>n.geomean=1</code> (i.e., individual observations). Note that all future geometric means must be based on the same sample size.
<code>k</code>	positive integer specifying the number of future observations or geometric means the prediction interval should contain with confidence level <code>conf.level</code> . The default value is <code>k=1</code> .
<code>method</code>	character string specifying the method to use if the number of future observations (k) is greater than 1. The possible values are <code>method="Bonferroni"</code> (approximate method based on Bonferroni inequality; the default), and <code>method="exact"</code> (exact method due to Dunnett, 1955). See the DETAILS section of predIntNormK for more information. This argument is ignored if <code>k=1</code> .
<code>pi.type</code>	character string indicating what kind of prediction interval to compute. The possible values are <code>pi.type="two-sided"</code> (the default), <code>pi.type="lower"</code> , and <code>pi.type="upper"</code> .
<code>conf.level</code>	a scalar between 0 and 1 indicating the confidence level of the prediction interval. The default value is <code>conf.level=0.95</code> .
<code>est.arg.list</code>	for <code>predIntLnormAlt</code> , a list containing arguments to pass to the function elnormAlt for estimating the mean and coefficient of variation. The default value is <code>est.arg.list=NULL</code> , which implies the default values will be used in the call to elnormAlt .

Details

The function `predIntLnorm` returns a prediction interval as well as estimates of the meanlog and sdlog parameters. The function `predIntLnormAlt` returns a prediction interval as well as estimates of the mean and coefficient of variation.

A prediction interval for a lognormal distribution is constructed by taking the natural logarithm of the observations and constructing a prediction interval based on the normal (Gaussian) distribution by calling [predIntNorm](#). These prediction limits are then exponentiated to produce a prediction interval on the original scale of the data.

Value

If `x` is a numeric vector, a list of class "estimate" containing the estimated parameters, the prediction interval, and other information. See the help file for [estimate.object](#) for details.

If `x` is the result of calling an estimation function, `predIntLnorm` returns a list whose class is the same as `x`. The list contains the same components as `x`, as well as a component called `interval`

containing the prediction interval information. If x already has a component called `interval`, this component is replaced with the prediction interval information.

Note

Prediction and tolerance intervals have long been applied to quality control and life testing problems (Hahn, 1970b,c; Hahn and Nelson, 1973; Krishnamoorthy and Mathew, 2009). In the context of environmental statistics, prediction intervals are useful for analyzing data from groundwater detection monitoring programs at hazardous and solid waste facilities (e.g., Gibbons et al., 2009; Millard and Neerchal, 2001; USEPA, 2009).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Berthouex, P.M., and L.C. Brown. (2002). *Statistics for Environmental Engineers*. Lewis Publishers, Boca Raton.
- Dunnett, C.W. (1955). A Multiple Comparisons Procedure for Comparing Several Treatments with a Control. *Journal of the American Statistical Association* **50**, 1096-1121.
- Dunnett, C.W. (1964). New Tables for Multiple Comparisons with a Control. *Biometrics* **20**, 482-491.
- Gibbons, R.D., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*, Second Edition. John Wiley & Sons, Hoboken.
- Hahn, G.J. (1969). Factors for Calculating Two-Sided Prediction Intervals for Samples from a Lognormal Distribution. *Journal of the American Statistical Association* **64**(327), 878-898.
- Hahn, G.J. (1970a). Additional Factors for Calculating Prediction Intervals for Samples from a Lognormal Distribution. *Journal of the American Statistical Association* **65**(332), 1668-1676.
- Hahn, G.J. (1970b). Statistical Intervals for a Lognormal Population, Part I: Tables, Examples and Applications. *Journal of Quality Technology* **2**(3), 115-125.
- Hahn, G.J. (1970c). Statistical Intervals for a Lognormal Population, Part II: Formulas, Assumptions, Some Derivations. *Journal of Quality Technology* **2**(4), 195-206.
- Hahn, G.J., and W.Q. Meeker. (1991). *Statistical Intervals: A Guide for Practitioners*. John Wiley and Sons, New York.
- Hahn, G., and W. Nelson. (1973). A Survey of Prediction Intervals and Their Applications. *Journal of Quality Technology* **5**, 178-188.
- Helsel, D.R., and R.M. Hirsch. (1992). *Statistical Methods in Water Resources Research*. Elsevier, New York.
- Helsel, D.R., and R.M. Hirsch. (2002). *Statistical Methods in Water Resources*. Techniques of Water Resources Investigations, Book 4, chapter A3. U.S. Geological Survey. (available on-line at: <http://pubs.usgs.gov/twri/twri4a3/>).
- Krishnamoorthy K., and T. Mathew. (2009). *Statistical Tolerance Regions: Theory, Applications, and Computation*. John Wiley and Sons, Hoboken.
- Millard, S.P., and Neerchal, N.K. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, Florida.
- Miller, R.G. (1981a). *Simultaneous Statistical Inference*. McGraw-Hill, New York.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery

Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

See Also

[elnorm](#), [elnormAlt](#), [predIntNorm](#), [predIntNormK](#), [predIntLnormSimultaneous](#), [predIntLnormAltSimultaneous](#), [tolIntLnorm](#), [tolIntLnormAlt](#), [Lognormal](#), [estimate.object](#).

Examples

```
# Generate 20 observations from a lognormal distribution with parameters
# meanlog=0 and sdlog=1. The exact two-sided 90% prediction interval for
# k=1 future observation is given by: [exp(-1.645), exp(1.645)] = [0.1930, 5.181].
# Use predIntLnorm to estimate the distribution parameters, and construct a
# two-sided 90% prediction interval.
# (Note: the call to set.seed simply allows you to reproduce this example.)
```

```
set.seed(47)
dat <- rlnorm(20, meanlog = 0, sdlog = 1)
predIntLnorm(dat, conf = 0.9)
```

```
#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Lognormal
#
#Estimated Parameter(s):      meanlog = -0.1035722
#                               sdlog   =  0.9106429
#
#Estimation Method:           mvue
#
#Data:                         dat
#
#Sample Size:                  20
#
#Prediction Interval Method:    exact
#
#Prediction Interval Type:      two-sided
#
#Confidence Level:              90%
#
#Number of Future Observations: 1
#
#Prediction Interval:          LPL = 0.1795898
#                               UPL = 4.5264399
```

```
#-----
```

```
# Repeat the above example, but do it in two steps.
# First create a list called est.list containing information about the
# estimated parameters, then create the prediction interval.
```

```
est.list <- elnorm(dat)
```

```

est.list
#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:      Lognormal
#
#Estimated Parameter(s):   meanlog = -0.1035722
#                           sdlog   =  0.9106429
#
#Estimation Method:        mvue
#
#Data:                     dat
#
#Sample Size:              20

predIntLnorm(est.list, conf = 0.9)
#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:      Lognormal
#
#Estimated Parameter(s):   meanlog = -0.1035722
#                           sdlog   =  0.9106429
#
#Estimation Method:        mvue
#
#Data:                     dat
#
#Sample Size:              20
#
#Prediction Interval Method: exact
#
#Prediction Interval Type:  two-sided
#
#Confidence Level:         90%
#
#Number of Future Observations: 1
#
#Prediction Interval:       LPL = 0.1795898
#                           UPL = 4.5264399
#-----

# Using the same data from the first example, create a one-sided
# upper 99% prediction limit for the next 3 geometric means of order 2
# (i.e., each of the 3 future geometric means is based on a sample size
# of 2 future observations).

predIntLnorm(dat, n.geomean = 2, k = 3, conf.level = 0.99,
  pi.type = "upper")

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:      Lognormal
#

```

```

#Estimated Parameter(s):      meanlog = -0.1035722
#                               sdlog   =  0.9106429
#
#Estimation Method:           mvue
#
#Data:                         dat
#
#Sample Size:                  20
#
#Prediction Interval Method:    Bonferroni
#
#Prediction Interval Type:      upper
#
#Confidence Level:             99%
#
#Number of Future
#Geometric Means:              3
#
#Sample Size for
#Geometric Means:              2
#
#Prediction Interval:           LPL = 0.000000
#                               UPL = 7.047571

#-----

# Compare the result above that is based on the Bonferroni method
# with the exact method

predIntLnorm(dat, n.geomean = 2, k = 3, conf.level = 0.99,
  pi.type = "upper", method = "exact")$interval$limits["UPL"]

#    UPL
#7.00316

#-----

# Clean up
rm(dat, est.list)

#-----

# Example 18-2 of USEPA (2009, p.18-15) shows how to construct a 99%
# upper prediction interval for the log-scale mean of 4 future observations
# (future mean of order 4) assuming a lognormal distribution based on
# chrysene concentrations (ppb) in groundwater at 2 background wells.
# Data were collected once per month over 4 months at the 2 background
# wells, and also at a compliance well.
# The question to be answered is whether there is evidence of
# contamination at the compliance well.

# Here we will follow the example, but look at the geometric mean
# instead of the log-scale mean.

#-----

# The data for this example are stored in EPA.09.Ex.18.2.chrysene.df.

```

```
EPA.09.Ex.18.2.chrysene.df
```

```
#   Month   Well Well.type Chrysene.ppb
#1     1   Well.1 Background         6.9
#2     2   Well.1 Background        27.3
#3     3   Well.1 Background        10.8
#4     4   Well.1 Background         8.9
#5     1   Well.2 Background        15.1
#6     2   Well.2 Background         7.2
#7     3   Well.2 Background        48.4
#8     4   Well.2 Background         7.8
#9     1   Well.3 Compliance        68.0
#10    2   Well.3 Compliance        48.9
#11    3   Well.3 Compliance        30.1
#12    4   Well.3 Compliance        38.1
```

```
Chrysene.bkgd <- with(EPA.09.Ex.18.2.chrysene.df,
  Chrysene.ppb[Well.type == "Background"])
Chrysene.cmpl <- with(EPA.09.Ex.18.2.chrysene.df,
  Chrysene.ppb[Well.type == "Compliance"])
```

```
#-----
```

```
# A Shapiro-Wilks goodness-of-fit test for normality indicates
# we should reject the assumption of normality and assume a
# lognormal distribution for the background well data:
```

```
gofTest(Chrysene.bkgd)
```

```
#Results of Goodness-of-Fit Test
```

```
#-----
```

```
#
#Test Method:                Shapiro-Wilk GOF
#
#Hypothesized Distribution:   Normal
#
#Estimated Parameter(s):     mean = 16.55000
#                             sd   = 14.54441
#
#Estimation Method:          mvue
#
#Data:                       Chrysene.bkgd
#
#Sample Size:                 8
#
#Test Statistic:              W = 0.7289006
#
#Test Statistic Parameter:    n = 8
#
#P-value:                     0.004759859
#
#Alternative Hypothesis:      True cdf does not equal the
#                             Normal Distribution.
```

```
gofTest(Chrysene.bkgd, dist = "lnorm")
```

```

#Results of Goodness-of-Fit Test
#-----
#
#Test Method:                Shapiro-Wilk GOF
#
#Hypothesized Distribution:   Lognormal
#
#Estimated Parameter(s):     meanlog = 2.5533006
#                             sdlog   = 0.7060038
#
#Estimation Method:          mvue
#
#Data:                       Chrysene.bkgd
#
#Sample Size:                8
#
#Test Statistic:             W = 0.8546352
#
#Test Statistic Parameter:   n = 8
#
#P-value:                    0.1061057
#
#Alternative Hypothesis:     True cdf does not equal the
#                             Lognormal Distribution.

#-----

# Here is the one-sided 99% upper prediction limit for
# a geometric mean based on 4 future observations:

predIntLnorm(Chrysene.bkgd, n.geomean = 4, k = 1,
  conf.level = 0.99, pi.type = "upper")

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:        Lognormal
#
#Estimated Parameter(s):     meanlog = 2.5533006
#                             sdlog   = 0.7060038
#
#Estimation Method:          mvue
#
#Data:                       Chrysene.bkgd
#
#Sample Size:                8
#
#Prediction Interval Method:  exact
#
#Prediction Interval Type:    upper
#
#Confidence Level:           99%
#
#Number of Future
#Geometric Means:            1
#

```

```

#Sample Size for
#Geometric Means:          4
#
#Prediction Interval:      LPL = 0.00000
#                          UPL = 46.96613

UPL <- predIntLnorm(Chrysene.bkgd, n.geomean = 4, k = 1,
  conf.level = 0.99, pi.type = "upper")$interval$limits["UPL"]

UPL
#      UPL
#46.96613

# Is there evidence of contamination at the compliance well?

geoMean(Chrysene.cmpl)
#[1] 44.19034

# Since the geometric mean at the compliance well is less than
# the upper prediction limit, there is no evidence of contamination.

#-----

# Cleanup
#-----

rm(Chrysene.bkgd, Chrysene.cmpl, UPL)

```

predIntLnormSimultaneous

Simultaneous Prediction Interval for a Lognormal Distribution

Description

Estimate the mean and standard deviation on the log-scale for a [lognormal distribution](#), or estimate the mean and coefficient of variation for a [lognormal distribution \(alternative parameterization\)](#), and construct a simultaneous prediction interval for the next r sampling occasions, based on one of three possible rules: k-of-m, California, or Modified California.

Usage

```

predIntLnormSimultaneous(x, n.geomean = 1, k = 1, m = 2, r = 1, rule = "k.of.m",
  delta.over.sigma = 0, pi.type = "upper", conf.level = 0.95,
  K.tol = .Machine$double.eps^0.5)

predIntLnormAltSimultaneous(x, n.geomean = 1, k = 1, m = 2, r = 1, rule = "k.of.m",
  delta.over.sigma = 0, pi.type = "upper", conf.level = 0.95,
  K.tol = .Machine$double.eps^0.5, est.arg.list = NULL)

```

Arguments

x For `predIntLnormSimultaneous`, **x** can be a numeric vector of positive observations, or an object resulting from a call to an estimating function that assumes

	<p>a lognormal distribution (i.e., <code>elnorm</code> or <code>elnormCensored</code>). You <i>cannot</i> supply objects resulting from a call to estimating functions that use the alternative parameterization such as <code>elnormAlt</code> or <code>elnormAltCensored</code>.</p> <p>For <code>predIntLnormAltSimultaneous</code>, a numeric vector of positive observations.</p> <p>If <code>x</code> is a numeric vector, missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.</p>
<code>n.geomean</code>	positive integer specifying the sample size associated with future geometric means. The default value is <code>n.geomean=1</code> (i.e., individual observations). Note that all future geometric means must be based on the same sample size.
<code>k</code>	for the <i>k</i> -of- <i>m</i> rule (<code>rule="k.of.m"</code>), a positive integer specifying the minimum number of observations (or geometric means) out of <i>m</i> observations (or geometric means) (all obtained on one future sampling “occasion”) the prediction interval should contain with confidence level <code>conf.level</code> . The default value is <code>k=1</code> . This argument is ignored when the argument <code>rule</code> is not equal to <code>"k.of.m"</code> .
<code>m</code>	positive integer specifying the maximum number of future observations (or geometric means) on one future sampling “occasion”. The default value is <code>m=2</code> , except when <code>rule="Modified.CA"</code> , in which case this argument is ignored and <code>m</code> is automatically set equal to 4.
<code>r</code>	positive integer specifying the number of future sampling “occasions”. The default value is <code>r=1</code> .
<code>rule</code>	character string specifying which rule to use. The possible values are <code>"k.of.m"</code> (<i>k</i> -of- <i>m</i> rule; the default), <code>"CA"</code> (California rule), and <code>"Modified.CA"</code> (modified California rule). See the DETAILS section below for more information.
<code>delta.over.sigma</code>	numeric scalar indicating the ratio Δ/σ . The quantity Δ (delta) denotes the difference between the mean of the population (on the log-scale) that was sampled to construct the prediction interval, and the mean of the population (on the log-scale) that will be sampled to produce the future observations. The quantity σ (sigma) denotes the population standard deviation (on the log-scale) for both populations. See the DETAILS section below for more information. The default value is <code>delta.over.sigma=0</code> .
<code>pi.type</code>	character string indicating what kind of prediction interval to compute. The possible values are <code>pi.type="upper"</code> (the default), and <code>pi.type="lower"</code> .
<code>conf.level</code>	a scalar between 0 and 1 indicating the confidence level of the prediction interval. The default value is <code>conf.level=0.95</code> .
<code>K.tol</code>	numeric scalar indicating the tolerance to use in the nonlinear search algorithm to compute <i>K</i> . The default value is <code>K.tol=.Machine\$double.eps^(1/2)</code> . For many applications, the value of <i>K</i> needs to be known only to the second decimal place, in which case setting <code>K.tol=1e-4</code> will speed up computation a bit.
<code>est.arg.list</code>	a list containing arguments to pass to the function <code>elnormAlt</code> for estimating the mean and coefficient of variation. The default value is <code>est.arg.list=NULL</code> , which implies the default values will be used in the call to <code>elnormAlt</code> .

Details

The function `predIntLnormSimultaneous` returns a simultaneous prediction interval as well as estimates of the `meanlog` and `sdlog` parameters. The function `predIntLnormAltSimultaneous` returns a prediction interval as well as estimates of the mean and coefficient of variation.

A simultaneous prediction interval for a lognormal distribution is constructed by taking the natural logarithm of the observations and constructing a prediction interval based on the normal (Gaussian) distribution by calling `predIntLnormSimultaneous`. These prediction limits are then exponentiated to produce a prediction interval on the original scale of the data.

Value

If `x` is a numeric vector, `predIntLnormSimultaneous` returns a list of class "estimate" containing the estimated parameters, the prediction interval, and other information. See the help file for `estimate.object` for details.

If `x` is the result of calling an estimation function, `predIntLnormSimultaneous` returns a list whose class is the same as `x`. The list contains the same components as `x`, as well as a component called `interval` containing the prediction interval information. If `x` already has a component called `interval`, this component is replaced with the prediction interval information.

Note

Motivation

Prediction and tolerance intervals have long been applied to quality control and life testing problems (Hahn, 1970b,c; Hahn and Nelson, 1973). In the context of environmental statistics, prediction intervals are useful for analyzing data from groundwater detection monitoring programs at hazardous and solid waste facilities.

One of the main statistical problems that plague groundwater monitoring programs at hazardous and solid waste facilities is the requirement of testing several wells and several constituents at each well on each sampling occasion. This is an obvious multiple comparisons problem, and the naive approach of using a standard t-test at a conventional α -level (e.g., 0.05 or 0.01) for each test leads to a very high probability of at least one significant result on each sampling occasion, when in fact no contamination has occurred. This problem was pointed out years ago by Millard (1987) and others.

Davis and McNichols (1987) proposed simultaneous prediction intervals as a way of controlling the facility-wide false positive rate (FWFPR) while maintaining adequate power to detect contamination in the groundwater. Because of the ubiquitous presence of spatial variability, it is usually best to use simultaneous prediction intervals at each well (Davis, 1998a). That is, by constructing prediction intervals based on background (pre-landfill) data on each well, and comparing future observations at a well to the prediction interval for that particular well. In each of these cases, the individual α -level at each well is equal to the FWFPR divided by the product of the number of wells and constituents.

Often, observations at downgradient wells are not available prior to the construction and operation of the landfill. In this case, upgradient well data can be combined to create a background prediction interval, and observations at each downgradient well can be compared to this prediction interval. If spatial variability is present and a major source of variation, however, this method is not really valid (Davis, 1994; Davis, 1998a).

Chapter 19 of USEPA (2009) contains an extensive discussion of using the 1-of- m rule and the Modified California rule.

Chapters 1 and 3 of Gibbons et al. (2009) discuss simultaneous prediction intervals for the normal and lognormal distributions, respectively.

The k-of-m Rule

For the k -of- m rule, Davis and McNichols (1987) give tables with "optimal" choices of k (in terms of best power for a given overall confidence level) for selected values of m , r , and n . They found that the optimal ratios of k to m (i.e., k/m) are generally small, in the range of 15-50%.

The California Rule

The California rule was mandated in that state for groundwater monitoring at waste disposal facilities when resampling verification is part of the statistical program (Barclay's Code of California Regulations, 1991). The California code mandates a "California" rule with $m \geq 3$. The motivation for this rule may have been a desire to have a majority of the observations in bounds (Davis, 1998a). For example, for a k -of- m rule with $k = 1$ and $m = 3$, a monitoring location will pass if the first observation is out of bounds, the second resample is out of bounds, but the last resample is in bounds, so that 2 out of 3 observations are out of bounds. For the California rule with $m = 3$, either the first observation must be in bounds, or the next 2 observations must be in bounds in order for the monitoring location to pass.

Davis (1998a) states that if the FWFPR is kept constant, then the California rule offers little increased power compared to the k -of- m rule, and can actually decrease the power of detecting contamination.

The Modified California Rule

The Modified California Rule was proposed as a compromise between a 1-of- m rule and the California rule. For a given FWFPR, the Modified California rule achieves better power than the California rule, and still requires at least as many observations in bounds as out of bounds, unlike a 1-of- m rule.

Different Notations Between Different References

For the k -of- m rule described in this help file, both Davis and McNichols (1987) and USEPA (2009, Chapter 19) use the variable p instead of k to represent the minimum number of future observations the interval should contain on each of the r sampling occasions.

Gibbons et al. (2009, Chapter 1) presents extensive lists of the value of K for both k -of- m rules and California rules. Gibbons et al.'s notation reverses the meaning of k and r compared to the notation used in this help file. That is, in Gibbons et al.'s notation, k represents the number of future sampling occasions or monitoring wells, and r represents the minimum number of observations the interval should contain on each sampling occasion.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Barclay's California Code of Regulations.** (1991). Title 22, Section 66264.97 [concerning hazardous waste facilities] and Title 23, Section 2550.7(e)(8) [concerning solid waste facilities]. Barclay's Law Publishers, San Francisco, CA.
- Davis, C.B. (1998a). *Ground-Water Statistics & Regulations: Principles, Progress and Problems*. Second Edition. Environmetrics & Statistics Limited, Henderson, NV.
- Davis, C.B. (1998b). Personal Communication, September 3, 1998.
- Davis, C.B., and R.J. McNichols. (1987). One-sided Intervals for at Least p of m Observations from a Lognormal Population on Each of r Future Occasions. *Technometrics* **29**, 359–370.
- Fertig, K.W., and N.R. Mann. (1977). One-Sided Prediction Intervals for at Least p Out of m Future Observations From a Lognormal Population. *Technometrics* **19**, 167–177.
- Gibbons, R.D., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*, Second Edition. John Wiley & Sons, Hoboken.
- Hahn, G.J. (1969). Factors for Calculating Two-Sided Prediction Intervals for Samples from a Lognormal Distribution. *Journal of the American Statistical Association* **64**(327), 878–898.

- Hahn, G.J. (1970a). Additional Factors for Calculating Prediction Intervals for Samples from a Lognormal Distribution. *Journal of the American Statistical Association* **65**(332), 1668-1676.
- Hahn, G.J. (1970b). Statistical Intervals for a Lognormal Population, Part I: Tables, Examples and Applications. *Journal of Quality Technology* **2**(3), 115-125.
- Hahn, G.J. (1970c). Statistical Intervals for a Lognormal Population, Part II: Formulas, Assumptions, Some Derivations. *Journal of Quality Technology* **2**(4), 195-206.
- Hahn, G.J., and W.Q. Meeker. (1991). *Statistical Intervals: A Guide for Practitioners*. John Wiley and Sons, New York.
- Hahn, G., and W. Nelson. (1973). A Survey of Prediction Intervals and Their Applications. *Journal of Quality Technology* **5**, 178-188.
- Hall, I.J., and R.R. Prairie. (1973). One-Sided Prediction Intervals to Contain at Least m Out of k Future Observations. *Technometrics* **15**, 897-914.
- Millard, S.P. (1987). Environmental Monitoring, Statistics, and the Law: Room for Improvement (with Comment). *The American Statistician* **41**(4), 249-259.
- Millard, S.P., and Neerchal, N.K. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, Florida.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.
- USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

See Also

[predIntLnormAltSimultaneousTestPower](#), [predIntNorm](#), [predIntNormSimultaneous](#), [predIntNormSimultaneous](#), [tolIntLnorm](#), [Lognormal](#), [LognormalAlt](#), [estimate.object](#), [elnorm](#), [elnormAlt](#).

Examples

```
# Generate 8 observations from a lognormal distribution with parameters
# mean=10 and cv=1, then use predIntLnormAltSimultaneous to estimate the
# mean and coefficient of variation of the true distribution and construct an
# upper 95% prediction interval to contain at least 1 out of the next
# 3 observations.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(479)
dat <- rlnormAlt(8, mean = 10, cv = 1)

predIntLnormAltSimultaneous(dat, k = 1, m = 3)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Lognormal
#
#Estimated Parameter(s):      mean = 11.2984322
#                             cv   = 0.9524441
#
#Estimation Method:           mvue
```

```

#
#Data:                                dat
#
#Sample Size:                          8
#
#Prediction Interval Method:           exact
#
#Prediction Interval Type:             upper
#
#Confidence Level:                     95%
#
#Minimum Number of
#Future Observations
#Interval Should Contain:              1
#
#Total Number of
#Future Observations:                  3
#
#Prediction Interval:                   LPL = 0.00000
#                                       UPL = 12.67555

#-----

# Compare the 95% 1-of-3 upper prediction limit to the California and
# Modified California upper prediction limits. Note that the upper
# prediction limit for the Modified California rule is between the limit
# for the 1-of-3 rule and the limit for the California rule.

predIntLnormAltSimultaneous(dat, k = 1, m = 3)$interval$limits["UPL"]
#      UPL
#12.67555

predIntLnormAltSimultaneous(dat, m = 3, rule = "CA")$interval$limits["UPL"]
#      UPL
#25.03606

predIntLnormAltSimultaneous(dat, rule = "Modified.CA")$interval$limits["UPL"]
#      UPL
#17.10475

#-----

# Show how the upper 95% simultaneous prediction limit increases
# as the number of future sampling occasions r increases.
# Here, we'll use the 1-of-3 rule.

predIntLnormAltSimultaneous(dat, k = 1, m = 3)$interval$limits["UPL"]
#      UPL
#12.67555

predIntLnormAltSimultaneous(dat, k = 1, m = 3, r = 10)$interval$limits["UPL"]
#      UPL
#27.72615

#-----

```

```
# Compare the upper simultaneous prediction limit for the 1-of-3 rule
# based on individual observations versus based on geometric means of
# order 4.
```

```
predIntLnormAltSimultaneous(dat, k = 1, m = 3)$interval$limits["UPL"]
#      UPL
#12.67555
```

```
predIntLnormAltSimultaneous(dat, n.geomean = 4, k = 1,
  m = 3)$interval$limits["UPL"]
#      UPL
#11.95532
```

```
#=====
```

```
# Example 19-1 of USEPA (2009, p. 19-17) shows how to compute an
# upper simultaneous prediction limit for the 1-of-3 rule for
# r = 2 future sampling occasions. The data for this example are
# stored in EPA.09.Ex.19.1.sulfate.df.
```

```
# We will pool data from 4 background wells that were sampled on
# a number of different occasions, giving us a sample size of
# n = 25 to use to construct the prediction limit.
```

```
# There are 50 compliance wells and we will monitor 10 different
# constituents at each well at each of the r=2 future sampling
# occasions. To determine the confidence level we require for
# the simultaneous prediction interval, USEPA (2009) recommends
# setting the individual Type I Error level at each well to
```

```
#  $1 - (1 - \text{SWFPR})^{(1 / (\text{Number of Constituents} * \text{Number of Wells}))}$ 
```

```
# which translates to setting the confidence limit to
```

```
#  $(1 - \text{SWFPR})^{(1 / (\text{Number of Constituents} * \text{Number of Wells}))}$ 
```

```
# where SWFPR = site-wide false positive rate. For this example, we
# will set SWFPR = 0.1. Thus, the confidence level is given by:
```

```
nc <- 10
nw <- 50
SWFPR <- 0.1
conf.level <- (1 - SWFPR)^(1 / (nc * nw))
```

```
conf.level
#[1] 0.9997893
```

```
#-----
```

```
# Look at the data:
```

```
names(EPA.09.Ex.19.1.sulfate.df)
#[1] "Well"           "Month"          "Day"
#[4] "Year"           "Date"           "Sulfate.mg.per.l"
#[7] "log.Sulfate.mg.per.l"
```

```
EPA.09.Ex.19.1.sulfate.df[,
```

```

c("Well", "Date", "Sulfate.mg.per.l", "log.Sulfate.mg.per.l")

#   Well      Date Sulfate.mg.per.l log.Sulfate.mg.per.l
#1  GW-01 1999-07-08          63.0          4.143135
#2  GW-01 1999-09-12          51.0          3.931826
#3  GW-01 1999-10-16          60.0          4.094345
#4  GW-01 1999-11-02          86.0          4.454347
#5  GW-04 1999-07-09         104.0          4.644391
#6  GW-04 1999-09-14         102.0          4.624973
#7  GW-04 1999-10-12          84.0          4.430817
#8  GW-04 1999-11-15          72.0          4.276666
#9  GW-08 1997-10-12          31.0          3.433987
#10 GW-08 1997-11-16          84.0          4.430817
#11 GW-08 1998-01-28          65.0          4.174387
#12 GW-08 1999-04-20          41.0          3.713572
#13 GW-08 2002-06-04          51.8          3.947390
#14 GW-08 2002-09-16          57.5          4.051785
#15 GW-08 2002-12-02          66.8          4.201703
#16 GW-08 2003-03-24          87.1          4.467057
#17 GW-09 1997-10-16          59.0          4.077537
#18 GW-09 1998-01-28          85.0          4.442651
#19 GW-09 1998-04-12          75.0          4.317488
#20 GW-09 1998-07-12          99.0          4.595120
#21 GW-09 2000-01-30          75.8          4.328098
#22 GW-09 2000-04-24          82.5          4.412798
#23 GW-09 2000-10-24          85.5          4.448516
#24 GW-09 2002-12-01         188.0          5.236442
#25 GW-09 2003-03-24         150.0          5.010635

# Construct the upper simultaneous prediction limit for the
# 1-of-3 plan assuming a lognormal distribution for the
# sulfate data

Sulfate <- EPA.09.Ex.19.1.sulfate.df$Sulfate.mg.per.l

predIntLnormSimultaneous(x = Sulfate, k = 1, m = 3, r = 2,
  rule = "k.of.m", pi.type = "upper", conf.level = conf.level)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Lognormal
#
#Estimated Parameter(s):      meanlog = 4.3156194
#                               sdlog   = 0.3756697
#
#Estimation Method:           mvue
#
#Data:                         Sulfate
#
#Sample Size:                  25
#
#Prediction Interval Method:   exact
#
#Prediction Interval Type:     upper
#
#Confidence Level:             99.97893%

```

```
#
#Minimum Number of
#Future Observations
#Interval Should Contain
#(per Sampling Occasion):      1
#
#Total Number of
#Future Observations
#(per Sampling Occasion):      3
#
#Number of Future
#Sampling Occasions:           2
#
#Prediction Interval:          LPL =  0.0000
#                               UPL = 159.5497

#=====

# Cleanup
#-----
rm(dat, nc, nw, SWFPR, conf.level, Sulfate)
```

predIntNorm	<i>Prediction Interval for a Normal Distribution</i>
-------------	--

Description

Estimate the mean and standard deviation of a [normal distribution](#), and construct a prediction interval for the next *k* observations or next set of *k* means.

Usage

```
predIntNorm(x, n.mean = 1, k = 1, method = "Bonferroni",
  pi.type = "two-sided", conf.level = 0.95)
```

Arguments

x	a numeric vector of observations, or an object resulting from a call to an estimating function that assumes a normal (Gaussian) distribution (e.g., enorm , eqnorm , enormCensored , etc.). If x is a numeric vector, missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
n.mean	positive integer specifying the sample size associated with the <i>k</i> future averages. The default value is n.mean=1 (i.e., individual observations). Note that all future averages must be based on the same sample size.
k	positive integer specifying the number of future observations or averages the prediction interval should contain with confidence level conf.level. The default value is k=1.
method	character string specifying the method to use if the number of future observations (<i>k</i>) is greater than 1. The possible values are method="Bonferroni" (approximate method based on Bonferonni inequality; the default), and method="exact" (exact method due to Dunnett, 1955). See the DETAILS section of predIntNormK for more information. This argument is ignored if k=1.

pi.type	character string indicating what kind of prediction interval to compute. The possible values are pi.type="two-sided" (the default), pi.type="lower", and pi.type="upper".
conf.level	a scalar between 0 and 1 indicating the confidence level of the prediction interval. The default value is conf.level=0.95.

Details

What is a Prediction Interval?

A prediction interval for some population is an interval on the real line constructed so that it will contain k future observations or averages from that population with some specified probability $(1 - \alpha)100\%$, where $0 < \alpha < 1$ and k is some pre-specified positive integer. The quantity $(1 - \alpha)100\%$ is called the confidence coefficient or confidence level associated with the prediction interval.

The Form of a Prediction Interval

Let $\underline{x} = x_1, x_2, \dots, x_n$ denote a vector of n observations from a [normal distribution](#) with parameters mean= μ and sd= σ . Also, let m denote the sample size associated with the k future averages (i.e., n.mean= m). When $m = 1$, each average is really just a single observation, so in the rest of this help file the term “averages” will replace the phrase “observations or averages”.

For a normal distribution, the form of a two-sided $(1 - \alpha)100\%$ prediction interval is:

$$[\bar{x} - Ks, \bar{x} + Ks] \quad (1)$$

where \bar{x} denotes the sample mean:

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i \quad (2)$$

s denotes the sample standard deviation:

$$s^2 = \frac{1}{n-1} \sum_{i=1}^n (x_i - \bar{x})^2 \quad (3)$$

and K denotes a constant that depends on the sample size n , the confidence level, the number of future averages k , and the sample size associated with the future averages, m . Do not confuse the constant K (uppercase K) with the number of future averages k (lowercase k). The symbol K is used here to be consistent with the notation used for tolerance intervals (see [tolIntNorm](#)).

Similarly, the form of a one-sided lower prediction interval is:

$$[\bar{x} - Ks, \infty] \quad (4)$$

and the form of a one-sided upper prediction interval is:

$$[-\infty, \bar{x} + Ks] \quad (5)$$

but K differs for one-sided versus two-sided prediction intervals. The derivation of the constant K is explained in the help file for [predIntNormK](#).

A Prediction Interval is a Random Interval

A prediction interval is a *random* interval; that is, the lower and/or upper bounds are random variables computed based on sample statistics in the baseline sample. Prior to taking one specific baseline sample, the probability that the prediction interval will contain the next k averages is

$(1 - \alpha)100\%$. Once a specific baseline sample is taken and the prediction interval based on that sample is computed, the probability that that prediction interval will contain the next k averages is not necessarily $(1 - \alpha)100\%$, but it should be close.

If an experiment is repeated N times, and for each experiment:

1. A sample is taken and a $(1 - \alpha)100\%$ prediction interval for $k = 1$ future observation is computed, and
2. One future observation is generated and compared to the prediction interval,

then the number of prediction intervals that actually contain the future observation generated in step 2 above is a [binomial random variable](#) with parameters $\text{size}=N$ and $\text{prob}=(1 - \alpha)100\%$.

If, on the other hand, only one baseline sample is taken and only one prediction interval for $k = 1$ future observation is computed, then the number of future observations out of a total of N future observations that will be contained in that one prediction interval is a binomial random variable with parameters $\text{size}=N$ and $\text{prob}=(1 - \alpha^*)100\%$, where α^* depends on the true population parameters and the computed bounds of the prediction interval.

Value

If x is a numeric vector, `predIntNorm` returns a list of class "estimate" containing the estimated parameters, the prediction interval, and other information. See the help file for [estimate.object](#) for details.

If x is the result of calling an estimation function, `predIntNorm` returns a list whose class is the same as x . The list contains the same components as x , as well as a component called `interval` containing the prediction interval information. If x already has a component called `interval`, this component is replaced with the prediction interval information.

Note

Prediction and tolerance intervals have long been applied to quality control and life testing problems (Hahn, 1970b,c; Hahn and Nelson, 1973; Krishnamoorthy and Mathew, 2009). In the context of environmental statistics, prediction intervals are useful for analyzing data from groundwater detection monitoring programs at hazardous and solid waste facilities (e.g., Gibbons et al., 2009; Millard and Neerchal, 2001; USEPA, 2009).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Berthouex, P.M., and L.C. Brown. (2002). *Statistics for Environmental Engineers*. Lewis Publishers, Boca Raton.
- Dunnett, C.W. (1955). A Multiple Comparisons Procedure for Comparing Several Treatments with a Control. *Journal of the American Statistical Association* **50**, 1096-1121.
- Dunnett, C.W. (1964). New Tables for Multiple Comparisons with a Control. *Biometrics* **20**, 482-491.
- Gibbons, R.D., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*, Second Edition. John Wiley & Sons, Hoboken.
- Hahn, G.J. (1969). Factors for Calculating Two-Sided Prediction Intervals for Samples from a Normal Distribution. *Journal of the American Statistical Association* **64**(327), 878-898.

- Hahn, G.J. (1970a). Additional Factors for Calculating Prediction Intervals for Samples from a Normal Distribution. *Journal of the American Statistical Association* **65**(332), 1668-1676.
- Hahn, G.J. (1970b). Statistical Intervals for a Normal Population, Part I: Tables, Examples and Applications. *Journal of Quality Technology* **2**(3), 115-125.
- Hahn, G.J. (1970c). Statistical Intervals for a Normal Population, Part II: Formulas, Assumptions, Some Derivations. *Journal of Quality Technology* **2**(4), 195-206.
- Hahn, G.J., and W.Q. Meeker. (1991). *Statistical Intervals: A Guide for Practitioners*. John Wiley and Sons, New York.
- Hahn, G., and W. Nelson. (1973). A Survey of Prediction Intervals and Their Applications. *Journal of Quality Technology* **5**, 178-188.
- Helsel, D.R., and R.M. Hirsch. (1992). *Statistical Methods in Water Resources Research*. Elsevier, New York.
- Helsel, D.R., and R.M. Hirsch. (2002). *Statistical Methods in Water Resources*. Techniques of Water Resources Investigations, Book 4, chapter A3. U.S. Geological Survey. (available on-line at: <http://pubs.usgs.gov/twri/twri4a3/>).
- Krishnamoorthy K., and T. Mathew. (2009). *Statistical Tolerance Regions: Theory, Applications, and Computation*. John Wiley and Sons, Hoboken.
- Millard, S.P., and Neerchal, N.K. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, Florida.
- Miller, R.G. (1981a). *Simultaneous Statistical Inference*. McGraw-Hill, New York.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.
- USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

See Also

[predIntNormK](#), [predIntNormSimultaneous](#), [predIntLnorm](#), [tolIntNorm](#), [Normal](#), [estimate.object](#), [enorm](#), [eqnorm](#).

Examples

```
# Generate 20 observations from a normal distribution with parameters
# mean=10 and sd=2, then create a two-sided 95% prediction interval for
# the next observation.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(47)
dat <- rnorm(20, mean = 10, sd = 2)
predIntNorm(dat)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Normal
#
#Estimated Parameter(s):      mean = 9.792856
#                               sd   = 1.821286
```

```

#
#Estimation Method:          mvue
#
#Data:                       dat
#
#Sample Size:                20
#
#Prediction Interval Method:  exact
#
#Prediction Interval Type:    two-sided
#
#Confidence Level:           95%
#
#Number of Future Observations:  1
#
#Prediction Interval:         LPL =  5.886723
#                             UPL = 13.698988

#-----

# Using the same data from the last example, create a one-sided
# upper 99% prediction limit for the next 3 averages of order 2
# (i.e., each of the 3 future averages is based on a sample size
# of 2 future observations).

predIntNorm(dat, n.mean = 2, k = 3, conf.level = 0.99,
  pi.type = "upper")

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:        Normal
#
#Estimated Parameter(s):     mean = 9.792856
#                             sd   = 1.821286
#
#Estimation Method:          mvue
#
#Data:                       dat
#
#Sample Size:                20
#
#Prediction Interval Method:  Bonferroni
#
#Prediction Interval Type:    upper
#
#Confidence Level:           99%
#
#Number of Future Averages:  3
#
#Sample Size for Averages:   2
#
#Prediction Interval:         LPL =    -Inf
#                             UPL = 13.90537

#-----

```

```

# Compare the result above that is based on the Bonferroni method
# with the exact method

predIntNorm(dat, n.mean = 2, k = 3, conf.level = 0.99,
  pi.type = "upper", method = "exact")$interval$limits["UPL"]

#      UPL
#13.89272

#-----

# Clean up
rm(dat)

#-----

# Example 18-1 of USEPA (2009, p.18-9) shows how to construct a 95%
# prediction interval for 4 future observations assuming a
# normal distribution based on arsenic concentrations (ppb) in
# groundwater at a solid waste landfill. There were 4 years of
# quarterly monitoring, and years 1-3 are considered background.
# The question to be answered is whether there is evidence of
# contamination in year 4.

# The data for this example is stored in EPA.09.Ex.18.1.arsenic.df.

EPA.09.Ex.18.1.arsenic.df

#   Year Sampling.Period Arsenic.ppb
#1    1      Background      12.6
#2    1      Background      30.8
#3    1      Background      52.0
#4    1      Background      28.1
#5    2      Background      33.3
#6    2      Background      44.0
#7    2      Background       3.0
#8    2      Background      12.8
#9    3      Background      58.1
#10   3      Background      12.6
#11   3      Background      17.6
#12   3      Background      25.3
#13   4      Compliance      48.0
#14   4      Compliance      30.3
#15   4      Compliance      42.5
#16   4      Compliance      15.0

As.bkgd <- with(EPA.09.Ex.18.1.arsenic.df,
  Arsenic.ppb[Sampling.Period == "Background"])
As.cmpl <- with(EPA.09.Ex.18.1.arsenic.df,
  Arsenic.ppb[Sampling.Period == "Compliance"])

# A Shapiro-Wilks goodness-of-fit test for normality indicates
# there is no evidence to reject the assumption of normality
# for the background data:

gofTest(As.bkgd)

```

```

#Results of Goodness-of-Fit Test
#-----
#
#Test Method:                Shapiro-Wilk GOF
#
#Hypothesized Distribution:   Normal
#
#Estimated Parameter(s):     mean = 27.51667
#                             sd   = 17.10119
#
#Estimation Method:          mvue
#
#Data:                        As.bkgd
#
#Sample Size:                 12
#
#Test Statistic:              W = 0.94695
#
#Test Statistic Parameter:    n = 12
#
#P-value:                     0.5929102
#
#Alternative Hypothesis:      True cdf does not equal the
#                             Normal Distribution.

# Here is the one-sided 95% upper prediction limit:

UPL <- predIntNorm(As.bkgd, k = 4,
  pi.type = "upper")$interval$limits["UPL"]
UPL
#      UPL
#73.67237

# Are any of the compliance observations above the prediction limit?

any(As.cmpl > UPL)
#[1] FALSE

#=====

# Cleanup
#-----

rm(As.bkgd, As.cmpl, UPL)

```

predIntNormK

Compute the Value of K for a Prediction Interval for a Normal Distribution

Description

Compute the value of K (the multiplier of estimated standard deviation) used to construct a prediction interval for the next k observations or next set of k means based on data from a [normal distribution](#). The function `predIntNormK` is called by `predIntNorm`.

Usage

```
predIntNormK(n, df = n - 1, n.mean = 1, k = 1,
  method = "Bonferroni", pi.type = "two-sided",
  conf.level = 0.95)
```

Arguments

<code>n</code>	a positive integer greater than 2 indicating the sample size upon which the prediction interval is based.
<code>df</code>	the degrees of freedom associated with the prediction interval. The default is <code>df=n-1</code> .
<code>n.mean</code>	positive integer specifying the sample size associated with the k future averages. The default value is <code>n.mean=1</code> (i.e., individual observations). Note that all future averages must be based on the same sample size.
<code>k</code>	positive integer specifying the number of future observations or averages the prediction interval should contain with confidence level <code>conf.level</code> . The default value is <code>k=1</code> .
<code>method</code>	character string specifying the method to use if the number of future observations (k) is greater than 1. The possible values are <code>method="Bonferroni"</code> (approximate method based on Bonferonni inequality; the default), and <code>method="exact"</code> (exact method due to Dunnett, 1955). See the DETAILS section for more information. This argument is ignored if <code>k=1</code> .
<code>pi.type</code>	character string indicating what kind of prediction interval to compute. The possible values are <code>pi.type="two-sided"</code> (the default), <code>pi.type="lower"</code> , and <code>pi.type="upper"</code> .
<code>conf.level</code>	a scalar between 0 and 1 indicating the confidence level of the prediction interval. The default value is <code>conf.level=0.95</code> .

Details

A prediction interval for some population is an interval on the real line constructed so that it will contain k future observations or averages from that population with some specified probability $(1 - \alpha)100\%$, where $0 < \alpha < 1$ and k is some pre-specified positive integer. The quantity $(1 - \alpha)100\%$ is called the confidence coefficient or confidence level associated with the prediction interval.

Let $\underline{x} = x_1, x_2, \dots, x_n$ denote a vector of n observations from a [normal distribution](#) with parameters $\text{mean}=\mu$ and $\text{sd}=\sigma$. Also, let m denote the sample size associated with the k future averages (i.e., `n.mean=m`). When $m = 1$, each average is really just a single observation, so in the rest of this help file the term “averages” will replace the phrase “observations or averages”.

For a normal distribution, the form of a two-sided $(1 - \alpha)100\%$ prediction interval is:

$$[\bar{x} - Ks, \bar{x} + Ks] \quad (1)$$

where \bar{x} denotes the sample mean:

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i \quad (2)$$

s denotes the sample standard deviation:

$$s^2 = \frac{1}{n-1} \sum_{i=1}^n (x_i - \bar{x})^2 \quad (3)$$

and K denotes a constant that depends on the sample size n , the confidence level, the number of future averages k , and the sample size associated with the future averages, m . Do not confuse the constant K (uppercase K) with the number of future averages k (lowercase k). The symbol K is used here to be consistent with the notation used for tolerance intervals (see [tolIntNorm](#)).

Similarly, the form of a one-sided lower prediction interval is:

$$[\bar{x} - Ks, \infty] \quad (4)$$

and the form of a one-sided upper prediction interval is:

$$[-\infty, \bar{x} + Ks] \quad (5)$$

but K differs for one-sided versus two-sided prediction intervals. The derivation of the constant K is explained below. The function `predIntNormK` computes the value of K and is called by [predIntNorm](#).

The Derivation of K for One Future Observation or Average ($k = 1$)

Let X denote a random variable from a [normal distribution](#) with parameters $\text{mean}=\mu$ and $\text{sd}=\sigma$, and let x_p denote the p 'th quantile of X .

A true two-sided $(1 - \alpha)100\%$ prediction interval for the next $k = 1$ observation of X is given by:

$$[x_{\alpha/2}, x_{1-\alpha/2}] = [\mu - z_{1-\alpha/2}\sigma, \mu + z_{1-\alpha/2}\sigma] \quad (6)$$

where z_p denotes the p 'th quantile of a standard normal distribution.

More generally, a true two-sided $(1 - \alpha)100\%$ prediction interval for the next $k = 1$ average based on a sample of size m is given by:

$$[\mu - z_{1-\alpha/2}\frac{\sigma}{\sqrt{m}}, \mu + z_{1-\alpha/2}\frac{\sigma}{\sqrt{m}}] \quad (7)$$

Because the values of μ and σ are unknown, they must be estimated, and a prediction interval then constructed based on the estimated values of μ and σ .

For a two-sided prediction interval (`pi.type="two-sided"`), the constant K for a $(1 - \alpha)100\%$ prediction interval for the next $k = 1$ average based on a sample size of m is computed as:

$$K = t_{n-1, 1-\alpha/2} \sqrt{\frac{1}{m} + \frac{1}{n}} \quad (8)$$

where $t_{\nu, p}$ denotes the p 'th quantile of the [Student's t-distribution](#) with ν degrees of freedom. For a one-sided prediction interval (`pi.type="lower"` or `pi.type="upper"`), the prediction interval is given by:

$$K = t_{n-1, 1-\alpha} \sqrt{\frac{1}{m} + \frac{1}{n}} \quad (9)$$

The formulas for these prediction intervals are derived as follows. Let \bar{y} denote the future average based on m observations. Then the quantity $\bar{y} - \bar{x}$ has a normal distribution with expectation and variance given by:

$$E(\bar{y} - \bar{x}) = 0 \quad (10)$$

$$Var(\bar{y} - \bar{x}) = Var(\bar{y}) + Var(\bar{x}) = \frac{\sigma^2}{m} + \frac{\sigma^2}{n} = \sigma^2 \left(\frac{1}{m} + \frac{1}{n} \right) \quad (11)$$

so the quantity

$$t = \frac{\bar{y} - \bar{x}}{s \sqrt{\frac{1}{m} + \frac{1}{n}}} \quad (12)$$

has a [Student's t-distribution](#) with $n - 1$ degrees of freedom.

The Derivation of K for More than One Future Observation or Average ($k > 1$)

When $k > 1$, the function predIntNormK allows for two ways to compute K : an exact method due to Dunnett (1955) (method="exact"), and an approximate (conservative) method based on the Bonferroni inequality (method="Bonferroni"; see Miller, 1981a, pp.8, 67-70; Gibbons et al., 2009, p.4). Each of these methods is explained below.

Exact Method Due to Dunnett (1955) (method="exact")

Dunnett (1955) derived the value of K in the context of the multiple comparisons problem of comparing several treatment means to one control mean. The value of K is computed as:

$$K = c \sqrt{\frac{1}{m} + \frac{1}{n}} \quad (13)$$

where c is a constant that depends on the sample size n , the number of future observations (averages) k , the sample size associated with the k future averages m , and the confidence level $(1 - \alpha)100\%$.

When pi.type="lower" or pi.type="upper", the value of c is the number that satisfies the following equation (Gupta and Sobel, 1957; Hahn, 1970a):

$$1 - \alpha = \int_0^\infty F_1(cs, k, \rho) h(s\sqrt{n-1}, n-1) \sqrt{n-1} ds \quad (14)$$

where

$$F_1(x, k, \rho) = \int_{-\infty}^{\infty} [\Phi(\frac{x + \rho^{1/2}y}{\sqrt{1-\rho}})]^k \phi(y) dy \quad (15)$$

$$\rho = 1/(\frac{n}{m} + 1) \quad (16)$$

$$h(x, \nu) = \frac{x^{\nu-1} e^{-x^2/2}}{2(\nu/2-1)\Gamma(\frac{\nu}{2})} \quad (17)$$

and $\Phi()$ and $\phi()$ denote the cumulative distribution function and probability density function, respectively, of the standard normal distribution. Note that the function $h(x, \nu)$ is the probability density function of a [chi random variable](#) with ν degrees of freedom.

When pi.type="two-sided", the value of c is the number that satisfies the following equation:

$$1 - \alpha = \int_0^\infty F_2(cs, k, \rho) h(s\sqrt{n-1}, n-1) \sqrt{n-1} ds \quad (18)$$

where

$$F_2(x, k, \rho) = \int_{-\infty}^{\infty} [\Phi(\frac{x + \rho^{1/2}y}{\sqrt{1-\rho}}) - \Phi(\frac{-x + \rho^{1/2}y}{\sqrt{1-\rho}})]^k \phi(y) dy \quad (19)$$

Approximate Method Based on the Bonferroni Inequality (method="Bonferroni")

As shown above, when $k = 1$, the value of K is given by Equation (8) or Equation (9) for two-sided or one-sided prediction intervals, respectively. When $k > 1$, a conservative way to construct a $(1 - \alpha^*)100\%$ prediction interval for the next k observations or averages is to use a Bonferroni correction (Miller, 1981a, p.8) and set $\alpha = \alpha^*/k$ in Equation (8) or (9) (Chew, 1968). This value of K will be conservative in that the computed prediction intervals will be wider than the exact predictions intervals. Hahn (1969, 1970a) compared the exact values of K with those based on the Bonferroni inequality for the case of $m = 1$ and found the approximation to be quite satisfactory except when n is small, k is large, and α is large. For example, Gibbons (1987a) notes that for a 99% prediction interval (i.e., $\alpha = 0.01$) for the next k observations, if $n > 4$, the bias of K is never greater than 1% no matter what the value of k .

Value

A numeric scalar equal to K , the multiplier of estimated standard deviation that is used to construct the prediction interval.

Note

Prediction and tolerance intervals have long been applied to quality control and life testing problems (Hahn, 1970b,c; Hahn and Nelson, 1973). In the context of environmental statistics, prediction intervals are useful for analyzing data from groundwater detection monitoring programs at hazardous and solid waste facilities (e.g., Gibbons et al., 2009; Millard and Neerchal, 2001; USEPA, 2009).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Berthouex, P.M., and L.C. Brown. (2002). *Statistics for Environmental Engineers*. Lewis Publishers, Boca Raton.
- Dunnett, C.W. (1955). A Multiple Comparisons Procedure for Comparing Several Treatments with a Control. *Journal of the American Statistical Association* **50**, 1096-1121.
- Dunnett, C.W. (1964). New Tables for Multiple Comparisons with a Control. *Biometrics* **20**, 482-491.
- Gibbons, R.D., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*, Second Edition. John Wiley & Sons, Hoboken.
- Hahn, G.J. (1969). Factors for Calculating Two-Sided Prediction Intervals for Samples from a Normal Distribution. *Journal of the American Statistical Association* **64**(327), 878-898.
- Hahn, G.J. (1970a). Additional Factors for Calculating Prediction Intervals for Samples from a Normal Distribution. *Journal of the American Statistical Association* **65**(332), 1668-1676.
- Hahn, G.J. (1970b). Statistical Intervals for a Normal Population, Part I: Tables, Examples and Applications. *Journal of Quality Technology* **2**(3), 115-125.
- Hahn, G.J. (1970c). Statistical Intervals for a Normal Population, Part II: Formulas, Assumptions, Some Derivations. *Journal of Quality Technology* **2**(4), 195-206.
- Hahn, G.J., and W.Q. Meeker. (1991). *Statistical Intervals: A Guide for Practitioners*. John Wiley and Sons, New York.
- Hahn, G., and W. Nelson. (1973). A Survey of Prediction Intervals and Their Applications. *Journal of Quality Technology* **5**, 178-188.
- Helsel, D.R., and R.M. Hirsch. (1992). *Statistical Methods in Water Resources Research*. Elsevier, New York.
- Helsel, D.R., and R.M. Hirsch. (2002). *Statistical Methods in Water Resources*. Techniques of Water Resources Investigations, Book 4, chapter A3. U.S. Geological Survey. (available on-line at: <http://pubs.usgs.gov/twri/twri4a3/>).
- Millard, S.P., and Neerchal, N.K. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, Florida.
- Miller, R.G. (1981a). *Simultaneous Statistical Inference*. McGraw-Hill, New York.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery

Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

See Also

[predIntNorm](#), [predIntNormSimultaneous](#), [predIntLnorm](#), [tolIntNorm](#), [Normal](#), [estimate.object](#), [enorm](#), [eqnorm](#).

Examples

```
# Compute the value of K for a two-sided 95% prediction interval
# for the next observation given a sample size of n=20.

predIntNormK(n = 20)
#[1] 2.144711

#-----

# Compute the value of K for a one-sided upper 99% prediction limit
# for the next 3 averages of order 2 (i.e., each of the 3 future
# averages is based on a sample size of 2 future observations) given a
# samle size of n=20.

predIntNormK(n = 20, n.mean = 2, k = 3, pi.type = "upper",
  conf.level = 0.99)
#[1] 2.258026

#-----

# Compare the result above that is based on the Bonferroni method
# with the exact method.

predIntNormK(n = 20, n.mean = 2, k = 3, method = "exact",
  pi.type = "upper", conf.level = 0.99)
#[1] 2.251084

#-----

# Example 18-1 of USEPA (2009, p.18-9) shows how to construct a 95%
# prediction interval for 4 future observations assuming a
# normal distribution based on arsenic concentrations (ppb) in
# groundwater at a solid waste landfill. There were 4 years of
# quarterly monitoring, and years 1-3 are considered background,

# So the sample size for the prediciton limit is n = 12,
# and the number of future samples is k = 4.

predIntNormK(n = 12, k = 4, pi.type = "upper")
#[1] 2.698976
```

predIntNormSimultaneous

Simultaneous Prediction Interval for a Normal Distribution

Description

Estimate the mean and standard deviation of a [normal distribution](#), and construct a simultaneous prediction interval for the next r sampling “occasions”, based on one of three possible rules: k -of- m , California, or Modified California.

Usage

```
predIntNormSimultaneous(x, n.mean = 1, k = 1, m = 2, r = 1, rule = "k.of.m",
  delta.over.sigma = 0, pi.type = "upper", conf.level = 0.95,
  K.tol = .Machine$double.eps^0.5)
```

Arguments

- | | |
|------------------|---|
| x | a numeric vector of observations, or an object resulting from a call to an estimating function that assumes a normal (Gaussian) distribution (e.g., enorm , eqnorm , enormCensored , etc.). If x is a numeric vector, missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed. |
| n.mean | positive integer specifying the sample size associated with the future averages. The default value is n.mean=1 (i.e., individual observations). Note that all future averages must be based on the same sample size. |
| k | for the k -of- m rule (rule="k.of.m"), a positive integer specifying the minimum number of observations (or averages) out of m observations (or averages) (all obtained on one future sampling “occasion”) the prediction interval should contain with confidence level conf.level. The default value is k=1. This argument is ignored when the argument rule is not equal to "k.of.m". |
| m | positive integer specifying the maximum number of future observations (or averages) on one future sampling “occasion”. The default value is m=2, except when rule="Modified.CA", in which case this argument is ignored and m is automatically set equal to 4. |
| r | positive integer specifying the number of future sampling “occasions”. The default value is r=1. |
| rule | character string specifying which rule to use. The possible values are "k.of.m" (k -of- m rule; the default), "CA" (California rule), and "Modified.CA" (modified California rule). See the DETAILS section below for more information. |
| delta.over.sigma | numeric scalar indicating the ratio Δ/σ . The quantity Δ (delta) denotes the difference between the mean of the population that was sampled to construct the prediction interval, and the mean of the population that will be sampled to produce the future observations. The quantity σ (sigma) denotes the population standard deviation for both populations. See the DETAILS section below for more information. The default value is delta.over.sigma=0. |
| pi.type | character string indicating what kind of prediction interval to compute. The possible values are pi.type="upper" (the default), and pi.type="lower". |

<code>conf.level</code>	a scalar between 0 and 1 indicating the confidence level of the prediction interval. The default value is <code>conf.level=0.95</code> .
<code>K.tol</code>	numeric scalar indicating the tolerance to use in the nonlinear search algorithm to compute K . The default value is <code>K.tol=.Machine\$double.eps^(1/2)</code> . For many applications, the value of K needs to be known only to the second decimal place, in which case setting <code>K.tol=1e-4</code> will speed up computation a bit.

Details

What is a Simultaneous Prediction Interval?

A prediction interval for some population is an interval on the real line constructed so that it will contain k future observations from that population with some specified probability $(1 - \alpha)100\%$, where $0 < \alpha < 1$ and k is some pre-specified positive integer. The quantity $(1 - \alpha)100\%$ is called the confidence coefficient or confidence level associated with the prediction interval. The function `predIntNorm` computes a standard prediction interval based on a sample from a [normal distribution](#).

The function `predIntNormSimultaneous` computes a simultaneous prediction interval that will contain a certain number of future observations with probability $(1 - \alpha)100\%$ for each of r future sampling “occasions”, where r is some pre-specified positive integer. The quantity r may refer to r distinct future sampling occasions in time, or it may for example refer to sampling at r distinct locations on one future sampling occasion, assuming that the population standard deviation is the same at all of the r distinct locations.

The function `predIntNormSimultaneous` computes a simultaneous prediction interval based on one of three possible rules:

- For the k -of- m rule (`rule="k.of.m"`), at least k of the next m future observations will fall in the prediction interval with probability $(1 - \alpha)100\%$ on each of the r future sampling occasions. If observations are being taken sequentially, for a particular sampling occasion, up to m observations may be taken, but once k of the observations fall within the prediction interval, sampling can stop. Note: When $k = m$ and $r = 1$, the results of `predIntNormSimultaneous` are equivalent to the results of `predIntNorm`.
- For the California rule (`rule="CA"`), with probability $(1 - \alpha)100\%$, for each of the r future sampling occasions, either the first observation will fall in the prediction interval, or else all of the next $m - 1$ observations will fall in the prediction interval. That is, if the first observation falls in the prediction interval then sampling can stop. Otherwise, $m - 1$ more observations must be taken.
- For the Modified California rule (`rule="Modified.CA"`), with probability $(1 - \alpha)100\%$, for each of the r future sampling occasions, either the first observation will fall in the prediction interval, or else at least 2 out of the next 3 observations will fall in the prediction interval. That is, if the first observation falls in the prediction interval then sampling can stop. Otherwise, up to 3 more observations must be taken.

Simultaneous prediction intervals can be extended to using averages (means) in place of single observations (USEPA, 2009, Chapter 19). That is, you can create a simultaneous prediction interval that will contain a specified number of averages (based on which rule you choose) on each of r future sampling occasions, where each average is based on w individual observations. For the function `predIntNormSimultaneous`, the argument `n.mean` corresponds to w .

The Form of a Prediction Interval

Let $\underline{x} = x_1, x_2, \dots, x_n$ denote a vector of n observations from a [normal distribution](#) with parameters $\text{mean}=\mu$ and $\text{sd}=\sigma$. Also, let w denote the sample size associated with the future averages (i.e., $\text{n.mean}=w$). When $w = 1$, each average is really just a single observation, so in the rest of this help file the term “averages” will replace the phrase “observations or averages”.

For a normal distribution, the form of a two-sided $(1 - \alpha)100\%$ prediction interval is:

$$[\bar{x} - Ks, \bar{x} + Ks] \quad (1)$$

where \bar{x} denotes the sample mean:

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i \quad (2)$$

s denotes the sample standard deviation:

$$s^2 = \frac{1}{n-1} \sum_{i=1}^n (x_i - \bar{x})^2 \quad (3)$$

and K denotes a constant that depends on the sample size n , the confidence level, the number of future sampling occasions r , and the sample size associated with the future averages, w . Do not confuse the constant K (uppercase K) with the number of future averages k (lowercase k) in the k -of- m rule. The symbol K is used here to be consistent with the notation used for tolerance intervals (see [tolIntNorm](#)).

Similarly, the form of a one-sided lower prediction interval is:

$$[\bar{x} - Ks, \infty] \quad (4)$$

and the form of a one-sided upper prediction interval is:

$$[-\infty, \bar{x} + Ks] \quad (5)$$

Note: For simultaneous prediction intervals, only lower (`pi.type="lower"`) and upper (`pi.type="upper"`) prediction intervals are available.

The derivation of the constant K is explained in the help file for [predIntNormSimultaneousK](#).

Prediction Intervals are Random Intervals

A prediction interval is a *random* interval; that is, the lower and/or upper bounds are random variables computed based on sample statistics in the baseline sample. Prior to taking one specific baseline sample, the probability that the prediction interval will perform according to the rule chosen is $(1 - \alpha)100\%$. Once a specific baseline sample is taken and the prediction interval based on that sample is computed, the probability that that prediction interval will perform according to the rule chosen is not necessarily $(1 - \alpha)100\%$, but it should be close. See the help file for [predIntNorm](#) for more information.

Value

If x is a numeric vector, `predIntNormSimultaneous` returns a list of class "estimate" containing the estimated parameters, the prediction interval, and other information. See the help file for [estimate.object](#) for details.

If x is the result of calling an estimation function, `predIntNormSimultaneous` returns a list whose class is the same as x . The list contains the same components as x , as well as a component called `interval` containing the prediction interval information. If x already has a component called `interval`, this component is replaced with the prediction interval information.

Note

Motivation

Prediction and tolerance intervals have long been applied to quality control and life testing problems (Hahn, 1970b,c; Hahn and Nelson, 1973). In the context of environmental statistics, prediction intervals are useful for analyzing data from groundwater detection monitoring programs at hazardous and solid waste facilities.

One of the main statistical problems that plague groundwater monitoring programs at hazardous and solid waste facilities is the requirement of testing several wells and several constituents at each well on each sampling occasion. This is an obvious multiple comparisons problem, and the naive approach of using a standard t-test at a conventional α -level (e.g., 0.05 or 0.01) for each test leads to a very high probability of at least one significant result on each sampling occasion, when in fact no contamination has occurred. This problem was pointed out years ago by Millard (1987) and others.

Davis and McNichols (1987) proposed simultaneous prediction intervals as a way of controlling the facility-wide false positive rate (FWFPR) while maintaining adequate power to detect contamination in the groundwater. Because of the ubiquitous presence of spatial variability, it is usually best to use simultaneous prediction intervals at each well (Davis, 1998a). That is, by constructing prediction intervals based on background (pre-landfill) data on each well, and comparing future observations at a well to the prediction interval for that particular well. In each of these cases, the individual α -level at each well is equal to the FWFPR divided by the product of the number of wells and constituents.

Often, observations at downgradient wells are not available prior to the construction and operation of the landfill. In this case, upgradient well data can be combined to create a background prediction interval, and observations at each downgradient well can be compared to this prediction interval. If spatial variability is present and a major source of variation, however, this method is not really valid (Davis, 1994; Davis, 1998a).

Chapter 19 of USEPA (2009) contains an extensive discussion of using the 1-of- m rule and the Modified California rule.

Chapters 1 and 3 of Gibbons et al. (2009) discuss simultaneous prediction intervals for the normal and lognormal distributions, respectively.

The k -of- m Rule

For the k -of- m rule, Davis and McNichols (1987) give tables with “optimal” choices of k (in terms of best power for a given overall confidence level) for selected values of m , r , and n . They found that the optimal ratios of k to m (i.e., k/m) are generally small, in the range of 15-50%.

The California Rule

The California rule was mandated in that state for groundwater monitoring at waste disposal facilities when resampling verification is part of the statistical program (Barclay’s Code of California Regulations, 1991). The California code mandates a “California” rule with $m \geq 3$. The motivation for this rule may have been a desire to have a majority of the observations in bounds (Davis, 1998a). For example, for a k -of- m rule with $k = 1$ and $m = 3$, a monitoring location will pass if the first observation is out of bounds, the second resample is out of bounds, but the last resample is in bounds, so that 2 out of 3 observations are out of bounds. For the California rule with $m = 3$, either the first observation must be in bounds, or the next 2 observations must be in bounds in order for the monitoring location to pass.

Davis (1998a) states that if the FWFPR is kept constant, then the California rule offers little increased power compared to the k -of- m rule, and can actually decrease the power of detecting contamination.

The Modified California Rule

The Modified California Rule was proposed as a compromise between a 1-of- m rule and the California rule. For a given FWPR, the Modified California rule achieves better power than the California rule, and still requires at least as many observations in bounds as out of bounds, unlike a 1-of- m rule.

Different Notations Between Different References

For the k -of- m rule described in this help file, both Davis and McNichols (1987) and USEPA (2009, Chapter 19) use the variable p instead of k to represent the minimum number of future observations the interval should contain on each of the r sampling occasions.

Gibbons et al. (2009, Chapter 1) presents extensive lists of the value of K for both k -of- m rules and California rules. Gibbons et al.'s notation reverses the meaning of k and r compared to the notation used in this help file. That is, in Gibbons et al.'s notation, k represents the number of future sampling occasions or monitoring wells, and r represents the minimum number of observations the interval should contain on each sampling occasion.

USEPA (2009, Chapter 19) uses p in place of k .

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Barclay's California Code of Regulations.** (1991). Title 22, Section 66264.97 [concerning hazardous waste facilities] and Title 23, Section 2550.7(e)(8) [concerning solid waste facilities]. Barclay's Law Publishers, San Francisco, CA.
- Davis, C.B. (1998a). *Ground-Water Statistics & Regulations: Principles, Progress and Problems*. Second Edition. Environmetrics & Statistics Limited, Henderson, NV.
- Davis, C.B. (1998b). Personal Communication, September 3, 1998.
- Davis, C.B., and R.J. McNichols. (1987). One-sided Intervals for at Least p of m Observations from a Normal Population on Each of r Future Occasions. *Technometrics* **29**, 359–370.
- Fertig, K.W., and N.R. Mann. (1977). One-Sided Prediction Intervals for at Least p Out of m Future Observations From a Normal Population. *Technometrics* **19**, 167–177.
- Gibbons, R.D., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*, Second Edition. John Wiley & Sons, Hoboken.
- Hahn, G.J. (1969). Factors for Calculating Two-Sided Prediction Intervals for Samples from a Normal Distribution. *Journal of the American Statistical Association* **64**(327), 878–898.
- Hahn, G.J. (1970a). Additional Factors for Calculating Prediction Intervals for Samples from a Normal Distribution. *Journal of the American Statistical Association* **65**(332), 1668–1676.
- Hahn, G.J. (1970b). Statistical Intervals for a Normal Population, Part I: Tables, Examples and Applications. *Journal of Quality Technology* **2**(3), 115–125.
- Hahn, G.J. (1970c). Statistical Intervals for a Normal Population, Part II: Formulas, Assumptions, Some Derivations. *Journal of Quality Technology* **2**(4), 195–206.
- Hahn, G.J., and W.Q. Meeker. (1991). *Statistical Intervals: A Guide for Practitioners*. John Wiley and Sons, New York.
- Hahn, G., and W. Nelson. (1973). A Survey of Prediction Intervals and Their Applications. *Journal of Quality Technology* **5**, 178–188.
- Hall, I.J., and R.R. Prairie. (1973). One-Sided Prediction Intervals to Contain at Least m Out of k Future Observations. *Technometrics* **15**, 897–914.

Millard, S.P. (1987). Environmental Monitoring, Statistics, and the Law: Room for Improvement (with Comment). *The American Statistician* **41**(4), 249–259.

Millard, S.P., and Neerchal, N.K. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, Florida.

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

See Also

[predIntNormSimultaneousK](#), [predIntNormSimultaneousTestPower](#), [predIntNorm](#), [predIntLnormSimultaneous](#), [tolIntNorm](#), [Normal](#), [estimate.object](#), [enorm](#)

Examples

```
# Generate 8 observations from a normal distribution with parameters
# mean=10 and sd=2, then use predIntNormSimultaneous to estimate the
# mean and standard deviation of the true distribution and construct an
# upper 95% prediction interval to contain at least 1 out of the next
# 3 observations.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(479)
dat <- rnorm(8, mean = 10, sd = 2)

predIntNormSimultaneous(dat, k = 1, m = 3)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Normal
#
#Estimated Parameter(s):      mean = 10.269773
#                               sd   =  2.210246
#
#Estimation Method:           mvue
#
#Data:                         dat
#
#Sample Size:                  8
#
#Prediction Interval Method:    exact
#
#Prediction Interval Type:      upper
#
#Confidence Level:             95%
#
#Minimum Number of
#Future Observations
#Interval Should Contain:      1
#
```

```

#Total Number of
#Future Observations:      3
#
#Prediction Interval:      LPL =  -Inf
#                          UPL = 11.4021

#-----

# Repeat the above example, but do it in two steps. First create a list called
# est.list containing information about the estimated parameters, then create the
# prediction interval.

est.list <- enorm(dat)
est.list

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:      Normal
#
#Estimated Parameter(s):   mean = 10.269773
#                          sd  = 2.210246
#
#Estimation Method:        mvue
#
#Data:                     dat
#
#Sample Size:              8

predIntNormSimultaneous(est.list, k = 1, m = 3)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:      Normal
#
#Estimated Parameter(s):   mean = 10.269773
#                          sd  = 2.210246
#
#Estimation Method:        mvue
#
#Data:                     dat
#
#Sample Size:              8
#
#Prediction Interval Method: exact
#
#Prediction Interval Type: upper
#
#Confidence Level:         95%
#
#Minimum Number of
#Future Observations
#Interval Should Contain:  1
#
#Total Number of

```



```

#Future Observations:          3
#
#Prediction Interval:          LPL =   -Inf
#                              UPL = 11.4021

#-----

# Compare the 95% 1-of-3 upper prediction interval to the California and
# Modified California prediction intervals. Note that the upper prediction
# bound for the Modified California rule is between the bound for the
# 1-of-3 rule bound and the bound for the California rule.

predIntNormSimultaneous(dat, k = 1, m = 3)$interval$limits["UPL"]
#    UPL
#11.4021

predIntNormSimultaneous(dat, m = 3, rule = "CA")$interval$limits["UPL"]
#    UPL
#13.03717

predIntNormSimultaneous(dat, rule = "Modified.CA")$interval$limits["UPL"]
#    UPL
#12.12201

#-----

# Show how the upper bound on an upper 95% simultaneous prediction limit increases
# as the number of future sampling occasions r increases. Here, we'll use the
# 1-of-3 rule.

predIntNormSimultaneous(dat, k = 1, m = 3)$interval$limits["UPL"]
#    UPL
#11.4021

predIntNormSimultaneous(dat, k = 1, m = 3, r = 10)$interval$limits["UPL"]
#    UPL
#13.28234

#-----

# Compare the upper simultaneous prediction limit for the 1-of-3 rule
# based on individual observations versus based on means of order 4.

predIntNormSimultaneous(dat, k = 1, m = 3)$interval$limits["UPL"]
#    UPL
#11.4021

predIntNormSimultaneous(dat, n.mean = 4, k = 1,
  m = 3)$interval$limits["UPL"]
#    UPL
#11.26157

#=====

# Example 19-1 of USEPA (2009, p. 19-17) shows how to compute an
# upper simultaneous prediction limit for the 1-of-3 rule for
# r = 2 future sampling occasions. The data for this example are

```

```

# stored in EPA.09.Ex.19.1.sulfate.df.

# We will pool data from 4 background wells that were sampled on
# a number of different occasions, giving us a sample size of
# n = 25 to use to construct the prediction limit.

# There are 50 compliance wells and we will monitor 10 different
# constituents at each well at each of the r=2 future sampling
# occasions. To determine the confidence level we require for
# the simultaneous prediction interval, USEPA (2009) recommends
# setting the individual Type I Error level at each well to

#  $1 - (1 - \text{SWFPR})^{(1 / (\text{Number of Constituents} * \text{Number of Wells}))}$ 

# which translates to setting the confidence limit to

#  $(1 - \text{SWFPR})^{(1 / (\text{Number of Constituents} * \text{Number of Wells}))}$ 

# where SWFPR = site-wide false positive rate. For this example, we
# will set SWFPR = 0.1. Thus, the confidence level is given by:

nc <- 10
nw <- 50
SWFPR <- 0.1
conf.level <- (1 - SWFPR)^(1 / (nc * nw))

conf.level
#[1] 0.9997893

#-----

# Look at the data:

names(EPA.09.Ex.19.1.sulfate.df)
#[1] "Well"           "Month"           "Day"
#[4] "Year"           "Date"            "Sulfate.mg.per.l"
#[7] "log.Sulfate.mg.per.l"

EPA.09.Ex.19.1.sulfate.df[,
  c("Well", "Date", "Sulfate.mg.per.l", "log.Sulfate.mg.per.l")]

#   Well      Date Sulfate.mg.per.l log.Sulfate.mg.per.l
#1  GW-01 1999-07-08          63.0          4.143135
#2  GW-01 1999-09-12          51.0          3.931826
#3  GW-01 1999-10-16          60.0          4.094345
#4  GW-01 1999-11-02          86.0          4.454347
#5  GW-04 1999-07-09         104.0          4.644391
#6  GW-04 1999-09-14         102.0          4.624973
#7  GW-04 1999-10-12          84.0          4.430817
#8  GW-04 1999-11-15          72.0          4.276666
#9  GW-08 1997-10-12          31.0          3.433987
#10 GW-08 1997-11-16          84.0          4.430817
#11 GW-08 1998-01-28          65.0          4.174387
#12 GW-08 1999-04-20          41.0          3.713572
#13 GW-08 2002-06-04          51.8          3.947390
#14 GW-08 2002-09-16          57.5          4.051785
#15 GW-08 2002-12-02          66.8          4.201703

```

#16	GW-08	2003-03-24	87.1	4.467057
#17	GW-09	1997-10-16	59.0	4.077537
#18	GW-09	1998-01-28	85.0	4.442651
#19	GW-09	1998-04-12	75.0	4.317488
#20	GW-09	1998-07-12	99.0	4.595120
#21	GW-09	2000-01-30	75.8	4.328098
#22	GW-09	2000-04-24	82.5	4.412798
#23	GW-09	2000-10-24	85.5	4.448516
#24	GW-09	2002-12-01	188.0	5.236442
#25	GW-09	2003-03-24	150.0	5.010635

```
# Construct the upper simultaneous prediction limit for the
# 1-of-3 plan based on the log-transformed sulfate data
```

```
log.Sulfate <- EPA.09.Ex.19.1.sulfate.df$log.Sulfate.mg.per.l
```

```
pred.int.list.log <-
  predIntNormSimultaneous(x = log.Sulfate, k = 1, m = 3, r = 2,
    rule = "k.of.m", pi.type = "upper", conf.level = conf.level)
```

```
pred.int.list.log
```

```
#Results of Distribution Parameter Estimation
```

```
#-----
```

```
#
#Assumed Distribution:      Normal
#
#Estimated Parameter(s):   mean = 4.3156194
#                           sd   = 0.3756697
#
#Estimation Method:        mvue
#
#Data:                     log.Sulfate
#
#Sample Size:              25
#
#Prediction Interval Method: exact
#
#Prediction Interval Type:  upper
#
#Confidence Level:         99.97893%
#
#Minimum Number of
#Future Observations
#Interval Should Contain
#(per Sampling Occasion):  1
#
#Total Number of
#Future Observations
#(per Sampling Occasion):  3
#
#Number of Future
#Sampling Occasions:       2
#
#Prediction Interval:      LPL =      -Inf
#                           UPL = 5.072355
```

```
# Now exponentiate the prediction interval to get the limit on
# the original scale

exp(pred.int.list.log$interval$limits["UPL"])
#      UPL
#159.5497

#=====

# Cleanup
#-----

rm(dat, est.list, nc, nw, SWFPR, conf.level, log.Sulfate,
    pred.int.list.log)
```

predIntNormSimultaneousK

Compute the Value of K for a Simultaneous Prediction Interval for a Normal Distribution

Description

Compute the value of K (the multiplier of estimated standard deviation) used to construct a simultaneous prediction interval based on data from a [normal distribution](#). The function `predIntNormSimultaneousK` is called by [predIntNormSimultaneous](#).

Usage

```
predIntNormSimultaneousK(n, df = n - 1, n.mean = 1, k = 1, m = 2, r = 1,
    rule = "k.of.m", delta.over.sigma = 0, pi.type = "upper", conf.level = 0.95,
    K.tol = .Machine$double.eps^0.5, integrate.args.list = NULL)
```

Arguments

- | | |
|--------|---|
| n | a positive integer greater than 2 indicating the sample size upon which the prediction interval is based. |
| df | the degrees of freedom associated with the prediction interval. The default is $df=n-1$. |
| n.mean | positive integer specifying the sample size associated with the future averages. The default value is $n.mean=1$ (i.e., individual observations). Note that all future averages must be based on the same sample size. |
| k | for the k -of- m rule ($rule="k.of.m"$), a positive integer specifying the minimum number of observations (or averages) out of m observations (or averages) (all obtained on one future sampling "occasion") the prediction interval should contain with confidence level <code>conf.level</code> . The default value is $k=1$. This argument is ignored when the argument <code>rule</code> is not equal to <code>"k.of.m"</code> . |
| m | positive integer specifying the maximum number of future observations (or averages) on one future sampling "occasion". The default value is $m=2$, except when $rule="Modified.CA"$, in which case this argument is ignored and m is automatically set equal to 4. |

<code>r</code>	positive integer specifying the number of future sampling “occasions”. The default value is <code>r=1</code> .
<code>rule</code>	character string specifying which rule to use. The possible values are “ <code>k.of.m</code> ” (<i>k</i> -of- <i>m</i> rule; the default), “ <code>CA</code> ” (California rule), and “ <code>Modified.CA</code> ” (modified California rule). See the DETAILS section below for more information.
<code>delta.over.sigma</code>	numeric scalar indicating the ratio Δ/σ . The quantity Δ (delta) denotes the difference between the mean of the population that was sampled to construct the prediction interval, and the mean of the population that will be sampled to produce the future observations. The quantity σ (sigma) denotes the population standard deviation for both populations. See the DETAILS section below for more information. The default value is <code>delta.over.sigma=0</code> .
<code>pi.type</code>	character string indicating what kind of prediction interval to compute. The possible values are <code>pi.type="upper"</code> (the default), and <code>pi.type="lower"</code> .
<code>conf.level</code>	a scalar between 0 and 1 indicating the confidence level of the prediction interval. The default value is <code>conf.level=0.95</code> .
<code>K.tol</code>	numeric scalar indicating the tolerance to use in the nonlinear search algorithm to compute K . The default value is <code>K.tol=Machine\$double.eps^(1/2)</code> . For many applications, the value of K needs to be known only to the second decimal place, in which case setting <code>K.tol=1e-4</code> will speed up computation a bit.
<code>integrate.args.list</code>	a list of arguments to supply to the integrate function. The default value is <code>integrate.args.list=NULL</code> which means that the default values of integrate are used.

Details

What is a Simultaneous Prediction Interval?

A prediction interval for some population is an interval on the real line constructed so that it will contain k future observations from that population with some specified probability $(1 - \alpha)100\%$, where $0 < \alpha < 1$ and k is some pre-specified positive integer. The quantity $(1 - \alpha)100\%$ is called the confidence coefficient or confidence level associated with the prediction interval. The function [predIntNorm](#) computes a standard prediction interval based on a sample from a [normal distribution](#).

The function `predIntNormSimultaneous` computes a simultaneous prediction interval that will contain a certain number of future observations with probability $(1 - \alpha)100\%$ for each of r future sampling “occasions”, where r is some pre-specified positive integer. The quantity r may refer to r distinct future sampling occasions in time, or it may for example refer to sampling at r distinct locations on one future sampling occasion, assuming that the population standard deviation is the same at all of the r distinct locations.

The function [predIntNormSimultaneous](#) computes a simultaneous prediction interval based on one of three possible rules:

- For the *k*-of-*m* rule (`rule="k.of.m"`), at least k of the next m future observations will fall in the prediction interval with probability $(1 - \alpha)100\%$ on each of the r future sampling occasions. If observations are being taken sequentially, for a particular sampling occasion, up to m observations may be taken, but once k of the observations fall within the prediction interval, sampling can stop. Note: When $k = m$ and $r = 1$, the results of `predIntNormSimultaneous` are equivalent to the results of [predIntNorm](#).
- For the California rule (`rule="CA"`), with probability $(1 - \alpha)100\%$, for each of the r future sampling occasions, either the first observation will fall in the prediction interval, or else all of the next $m - 1$ observations will fall in the prediction interval. That is, if the first observation

falls in the prediction interval then sampling can stop. Otherwise, $m - 1$ more observations must be taken.

- For the Modified California rule (rule="Modified.CA"), with probability $(1 - \alpha)100\%$, for each of the r future sampling occasions, either the first observation will fall in the prediction interval, or else at least 2 out of the next 3 observations will fall in the prediction interval. That is, if the first observation falls in the prediction interval then sampling can stop. Otherwise, up to 3 more observations must be taken.

Simultaneous prediction intervals can be extended to using averages (means) in place of single observations (USEPA, 2009, Chapter 19). That is, you can create a simultaneous prediction interval that will contain a specified number of averages (based on which rule you choose) on each of r future sampling occasions, where each average is based on w individual observations. For the functions `predIntNormSimultaneous` and `predIntNormSimultaneousK`, the argument `n.mean` corresponds to w .

The Form of a Prediction Interval

Let $\underline{x} = x_1, x_2, \dots, x_n$ denote a vector of n observations from a [normal distribution](#) with parameters `mean= μ` and `sd= σ` . Also, let w denote the sample size associated with the future averages (i.e., `n.mean= w`). When $w = 1$, each average is really just a single observation, so in the rest of this help file the term "averages" will sometimes replace the phrase "observations or averages".

For a normal distribution, the form of a two-sided $(1 - \alpha)100\%$ simultaneous prediction interval is:

$$[\bar{x} - Ks, \bar{x} + Ks] \quad (1)$$

where \bar{x} denotes the sample mean:

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i \quad (2)$$

s denotes the sample standard deviation:

$$s^2 = \frac{1}{n-1} \sum_{i=1}^n (x_i - \bar{x})^2 \quad (3)$$

and K denotes a constant that depends on the sample size n , the confidence level, the number of future sampling occasions r , and the sample size associated with the future averages, w . Do not confuse the constant K (uppercase K) with the number of future averages k (lowercase k) in the k -of- m rule. The symbol K is used here to be consistent with the notation used for tolerance intervals (see [tolIntNorm](#)).

Similarly, the form of a one-sided lower prediction interval is:

$$[\bar{x} - Ks, \infty] \quad (4)$$

and the form of a one-sided upper prediction interval is:

$$[-\infty, \bar{x} + Ks] \quad (5)$$

Note: For simultaneous prediction intervals, only lower (`pi.type="lower"`) and upper (`pi.type="upper"`) prediction intervals are available.

The derivation of the constant K is explained below.

The Derivation of K for Future Observations

First we will show the derivation based on future observations (i.e., $w = 1$, $n.mean=1$), and then extend the formulas to future averages.

The Derivation of K for the k-of-m Rule (rule="k.of.m")

For the k -of- m rule (rule="k.of.m") with $w = 1$ (i.e., $n.mean=1$), at least k of the next m future observations will fall in the prediction interval with probability $(1 - \alpha)100\%$ on each of the r future sampling occasions. If observations are being taken sequentially, for a particular sampling occasion, up to m observations may be taken, but once k of the observations fall within the prediction interval, sampling can stop. Note: When $k = m$ and $r = 1$, this kind of simultaneous prediction interval becomes the same as a standard prediction interval for the next k observations (see [predIntNorm](#)).

For the case when $r = 1$ future sampling occasion, both Hall and Prairie (1973) and Fertig and Mann (1977) discuss the derivation of K . Davis and McNichols (1987) extend the derivation to the case where r is a positive integer. They show that for a one-sided upper prediction interval ($pi.type="upper"$), the probability p that at least k of the next m future observations will be contained in the interval given in Equation (5) above, for each of r future sampling occasions, is given by:

$$p = \int_0^1 T(\sqrt{n}K; n-1, \sqrt{n}[\Phi^{-1}(v) + \frac{\Delta}{\sigma}]) r [I(v; k, m+1-k)]^{r-1} \left[\frac{v^{k-1}(1-v)^{m-k}}{B(k, m+1-k)} \right] dv \quad (6)$$

where $T(x; \nu, \delta)$ denotes the cdf of the [non-central Student's t-distribution](#) with parameters $df=\nu$ and $ncp=\delta$ evaluated at x ; $\Phi(x)$ denotes the cdf of the standard [normal distribution](#) evaluated at x ; $I(x; \nu, \omega)$ denotes the cdf of the [beta distribution](#) with parameters $shape1=\nu$ and $shape2=\omega$; and $B(\nu, \omega)$ denotes the value of the [beta function](#) with parameters $a=\nu$ and $b=\omega$.

The quantity Δ (upper case delta) denotes the difference between the mean of the population that was sampled to construct the prediction interval, and the mean of the population that will be sampled to produce the future observations. The quantity σ (sigma) denotes the population standard deviation of both of these populations. Usually you assume $\Delta = 0$ unless you are interested in computing the power of the rule to detect a change in means between the populations (see [predIntNormSimultaneousTestPower](#)).

For given values of the confidence level (p), sample size (n), minimum number of future observations to be contained in the interval per sampling occasion (k), number of future observations per sampling occasion (m), and number of future sampling occasions (r), Equation (6) can be solved for K . The function `predIntNormSimultaneousK` uses the R function `nlfminb` to solve Equation (6) for K .

When $pi.type="lower"$, the same value of K is used as when $pi.type="upper"$, but Equation (4) is used to construct the prediction interval.

The Derivation of K for the California Rule (rule="CA")

For the California rule (rule="CA"), with probability $(1 - \alpha)100\%$, for each of the r future sampling occasions, either the first observation will fall in the prediction interval, or else all of the next $m - 1$ observations will fall in the prediction interval. That is, if the first observation falls in the prediction interval then sampling can stop. Otherwise, $m - 1$ more observations must be taken.

The formula for K is the same as for the k -of- m rule, except that Equation (6) becomes the following (Davis, 1998b):

$$p = \int_0^1 T(\sqrt{n}K; n-1, \sqrt{n}[\Phi^{-1}(v) + \frac{\Delta}{\sigma}]) r \{v[1+v^{m-2}(1-v)]\}^{r-1} [1+v^{m-2}(m-1-mv)] dv \quad (7)$$

The Derivation of K for the Modified California Rule (rule="Modified.CA")

For the Modified California rule (rule="Modified.CA"), with probability $(1 - \alpha)100\%$, for each of the r future sampling occasions, either the first observation will fall in the prediction interval, or else at least 2 out of the next 3 observations will fall in the prediction interval. That is, if the first observation falls in the prediction interval then sampling can stop. Otherwise, up to 3 more observations must be taken.

The formula for K is the same as for the k -of- m rule, except that Equation (6) becomes the following (Davis, 1998b):

$$p = \int_0^1 T(\sqrt{n}K; n-1, \sqrt{n}[\Phi^{-1}(v) + \frac{\Delta}{\sigma}]) r\{v[1+v(3-v[5-2v])]\}^{r-1} \{1+v[6-v(15-8v)]\} dv \quad (8)$$

The Derivation of K for Future Means

For each of the above rules, if we are interested in using averages instead of single observations, with $w \geq 1$ (i.e., $n \cdot \text{mean} \geq 1$), the first term in the integral in Equations (6)-(8) that involves the cdf of the [non-central Student's t-distribution](#) becomes:

$$T(\sqrt{n}K; n-1, \frac{\sqrt{n}}{\sqrt{w}}[\Phi^{-1}(v) + \frac{\sqrt{w}\Delta}{\sigma}]) \quad (9)$$

Value

A numeric scalar equal to K , the multiplier of estimated standard deviation that is used to construct the simultaneous prediction interval.

Note

Motivation

Prediction and tolerance intervals have long been applied to quality control and life testing problems (Hahn, 1970b,c; Hahn and Nelson, 1973). In the context of environmental statistics, prediction intervals are useful for analyzing data from groundwater detection monitoring programs at hazardous and solid waste facilities.

One of the main statistical problems that plague groundwater monitoring programs at hazardous and solid waste facilities is the requirement of testing several wells and several constituents at each well on each sampling occasion. This is an obvious multiple comparisons problem, and the naive approach of using a standard t-test at a conventional α -level (e.g., 0.05 or 0.01) for each test leads to a very high probability of at least one significant result on each sampling occasion, when in fact no contamination has occurred. This problem was pointed out years ago by Millard (1987) and others.

Davis and McNichols (1987) proposed simultaneous prediction intervals as a way of controlling the facility-wide false positive rate (FWFPR) while maintaining adequate power to detect contamination in the groundwater. Because of the ubiquitous presence of spatial variability, it is usually best to use simultaneous prediction intervals at each well (Davis, 1998a). That is, by constructing prediction intervals based on background (pre-landfill) data on each well, and comparing future observations at a well to the prediction interval for that particular well. In each of these cases, the individual α -level at each well is equal to the FWFPR divided by the product of the number of wells and constituents.

Often, observations at downgradient wells are not available prior to the construction and operation of the landfill. In this case, upgradient well data can be combined to create a background prediction interval, and observations at each downgradient well can be compared to this prediction interval. If spatial variability is present and a major source of variation, however, this method is not really valid (Davis, 1994; Davis, 1998a).

Chapter 19 of USEPA (2009) contains an extensive discussion of using the 1-of- m rule and the Modified California rule.

Chapters 1 and 3 of Gibbons et al. (2009) discuss simultaneous prediction intervals for the normal and lognormal distributions, respectively.

The k-of-m Rule

For the k -of- m rule, Davis and McNichols (1987) give tables with “optimal” choices of k (in terms of best power for a given overall confidence level) for selected values of m , r , and n . They found that the optimal ratios of k to m (i.e., k/m) are generally small, in the range of 15-50%.

The California Rule

The California rule was mandated in that state for groundwater monitoring at waste disposal facilities when resampling verification is part of the statistical program (Barclay’s Code of California Regulations, 1991). The California code mandates a “California” rule with $m \geq 3$. The motivation for this rule may have been a desire to have a majority of the observations in bounds (Davis, 1998a). For example, for a k -of- m rule with $k = 1$ and $m = 3$, a monitoring location will pass if the first observation is out of bounds, the second resample is out of bounds, but the last resample is in bounds, so that 2 out of 3 observations are out of bounds. For the California rule with $m = 3$, either the first observation must be in bounds, or the next 2 observations must be in bounds in order for the monitoring location to pass.

Davis (1998a) states that if the FWFPR is kept constant, then the California rule offers little increased power compared to the k -of- m rule, and can actually decrease the power of detecting contamination.

The Modified California Rule

The Modified California Rule was proposed as a compromise between a 1-of- m rule and the California rule. For a given FWFPR, the Modified California rule achieves better power than the California rule, and still requires at least as many observations in bounds as out of bounds, unlike a 1-of- m rule.

Different Notations Between Different References

For the k -of- m rule described in this help file, both Davis and McNichols (1987) and USEPA (2009, Chapter 19) use the variable p instead of k to represent the minimum number of future observations the interval should contain on each of the r sampling occasions.

Gibbons et al. (2009, Chapter 1) presents extensive lists of the value of K for both k -of- m rules and California rules. Gibbons et al.’s notation reverses the meaning of k and r compared to the notation used in this help file. That is, in Gibbons et al.’s notation, k represents the number of future sampling occasions or monitoring wells, and r represents the minimum number of observations the interval should contain on each sampling occasion.

USEPA (2009, Chapter 19) uses p in place of k .

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Barclay’s California Code of Regulations. (1991). Title 22, Section 66264.97 [concerning hazardous waste facilities] and Title 23, Section 2550.7(e)(8) [concerning solid waste facilities]. Barclay’s Law Publishers, San Francisco, CA.

- Davis, C.B. (1998a). *Ground-Water Statistics & Regulations: Principles, Progress and Problems*. Second Edition. Environmetrics & Statistics Limited, Henderson, NV.
- Davis, C.B. (1998b). Personal Communication, September 3, 1998.
- Davis, C.B., and R.J. McNichols. (1987). One-sided Intervals for at Least p of m Observations from a Normal Population on Each of r Future Occasions. *Technometrics* **29**, 359–370.
- Fertig, K.W., and N.R. Mann. (1977). One-Sided Prediction Intervals for at Least p Out of m Future Observations From a Normal Population. *Technometrics* **19**, 167–177.
- Gibbons, R.D., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*, Second Edition. John Wiley & Sons, Hoboken.
- Hahn, G.J. (1969). Factors for Calculating Two-Sided Prediction Intervals for Samples from a Normal Distribution. *Journal of the American Statistical Association* **64**(327), 878–898.
- Hahn, G.J. (1970a). Additional Factors for Calculating Prediction Intervals for Samples from a Normal Distribution. *Journal of the American Statistical Association* **65**(332), 1668–1676.
- Hahn, G.J. (1970b). Statistical Intervals for a Normal Population, Part I: Tables, Examples and Applications. *Journal of Quality Technology* **2**(3), 115–125.
- Hahn, G.J. (1970c). Statistical Intervals for a Normal Population, Part II: Formulas, Assumptions, Some Derivations. *Journal of Quality Technology* **2**(4), 195–206.
- Hahn, G.J., and W.Q. Meeker. (1991). *Statistical Intervals: A Guide for Practitioners*. John Wiley and Sons, New York.
- Hahn, G., and W. Nelson. (1973). A Survey of Prediction Intervals and Their Applications. *Journal of Quality Technology* **5**, 178–188.
- Hall, I.J., and R.R. Prairie. (1973). One-Sided Prediction Intervals to Contain at Least m Out of k Future Observations. *Technometrics* **15**, 897–914.
- Millard, S.P. (1987). Environmental Monitoring, Statistics, and the Law: Room for Improvement (with Comment). *The American Statistician* **41**(4), 249–259.
- Millard, S.P., and Neerchal, N.K. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, Florida.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.
- USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

See Also

[predIntNormSimultaneous](#), [predIntNormSimultaneousTestPower](#), [predIntNorm](#), [predIntNormK](#), [predIntLnormSimultaneous](#), [tolIntNorm](#), [Normal](#), [estimate.object](#), [enorm](#)

Examples

```
# Compute the value of K for an upper 95% simultaneous prediction
# interval to contain at least 1 out of the next 3 observations
# given a background sample size of n=8.

predIntNormSimultaneousK(n = 8, k = 1, m = 3)
#[1] 0.5123091
```

```

#-----

# Compare the value of K for a 95% 1-of-3 upper prediction interval to
# the value for the California and Modified California rules.
# Note that the value of K for the Modified California rule is between
# the value of K for the 1-of-3 rule and the California rule.

predIntNormSimultaneousK(n = 8, k = 1, m = 3)
#[1] 0.5123091

predIntNormSimultaneousK(n = 8, m = 3, rule = "CA")
#[1] 1.252077

predIntNormSimultaneousK(n = 8, rule = "Modified.CA")
#[1] 0.8380233

#-----

# Show how the value of K for an upper 95% simultaneous prediction
# limit increases as the number of future sampling occasions r increases.
# Here, we'll use the 1-of-3 rule.

predIntNormSimultaneousK(n = 8, k = 1, m = 3)
#[1] 0.5123091

predIntNormSimultaneousK(n = 8, k = 1, m = 3, r = 10)
#[1] 1.363002

#=====

# Example 19-1 of USEPA (2009, p. 19-17) shows how to compute an
# upper simultaneous prediction limit for the 1-of-3 rule for
# r = 2 future sampling occasions. The data for this example are
# stored in EPA.09.Ex.19.1.sulfate.df.

# We will pool data from 4 background wells that were sampled on
# a number of different occasions, giving us a sample size of
# n = 25 to use to construct the prediction limit.

# There are 50 compliance wells and we will monitor 10 different
# constituents at each well at each of the r=2 future sampling
# occasions. To determine the confidence level we require for
# the simultaneous prediction interval, USEPA (2009) recommends
# setting the individual Type I Error level at each well to

#  $1 - (1 - \text{SWFPR})^{(1 / (\text{Number of Constituents} * \text{Number of Wells}))}$ 

# which translates to setting the confidence limit to

#  $(1 - \text{SWFPR})^{(1 / (\text{Number of Constituents} * \text{Number of Wells}))}$ 

# where SWFPR = site-wide false positive rate. For this example, we
# will set SWFPR = 0.1. Thus, the confidence level is given by:

nc <- 10
nw <- 50

```

```

SWFPR <- 0.1
conf.level <- (1 - SWFPR)^(1 / (nc * nw))

conf.level
#[1] 0.9997893

#-----

# Compute the value of K for the upper simultaneous prediction
# limit for the 1-of-3 plan.

predIntNormSimultaneousK(n = 25, k = 1, m = 3, r = 2,
  rule = "k.of.m", pi.type = "upper", conf.level = conf.level)
#[1] 2.014365

#=====

# Cleanup
#-----

rm(nc, nw, SWFPR, conf.level)

```

predIntNpar

Nonparametric Prediction Interval for a Continuous Distribution

Description

Construct a nonparametric prediction interval to contain at least k out of the next m future observations with probability $(1 - \alpha)100\%$ for a continuous distribution.

Usage

```

predIntNpar(x, k = m, m = 1, lp1.rank = 1, n.plus.one.minus.up1.rank = 1,
  lb = -Inf, ub = Inf, pi.type = "two-sided")

```

Arguments

- | | |
|-----------------------|---|
| <code>x</code> | a numeric vector of observations. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed. |
| <code>k</code> | positive integer specifying the minimum number of future observations out of m that should be contained in the prediction interval. The default value is $k=m$. |
| <code>m</code> | positive integer specifying the number of future observations. The default value is $m=1$. |
| <code>lp1.rank</code> | positive integer indicating the rank of the order statistic to use for the lower bound of the prediction interval. If <code>pi.type="two-sided"</code> or <code>pi.type="lower"</code> , the default value is <code>lp1.rank=1</code> (implying the minimum value of x is used as the lower bound of the prediction interval). If <code>pi.type="upper"</code> , this argument is set equal to 0 and the value of <code>lb</code> is used as the lower bound of the tolerance interval. |

<code>n.plus.one.minus.upl.rank</code>	positive integer related to the rank of the order statistic to use for the upper bound of the prediction interval. Set $v = n.plus.one.minus.upl.rank$, let w denote the rank of the order statistic used as the upper bound (<code>upl.rank</code>), and let n denote the sample size. The relationship between v and w is given by: $v = n + 1 - w$. If <code>pi.type="two-sided"</code> or <code>pi.type="upper"</code> , the default value is <code>n.plus.one.minus.upl.rank=1</code> (so $v = n$, implying the maximum value of x is used as the upper bound of the prediction interval). If <code>pi.type="lower"</code> , this argument is set equal to 0 and the value of <code>ub</code> is used as the upper bound of the prediction interval.
<code>lb, ub</code>	scalars indicating lower and upper bounds on the distribution. By default, <code>lb=-Inf</code> and <code>ub=Inf</code> . If you are constructing a prediction interval for a distribution that you know has a lower bound other than <code>-Inf</code> (e.g., 0), set <code>lb</code> to this value. Similarly, if you know the distribution has an upper bound other than <code>Inf</code> , set <code>ub</code> to this value. The argument <code>lb</code> is ignored if <code>pi.type="two-sided"</code> or <code>pi.type="lower"</code> . The argument <code>ub</code> is ignored if <code>pi.type="two-sided"</code> or <code>pi.type="upper"</code> .
<code>pi.type</code>	character string indicating what kind of prediction interval to compute. The possible values are "two-sided" (the default), "lower", and "upper".

Details

What is a Nonparametric Prediction Interval?

A nonparametric prediction interval for some population is an interval on the real line constructed so that it will contain at least k of m future observations from that population with some specified probability $(1 - \alpha)100\%$, where $0 < \alpha < 1$ and k and m are pre-specified positive integer where $k \leq m$. The quantity $(1 - \alpha)100\%$ is called the confidence coefficient or confidence level associated with the prediction interval.

The Form of a Nonparametric Prediction Interval

Let $\underline{x} = x_1, x_2, \dots, x_n$ denote a vector of n independent observations from some continuous distribution, and let $x_{(i)}$ denote the i 'th order statistics in \underline{x} . A two-sided nonparametric prediction interval is constructed as:

$$[x_{(u)}, x_{(v)}] \quad (1)$$

where u and v are positive integers between 1 and n , and $u < v$. That is, u denotes the rank of the lower prediction limit, and v denotes the rank of the upper prediction limit. To make it easier to write some equations later on, we can also write the prediction interval (1) in a slightly different way as:

$$[x_{(u)}, x_{(n+1-w)}] \quad (2)$$

where

$$w = n + 1 - v \quad (3)$$

so that w is a positive integer between 1 and $n - 1$, and $u < n + 1 - w$. In terms of the arguments to the function `predIntNpar`, the argument `lpl.rank` corresponds to u , and the argument `n.plus.one.minus.upl.rank` corresponds to w .

If we allow $u = 0$ and $w = 0$ and define lower and upper bounds as:

$$x_{(0)} = lb \quad (4)$$

$$x_{(n+1)} = ub \quad (5)$$

then Equation (2) above can also represent a one-sided lower or one-sided upper prediction interval as well. That is, a one-sided lower nonparametric prediction interval is constructed as:

$$[x_{(u)}, x_{(n+1)}] = [x_{(u)}, ub] \quad (6)$$

and a one-sided upper nonparametric prediction interval is constructed as:

$$[x_{(0)}, x_{(n+1-w)}] = [lb, x_{(n+1-w)}] \quad (7)$$

Usually, $lb = -\infty$ or $lb = 0$ and $ub = \infty$.

Constructing Nonparametric Prediction Intervals for Future Observations

Danziger and Davis (1964) show that the probability that at least k out of the next m observations will fall in the interval defined in Equation (2) is given by:

$$(1 - \alpha) = \left[\sum_{i=k}^m \binom{m-i+u+w-1}{m-i} \binom{i+n-u-w}{i} \right] / \binom{n+m}{m} \quad (8)$$

(Note that computing a nonparametric prediction interval for the case $k = m = 1$ is equivalent to computing a nonparametric β -expectation tolerance interval with coverage $(1 - \alpha)100\%$; see [tolIntNpar](#)).

The Special Case of Using the Minimum and the Maximum

Setting $u = w = 1$ implies using the smallest and largest observed values as the prediction limits. In this case, it can be shown that the probability that at least k out of the next m observations will fall in the interval

$$[x_{(1)}, x_{(n)}] \quad (9)$$

is given by:

$$(1 - \alpha) = \left[\sum_{i=k}^m (m-i-1) \binom{n+i-2}{i} \right] / \binom{n+m}{m} \quad (10)$$

Setting $k = m$ in Equation (10), the probability that all of the next m observations will fall in the interval defined in Equation (9) is given by:

$$(1 - \alpha) = \frac{n(n-1)}{(n+m)(n+m-1)} \quad (11)$$

For one-sided prediction limits, the probability that all m future observations will fall below $x_{(n)}$ (upper prediction limit; `pi.type="upper"`) and the probability that all m future observations will fall above $x_{(1)}$ (lower prediction limit; `pi.type="lower"`) are both given by:

$$(1 - \alpha) = \frac{n}{n+m} \quad (12)$$

Constructing Nonparametric Prediction Intervals for Future Medians

To construct a nonparametric prediction interval for a future median based on s future observations, where s is odd, note that this is equivalent to constructing a nonparametric prediction interval that must hold at least $k = (s+1)/2$ of the next $m = s$ future observations.

Value

a list of class "estimate" containing the prediction interval and other information. See the help file for [estimate.object](#) for details.

Note

Prediction and tolerance intervals have long been applied to quality control and life testing problems (Hahn, 1970b,c; Hahn and Nelson, 1973; Krishnamoorthy and Mathew, 2009). In the context of environmental statistics, prediction intervals are useful for analyzing data from groundwater detection monitoring programs at hazardous and solid waste facilities (e.g., Gibbons et al., 2009; Millard and Neerchal, 2001; USEPA, 2009).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Danziger, L., and S. Davis. (1964). Tables of Distribution-Free Tolerance Limits. *Annals of Mathematical Statistics* **35**(5), 1361–1365.
- Davis, C.B. (1994). Environmental Regulatory Statistics. In Patil, G.P., and C.R. Rao, eds., *Handbook of Statistics, Vol. 12: Environmental Statistics*. North-Holland, Amsterdam, a division of Elsevier, New York, NY, Chapter 26, 817–865.
- Davis, C.B., and R.J. McNichols. (1987). One-sided Intervals for at Least p of m Observations from a Normal Population on Each of r Future Occasions. *Technometrics* **29**, 359–370.
- Davis, C.B., and R.J. McNichols. (1994a). Ground Water Monitoring Statistics Update: Part I: Progress Since 1988. *Ground Water Monitoring and Remediation* **14**(4), 148–158.
- Davis, C.B., and R.J. McNichols. (1994b). Ground Water Monitoring Statistics Update: Part II: Nonparametric Prediction Limits. *Ground Water Monitoring and Remediation* **14**(4), 159–175.
- Davis, C.B., and R.J. McNichols. (1999). Simultaneous Nonparametric Prediction Limits (with Discussion). *Technometrics* **41**(2), 89–112.
- Gibbons, R.D. (1987a). Statistical Prediction Intervals for the Evaluation of Ground-Water Quality. *Ground Water* **25**, 455–465.
- Gibbons, R.D. (1991b). Statistical Tolerance Limits for Ground-Water Monitoring. *Ground Water* **29**, 563–570.
- Gibbons, R.D., and J. Baker. (1991). The Properties of Various Statistical Prediction Intervals for Ground-Water Detection Monitoring. *Journal of Environmental Science and Health* **A26**(4), 535–553.
- Gibbons, R.D., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*, Second Edition. John Wiley & Sons, Hoboken.
- Hahn, G.J., and W.Q. Meeker. (1991). *Statistical Intervals: A Guide for Practitioners*. John Wiley and Sons, New York, 392pp.
- Hahn, G., and W. Nelson. (1973). A Survey of Prediction Intervals and Their Applications. *Journal of Quality Technology* **5**, 178–188.
- Hall, I.J., R.R. Prairie, and C.K. Motlagh. (1975). Non-Parametric Prediction Intervals. *Journal of Quality Technology* **7**(3), 109–114.
- Millard, S.P., and Neerchal, N.K. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, Florida.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

See Also

[estimate.object](#), [predIntNparN](#), [predIntNparConfLevel](#), [plotPredIntNparDesign](#).

Examples

```
# Generate 20 observations from a lognormal mixture distribution with
# parameters mean1=1, cv1=0.5, mean2=5, cv2=1, and p.mix=0.1. Use
# predIntNpar to construct a two-sided prediction interval using the
# minimum and maximum observed values. Note that the associated confidence
# level is 90%. A larger sample size is required to obtain a larger
# confidence level (see the help file for predIntNparN).
# (Note: the call to set.seed simply allows you to reproduce this example.)
```

```
set.seed(250)
dat <- rlnormMixAlt(n = 20, mean1 = 1, cv1 = 0.5,
  mean2 = 5, cv2 = 1, p.mix = 0.1)
```

```
predIntNpar(dat)
```

```
#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:      None
#
#Data:                     dat
#
#Sample Size:              20
#
#Prediction Interval Method: Exact
#
#Prediction Interval Type:  two-sided
#
#Confidence Level:         90.47619%
#
#Prediction Limit Rank(s):  1 20
#
#Number of Future Observations: 1
#
#Prediction Interval:      LPL = 0.3647875
#                          UPL = 1.8173115
```

```
#-----
```

```
# Repeat the above example, but specify m=5 future observations should be
# contained in the prediction interval. Note that the confidence level is
# now only 63%.
```

```
predIntNpar(dat, m = 5)
```

```
#Results of Distribution Parameter Estimation
#-----
#
```



```

#Assumed Distribution:      None
#
#Data:                     dat
#
#Sample Size:              20
#
#Prediction Interval Method: Exact
#
#Prediction Interval Type:  two-sided
#
#Confidence Level:         63.33333%
#
#Prediction Limit Rank(s):  1 20
#
#Number of Future Observations: 5
#
#Prediction Interval:       LPL = 0.3647875
#                           UPL = 1.8173115

#-----

# Repeat the above example, but specify that a minimum of k=3 observations
# out of a total of m=5 future observations should be contained in the
# prediction interval. Note that the confidence level is now 98%.

predIntNpar(dat, k = 3, m = 5)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:      None
#
#Data:                     dat
#
#Sample Size:              20
#
#Prediction Interval Method: Exact
#
#Prediction Interval Type:  two-sided
#
#Confidence Level:         98.37945%
#
#Prediction Limit Rank(s):  1 20
#
#Minimum Number of
#Future Observations
#Interval Should Contain:   3
#
#Total Number of
#Future Observations:      5
#
#Prediction Interval:       LPL = 0.3647875
#                           UPL = 1.8173115

#=====

# Example 18-3 of USEPA (2009, p.18-19) shows how to construct

```

```
# a one-sided upper nonparametric prediction interval for the next
# 4 future observations of trichloroethylene (TCE) at a downgradient well.
# The data for this example are stored in EPA.09.Ex.18.3.TCE.df.
# There are 6 monthly observations of TCE (ppb) at 3 background wells,
# and 4 montly observations of TCE at a compliance well.
```

```
# Look at the data
#-----
```

```
EPA.09.Ex.18.3.TCE.df
```

```
#   Month Well Well.type TCE.ppb.orig TCE.ppb Censored
#1      1 BW-1 Background          <5    5.0      TRUE
#2      2 BW-1 Background          <5    5.0      TRUE
#3      3 BW-1 Background           8    8.0     FALSE
#...
#22     4 CW-4 Compliance          <5    5.0      TRUE
#23     5 CW-4 Compliance           8    8.0     FALSE
#24     6 CW-4 Compliance          14   14.0     FALSE
```

```
longToWide(EPA.09.Ex.18.3.TCE.df, "TCE.ppb.orig", "Month", "Well",
  paste.row.name = TRUE)
```

```
#           BW-1 BW-2 BW-3 CW-4
#Month.1    <5    7   <5
#Month.2    <5  6.5   <5
#Month.3     8   <5 10.5  7.5
#Month.4    <5    6   <5   <5
#Month.5     9   12   <5    8
#Month.6    10   <5    9   14
```

```
# Construct the prediction limit based on the background well data
# using the maximum value as the upper prediction limit.
# Note that since all censored observations are censored at one
# censoring level and the censoring level is less than all of the
# uncensored observations, we can just supply the censoring level
# to predIntNpar.
```

```
#-----
```

```
with(EPA.09.Ex.18.3.TCE.df,
  predIntNpar(TCE.ppb[Well.type == "Background"],
    m = 4, pi.type = "upper", lb = 0))
```

```
#Results of Distribution Parameter Estimation
```

```
#-----
```

```
#
#Assumed Distribution:          None
#
#Data:                        TCE.ppb[Well.type == "Background"]
#
#Sample Size:                  18
#
#Prediction Interval Method:    Exact
#
#Prediction Interval Type:      upper
```

```

#
#Confidence Level:          81.81818%
#
#Prediction Limit Rank(s):    18
#
#Number of Future Observations:  4
#
#Prediction Interval:         LPL = 0
#                               UPL = 12

# Since the value of 14 ppb for Month 6 at the compliance well exceeds
# the upper prediction limit of 12, we might conclude that there is
# statistically significant evidence of an increase over background
# at CW-4. However, the confidence level associated with this
# prediction limit is about 82%, which implies a Type I error level of
# 18%. This means there is nearly a one in five chance of a false positive.
# Only additional background data and/or use of a retesting strategy
# (see predIntNparSimultaneous) would lower the false positive rate.

#=====

# Example 18-4 of USEPA (2009, p.18-19) shows how to construct
# a one-sided upper nonparametric prediction interval for the next
# median of order 3 of xylene at a downgradient well.
# The data for this example are stored in EPA.09.Ex.18.4.xylene.df.
# There are 8 monthly observations of xylene (ppb) at 3 background wells,
# and 3 montly observations of TCE at a compliance well.

# Look at the data
#-----

EPA.09.Ex.18.4.xylene.df

#   Month  Well  Well.type  Xylene.ppb.orig  Xylene.ppb  Censored
#1      1  Well.1 Background                <5         5.0      TRUE
#2      2  Well.1 Background                <5         5.0      TRUE
#3      3  Well.1 Background                7.5         7.5     FALSE
#...
#30     6  Well.4 Compliance                <5         5.0      TRUE
#31     7  Well.4 Compliance                7.8         7.8     FALSE
#32     8  Well.4 Compliance               10.4        10.4     FALSE

longToWide(EPA.09.Ex.18.4.xylene.df, "Xylene.ppb.orig", "Month", "Well",
  paste.row.name = TRUE)

#      Well.1 Well.2 Well.3 Well.4
#Month.1    <5    9.2    <5
#Month.2    <5     <5    5.4
#Month.3    7.5     <5    6.7
#Month.4    <5    6.1    <5
#Month.5    <5     8     <5
#Month.6    <5    5.9    <5    <5
#Month.7    6.4     <5    <5    7.8
#Month.8     6     <5    <5   10.4

# Construct the prediction limit based on the background well data
# using the maximum value as the upper prediction limit.

```

```

# Note that since all censored observations are censored at one
# censoring level and the censoring level is less than all of the
# uncensored observations, we can just supply the censoring level
# to predIntNpar.
#
# To compute a prediction interval for a median of order 3 (i.e.,
# a median based on 3 observations), this is equivalent to
# constructing a nonparametric prediction interval that must hold
# at least 2 of the next 3 future observations.
#-----

with(EPA.09.Ex.18.4.xylene.df,
  predIntNpar(Xylene.ppb[Well.type == "Background"],
    k = 2, m = 3, pi.type = "upper", lb = 0))

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          None
#
#Data:                         Xylene.ppb[Well.type == "Background"]
#
#Sample Size:                  24
#
#Prediction Interval Method:    Exact
#
#Prediction Interval Type:      upper
#
#Confidence Level:              99.1453%
#
#Prediction Limit Rank(s):      24
#
#Minimum Number of
#Future Observations
#Interval Should Contain:       2
#
#Total Number of
#Future Observations:           3
#
#Prediction Interval:           LPL = 0.0
#                               UPL = 9.2

# The Month 8 observation at the Compliance well is 10.4 ppb of Xylene,
# which is greater than the upper prediction limit of 9.2 ppb, so
# conclude there is evidence of contamination at the
# 100% - 99% = 1% Type I Error Level

#=====

# Cleanup
#-----

rm(dat)

```

predIntNparSimultaneous

*Nonparametric Simultaneous Prediction Interval for a Continuous Distribution***Description**

Construct a nonparametric simultaneous prediction interval for the next r sampling “occasions” based on one of three possible rules: k -of- m , California, or Modified California. The simultaneous prediction interval assumes the observations from from a continuous distribution.

Usage

```
predIntNparSimultaneous(x, n.median = 1, k = 1, m = 2, r = 1, rule = "k.of.m",
  lpl.rank = 1, n.plus.one.minus.upl.rank = 1, lb = -Inf, ub = Inf,
  pi.type = "upper", integrate.args.list = NULL)
```

Arguments

<code>x</code>	a numeric vector of observations. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
<code>n.median</code>	positive odd integer specifying the sample size associated with the future medians. The default value is <code>n.median=1</code> (i.e., individual observations). Note that all future medians must be based on the same sample size.
<code>k</code>	for the k -of- m rule (<code>rule="k.of.m"</code>), a positive integer specifying the minimum number of observations (or medians) out of m observations (or medians) (all obtained on one future sampling “occassion”) the prediction interval should contain. The default value is <code>k=1</code> . This argument is ignored when the argument rule is not equal to “ <code>k.of.m</code> ”.
<code>m</code>	positive integer specifying the maximum number of future observations (or medians) on one future sampling “occasion”. The default value is <code>m=2</code> , except when <code>rule="Modified.CA"</code> , in which case this argument is ignored and <code>m</code> is automatically set equal to 4.
<code>r</code>	positive integer specifying the number of future sampling “occasions”. The default value is <code>r=1</code> .
<code>rule</code>	character string specifying which rule to use. The possible values are “ <code>k.of.m</code> ” (k -of- m rule; the default), “ <code>CA</code> ” (California rule), and “ <code>Modified.CA</code> ” (modified California rule). See the DETAILS section below for more information.
<code>lpl.rank</code>	positive integer indicating the rank of the order statistic to use for the lower bound of the prediction interval. When <code>pi.type="lower"</code> , the default value is <code>lpl.rank=1</code> (implying the minimum value of <code>x</code> is used as the lower bound of the prediction interval). When <code>pi.type="upper"</code> , the argument <code>lpl.rank</code> is set equal to 0 and the value of <code>lb</code> is used as the lower bound of the tolerance interval.
<code>n.plus.one.minus.upl.rank</code>	positive integer related to the rank of the order statistic to use for the upper bound of the prediction interval. Set $v = n.plus.one.minus.upl.rank$, let w denote the rank of the order statistic used as the upper bound (<code>upl.rank</code>), and let n denote the sample size. The relationship between v and w is given by: $v = n + 1 - w$. When <code>pi.type="upper"</code> , the default value is <code>n.plus.one.minus.upl.rank=1</code> (so $v = n$, implying the maximum value of <code>x</code> is used as the upper bound of the

	prediction interval). When <code>pi.type="lower"</code> , the argument <code>n.plus.one.minus.upl.rank</code> is set equal to 0 and the value of <code>ub</code> is used as the upper bound of the prediction interval.
<code>lb, ub</code>	scalars indicating lower and upper bounds on the distribution. By default, <code>lb=-Inf</code> and <code>ub=Inf</code> . If you are constructing a prediction interval for a distribution that you know has a lower bound other than <code>-Inf</code> (e.g., 0), set <code>lb</code> to this value. Similarly, if you know the distribution has an upper bound other than <code>Inf</code> , set <code>ub</code> to this value. The argument <code>lb</code> is ignored if <code>pi.type="two-sided"</code> or <code>pi.type="lower"</code> . The argument <code>ub</code> is ignored if <code>pi.type="two-sided"</code> or <code>pi.type="upper"</code> .
<code>pi.type</code>	character string indicating what kind of prediction interval to compute. The possible values are "upper" (the default) and "lower".
<code>integrate.args.list</code>	a list of arguments to supply to the <code>integrate</code> function. The default value is <code>integrate.args.list=NULL</code> which means that the default values of <code>integrate</code> are used.

Details

What is a Nonparametric Simultaneous Prediction Interval?

A nonparametric prediction interval for some population is an interval on the real line constructed so that it will contain at least k of m future observations from that population with some specified probability $(1 - \alpha)100\%$, where $0 < \alpha < 1$ and k and m are some pre-specified positive integers and $k \leq m$. The quantity $(1 - \alpha)100\%$ is called the confidence coefficient or confidence level associated with the prediction interval. The function `predIntNpar` computes a standard nonparametric prediction interval.

The function `predIntNparSimultaneous` computes a nonparametric simultaneous prediction interval that will contain a certain number of future observations with probability $(1 - \alpha)100\%$ for each of r future sampling "occasions", where r is some pre-specified positive integer. The quantity r may refer to r distinct future sampling occasions in time, or it may for example refer to sampling at r distinct locations on one future sampling occasion, assuming that the population standard deviation is the same at all of the r distinct locations.

The function `predIntNparSimultaneous` computes a nonparametric simultaneous prediction interval based on one of three possible rules:

- For the k -of- m rule (`rule="k.of.m"`), at least k of the next m future observations will fall in the prediction interval with probability $(1 - \alpha)100\%$ on each of the r future sampling occasions. If observations are being taken sequentially, for a particular sampling occasion, up to m observations may be taken, but once k of the observations fall within the prediction interval, sampling can stop. Note: For this rule, when $r = 1$, the results of `predIntNparSimultaneous` are equivalent to the results of `predIntNpar`.
- For the California rule (`rule="CA"`), with probability $(1 - \alpha)100\%$, for each of the r future sampling occasions, either the first observation will fall in the prediction interval, or else all of the next $m - 1$ observations will fall in the prediction interval. That is, if the first observation falls in the prediction interval then sampling can stop. Otherwise, $m - 1$ more observations must be taken.
- For the Modified California rule (`rule="Modified.CA"`), with probability $(1 - \alpha)100\%$, for each of the r future sampling occasions, either the first observation will fall in the prediction interval, or else at least 2 out of the next 3 observations will fall in the prediction interval. That is, if the first observation falls in the prediction interval then sampling can stop. Otherwise, up to 3 more observations must be taken.

Nonparametric simultaneous prediction intervals can be extended to using medians in place of single observations (USEPA, 2009, Chapter 19). That is, you can create a nonparametric simultaneous prediction interval that will contain a specified number of medians (based on which rule you choose) on each of r future sampling occasions, where each median is based on b individual observations. For the function `predIntNparSimultaneous`, the argument `n.median` corresponds to b .

The Form of a Nonparametric Prediction Interval

Let $\underline{x} = x_1, x_2, \dots, x_n$ denote a vector of n independent observations from some continuous distribution, and let $x_{(i)}$ denote the i 'th order statistics in \underline{x} . A two-sided nonparametric prediction interval is constructed as:

$$[x_{(u)}, x_{(v)}] \quad (1)$$

where u and v are positive integers between 1 and n , and $u < v$. That is, u denotes the rank of the lower prediction limit, and v denotes the rank of the upper prediction limit. To make it easier to write some equations later on, we can also write the prediction interval (1) in a slightly different way as:

$$[x_{(u)}, x_{(n+1-w)}] \quad (2)$$

where

$$w = n + 1 - v \quad (3)$$

so that w is a positive integer between 1 and $n - 1$, and $u < n + 1 - w$. In terms of the arguments to the function `predIntNparSimultaneous`, the argument `lp1.rank` corresponds to u , and the argument `n.plus.one.minus.up1.rank` corresponds to w .

If we allow $u = 0$ and $w = 0$ and define lower and upper bounds as:

$$x_{(0)} = lb \quad (4)$$

$$x_{(n+1)} = ub \quad (5)$$

then Equation (2) above can also represent a one-sided lower or one-sided upper prediction interval as well. That is, a one-sided lower nonparametric prediction interval is constructed as:

$$[x_{(u)}, x_{(n+1)}] = [x_{(u)}, ub] \quad (6)$$

and a one-sided upper nonparametric prediction interval is constructed as:

$$[x_{(0)}, x_{(n+1-w)}] = [lb, x_{(n+1-w)}] \quad (7)$$

Usually, $lb = -\infty$ or $lb = 0$ and $ub = \infty$.

Note: For nonparametric simultaneous prediction intervals, only lower (`pi.type="lower"`) and upper (`pi.type="upper"`) prediction intervals are available.

Constructing Nonparametric Simultaneous Prediction Intervals for Future Observations

First we will show how to construct a nonparametric simultaneous prediction interval based on future observations (i.e., $b = 1$, `n.median=1`), and then extend the formulas to future medians.

Simultaneous Prediction Intervals for the k-of-m Rule (rule="k.of.m")

For the k -of- m rule (`rule="k.of.m"`) with $w = 1$ (i.e., `n.median=1`), at least k of the next m future observations will fall in the prediction interval with probability $(1 - \alpha)100\%$ on each of the r future sampling occasions. If observations are being taken sequentially, for a particular sampling occasion, up to m observations may be taken, but once k of the observations fall within the prediction interval, sampling can stop. Note: When $r = 1$, this kind of simultaneous prediction interval becomes the same as a standard nonparametric prediction interval (see [predIntNpar](#)).

Chou and Owen (1986) developed the theory for nonparametric simultaneous prediction limits for various rules, including the 1-of- m rule. Their theory, however, does not cover the California or Modified California rules, and uses an r -fold summation involving a minimum of 2^r terms. Davis and McNichols (1994b; 1999) extended the results of Chou and Owen (1986) to include the California and Modified California rule, and developed algorithms that involve summing far fewer terms.

Davis and McNichols (1999) give formulas for the probabilities associated with the one-sided upper simultaneous prediction interval shown in Equation (7). For the k -of- m rule, the probability that at least k of the next m future observations will be contained in the interval given in Equation (7) for each of r future sampling occasions is given by:

$$\begin{aligned} 1 - \alpha &= E\left[\sum_{i=0}^{m-k} \binom{k-1+i}{k-1} Y^k (1-Y)^i\right]^r \\ &= \int_0^1 \left[\sum_{i=0}^{m-k} \binom{k-1+i}{k-1} y^k (1-y)^i\right]^r f(y) dy \quad (8) \end{aligned}$$

where Y denotes a random variable with a [beta distribution](#) with parameters v and $n+1-v$, and $f(\cdot)$ denotes the pdf of this distribution. Note that v denotes the rank of the order statistic used as the upper prediction limit (i.e., `n.plus.one.minus.upl.rank`), and that v is usually equal to n .

Also note that the summation term in Equation (8) corresponds to the cumulative distribution function of a [Negative Binomial distribution](#) with parameters `size`= k and `prob`= y evaluated at `q`= $m-k$.

When `pi.type`="lower", Y denotes a random variable with a [beta distribution](#) with parameters $n+1-u$ and u . Note that u denotes the rank of the order statistic used as the lower prediction limit (i.e., `lpl.rank`= u), and that u is usually equal to 1.

Simultaneous Prediction Intervals for the California Rule (rule="CA")

For the California rule (rule="CA"), with probability $(1-\alpha)100\%$, for each of the r future sampling occasions, either the first observation will fall in the prediction interval, or else all of the next $m-1$ observations will fall in the prediction interval. That is, if the first observation falls in the prediction interval then sampling can stop. Otherwise, $m-1$ more observations must be taken.

In this case, the probability is given by:

$$\begin{aligned} 1 - \alpha &= E\left[\sum_{i=0}^r \binom{r}{i} Y^{r-i+(m-1)i} (1-Y)^i\right] \\ &= \int_0^1 \left[\sum_{i=0}^r \binom{r}{i} y^{r-i+(m-1)i} (1-y)^i\right] f(y) dy \quad (9) \end{aligned}$$

Simultaneous Prediction Intervals for the Modified California Rule (rule="Modified.CA")

For the Modified California rule (rule="Modified.CA"), with probability $(1-\alpha)100\%$, for each of the r future sampling occasions, either the first observation will fall in the prediction interval, or else at least 2 out of the next 3 observations will fall in the prediction interval. That is, if the first observation falls in the prediction interval then sampling can stop. Otherwise, up to 3 more observations must be taken.

In this case, the probability is given by:

$$\begin{aligned} 1 - \alpha &= E[Y^r (1+Q+Q^2-2Q^3)^r] \\ &= \int_0^1 [y^r (1+q+q^2-2q^3)^r] f(y) dy \quad (10) \end{aligned}$$

where $Q = 1 - Y$ and $q = 1 - y$.

Davis and McNichols (1999) provide algorithms for computing the probabilities based on expanding polynomials and the formula for the expected value of a beta random variable. In the discussion section of Davis and McNichols (1999), however, Vangel points out that numerical integration is adequate, and this is how these probabilities are computed in the function `predIntNparSimultaneous`.

Constructing Nonparametric Simultaneous Prediction Intervals for Future Medians

USEPA (2009, Chapter 19; Cameron, 2011) extends nonparametric simultaneous prediction intervals to testing future medians for the case of the 1-of-1 and 1-of-2 plans for medians of order 3. In general, each of the rules (k -of- m , California, and Modified California) can be easily extended to the case of using medians as long as the medians are based on an odd (as opposed to even) sample size.

For each of the above rules, if we are interested in using medians instead of single observations (i.e., $b \geq 1$; $n.\text{median} \geq 1$), and we force b to be odd, then a median will be less than a prediction limit once $(b + 1)/2$ observations are less than the prediction limit. Thus, Equations (8) - (10) are modified by replacing y with the term:

$$\sum_{i=0}^{b-b'} \binom{b'-1+i}{b'-1} y^{b'} (1-y)^i \quad (11)$$

where

$$b' = \frac{b+1}{2} \quad (12)$$

Value

a list of class "estimate" containing the simultaneous prediction interval and other information. See the help file for `estimate.object` for details.

Note

Prediction and tolerance intervals have long been applied to quality control and life testing problems (Hahn, 1970b,c; Hahn and Nelson, 1973; Krishnamoorthy and Mathew, 2009). In the context of environmental statistics, prediction intervals are useful for analyzing data from groundwater detection monitoring programs at hazardous and solid waste facilities (e.g., Gibbons et al., 2009; Millard and Neerchal, 2001; USEPA, 2009).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Cameron, Kirk. (2011). Personal communication, February 16, 2011. MacStat Consulting, Ltd., Colorado Springs, Colorado.
- Chew, V. (1968). Simultaneous Prediction Intervals. *Technometrics* **10**(2), 323–331.
- Danziger, L., and S. Davis. (1964). Tables of Distribution-Free Tolerance Limits. *Annals of Mathematical Statistics* **35**(5), 1361–1365.
- Davis, C.B. (1994). Environmental Regulatory Statistics. In Patil, G.P., and C.R. Rao, eds., *Handbook of Statistics, Vol. 12: Environmental Statistics*. North-Holland, Amsterdam, a division of Elsevier, New York, NY, Chapter 26, 817–865.
- Davis, C.B., and R.J. McNichols. (1987). One-sided Intervals for at Least p of m Observations from a Normal Population on Each of r Future Occasions. *Technometrics* **29**, 359–370.

- Davis, C.B., and R.J. McNichols. (1994a). Ground Water Monitoring Statistics Update: Part I: Progress Since 1988. *Ground Water Monitoring and Remediation* **14**(4), 148–158.
- Davis, C.B., and R.J. McNichols. (1994b). Ground Water Monitoring Statistics Update: Part II: Nonparametric Prediction Limits. *Ground Water Monitoring and Remediation* **14**(4), 159–175.
- Davis, C.B., and R.J. McNichols. (1999). Simultaneous Nonparametric Prediction Limits (with Discussion). *Technometrics* **41**(2), 89–112.
- Gibbons, R.D. (1987a). Statistical Prediction Intervals for the Evaluation of Ground-Water Quality. *Ground Water* **25**, 455–465.
- Gibbons, R.D. (1991b). Statistical Tolerance Limits for Ground-Water Monitoring. *Ground Water* **29**, 563–570.
- Gibbons, R.D., and J. Baker. (1991). The Properties of Various Statistical Prediction Intervals for Ground-Water Detection Monitoring. *Journal of Environmental Science and Health* **A26**(4), 535–553.
- Gibbons, R.D., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*, Second Edition. John Wiley & Sons, Hoboken.
- Hahn, G.J., and W.Q. Meeker. (1991). *Statistical Intervals: A Guide for Practitioners*. John Wiley and Sons, New York, 392pp.
- Hahn, G., and W. Nelson. (1973). A Survey of Prediction Intervals and Their Applications. *Journal of Quality Technology* **5**, 178–188.
- Hall, I.J., R.R. Prairie, and C.K. Motlagh. (1975). Non-Parametric Prediction Intervals. *Journal of Quality Technology* **7**(3), 109–114.
- Millard, S.P., and Neerchal, N.K. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, Florida.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.
- USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

See Also

[predIntNparSimultaneousConfLevel](#), [predIntNparSimultaneousN](#), [predIntNparSimultaneousTestPower](#), [predIntNpar](#), [tolIntNpar](#), [estimate.object](#).

Examples

```
# Generate 20 observations from a lognormal mixture distribution with
# parameters mean1=1, cv1=0.5, mean2=5, cv2=1, and p.mix=0.1. Use
# predIntNparSimultaneous to construct an upper one-sided prediction interval
# using the maximum observed value using the 1-of-3 rule.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
dat <- rlnormMixAlt(n = 20, mean1 = 1, cv1 = 0.5,
  mean2 = 5, cv2 = 1, p.mix = 0.1)

predIntNparSimultaneous(dat, k = 1, m = 3, lb = 0)
```

```

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          None
#
#Data:                          dat
#
#Sample Size:                  20
#
#Prediction Interval Method:    exact
#
#Prediction Interval Type:      upper
#
#Confidence Level:             99.94353%
#
#Prediction Limit Rank(s):      20
#
#Minimum Number of
#Future Observations
#Interval Should Contain:      1
#
#Total Number of
#Future Observations:          3
#
#Prediction Interval:           LPL = 0.000000
#                               UPL = 1.817311
#-----

# Compare the confidence levels for the 1-of-3 rule, California Rule, and
# Modified California Rule.

predIntNparSimultaneous(dat, k = 1, m = 3, lb = 0)$interval$conf.level
#[1] 0.9994353

predIntNparSimultaneous(dat, m = 3, rule = "CA", lb = 0)$interval$conf.level
#[1] 0.9919066

predIntNparSimultaneous(dat, rule = "Modified.CA", lb = 0)$interval$conf.level
#[1] 0.9984943

#=====

# Repeat the above example, but create the baseline data using just
# n=8 observations and set r to 4 future sampling occasions

set.seed(598)
dat <- rlnormMixAlt(n = 8, mean1 = 1, cv1 = 0.5,
  mean2 = 5, cv2 = 1, p.mix = 0.1)

predIntNparSimultaneous(dat, k = 1, m = 3, r = 4, lb = 0)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          None
#

```

```

#Data:                                dat
#
#Sample Size:                          8
#
#Prediction Interval Method:           exact
#
#Prediction Interval Type:             upper
#
#Confidence Level:                     97.7599%
#
#Prediction Limit Rank(s):             8
#
#Minimum Number of
#Future Observations
#Interval Should Contain
#(per Sampling Occasion):             1
#
#Total Number of
#Future Observations
#(per Sampling Occasion):             3
#
#Number of Future
#Sampling Occasions:                   4
#
#Prediction Interval:                  LPL = 0.000000
#                                      UPL = 5.683453

#-----

# Compare the confidence levels for the 1-of-3 rule, California Rule, and
# Modified California Rule.

predIntNparSimultaneous(dat, k = 1, m = 3, r = 4, lb = 0)$interval$conf.level
#[1] 0.977599

predIntNparSimultaneous(dat, m = 3, r = 4, rule = "CA", lb = 0)$interval$conf.level
#[1] 0.8737798

predIntNparSimultaneous(dat, r = 4, rule = "Modified.CA", lb = 0)$interval$conf.level
#[1] 0.9510178

#=====

# Example 19-5 of USEPA (2009, p. 19-33) shows how to compute nonparametric upper
# simultaneous prediction limits for various rules based on trace mercury data (ppb)
# collected in the past year from a site with four background wells and 10 compliance
# wells (data for two of the compliance wells are shown in the guidance document).
# The facility must monitor the 10 compliance wells for five constituents
# (including mercury) annually.

# Here we will compute the confidence level associated with two different sampling plans:
# 1) the 1-of-2 retesting plan for a median of order 3 using the background maximum and
# 2) the 1-of-4 plan on individual observations using the 3rd highest background value.
# The data for this example are stored in EPA.09.Ex.19.5.mercury.df.

# We will pool data from 4 background wells that were sampled on
# a number of different occasions, giving us a sample size of

```

```

# n = 20 to use to construct the prediction limit.

# There are 10 compliance wells and we will monitor 5 different
# constituents at each well annually. For this example, USEPA (2009)
# recommends setting r to the product of the number of compliance wells and
# the number of evaluations per year.

# To determine the minimum confidence level we require for
# the simultaneous prediction interval, USEPA (2009) recommends
# setting the maximum allowed individual Type I Error level per constituent to:

#  $1 - (1 - \text{SWFPR})^{(1 / \text{Number of Constituents})}$ 

# which translates to setting the confidence limit to

#  $(1 - \text{SWFPR})^{(1 / \text{Number of Constituents})}$ 

# where SWFPR = site-wide false positive rate. For this example, we
# will set SWFPR = 0.1. Thus, the required individual Type I Error level
# and confidence level per constituent are given as follows:

# n = 20 based on 4 Background Wells
# nw = 10 Compliance Wells
# nc = 5 Constituents
# ne = 1 Evaluation per year

n <- 20
nw <- 10
nc <- 5
ne <- 1

# Set number of future sampling occasions r to
# Number Compliance Wells x Number Evaluations per Year
r <- nw * ne

conf.level <- (1 - 0.1)^(1 / nc)
conf.level
#[1] 0.9791484

alpha <- 1 - conf.level
alpha
#[1] 0.02085164

#-----

# Look at the data:

head(EPA.09.Ex.19.5.mercury.df)
# Event Well Well.type Mercury.ppb.orig Mercury.ppb Censored
#1 1 BG-1 Background 0.21 0.21 FALSE
#2 2 BG-1 Background <.2 0.20 TRUE
#3 3 BG-1 Background <.2 0.20 TRUE
#4 4 BG-1 Background <.2 0.20 TRUE
#5 5 BG-1 Background <.2 0.20 TRUE
#6 6 BG-1 Background NA FALSE

longToWide(EPA.09.Ex.19.5.mercury.df, "Mercury.ppb.orig",

```

```

      "Event", "Well", paste.row.name = TRUE)
#           BG-1 BG-2 BG-3 BG-4 CW-1 CW-2
#Event.1 0.21  <.2  <.2  <.2 0.22 0.36
#Event.2  <.2  <.2 0.23 0.25  0.2 0.41
#Event.3  <.2  <.2  <.2 0.28  <.2 0.28
#Event.4  <.2 0.21 0.23  <.2 0.25 0.45
#Event.5  <.2  <.2 0.24  <.2 0.24 0.43
#Event.6                <.2 0.54

# Construct the upper simultaneous prediction limit using the 1-of-2
# retesting plan for a median of order 3 based on the background maximum

Hg.Back <- with(EPA.09.Ex.19.5.mercury.df,
  Mercury.ppb[Well.type == "Background"])

pred.int.1.of.2.med.3 <- predIntNparSimultaneous(Hg.Back, n.median = 3,
  k = 1, m = 2, r = r, lb = 0)

pred.int.1.of.2.med.3

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          None
#
#Data:                         Hg.Back
#
#Sample Size:                  20
#
#Number NA/NaN/Inf's:         4
#
#Prediction Interval Method:    exact
#
#Prediction Interval Type:      upper
#
#Confidence Level:             99.40354%
#
#Prediction Limit Rank(s):      20
#
#Minimum Number of
#Future Medians
#Interval Should Contain
#(per Sampling Occasion):      1
#
#Total Number of
#Future Medians
#(per Sampling Occasion):      2
#
#Number of Future
#Sampling Occasions:           10
#
#Sample Size for Medians:      3
#
#Prediction Interval:           LPL = 0.00
#                               UPL = 0.28

```

```

# Note that the achieved confidence level of 99.4% is greater than the
# required confidence level of 97.9%.

# Now determine whether either compliance well indicates evidence of
# Mercury contamination.

# Compliance Well 1
#-----
Hg.CW.1 <- with(EPA.09.Ex.19.5.mercury.df, Mercury.ppb.orig[Well == "CW-1"])

Hg.CW.1
#[1] "0.22" "0.2" "<.2" "0.25" "0.24" "<.2"

# The median of the first 3 observations is 0.2, which is less than
# the UPL of 0.28, so there is no evidence of contamination.

# Compliance Well 2
#-----
Hg.CW.2 <- with(EPA.09.Ex.19.5.mercury.df, Mercury.ppb.orig[Well == "CW-2"])

Hg.CW.2
#[1] "0.36" "0.41" "0.28" "0.45" "0.43" "0.54"

# The median of the first 3 observations is 0.36, so 3 more observations have to
# be looked at. The median of the second 3 observations is 0.45, which is
# larger than the UPL of 0.28, so there is evidence of contamination.

#-----

# Now create the upper simultaneous prediction limit using the 1-of-4 plan
# on individual observations using the 3rd highest background value.

pred.int.1.of.4.3rd <- predIntNparSimultaneous(Hg.Back, k = 1, m = 4,
  r = r, lb = 0, n.plus.one.minus.upl.rank = 3)

pred.int.1.of.4.3rd

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:      None
#
#Data:                     Hg.Back
#
#Sample Size:              20
#
#Number NA/NaN/Inf's:      4
#
#Prediction Interval Method: exact
#
#Prediction Interval Type:  upper
#
#Confidence Level:         98.64909%
#
#Prediction Limit Rank(s):  18
#
#Minimum Number of

```

```

#Future Observations
#Interval Should Contain
#(per Sampling Occasion):      1
#
#Total Number of
#Future Observations
#(per Sampling Occasion):      4
#
#Number of Future
#Sampling Occasions:           10
#
#Prediction Interval:           LPL = 0.00
#                               UPL = 0.24

# Note that the achieved confidence level of 98.6% is greater than the
# required confidence level of 97.9%.

# Now determine whether either compliance well indicates evidence of
# Mercury contamination.

# Compliance Well 1
#-----
Hg.CW.1 <- with(EPA.09.Ex.19.5.mercury.df, Mercury.ppb.orig[Well == "CW-1"])

Hg.CW.1
#[1] "0.22" "0.2"  "<.2"  "0.25" "0.24" "<.2"

# The first observation is less than the UPL of 0.24, which is less than
# the UPL of 0.28, so there is no evidence of contamination.

# Compliance Well 2
#-----
Hg.CW.2 <- with(EPA.09.Ex.19.5.mercury.df, Mercury.ppb.orig[Well == "CW-2"])

Hg.CW.2
#[1] "0.36" "0.41" "0.28" "0.45" "0.43" "0.54"

# All of the first 4 observations are greater than the UPL of 0.24, so there
# is evidence of contamination.

#=====

# Cleanup
#-----
rm(dat, n, nw, nc, ne, r, conf.level, alpha, Hg.Back, pred.int.1.of.2.med.3,
  pred.int.1.of.4.3rd, Hg.CW.1, Hg.CW.2)

```

predIntPois

Prediction Interval for a Poisson Distribution

Description

Estimate the mean of a [Poisson distribution](#), and construct a prediction interval for the next k observations or next set of k sums.

Usage

```
predIntPois(x, k = 1, n.sum = 1, method = "conditional",
  pi.type = "two-sided", conf.level = 0.95, round.limits = TRUE)
```

Arguments

<code>x</code>	numeric vector of observations, or an object resulting from a call to an estimating function that assumes a Poisson distribution (i.e., <code>epois</code> or <code>epoisCensored</code>). If <code>x</code> is a numeric vector, missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
<code>k</code>	positive integer specifying the number of future observations or sums the prediction interval should contain with confidence level <code>conf.level</code> . The default value is <code>k=1</code> .
<code>n.sum</code>	positive integer specifying the sample size associated with the k future sums. The default value is <code>n.sum=1</code> (i.e., individual observations). Note that all future sums must be based on the same sample size.
<code>method</code>	character string specifying the method to use. The possible values are "conditional" (based on a conditional distribution; the default), "conditional.approx.normal" (method based on approximating a conditional distribution with the standard normal distribution), "conditional.approx.t" (method based on approximating a conditional distribution with Student's t-distribution), and "normal.approx" (approximate method based on the fact that the mean and variance of a Poisson distribution are the same). See the DETAILS section for more information on these methods. The "conditional" method is only implemented for $k=1$; when k is bigger than 1, the value of <code>method</code> cannot be "conditional".
<code>pi.type</code>	character string indicating what kind of prediction interval to compute. The possible values are <code>pi.type="two-sided"</code> (the default), <code>pi.type="lower"</code> , and <code>pi.type="upper"</code> .
<code>conf.level</code>	a scalar between 0 and 1 indicating the confidence level of the prediction interval. The default value is <code>conf.level=0.95</code> .
<code>round.limits</code>	logical scalar indicating whether to round the computed prediction limits to the nearest integer. The default value is <code>round.limits=TRUE</code> .

Details

A prediction interval for some population is an interval on the real line constructed so that it will contain k future observations or averages from that population with some specified probability $(1 - \alpha)100\%$, where $0 < \alpha < 1$ and k is some pre-specified positive integer. The quantity $(1 - \alpha)100\%$ is called the confidence coefficient or confidence level associated with the prediction interval.

In the case of a [Poisson distribution](#), we have modified the usual meaning of a prediction interval and instead construct an interval that will contain k future observations or k future *sums* with a certain confidence level.

A prediction interval is a *random* interval; that is, the lower and/or upper bounds are random variables computed based on sample statistics in the baseline sample. Prior to taking one specific baseline sample, the probability that the prediction interval will contain the next k averages is $(1 - \alpha)100\%$. Once a specific baseline sample is taken and the prediction interval based on that sample is computed, the probability that that prediction interval will contain the next k averages is not necessarily $(1 - \alpha)100\%$, but it should be close.

If an experiment is repeated N times, and for each experiment:

1. A sample is taken and a $(1 - \alpha)100\%$ prediction interval for $k = 1$ future observation is computed, and
2. One future observation is generated and compared to the prediction interval,

then the number of prediction intervals that actually contain the future observation generated in step 2 above is a binomial random variable with parameters $\text{size}=N$ and $\text{prob}=(1 - \alpha)100\%$ (see [Binomial](#)).

If, on the other hand, only one baseline sample is taken and only one prediction interval for $k = 1$ future observation is computed, then the number of future observations out of a total of N future observations that will be contained in that one prediction interval is a binomial random variable with parameters $\text{size}=N$ and $\text{prob}=(1 - \alpha^*)100\%$, where α^* depends on the true population parameters and the computed bounds of the prediction interval.

Because of the discrete nature of the [Poisson distribution](#), even if the true mean of the distribution λ were known exactly, the actual confidence level associated with a prediction limit will usually not be exactly equal to $(1 - \alpha)100\%$. For example, for the Poisson distribution with parameter $\text{lambda}=2$, the interval $[0, 4]$ contains 94.7% of this distribution and the interval $[0,5]$ contains 98.3% of this distribution. Thus, no interval can contain exactly 95% of this distribution, so it is impossible to construct an exact 95% prediction interval for the next $k = 1$ observation for a Poisson distribution with parameter $\text{lambda}=2$.

The Form of a Poisson Prediction Interval

Let $\underline{x} = x_1, x_2, \dots, x_n$ denote a vector of n observations from a [Poisson distribution](#) with parameter $\text{lambda}=\lambda$. Also, let X denote the sum of these n random variables, i.e.,

$$X = \sum_{i=1}^n x_i \quad (1)$$

Finally, let m denote the sample size associated with the k future sums (i.e., $n.\text{sum}=m$). When $m = 1$, each sum is really just a single observation, so in the rest of this help file the term “sums” replaces the phrase “observations or sums”.

Let $\underline{y} = y_1, y_2, \dots, y_m$ denote a vector of m future observations from a Poisson distribution with parameter $\text{lambda}=\lambda^*$, and set Y equal to the sum of these m random variables, i.e.,

$$Y = \sum_{i=1}^m y_i \quad (2)$$

Then Y has a Poisson distribution with parameter $\text{lambda}=m\lambda^*$ (Johnson et al., 1992, p.160). We are interested in constructing a prediction limit for the next value of Y , or else the next k sums of m Poisson random variables, based on the observed value of X and assuming $\lambda^* = \lambda$.

For a Poisson distribution, the form of a two-sided prediction interval is:

$$[m\bar{x} - K, m\bar{x} + K] = [cX - K, cX + K] \quad (3)$$

where

$$\bar{x} = \frac{X}{n} = \sum_{i=1}^n x_i \quad (4)$$

$$c = \frac{m}{n} \quad (5)$$

and K is a constant that depends on the sample size n , the confidence level $(1 - \alpha)100\%$, the number of future sums k , and the sample size associated with the future sums m . Do not confuse the constant K (uppercase K) with the number of future sums k (lowercase k). The symbol K is

used here to be consistent with the notation used for prediction intervals for the normal distribution (see [predIntNorm](#)).

Similarly, the form of a one-sided lower prediction interval is:

$$[m\bar{x} - K, \infty] = [cX - K, \infty] \quad (6)$$

and the form of a one-sided upper prediction interval is:

$$[0, m\bar{x} + K] = [0, cX + K] \quad (7)$$

The derivation of the constant K is explained below.

Conditional Distribution (method="conditional")

Nelson (1970) derives a prediction interval for the case $k = 1$ based on the conditional distribution of Y given $X + Y$. He notes that the conditional distribution of Y given the quantity $X + Y = w$ is [binomial](#) with parameters $\text{size}=w$ and $\text{prob}=[m\lambda^*/(m\lambda^* + n\lambda)]$ (Johnson et al., 1992, p.161). When $k = 1$, the prediction limits are computed as those most extreme values of Y that still yield a non-significant test of the hypothesis $H_0 : \lambda^* = \lambda$, which for the conditional distribution of Y is equivalent to the hypothesis $H_0 : \text{prob}=[m/(m + n)]$.

Using the relationship between the [binomial](#) and [F-distribution](#) (see the explanation of exact confidence intervals in the help file for [ebinom](#)), Nelson (1982, p. 203) states that exact two-sided $(1 - \alpha)100\%$ prediction limits [LPL, UPL] are the closest integer solutions to the following equations:

$$\frac{m}{LPL + 1} = \frac{n}{X} F(2LPL + 2, 2X, 1 - \alpha/2) \quad (8)$$

$$\frac{UPL}{n} = \frac{X + 1}{n} F(2X + 2, 2UPL, 1 - \alpha/2) \quad (9)$$

where $F(\nu_1, \nu_2, p)$ denotes the p 'th quantile of the [F-distribution](#) with ν_1 and ν_2 degrees of freedom. If `ci.type="lower"`, $\alpha/2$ is replaced with α in Equation (8) above for LPL , and UPL is set to ∞ .

If `ci.type="upper"`, $\alpha/2$ is replaced with α in Equation (9) above for UPL , and LPL is set to 0.

NOTE: This method is not extended to the case $k > 1$.

Conditional Distribution Approximation Based on Normal Distribution

(method="conditional.approx.normal")

Cox and Hinkley (1974, p.245) derive an approximate prediction interval for the case $k = 1$. Like Nelson (1970), they note that the conditional distribution of Y given the quantity $X + Y = w$ is [binomial](#) with parameters $\text{size}=w$ and $\text{prob}=[m\lambda^*/(m\lambda^* + n\lambda)]$, and that the hypothesis $H_0 : \lambda^* = \lambda$ is equivalent to the hypothesis $H_0 : \text{prob}=[m/(m + n)]$.

Cox and Hinkley (1974, p.245) suggest using the normal approximation to the binomial distribution (in this case, without the continuity correction; see Zar, 2010, pp.534-536 for information on the continuity correction associated with the normal approximation to the binomial distribution). Under the null hypothesis $H_0 : \lambda^* = \lambda$, the quantity

$$z = [Y - \frac{c(X + Y)}{1 + c}] / \{[\frac{c(X + Y)}{(1 + c)^2}]^{1/2}\} \quad (10)$$

is approximately distributed as a standard normal random variable.

The Case When $k = 1$

When $k = 1$ and `pi.type="two-sided"`, the prediction limits are computed by solving the equation

$$z^2 \leq z_{1-\alpha/2}^2 \quad (11)$$

where z_p denotes the p 'th quantile of the standard normal distribution. In this case, Gibbons (1987b) notes that the quantity K in Equation (3) above is given by:

$$K = \frac{t^2 c}{2} t c [X(1 + \frac{1}{c}) + \frac{t^2}{4}]^{1/2} \quad (12)$$

where $t = z_{1-\alpha/2}$.

When `pi.type="lower"` or `pi.type="upper"`, K is computed exactly as above, except t is set to $t = z_{1-\alpha}$.

The Case When $k > 1$

When $k > 1$, Gibbons (1987b) suggests using the Bonferroni inequality. That is, the value of K is computed exactly as for the case $k = 1$ described above, except that the Bonferroni value of t is used in place of the usual value of t :

When `pi.type="two-side"`, $t = z_{1-(\alpha/k)/2}$.

When `pi.type="lower"` or `pi.type="upper"`, $t = z_{1-\alpha/k}$.

Conditional Distribution Approximation Based on Student's t-Distribution

(`method="conditional.approx.t"`)

When `method="conditional.approx.t"`, the exact same procedure is used as when `method="conditional.approx.n"` except that the quantity in Equation (10) is assumed to follow a Student's t-distribution with $n - 1$ degrees of freedom. Thus, all occurrences of z_p are replaced with $t_{n-1,p}$, where $t_{\nu,p}$ denotes the p 'th quantile of [Student's t-distribution](#) with ν degrees of freedom.

Normal Approximation (`method="normal.approx"`)

The normal approximation for Poisson prediction limits was given by Nelson (1970; 1982, p.203) and is based on the fact that the mean and variance of a Poisson distribution are the same (Johnson et al, 1992, p.157), and for "large" values of n and m , both X and Y are approximately normally distributed.

The Case When $k = 1$

The quantity $Y - cX$ is approximately normally distributed with expectation and variance given by:

$$E(Y - cX) = E(Y) - cE(X) = m\lambda - cn\lambda = 0 \quad (13)$$

$$Var(Y - cX) = Var(Y) + c^2 Var(X) = m\lambda + c^2 n\lambda = m\lambda(1 + \frac{m}{n}) \quad (14)$$

so the quantity

$$z = \frac{Y - cX}{\sqrt{m\hat{\lambda}(1 + \frac{m}{n})}} = \frac{Y - cX}{\sqrt{m\bar{x}(1 + \frac{m}{n})}} \quad (15)$$

is approximately distributed as a standard normal random variable. The function `predIntPois`, however, assumes this quantity is distributed as approximately a [Student's t-distribution](#) with $n - 1$ degrees of freedom.

Thus, following the idea of prediction intervals for a normal distribution (see `predIntNorm`), when `pi.type="two-sided"`, the constant K for a $(1 - \alpha)100\%$ prediction interval for the next $k = 1$ sum of m observations is computed as:

$$K = t_{n-1,1-\alpha/2} \sqrt{m\bar{x}(1 + \frac{m}{n})} \quad (16)$$

where $t_{\nu,p}$ denotes the p 'th quantile of a [Student's t-distribution](#) with ν degrees of freedom.

Similarly, when `pi.type="lower"` or `pi.type="upper"`, the constant K is computed as:

$$K = t_{n-1, 1-\alpha} \sqrt{m\bar{x}(1 + \frac{m}{n})} \quad (17)$$

The Case When $k > 1$

When $k > 1$, the value of K is computed exactly as for the case $k = 1$ described above, except that the Bonferroni value of t is used in place of the usual value of t :

When `pi.type="two-sided"`,

$$K = t_{n-1, 1-(\alpha/k)/2} \sqrt{m\bar{x}(1 + \frac{m}{n})} \quad (18)$$

When `pi.type="lower"` or `pi.type="upper"`,

$$K = t_{n-1, 1-(\alpha/k)} \sqrt{m\bar{x}(1 + \frac{m}{n})} \quad (19)$$

Hahn and Nelson (1973, p.182) discuss another method of computing K when $k > 1$, but this method is not implemented here.

Value

If `x` is a numeric vector, `predIntPois` returns a list of class "estimate" containing the estimated parameter, the prediction interval, and other information. See the help file for [estimate.object](#) for details.

If `x` is the result of calling an estimation function, `predIntPois` returns a list whose class is the same as `x`. The list contains the same components as `x`, as well as a component called `interval` containing the prediction interval information. If `x` already has a component called `interval`, this component is replaced with the prediction interval information.

Note

Prediction and tolerance intervals have long been applied to quality control and life testing problems. Nelson (1970) notes that his development of confidence and prediction limits for the Poisson distribution is based on well-known results dating back to the 1950's. Hahn and Nelson (1973) review prediction intervals for several distributions, including Poisson prediction intervals. The monograph by Hahn and Meeker (1991) includes a discussion of Poisson prediction intervals.

Gibbons (1987b) uses the Poisson distribution to model the number of detected compounds per scan of the 32 volatile organic priority pollutants (VOC), and also to model the distribution of chemical concentration (in ppb), and presents formulas for prediction and tolerance intervals. The formulas for prediction intervals are based on Cox and Hinkley (1974, p.245). Gibbons (1987b) only deals with the case where `n.sum=1`.

Gibbons et al. (2009, pp. 72–76) discuss methods for Poisson prediction limits.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Cox, D.R., and D.V. Hinkley. (1974). *Theoretical Statistics*. Chapman and Hall, New York, pp.242–245.
- Gibbons, R.D. (1987b). Statistical Models for the Analysis of Volatile Organic Compounds in Waste Disposal Sites. *Ground Water* **25**, 572–580.
- Gibbons, R.D., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*, Second Edition. John Wiley & Sons, Hoboken, pp. 72–76.
- Hahn, G.J., and W.Q. Meeker. (1991). *Statistical Intervals: A Guide for Practitioners*. John Wiley and Sons, New York.
- Hahn, G., and W. Nelson. (1973). A Survey of Prediction Intervals and Their Applications. *Journal of Quality Technology* **5**, 178–188.
- Johnson, N. L., S. Kotz, and A. Kemp. (1992). *Univariate Discrete Distributions*. Second Edition. John Wiley and Sons, New York, Chapter 4.
- Millard, S.P., and N.K. Neerchal. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton.
- Miller, R.G. (1981a). *Simultaneous Statistical Inference*. McGraw-Hill, New York, pp.8, 76–81.
- Nelson, W.R. (1970). Confidence Intervals for the Ratio of Two Poisson Means and Poisson Predictor Intervals. *IEEE Transactions of Reliability* **R-19**, 42–49.
- Nelson, W.R. (1982). *Applied Life Data Analysis*. John Wiley and Sons, New York, pp.200–204.
- Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ, pp. 585–586.

See Also

[Poisson](#), [epois](#), [estimate.object](#), [Prediction Intervals](#), [tolIntPois](#), [Estimating Distribution Parameters](#).

Examples

```
# Generate 20 observations from a Poisson distribution with parameter
# lambda=2. The interval [0, 4] contains 94.7% of this distribution and
# the interval [0,5] contains 98.3% of this distribution. Thus, because
# of the discrete nature of the Poisson distribution, no interval contains
# exactly 95% of this distribution. Use predIntPois to estimate the mean
# parameter of the true distribution, and construct a one-sided upper
# 95% prediction interval for the next single observation from this distribution.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
dat <- rpois(20, lambda = 2)

predIntPois(dat, pi.type = "upper")

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Poisson
#
#Estimated Parameter(s):      lambda = 1.8
#
#Estimation Method:           mle/mme/mvue
```

```

#
#Data:                                dat
#
#Sample Size:                          20
#
#Prediction Interval Method:           conditional
#
#Prediction Interval Type:             upper
#
#Confidence Level:                     95%
#
#Number of Future Observations:       1
#
#Prediction Interval:                  LPL = 0
#                                      UPL = 5

#-----

# Compare results above with the other approximation methods:

predIntPois(dat, method = "conditional.approx.normal",
  pi.type = "upper")$interval$limits
#LPL UPL
# 0    4

predIntPois(dat, method = "conditional.approx.t",
  pi.type = "upper")$interval$limits
#LPL UPL
# 0    4

predIntPois(dat, method = "normal.approx",
  pi.type = "upper")$interval$limits
#LPL UPL
# 0    4
#Warning message:
#In predIntPois(dat, method = "normal.approx", pi.type = "upper") :
# Estimated value of 'lambda' and/or number of future observations
# is/are probably too small for the normal approximation to work well.

#=====

# Using the same data as in the previous example, compute a one-sided
# upper 95% prediction limit for k=10 future observations.

# Using conditional approximation method based on the normal distribution.

predIntPois(dat, k = 10, method = "conditional.approx.normal",
  pi.type = "upper")

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:                 Poisson
#
#Estimated Parameter(s):              lambda = 1.8

```

```

#
#Estimation Method:          mle/mme/mvue
#
#Data:                       dat
#
#Sample Size:                20
#
#Prediction Interval Method:  conditional.approx.normal
#
#Prediction Interval Type:    upper
#
#Confidence Level:           95%
#
#Number of Future Observations: 10
#
#Prediction Interval:        LPL = 0
#                             UPL = 6

# Using method based on approximating conditional distribution with
# Student's t-distribution

predIntPois(dat, k = 10, method = "conditional.approx.t",
  pi.type = "upper")$interval$limits
#LPL UPL
# 0 6

#=====

# Repeat the above example, but set k=5 and n.sum=3. Thus, we want a
# 95% upper prediction limit for the next 5 sets of sums of 3 observations.

predIntPois(dat, k = 5, n.sum = 3, method = "conditional.approx.t",
  pi.type = "upper")$interval$limits
#LPL UPL
# 0 12

#=====

# Reproduce Example 3.6 in Gibbons et al. (2009, p. 75)
# A 32-constituent VOC scan was performed for n=16 upgradient
# samples and there were 5 detections out of these 16. We
# want to construct a one-sided upper 95% prediction limit
# for 20 monitoring wells (so k=20 future observations) based
# on these data.

# First we need to create a data set that will yield a mean
# of 5/16 based on a sample size of 16. Any number of data
# sets will do. Here are two possible ones:

dat <- c(rep(1, 5), rep(0, 11))
dat <- c(2, rep(1, 3), rep(0, 12))

# Now call predIntPois. Don't round the limits so we can
# compare to the example in Gibbons et al. (2009).

predIntPois(dat, k = 20, method = "conditional.approx.t",

```



```

    pi.type = "upper", round.limits = FALSE)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:      Poisson
#
#Estimated Parameter(s):   lambda = 0.3125
#
#Estimation Method:        mle/mme/mvue
#
#Data:                     dat
#
#Sample Size:              16
#
#Prediction Interval Method: conditional.approx.t
#
#Prediction Interval Type:  upper
#
#Confidence Level:         95%
#
#Number of Future Observations: 20
#
#Prediction Interval:       LPL = 0.000000
#                           UPL = 2.573258

#=====

# Cleanup
#-----
rm(dat)

```

print.boxcox

Print Output of Objective for Box-Cox Power Transformations

Description

Formats and prints the results of calling the function [boxcox](#). This method is automatically called by [print](#) when given an object of class "boxcox". The names of other functions involved in Box-Cox transformations are listed under [Data Transformations](#).

Usage

```
## S3 method for class 'boxcox'
print(x, ...)
```

Arguments

x an object of class "boxcox". See [boxcox.object](#) for details.

... arguments that can be supplied to the [format](#) function.

Details

This is the "boxcox" method for the generic function [print](#). Prints the objective name, the name of the data object used, the sample size, the values of the powers, and the values of the objective. In the case of optimization, also prints the range of powers over which the optimization took place.

Value

Invisibly returns the input x.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Chambers, J. M. and Hastie, T. J. (1992). *Statistical Models in S*. Wadsworth & Brooks/Cole.

See Also

[boxcox](#), [boxcox.object](#), [plot.boxcox](#), [Data Transformations](#), [print](#).

print.boxcoxCensored	<i>Print Output of Objective for Box-Cox Power Transformations Based on Type I Censored Data</i>
----------------------	--

Description

Formats and prints the results of calling the function [boxcoxCensored](#). This method is automatically called by [print](#) when given an object of class "boxcoxCensored".

Usage

```
## S3 method for class 'boxcoxCensored'
print(x, ...)
```

Arguments

x	an object of class "boxcoxCensored". See boxcoxCensored.object for details.
...	arguments that can be supplied to the format function.

Details

This is the "boxcoxCensored" method for the generic function [print](#). Prints the objective name, the name of the data object used, the sample size, the percentage of censored observations, the values of the powers, and the values of the objective. In the case of optimization, also prints the range of powers over which the optimization took place.

Value

Invisibly returns the input x.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Chambers, J. M. and Hastie, T. J. (1992). *Statistical Models in S*. Wadsworth & Brooks/Cole.

See Also

[boxcoxCensored](#), [boxcoxCensored.object](#), [plot.boxcoxCensored](#), [Data Transformations](#), [print](#).

print.boxcoxLm	<i>Print Output of Objective for Box-Cox Power Transformations for an "lm" Object</i>
----------------	---

Description

Formats and prints the results of calling the function [boxcox](#) when the argument `x` supplied to [boxcox](#) is an object of class "lm". This method is automatically called by [print](#) when given an object of class "boxcoxLm". The names of other functions involved in Box-Cox transformations are listed under [Data Transformations](#).

Usage

```
## S3 method for class 'boxcoxLm'
print(x, ...)
```

Arguments

`x` an object of class "boxcoxLm". See [boxcoxLm.object](#) for details.

`...` arguments that can be supplied to the [format](#) function.

Details

This is the "boxcoxLm" method for the generic function [print](#). Prints the objective name, the details of the "lm" object used, the sample size, the values of the powers, and the values of the objective. In the case of optimization, also prints the range of powers over which the optimization took place.

Value

Invisibly returns the input `x`.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Chambers, J. M. and Hastie, T. J. (1992). *Statistical Models in S*. Wadsworth & Brooks/Cole.

See Also

[boxcox](#), [boxcoxLm.object](#), [plot.boxcoxLm](#), [Data Transformations](#), [print](#).

print.gof

Print Output of Goodness-of-Fit Tests

Description

Formats and prints the results of performing a goodness-of-fit test. This method is automatically called by [print](#) when given an object of class "gof". The names of the functions that perform goodness-of-fit tests and that produce objects of class "gof" are listed under [Goodness-of-Fit Tests](#).

Usage

```
## S3 method for class 'gof'
print(x, ...)
```

Arguments

`x` an object of class "gof". See [gof.object](#) for details.
`...` arguments that can be supplied to the [format](#) function.

Details

This is the "gof" method for the generic function [print](#). Prints name of the test, hypothesized distribution, estimated population parameter(s), estimation method, data name, sample size, value of the test statistic, parameters associated with the null distribution of the test statistic, p-value associated with the test statistic, and the alternative hypothesis.

Value

Invisibly returns the input `x`.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Chambers, J. M. and Hastie, T. J. (1992). *Statistical Models in S*. Wadsworth & Brooks/Cole.

See Also

[Goodness-of-Fit Tests](#), [gof.object](#), [print](#).

`print.gofGroup`*Print Output of Group Goodness-of-Fit Tests*

Description

Formats and prints the results of performing a group goodness-of-fit test. This method is automatically called by `print` when given an object of class `"gofGroup"`. Currently, the only **EnvStats** function that performs a group goodness-of-fit test that produces an object of class `"gofGroup"` is `gofGroupTest`.

Usage

```
## S3 method for class 'gofGroup'
print(x, ...)
```

Arguments

<code>x</code>	an object of class <code>"gofGroup"</code> . See <code>gofGroup.object</code> for details.
<code>...</code>	arguments that can be supplied to the <code>format</code> function.

Details

This is the `"gofGroup"` method for the generic function `print`. See the help file for `gofGroup.object` for information on the information contained in this kind of object.

Value

Invisibly returns the input `x`.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Chambers, J. M. and Hastie, T. J. (1992). *Statistical Models in S*. Wadsworth & Brooks/Cole.

See Also

[Goodness-of-Fit Tests](#), `gofGroup.object`, `print`.

print.gofTwoSample	<i>Print Output of Two-Sample Goodness-of-Fit Tests</i>
--------------------	---

Description

Formats and prints the results of performing a two-sample goodness-of-fit test. This method is automatically called by [print](#) when given an object of class "[gofTwoSample](#)". Currently, the only **EnvStats** function that performs a two-sample goodness-of-fit test that produces an object of class "[gofTwoSample](#)" is [gofTest](#).

Usage

```
## S3 method for class 'gofTwoSample'  
print(x, ...)
```

Arguments

x	an object of class " gofTwoSample ". See gofTwoSample.object for details.
...	arguments that can be supplied to the format function.

Details

This is the "[gofTwoSample](#)" method for the generic function [print](#). See the help file for [gofTwoSample.object](#) for information on the information contained in this kind of object.

Value

Invisibly returns the input x.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Chambers, J. M. and Hastie, T. J. (1992). *Statistical Models in S*. Wadsworth & Brooks/Cole.

See Also

[Goodness-of-Fit Tests](#), [gofTwoSample.object](#), [print](#).

propTestN	<i>Compute Sample Size Necessary to Achieve a Specified Power for a One- or Two-Sample Proportion Test</i>
-----------	--

Description

Compute the sample size necessary to achieve a specified power for a one- or two-sample proportion test, given the true proportion(s) and significance level.

Usage

```
propTestN(p.or.p1, p0.or.p2, alpha = 0.05, power = 0.95,
  sample.type = "one.sample", alternative = "two.sided",
  ratio = 1, approx = TRUE,
  correct = sample.type == "two.sample",
  round.up = TRUE, warn = TRUE, return.exact.list = TRUE,
  n.min = 2, n.max = 10000, tol.alpha = 0.1 * alpha,
  tol = 1e-7, maxiter = 1000)
```

Arguments

p.or.p1	numeric vector of proportions. When sample.type="one.sample", this argument denotes the true value of p , the probability of dQuotesuccess. When sample.type="two.sample", this argument denotes the value of p_1 , the probability of dQuotesuccess in group 1. The default value is p.or.p1=0.5. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed.
p0.or.p2	numeric vector of proportions. When sample.type="one.sample", this argument denotes the hypothesized value of p , the probability of dQuotesuccess. When sample.type="two.sample", this argument denotes the value of p_2 , the probability of dQuotesuccess in group 2. The default value is p0.or.p2=0.5. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed.
alpha	numeric vector of numbers between 0 and 1 indicating the Type I error level associated with the hypothesis test. The default value is alpha=0.05.
power	numeric vector of numbers between 0 and 1 indicating the power associated with the hypothesis test. The default value is power=0.95.
sample.type	character string indicating whether to compute sample size based on a one-sample or two-sample hypothesis test. When sample.type="one.sample", the computed sample size is based on a hypothesis test for a single proportion. When sample.type="two.sample", the computed sample size is based on a hypothesis test for the difference between two proportions. The default value is sample.type="one.sample".
alternative	character string indicating the kind of alternative hypothesis. The possible values are "two.sided" (the default), "less", and "greater".
ratio	numeric vector indicating the ratio of sample size in group 2 to sample size in group 1 (n_2/n_1). The default value is ratio=1. All values of ratio must be greater than or equal to 1. This argument is ignored if sample.type="one.sample".

approx	logical scalar indicating whether to compute the sample size based on the normal approximation to the binomial distribution. The default value is approx=TRUE. Currently, the exact method (approx=FALSE) is only available for the one-sample case (i.e., sample.type="one.sample").
correct	logical scalar indicating whether to use the continuity correction when approx=TRUE. The default value is approx=TRUE when sample.type="two.sample" and approx=FALSE when sample.type="one.sample". This argument is ignored when approx=FALSE.
round.up	logical scalar indicating whether to round up the values of the computed sample size(s) to the next smallest integer. The default value is round.up=TRUE.
warn	logical scalar indicating whether to issue a warning. The default value is warn=TRUE. When approx=TRUE (sample size based on the normal approximation) and warn=T, a warning is issued for cases when the normal approximation to the binomial distribution probably is not accurate. When approx=FALSE (sample size based on the exact test) and warn=TRUE, a warning is issued when the user-supplied sample size is too small to yield a significance level less than or equal to the user-supplied value of alpha.
return.exact.list	logical scalar relevant to the case when approx=FALSE (i.e., when the power is based on the exact test). This argument indicates whether to return a list containing extra information about the exact test in addition to the power of the exact test. By default, propTestN returns only a vector containing the computed sample size(s) (see the VALUE section below). When return.exact.list=TRUE (the default) and approx=FALSE, propTestN returns a list with components indicating the required sample size, power of the exact test, the true significance level associated with the exact test, and the critical values associated with the exact test (see the DETAILS section for more information).
n.min	integer relevant to the case when approx=FALSE (i.e., when the power is based on the exact test). This argument indicates the minimum allowed value for n to use in the search algorithm. The default value is n.min=2.
n.max	integer relevant to the case when approx=FALSE (i.e., when the power is based on the exact test). This argument indicates the maximum allowed value for n to use in the search algorithm. The default value is n.max=10000.
tol.alpha	numeric vector relevant to the case when approx=FALSE (i.e., when the power is based on the exact test). This argument indicates the tolerance on alpha to use in the search algorithm (i.e., how close the actual Type I error level is to the value prescribed by alpha. The default value is tol.alpha=0.1*alpha.
tol	numeric scalar relevant to the case when approx=FALSE (i.e., when the power is based on the exact test). This argument is passed to the uniroot function and indicates the tolerance to use in the search algorithm. The default value is tol=1e-7.
maxiter	integer relevant to the case when approx=FALSE (i.e., when the power is based on the exact test). This argument is passed to the uniroot function and indicates the maximum number of iterations to use in the search algorithm. The default value is maxiter=1000.

Details

If the arguments `p.or.p1`, `p0.or.p2`, `alpha`, `power`, `ratio`, and `tol.alpha` are not all the same length, they are replicated to be the same length as the length of the longest argument.

One-Sample Case (`sample.type="one.sample"`).

`approx=TRUE` When `sample.type="one.sample"` and `approx=TRUE`, sample size is computed based on the test that uses the normal approximation to the binomial distribution; see the help file for [prop.test](#). The formula for this test and the associated power is presented in standard statistics texts, including Zar (2010, pp. 534-537, 539-541). These equations can be inverted to solve for the sample size, given a specified power, significance level, hypothesized proportion, and true proportion.

`approx=FALSE` When `sample.type="one.sample"` and `approx=FALSE`, sample size is computed based on the exact binomial test; see the help file for [binom.test](#). The formula for this test and its associated power is presented in standard statistics texts, including Zar (2010, pp. 532-534, 539) and Millard and Neerchal (2001, pp. 385-386, 504-506). The formula for the power involves five quantities: the hypothesized proportion (p_0), the true proportion (p), the significance level (α), the power, and the sample size (n). In this case the function `propTestN` uses a search algorithm to determine the required sample size to attain a specified power, given the values of the hypothesized and true proportions and the significance level.

Two-Sample Case (`sample.type="two.sample"`).

When `sample.type="two.sample"`, sample size is computed based on the test that uses the normal approximation to the binomial distribution; see the help file for [prop.test](#). The formula for this test and its associated power is presented in standard statistics texts, including Zar (2010, pp. 549-550, 552-553) and Millard and Neerchal (2001, pp. 443-445, 508-510). These equations can be inverted to solve for the sample size, given a specified power, significance level, true proportions, and ratio of sample size in group 2 to sample size in group 1.

Value

Approximate Test (`approx=TRUE`).

When `sample.type="one.sample"`, or `sample.type="two.sample"` and `ratio=1` (i.e., equal sample sizes for each group), `propTestN` returns a numeric vector of sample sizes. When `sample.type="two.sample"` and at least one element of `ratio` is greater than 1, `propTestN` returns a list with two components called `n1` and `n2`, specifying the sample sizes for each group.

Exact Test (`approx=FALSE`).

If `return.exact.list=FALSE`, `propTestN` returns a numeric vector of sample sizes.

If `return.exact.list=TRUE`, `propTestN` returns a list with the following components:

<code>n</code>	numeric vector of sample sizes.
<code>power</code>	numeric vector of powers.
<code>alpha</code>	numeric vector containing the true significance levels. Because of the discrete nature of the binomial distribution, the true significance levels usually do not equal the significance level supplied by the user in the argument <code>alpha</code> .
<code>q.critical.lower</code>	numeric vector of lower critical values for rejecting the null hypothesis. If the observed number of "successes" is <i>less than or equal to</i> these values, the null hypothesis is rejected. (Not present if <code>alternative="greater"</code> .)

`q.critical.upper`

numeric vector of upper critical values for rejecting the null hypothesis. If the observed number of "successes" is *greater than* these values, the null hypothesis is rejected. (Not present if `alternative="less"`.)

Note

The binomial distribution is used to model processes with binary (Yes-No, Success-Failure, Heads-Tails, etc.) outcomes. It is assumed that the outcome of any one trial is independent of any other trial, and that the probability of success, p , is the same on each trial. A binomial discrete random variable X is the number of "successes" in n independent trials. A special case of the binomial distribution occurs when $n = 1$, in which case X is also called a Bernoulli random variable.

In the context of environmental statistics, the binomial distribution is sometimes used to model the proportion of times a chemical concentration exceeds a set standard in a given period of time (e.g., Gilbert, 1987, p.143), or to compare the proportion of detects in a compliance well vs. a background well (e.g., USEPA, 1989b, Chapter 8, p.3-7).

In the course of designing a sampling program, an environmental scientist may wish to determine the relationship between sample size, power, significance level, and the difference between the hypothesized and true proportions if one of the objectives of the sampling program is to determine whether a proportion differs from a specified level or two proportions differ from each other. The functions [propTestPower](#), [propTestN](#), [propTestMdd](#), and [plotPropTestDesign](#) can be used to investigate these relationships for the case of binomial proportions.

Studying the two-sample proportion test, Haseman (1978) found that the formulas used to estimate the power that do not incorporate the continuity correction tend to underestimate the power. Casagrande, Pike, and Smith (1978) found that the formulas that do incorporate the continuity correction provide an excellent approximation.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Berthouex, P.M., and L.C. Brown. (1994). *Statistics for Environmental Engineers*. Lewis Publishers, Boca Raton, FL, Chapter 15.
- Casagrande, J.T., M.C. Pike, and P.G. Smith. (1978). An Improved Approximation Formula for Calculating Sample Sizes for Comparing Two Binomial Distributions. *Biometrics* **34**, 483-486.
- Fleiss, J. L. (1981). *Statistical Methods for Rates and Proportions*. Second Edition. John Wiley and Sons, New York, Chapters 1-2.
- Gilbert, R.O. (1987). *Statistical Methods for Environmental Pollution Monitoring*. Van Nostrand Reinhold, New York, NY.
- Haseman, J.K. (1978). Exact Sample Sizes for Use with the Fisher-Irwin Test for 2x2 Tables. *Biometrics* **34**, 106-109.
- Millard, S.P., and N. Neerchal. (2001). *Environmental Statistics with S-Plus*. CRC Press, Boca Raton, FL.
- Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ.

See Also

[propTestPower](#), [propTestMdd](#), [plotPropTestDesign](#), [prop.test](#), [binom.test](#).

Examples

```

# Look at how the required sample size of the one-sample
# proportion test with a two-sided alternative and Type I error
# set to 5% increases with increasing power:

seq(0.5, 0.9, by=0.1)
#[1] 0.5 0.6 0.7 0.8 0.9

propTestN(p.or.p1 = 0.7, p0.or.p2 = 0.5,
  power = seq(0.5, 0.9, by=0.1))
#[1] 25 31 38 47 62

#-----

# Repeat the last example, but compute the sample size based on
# the exact test instead of the approximation. Note that because
# we require the actual Type I error (alpha) to be within
# 10% of the supplied value of alpha (which is 0.05 by default),
# due to the discrete nature of the exact binomial test
# we end up with more power than we specified.

n.list <- propTestN(p.or.p1 = 0.7, p0.or.p2 = 0.5,
  power = seq(0.5, 0.9, by=0.1), approx=FALSE)

lapply(n.list, round, 3)
#$n
#[1] 37 37 44 51 65
#
#$power
#[1] 0.698 0.698 0.778 0.836 0.910
#
#$alpha
#[1] 0.047 0.047 0.049 0.049 0.046
#
#$q.critical.lower
#[1] 12 12 15 18 24
#
#$q.critical.upper
#[1] 24 24 28 32 40

#-----

# Using the example above, see how the sample size changes
# if we allow the Type I error to deviate by more than 10 percent
# of the value of alpha (i.e., by more than 0.005).

n.list <- propTestN(p.or.p1 = 0.7, p0.or.p2 = 0.5,
  power = seq(0.5, 0.9, by=0.1), approx=FALSE, tol.alpha = 0.01)

lapply(n.list, round, 3)
#$n
#[1] 25 35 42 49 65
#
#$power
#[1] 0.512 0.652 0.743 0.810 0.910
#

```

```

#$alpha
#[1] 0.043 0.041 0.044 0.044 0.046
#
#$q.critical.lower
#[1] 7 11 14 17 24
#
#$q.critical.upper
#[1] 17 23 27 31 40

rm(n.list)

#-----

# Look at how the required sample size for the two-sample
# proportion test decreases with increasing difference between
# the two population proportions:

seq(0.4, 0.1, by=-0.1)
#[1] 0.4 0.3 0.2 0.1

propTestN(p.or.p1 = seq(0.4, 0.1, by=-0.1),
  p0.or.p2 = 0.5, sample.type = "two")
#[1] 661 163 70 36
#Warning message:
#In propTestN(p.or.p1 = seq(0.4, 0.1, by = -0.1), p0.or.p2 = 0.5, :
# The computed sample sizes 'n1' and 'n2' are too small,
# relative to the given values of 'p1' and 'p2', for the normal
# approximation to work well for the following element indices:
#      4

#-----

# Look at how the required sample size for the two-sample
# proportion test decreases with increasing values of Type I error:

propTestN(p.or.p1 = 0.7, p0.or.p2 = 0.5,
  sample.type = "two",
  alpha = c(0.001, 0.01, 0.05, 0.1))
#[1] 299 221 163 137

#-----

# Modifying the example on pages 8-5 to 8-7 of USEPA (1989b),
# determine the required sample size to detect a difference in the
# proportion of detects of cadmium between the background and
# compliance wells. Set the compliance well to "group 1" and
# the background well to "group 2". Assume the true probability
# of a "detect" at the background well is 1/3, set the probability
# of a "detect" at the compliance well to 0.4 and 0.5, use a 5%
# significance level and 95% power, and use the upper
# one-sided alternative (probability of a "detect" at the compliance
# well is greater than the probability of a "detect" at the background
# well). (The original data are stored in EPA.89b.cadmium.df.)
#
# Note that the required sample size decreases from about
# 1160 at each well to about 200 at each well as the difference in
# proportions changes from (0.4 - 1/3) to (0.5 - 1/3), but both of

```

```
# these sample sizes are enormous compared to the number of samples
# usually collected in the field.
```

```
EPA.89b.cadmium.df
#   Cadmium.orig Cadmium Censored Well.type
#1          0.1   0.100    FALSE Background
#2          0.12  0.120    FALSE Background
#3           BDL   0.000     TRUE Background
# .....
#86          BDL   0.000     TRUE Compliance
#87          BDL   0.000     TRUE Compliance
#88          BDL   0.000     TRUE Compliance
```

```
attach(EPA.89b.cadmium.df)
```

```
p.hat.back <- mean(!Censored[Well.type=="Background"])
```

```
p.hat.back
#[1] 0.3333333
```

```
p.hat.comp <- mean(!Censored[Well.type=="Compliance"])
```

```
p.hat.comp
#[1] 0.375
```

```
n.back <- sum(Well.type == "Background")
```

```
n.back
#[1] 24
```

```
n.comp <- sum(Well.type == "Compliance")
```

```
n.comp
#[1] 64
```

```
propTestN(p.or.p1 = c(0.4, 0.50), p0.or.p2 = p.hat.back,
  alt="greater", sample.type="two")
#[1] 1159 199
```

```
rm(p.hat.back, p.hat.comp, n.back, n.comp)
detach("EPA.89b.cadmium.df")
```

propTestPower

Compute the Power of a One- or Two-Sample Proportion Test

Description

Compute the power of a one- or two-sample proportion test, given the sample size(s), true proportion(s), and significance level.

Usage

```
propTestPower(n.or.n1, p.or.p1 = 0.5, n2 = n.or.n1,
  p0.or.p2 = 0.5, alpha = 0.05, sample.type = "one.sample",
```

```

alternative = "two.sided", approx = TRUE,
correct = sample.type == "two.sample", warn = TRUE,
return.exact.list = TRUE)

```

Arguments

n.or.n1	numeric vector of sample sizes. When <code>sample.type="one.sample"</code> , this argument denotes n , the number of observations in the single sample. When <code>sample.type="two.sample"</code> , this argument denotes n_1 , the number of observations from group 1. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed.
p.or.p1	numeric vector of proportions. When <code>sample.type="one.sample"</code> , this argument denotes the true value of p , the probability of dQuotesuccess. When <code>sample.type="two.sample"</code> , this argument denotes the value of p_1 , the probability of dQuotesuccess in group 1. The default value is <code>p.or.p1=0.5</code> . Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed.
n2	numeric vector of sample sizes for group 2. The default value is <code>n2=n.or.n1</code> . This argument is ignored when <code>sample.type="one.sample"</code> . Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed.
p0.or.p2	numeric vector of proportions. When <code>sample.type="one.sample"</code> , this argument denotes the hypothesized value of p , the probability of dQuotesuccess. When <code>sample.type="two.sample"</code> , this argument denotes the value of p_2 , the probability of dQuotesuccess in group 2. The default value is <code>p0.or.p2=0.5</code> . Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are not allowed.
alpha	numeric vector of numbers between 0 and 1 indicating the Type I error level associated with the hypothesis test. The default value is <code>alpha=0.05</code> .
sample.type	character string indicating whether to compute power based on a one-sample or two-sample hypothesis test. When <code>sample.type="one.sample"</code> , the computed power is based on a hypothesis test for a single proportion. When <code>sample.type="two.sample"</code> , the computed power is based on a hypothesis test for the difference between two proportions. The default value is <code>sample.type="one.sample"</code> .
alternative	character string indicating the kind of alternative hypothesis. The possible values are "two.sided" (the default), "less", and "greater".
approx	logical scalar indicating whether to compute the power based on the normal approximation to the binomial distribution. The default value is <code>approx=TRUE</code> . Currently, the exact method (<code>approx=FALSE</code>) is only available for the one-sample case (i.e., <code>sample.type="one.sample"</code>).
correct	logical scalar indicating whether to use the continuity correction when <code>approx=TRUE</code> . The default value is <code>approx=TRUE</code> when <code>sample.type="two.sample"</code> and <code>approx=FALSE</code> when <code>sample.type="one.sample"</code> . This argument is ignored when <code>approx=FALSE</code> .
warn	logical scalar indicating whether to issue a warning. The default value is <code>warn=TRUE</code> . When <code>approx=TRUE</code> (power based on the normal approximation) and <code>warn=TRUE</code> , a warning is issued for cases when the normal approximation to the binomial distribution probably is not accurate. When <code>approx=FALSE</code> (power based on the exact test) and <code>warn=TRUE</code> , a warning is issued when the user-supplied sample size is too small to yield a significance level less than or equal to the user-supplied value of <code>alpha</code> .

`return.exact.list`

logical scalar relevant to the case when `approx=FALSE` (i.e., when the power is based on the exact test). This argument indicates whether to return a list containing extra information about the exact test in addition to the power of the exact test. By default, `propTestPower` returns only a vector containing the computed power(s) (see the `VALUE` section below). When `return.exact.list=TRUE` (the default) and `approx=FALSE`, `propTestPower` returns a list with components indicating the power of the exact test, the true significance level associated with the exact test, and the critical values associated with the exact test (see the `DETAILS` section for more information).

Details

If the arguments `n` or `n1`, `p` or `p1`, `n2`, `p0` or `p2`, and `alpha` are not all the same length, they are replicated to be the same length as the length of the longest argument.

One-Sample Case (`sample.type="one.sample"`).

`approx=TRUE` When `sample.type="one.sample"` and `approx=TRUE`, power is computed based on the test that uses the normal approximation to the binomial distribution; see the help file for [prop.test](#). The formula for this test and its associated power is presented in most standard statistics texts, including Zar (2010, pp. 534-537, 539-541).

`approx=FALSE` When `sample.type="one.sample"` and `approx=FALSE`, power is computed based on the exact binomial test; see the help file for [binom.test](#). The formula for this test and its associated power is presented in most standard statistics texts, including Zar (2010, pp. 532-534, 539) and Millard and Neerchal (2001, pp. 385-386, 504-506).

Two-Sample Case (`sample.type="two.sample"`).

When `sample.type="two.sample"`, power is computed based on the test that uses the normal approximation to the binomial distribution; see the help file for [prop.test](#). The formula for this test and its associated power is presented in standard statistics texts, including Zar (2010, pp. 549-550, 552-553) and Millard and Neerchal (2001, pp. 443-445, 508-510).

Value

By default, `propTestPower` returns a numeric vector of powers. For the one-sample proportion test (`sample.type="one.sample"`), when `approx=FALSE` and `return.exact.list=TRUE`, `propTestPower` returns a list with the following components:

<code>power</code>	numeric vector of powers.
<code>alpha</code>	numeric vector containing the true significance levels. Because of the discrete nature of the binomial distribution, the true significance levels usually do not equal the significance level supplied by the user in the argument <code>alpha</code> .
<code>q.critical.lower</code>	numeric vector of lower critical values for rejecting the null hypothesis. If the observed number of "successes" is <i>less than or equal to</i> these values, the null hypothesis is rejected. (Not present if <code>alternative="greater"</code> .)
<code>q.critical.upper</code>	numeric vector of upper critical values for rejecting the null hypothesis. If the observed number of "successes" is <i>greater than</i> these values, the null hypothesis is rejected. (Not present if <code>alternative="less"</code> .)

Note

The binomial distribution is used to model processes with binary (Yes-No, Success-Failure, Heads-Tails, etc.) outcomes. It is assumed that the outcome of any one trial is independent of any other trial, and that the probability of success, p , is the same on each trial. A binomial discrete random variable X is the number of "successes" in n independent trials. A special case of the binomial distribution occurs when $n = 1$, in which case X is also called a Bernoulli random variable.

In the context of environmental statistics, the binomial distribution is sometimes used to model the proportion of times a chemical concentration exceeds a set standard in a given period of time (e.g., Gilbert, 1987, p.143), or to compare the proportion of detects in a compliance well vs. a background well (e.g., USEPA, 1989b, Chapter 8, p.3-7).

In the course of designing a sampling program, an environmental scientist may wish to determine the relationship between sample size, power, significance level, and the difference between the hypothesized and true proportions if one of the objectives of the sampling program is to determine whether a proportion differs from a specified level or two proportions differ from each other. The functions `propTestPower`, `propTestN`, `propTestMdd`, and `plotPropTestDesign` can be used to investigate these relationships for the case of binomial proportions.

Studying the two-sample proportion test, Haseman (1978) found that the formulas used to estimate the power that do not incorporate the continuity correction tend to underestimate the power. Casagrande, Pike, and Smith (1978) found that the formulas that do incorporate the continuity correction provide an excellent approximation.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Berthouex, P.M., and L.C. Brown. (1994). *Statistics for Environmental Engineers*. Lewis Publishers, Boca Raton, FL, Chapter 15.
- Casagrande, J.T., M.C. Pike, and P.G. Smith. (1978). An Improved Approximation Formula for Calculating Sample Sizes for Comparing Two Binomial Distributions. *Biometrics* **34**, 483-486.
- Fleiss, J. L. (1981). *Statistical Methods for Rates and Proportions*. Second Edition. John Wiley and Sons, New York, Chapters 1-2.
- Gilbert, R.O. (1987). *Statistical Methods for Environmental Pollution Monitoring*. Van Nostrand Reinhold, New York, NY.
- Haseman, J.K. (1978). Exact Sample Sizes for Use with the Fisher-Irwin Test for 2x2 Tables. *Biometrics* **34**, 106-109.
- Millard, S.P., and N. Neerchal. (2001). *Environmental Statistics with S-Plus*. CRC Press, Boca Raton, FL.
- Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ.

See Also

`propTestN`, `propTestMdd`, `plotPropTestDesign`, `prop.test`, `binom.test`.

Examples

```

# Look at how the power of the one-sample proportion test
# increases with increasing sample size:

seq(20, 50, by=10)
#[1] 20 30 40 50

power <- propTestPower(n.or.n1 = seq(20, 50, by=10),
  p.or.p1 = 0.7, p0 = 0.5)

round(power, 2)
#[1] 0.43 0.60 0.73 0.83

rm(power)

#-----

# Repeat the last example, but compute the power based on
# the exact test instead of the approximation.
# Note that the significance level varies with sample size and
# never attains the requested level of 0.05.

prop.test.list <- propTestPower(n.or.n1 = seq(20, 50, by=10),
  p.or.p1 = 0.7, p0 = 0.5, approx=FALSE)

lapply(prop.test.list, round, 2)
#$power:
#[1] 0.42 0.59 0.70 0.78
#
#$alpha:
#[1] 0.04 0.04 0.04 0.03
#
#$q.critical.lower:
#[1] 5 9 13 17
#
#$q.critical.upper:
#[1] 14 20 26 32

rm(prop.test.list)

#-----

# Look at how the power of the two-sample proportion test
# increases with increasing difference between the two
# population proportions:

seq(0.5, 0.1, by=-0.1)
#[1] 0.5 0.4 0.3 0.2 0.1

power <- propTestPower(30, sample.type = "two",
  p.or.p1 = seq(0.5, 0.1, by=-0.1))
#Warning message:
#In propTestPower(30, sample.type = "two", p.or.p1 = seq(0.5, 0.1, :
#The sample sizes 'n1' and 'n2' are too small, relative to the given
# values of 'p1' and 'p2', for the normal approximation to work well
# for the following element indices:

```

```

#           5

round(power, 2)
#[1] 0.05 0.08 0.26 0.59 0.90

rm(power)

#-----

# Look at how the power of the two-sample proportion test
# increases with increasing values of Type I error:

power <- propTestPower(30, sample.type = "two",
  p.or.p1 = 0.7,
  alpha = c(0.001, 0.01, 0.05, 0.1))

round(power, 2)
#[1] 0.02 0.10 0.26 0.37

rm(power)

#-----

# Modifying the example on pages 8-5 to 8-7 of USEPA (1989b),
# determine how adding another 20 observations to the background
# well to increase the sample size from 24 to 44 will affect the
# power of detecting a difference in the proportion of detects of
# cadmium between the background and compliance wells. Set the
# compliance well to "group 1" and set the background well to
# "group 2". Assume the true probability of a "detect" at the
# background well is 1/3, set the probability of a "detect" at the
# compliance well to 0.4, use a 5% significance level, and use the
# upper one-sided alternative (probability of a "detect" at the
# compliance well is greater than the probability of a "detect" at
# the background well).
# (The original data are stored in EPA.89b.cadmium.df.)
#
# Note that the power does increase (from 9% to 12%), but is relatively
# very small.

EPA.89b.cadmium.df
#   Cadmium.orig Cadmium Censored Well.type
#1          0.1   0.100   FALSE Background
#2          0.12  0.120   FALSE Background
#3          BDL   0.000    TRUE Background
# .....
#86          BDL   0.000    TRUE Compliance
#87          BDL   0.000    TRUE Compliance
#88          BDL   0.000    TRUE Compliance

attach(EPA.89b.cadmium.df)

p.hat.back <- mean(!Censored[Well.type=="Background"])

p.hat.back
#[1] 0.3333333

```

```

p.hat.comp <- mean(!Censored[Well.type=="Compliance"])

p.hat.comp
#[1] 0.375

n.back <- sum(Well.type == "Background")

n.back
#[1] 24

n.comp <- sum(Well.type == "Compliance")

n.comp
#[1] 64

propTestPower(n.or.n1 = n.comp,
  p.or.p1 = 0.4,
  n2 = c(n.back, 44), p0.or.p2 = p.hat.back,
  alt="greater", sample.type="two")
#[1] 0.08953013 0.12421135

rm(p.hat.back, p.hat.comp, n.back, n.comp)
detach("EPA.89b.cadmium.df")

```

pwMoment

*Estimate Probability-Weighted Moments***Description**

Estimate the $1jk$ 'th probability-weighted moment from a random sample, where either $j = 0$, $k = 0$, or both.

Usage

```

pwMoment(x, j = 0, k = 0, method = "unbiased",
  plot.pos.cons = c(a = 0.35, b = 0), na.rm = FALSE)

```

Arguments

<code>x</code>	numeric vector of observations.
<code>j, k</code>	non-negative integers specifying the order of the moment.
<code>method</code>	character string specifying what method to use to compute the probability-weighted moment. The possible values are "unbiased" (method based on the U-statistic; the default), or "plotting.position" (method based on the plotting position formula). See the DETAILS section for more information.
<code>plot.pos.cons</code>	numeric vector of length 2 specifying the constants used in the formula for the plotting positions when <code>method="plotting.position"</code> . The default value is <code>plot.pos.cons=c(a=0.35, b=0)</code> . If this vector has a names attribute with the value <code>c("a","b")</code> or <code>c("b","a")</code> , then the elements will be matched by name in the formula for computing the plotting positions. Otherwise, the first element is mapped to the name "a" and the second element to the name "b".

	See the DETAILS section for more information. This argument is ignored if <code>method="unbiased"</code> .
<code>na.rm</code>	logical scalar indicating whether to remove missing values from <code>x</code> . If <code>na.rm=FALSE</code> (the default) and <code>x</code> contains missing values, then a missing value (NA) is returned. If <code>na.rm=TRUE</code> , missing values are removed from <code>x</code> prior to computing the probability-weighted moment.

Details

The definition of a probability-weighted moment, introduced by Greenwood et al. (1979), is as follows. Let X denote a random variable with cdf F , and let $x(p)$ denote the p 'th quantile of the distribution. Then the ijk 'th probability-weighted moment is given by:

$$M(i, j, k) = E[X^i F^j (1 - F)^k] = \int_0^1 [x(F)]^i F^j (1 - F)^k dF$$

where i , j , and k are real numbers. Note that if i is a nonnegative integer, then $M(i, 0, 0)$ is the conventional i 'th moment about the origin.

Greenwood et al. (1979) state that in the special case where i , j , and k are nonnegative integers:

$$M(i, j, k) = B(j + 1, k + 1) E[X_{j+1, j+k+1}^i]$$

where $B(a, b)$ denotes the [beta function](#) evaluated at a and b , and

$$E[X_{j+1, j+k+1}^i]$$

denotes the i 'th moment about the origin of the $(j + 1)$ 'th order statistic for a sample of size $(j + k + 1)$. In particular,

$$M(1, 0, k) = \frac{1}{k + 1} E[X_{1, k+1}]$$

$$M(1, j, 0) = \frac{1}{j + 1} E[X_{j+1, j+1}]$$

where

$$E[X_{1, k+1}]$$

denotes the expected value of the first order statistic (i.e., the minimum) in a sample of size $(k + 1)$, and

$$E[X_{j+1, j+1}]$$

denotes the expected value of the $(j + 1)$ 'th order statistic (i.e., the maximum) in a sample of size $(j + 1)$.

Unbiased Estimators (`method="unbiased"`)

Landwehr et al. (1979) show that, given a random sample of n values from some arbitrary distribution, an unbiased, distribution-free, and parameter-free estimator of $M(1, 0, k)$ is given by:

$$\hat{M}(1, 0, k) = \frac{1}{n} \sum_{i=1}^{n-k} x_{i,n} \frac{\binom{n-i}{k}}{\binom{n-1}{k}}$$

where the quantity $x_{i,n}$ denotes the i 'th order statistic in the random sample of size n . Hosking et al. (1985) note that this estimator is closely related to U-statistics (Hoeffding, 1948; Lehmann, 1975, pp. 362-371). Hosking et al. (1985) note that an unbiased, distribution-free, and parameter-free estimator of $M(1, j, 0)$ is given by:

$$\hat{M}(1, j, 0) = \frac{1}{n} \sum_{i=j+1}^n x_{i,n} \frac{\binom{i-1}{j}}{\binom{n-1}{j}}$$

Plotting-Position Estimators (method="plotting.position")

Hosking et al. (1985) propose alternative estimators of $M(1, 0, k)$ and $M(1, j, 0)$ based on plotting positions:

$$\hat{M}(1, 0, k) = \frac{1}{n} \sum_{i=1}^n (1 - p_{i,n})^k x_{i,n}$$

$$\hat{M}(1, j, 0) = \frac{1}{n} \sum_{i=1}^n p_{i,n}^j x_{i,n}$$

where

$$p_{i,n} = \hat{F}(x_{i,n})$$

denotes the plotting position of the i 'th order statistic in the random sample of size n , that is, a distribution-free estimate of the cdf of X evaluated at the i 'th order statistic. Typically, plotting positions have the form:

$$p_{i,n} = \frac{i - a}{n + b}$$

where $b > -a > -1$. For this form of plotting position, the plotting-position estimators are asymptotically equivalent to the U-statistic estimators.

Value

A numeric scalar—the value of the $1jk$ 'th probability-weighted moment as defined by Greenwood et al. (1979).

Note

Greenwood et al. (1979) introduced the concept of probability-weighted moments as a tool to derive estimates of distribution parameters for distributions that can be (perhaps only be) expressed in inverse form. The term “inverse form” simply means that instead of characterizing the distribution by the formula for its cumulative distribution function (cdf), the distribution is characterized by the formula for the p 'th quantile ($0 \leq p \leq 1$).

For distributions that can only be expressed in inverse form, moment estimates of their parameters are not available, and maximum likelihood estimates are not easy to compute. Greenwood et al. (1979) show that in these cases, it is often possible to derive expressions for the distribution parameters in terms of probability-weighted moments. Thus, for these cases the distribution parameters can be estimated based on the sample probability-weighted moments, which are fairly easy to compute. Furthermore, for distributions whose parameters can be expressed as functions of conventional moments, the method of probability-weighted moments provides an alternative to method of moments and maximum likelihood estimators.

Landwehr et al. (1979) use the method of probability-weighted moments to estimate the parameters of the [Type I Extreme Value \(Gumbel\) distribution](#).

Hosking et al. (1985) use the method of probability-weighted moments to estimate the parameters of the [generalized extreme value distribution](#).

Hosking (1990) and Hosking and Wallis (1995) show the relationship between probability-weighted moments and [L-moments](#).

Hosking and Wallis (1995) recommend using the unbiased estimators of probability-weighted moments for almost all applications.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Greenwood, J.A., J.M. Landwehr, N.C. Matalas, and J.R. Wallis. (1979). Probability Weighted Moments: Definition and Relation to Parameters of Several Distributions Expressible in Inverse Form. *Water Resources Research* **15**(5), 1049–1054.
- Hoeffding, W. (1948). A Class of Statistics with Asymptotically Normal Distribution. *Annals of Mathematical Statistics* **19**, 293–325.
- Hosking, J.R.M. (1990). L-Moments: Analysis and Estimation of Distributions Using Linear Combinations of Order Statistics. *Journal of the Royal Statistical Society, Series B* **52**(1), 105–124.
- Hosking, J.R.M., and J.R. Wallis (1995). A Comparison of Unbiased and Plotting-Position Estimators of L Moments. *Water Resources Research* **31**(8), 2019–2025.
- Hosking, J.R.M., J.R. Wallis, and E.F. Wood. (1985). Estimation of the Generalized Extreme-Value Distribution by the Method of Probability-Weighted Moments. *Technometrics* **27**(3), 251–261.
- Landwehr, J.M., N.C. Matalas, and J.R. Wallis. (1979). Probability Weighted Moments Compared With Some Traditional Techniques in Estimating Gumbel Parameters and Quantiles. *Water Resources Research* **15**(5), 1055–1064.
- Lehmann, E.L. (1975). *Nonparametrics: Statistical Methods Based on Ranks*. Holden-Day, Oakland, CA, pp.362-371.

See Also

[eevd](#), [egevd](#), [lMoment](#).

Examples

```
# Generate 20 observations from a generalized extreme value distribution
# with parameters location=10, scale=2, and shape=.25, then compute the
# 0'th, 1'st and 2'nd probability-weighted moments.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
dat <- rgev(20, location = 10, scale = 2, shape = 0.25)

pwMoment(dat)
#[1] 10.59556

pwMoment(dat, 1)
#[1] 5.798481

pwMoment(dat, 2)
#[1] 4.060574

pwMoment(dat, k = 1)
#[1] 4.797081

pwMoment(dat, k = 2)
#[1] 3.059173

pwMoment(dat, 1, method = "plotting.position")
# [1] 5.852913
```

```

pwMoment(dat, 1, method = "plotting.position",
  plot.pos = c(.325, 1))
#[1] 5.586817

#-----

# Clean Up
#-----
rm(dat)

```

qqPlot

Quantile-Quantile (Q-Q) Plot

Description

Produces a quantile-quantile (Q-Q) plot, also called a probability plot. The qqPlot function is a modified version of the R functions [qqnorm](#) and [qqplot](#). The **EnvStats** function qqPlot allows the user to specify a number of different distributions in addition to the normal distribution, and to optionally estimate the distribution parameters of the fitted distribution.

Usage

```

qqPlot(x, y = NULL, distribution = "norm", param.list = list(mean = 0, sd = 1),
  estimate.params = plot.type == "Tukey Mean-Difference Q-Q",
  est.arg.list = NULL, plot.type = "Q-Q", plot.pos.con = NULL, plot.it = TRUE,
  equal.axes = qq.line.type == "0-1" || estimate.params, add.line = FALSE,
  qq.line.type = "least squares", duplicate.points.method = "standard",
  points.col = 1, line.col = 1, line.lwd = par("cex"), line.lty = 1,
  digits = .Options$digits, ..., main = NULL, xlab = NULL, ylab = NULL,
  xlim = NULL, ylim = NULL)

```

Arguments

- | | |
|--------------|--|
| x | numeric vector of observations. When y is not supplied, x represents a sample from the hypothesized distribution specified by distribution. When y is supplied, the distribution of x is compared with the distribution of y. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed. |
| y | optional numeric vector of observations (not necessarily the same length as x). Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed. |
| distribution | when y is not supplied, a character string denoting the distribution abbreviation. The default value is distribution="norm". See the help file for Distribution.df for a list of possible distribution abbreviations. This argument is ignored if y is supplied. |
| param.list | when y is not supplied, a list with values for the parameters of the distribution. The default value is param.list=list(mean=0, sd=1). See the help file for Distribution.df for the names and possible values of the parameters associated with each distribution. This argument is ignored if y is supplied or estimate.params=TRUE. |

estimate.params

when *y* is not supplied, a logical scalar indicating whether to compute quantiles based on estimating the distribution parameters (`estimate.params=TRUE`) or using the known distribution parameters specified in `param.list` (`estimate.params=FALSE`). The default value of `estimate.params` is `FALSE` if `plot.type="Q-Q"` because the default configuration is a standard normal (`mean=0`, `sd=1`) Q-Q plot, which will yield roughly a straight line if the observations in *x* are from any normal distribution. The default value of `estimate.params` is `TRUE` if `plot.type="Tukey Mean-Difference Q-Q"`. The argument `estimate.params` is ignored if *y* is supplied.

est.arg.list

when *y* is not supplied and `estimate.params=TRUE`, a list whose components are optional arguments associated with the function used to estimate the parameters of the assumed distribution (see the help file [Estimating Distribution Parameters](#)). For example, all functions used to estimate distribution parameters have an optional argument called `method` that specifies the method to use to estimate the parameters. (See the help file for [Distribution.df](#) for a list of available estimation methods for each distribution.) To override the default estimation method, supply the argument `est.arg.list` with a component called `method`; for example `est.arg.list=list(method="mle")`. The default value is `est.arg.list=NULL` so that all default values for the estimating function are used. This argument is ignored if `estimate.params=FALSE` or *y* is supplied.

plot.type

a character string denoting the kind of plot. Possible values are "Q-Q" (Quantile-Quantile plot, the default) and "Tukey Mean-Difference Q-Q" (Tukey mean-difference Q-Q plot). This argument may be abbreviated (e.g., `plot.type="T"` to indicate a Tukey mean-difference Q-Q plot).

plot.pos.con

numeric scalar between 0 and 1 containing the value of the plotting position constant. The default value of `plot.pos.con` depends on whether the argument *y* is supplied, and if not the value of the argument distribution. When *y* is supplied, the default value is `plot.pos.con=0.5`, corresponding to Hazen plotting positions. When *y* is not supplied, for the normal, lognormal, three-parameter lognormal, zero-modified normal, and zero-modified lognormal distributions, the default value is `plot.pos.con=0.375`. For the Type I extreme value (Gumbel) distribution (`distribution="evd"`), the default value is `plot.pos.con=0.44`. For all other distributions, the default value is `plot.pos.con=0.4`.

plot.it

a logical scalar indicating whether to create a plot on the current graphics device. The default value is `plot.it=TRUE`.

equal.axes

a logical scalar indicating whether to use the same range on the *x*- and *y*-axes when `plot.type="Q-Q"`. The default value is `TRUE` if `qq.line.type="0-1"` or `estimate.params=TRUE`, otherwise it is `FALSE`. This argument is ignored if `plot.type="Tukey Mean-Difference Q-Q"`.

add.line

a logical scalar indicating whether to add a line to the plot. If `add.line=TRUE` and `plot.type="Q-Q"`, a line determined by the value of `qq.line.type` is added to the plot. If `add.line=TRUE` and `plot.type="Tukey Mean-Difference Q-Q"`, a horizontal line at $y = 0$ is added to the plot. The default value is `add.line=FALSE`.

qq.line.type

character string determining what kind of line to add to the Q-Q plot. Possible values are "least squares" (the default), "0-1" and "robust". For the value "least squares", a least squares line is fit and added. For the value "0-1", a line with intercept 0 and slope 1 is added. For the value "robust", a line is fit through the first and third quartiles of the *x* and *y* data. This argument is ignored if `add.line=FALSE` or `plot.type="Tukey Mean-Difference Q-Q"`.

<code>duplicate.points.method</code>	a character string denoting how to plot points with duplicate (x, y) values. Possible values are "standard" (the default), "jitter", and "number". For the value "standard", a single plotting symbol is plotted (this is the default behavior of R). For the value "jitter", a separate plotting symbol is plotted for each duplicate point, where the plotting symbols cluster around the true value of x and y . For the value "number", a single number is plotted at (x, y) that represents how many duplicate points are at that (x, y) coordinate.
<code>points.col</code>	a numeric scalar or character string determining the color of the points in the plot. The default value is <code>points.col=1</code> . See the entry for <code>col</code> in the help file for par for more information.
<code>line.col</code>	a numeric scalar or character string determining the color of the line in the plot. The default value is <code>points.col=1</code> . See the entry for <code>col</code> in the help file for par for more information. This argument is ignored if <code>add.line=FALSE</code> .
<code>line.lwd</code>	a numeric scalar determining the width of the line in the plot. The default value is <code>line.lwd=par("cex")</code> . See the entry for <code>lwd</code> in the help file for par for more information. This argument is ignored if <code>add.line=FALSE</code> .
<code>line.lty</code>	a numeric scalar determining the line type of the line in the plot. The default value is <code>line.lty=1</code> . See the entry for <code>lty</code> in the help file for par for more information. This argument is ignored if <code>add.line=FALSE</code> .
<code>digits</code>	a scalar indicating how many significant digits to print for the distribution parameters. The default value is <code>digits=Options\$digits</code> .
<code>main, xlab, ylab, xlim, ylim, ...</code>	additional graphical parameters (see par).

Details

If y is not supplied, the vector x is assumed to be a sample from the probability distribution specified by the argument `distribution` (and `param.list` if `estimate.params=FALSE`). When `plot.type="Q-Q"`, the quantiles of x are plotted on the y -axis against the quantiles of the assumed distribution on the x -axis.

If y is supplied and `plot.type="Q-Q"`, the empirical quantiles of y are plotted against the empirical quantiles of x .

When `plot.type="Tukey Mean-Difference Q-Q"`, the difference of the quantiles is plotted on the y -axis against the mean of the quantiles on the x -axis.

Special Distributions

When y is not supplied and the argument `distribution` specifies one of the following distributions, the function `qqPlot` behaves in the manner described below.

"lnorm" *Lognormal Distribution*. The log-transformed quantiles are plotted against quantiles from a Normal (Gaussian) distribution.

"lnormAlt" *Lognormal Distribution (alternative parameterization)*. The untransformed quantiles are plotted against quantiles from a Lognormal distribution.

"lnorm3" *Three-Parameter Lognormal Distribution*. The quantiles of $\log(x - \text{threshold})$ are plotted against quantiles from a Normal (Gaussian) distribution. The value of `threshold` is either specified in the argument `param.list`, or, if `estimate.params=TRUE`, then it is estimated.

"zmnorm" *Zero-Modified Normal Distribution*. The quantiles of the non-zero values (i.e., $x[x \neq 0]$) are plotted against quantiles from a Normal (Gaussian) distribution.

"zmlnorm" *Zero-Modified Lognormal Distribution*. The quantiles of the log-transformed positive values (i.e., $\log(x[x>0])$) are plotted against quantiles from a Normal (Gaussian) distribution.

"zmlnormAlt" *Lognormal Distribution (alternative parameterization)*. The quantiles of the untransformed positive values (i.e., $x[x>0]$) are plotted against quantiles from a Lognormal distribution.

Explanation of Q-Q Plots

A **probability plot** or **quantile-quantile (Q-Q) plot** is a graphical display invented by Wilk and Gnanadesikan (1968) to compare a data set to a particular probability distribution or to compare it to another data set. The idea is that if two population distributions are exactly the same, then they have the same quantiles (percentiles), so a plot of the quantiles for the first distribution vs. the quantiles for the second distribution will fall on the 0-1 line (i.e., the straight line $y = x$ with intercept 0 and slope 1). If the two distributions have the same shape and spread but different locations, then the plot of the quantiles will fall on the line $y = x + b$ (parallel to the 0-1 line) where b denotes the difference in locations. If the distributions have different locations and differ by a multiplicative constant m , then the plot of the quantiles will fall on the line $y = mx + b$ (D'Agostino, 1986a, p. 25; Helsel and Hirsch, 1986, p. 42). Various kinds of differences between distributions will yield various kinds of deviations from a straight line.

Comparing Observations to a Hypothesized Distribution

Let $\underline{x} = x_1, x_2, \dots, x_n$ denote the observations in a random sample of size n from some unknown distribution with cumulative distribution function $F()$, and let $x_{(1)}, x_{(2)}, \dots, x_{(n)}$ denote the ordered observations. Depending on the particular formula used for the empirical cdf (see [ecdfPlot](#)), the i 'th order statistic is an estimate of the $i/(n+1)$ 'th, $(i-0.5)/n$ 'th, etc., quantile. For the moment, assume the i 'th order statistic is an estimate of the $i/(n+1)$ 'th quantile, that is:

$$\hat{F}[x_{(i)}] = \hat{p}_i = \frac{i}{n+1} \quad (1)$$

so

$$x_{(i)} \approx F^{-1}(\hat{p}_i) \quad (2)$$

If we knew the form of the true cdf F , then the plot of $x_{(i)}$ vs. $F^{-1}(\hat{p}_i)$ would form approximately a straight line based on Equation (2) above. A probability plot is a plot of $x_{(i)}$ vs. $F_0^{-1}(\hat{p}_i)$, where F_0 denotes the cdf associated with the hypothesized distribution. The probability plot should fall roughly on the line $y = x$ if $F = F_0$. If F and F_0 merely differ by a shift in location and scale, that is, if $F[(x - \mu)/\sigma] = F_0(x)$, then the plot should fall roughly on the line $y = \sigma x + \mu$.

The quantity $\hat{p}_i = i/(n+1)$ in Equation (1) above is called the **plotting position** for the probability plot. This particular formula for the plotting position is appealing because it can be shown that for any continuous distribution

$$E\{F[x_{(i)}]\} = \frac{i}{n+1} \quad (3)$$

(Nelson, 1982, pp. 299-300; Stedinger et al., 1993). That is, the i 'th plotting position defined as in Equation (1) is the expected value of the true cdf evaluated at the i 'th order statistic. Many authors and practitioners, however, prefer to use a plotting position that satisfies:

$$F^{-1}(\hat{p}_i) = E[x_{(i)}] \quad (4)$$

or one that satisfies

$$F^{-1}(\hat{p}_i) = M[x_{(i)}] = F^{-1}\{M[u_{(i)}]\} \quad (5)$$

where $M[x_{(i)}]$ denotes the median of the distribution of the i 'th order statistic, and $u_{(i)}$ denotes the i 'th order statistic in a random sample of n [uniform \(0,1\)](#) random variates.

The plotting positions in Equation (4) are often approximated since the expected value of the i 'th order statistic is often difficult and time-consuming to compute. Note that these plotting positions will differ for different distributions.

The plotting positions in Equation (5) were recommended by Filliben (1975) because they require computing or approximating only the medians of [uniform \(0,1\)](#) order statistics, no matter what the form of the assumed cdf F_0 . Also, the median may be preferred as a measure of central tendency because the distributions of most order statistics are skewed.

Most plotting positions can be written as:

$$\hat{p}_i = \frac{i - a}{n - 2a + 1} \quad (6)$$

where $0 \leq a \leq 1$ (D'Agostino, 1986a, p.25; Stedinger et al., 1993). The quantity a is sometimes called the “plotting position constant”, and is determined by the argument `plot.pos.con` in the function `qqPlot`. The table below, adapted from Stedinger et al. (1993), displays commonly used plotting positions based on equation (6) for several distributions.

Name	a	Distribution Often Used With	References
Weibull	0	Weibull, Uniform	Weibull (1939), Stedinger et al. (1993)
Median	0.3175	Several	Filliben (1975), Vogel (1986)
Blom	0.375	Normal and Others	Blom (1958), Looney and Gullledge (1985)
Cunnane	0.4	Several	Cunnane (1978), Chowdhury et al. (1991)
Gringorten	0.44	Gumbel	Gringorton (1963), Vogel (1986)
Hazen	0.5	Several	Hazen (1914), Chambers et al. (1983), Cleveland (1993)

For moderate and large sample sizes, there is very little difference in visual appearance of the Q-Q plot for different choices of plotting positions.

Comparing Two Data Sets

Let $\underline{x} = x_1, x_2, \dots, x_n$ denote the observations in a random sample of size n from some unknown distribution with cumulative distribution function $F(\cdot)$, and let $x_{(1)}, x_{(2)}, \dots, x_{(n)}$ denote the ordered observations. Similarly, let $\underline{y} = y_1, y_2, \dots, y_m$ denote the observations in a random sample of size m from some unknown distribution with cumulative distribution function $G(\cdot)$, and let $y_{(1)}, y_{(2)}, \dots, y_{(m)}$ denote the ordered observations. Suppose we are interested in investigating whether the shape of the distribution with cdf F is the same as the shape of the distribution with cdf G (e.g., F and G may both be normal distributions but differ in mean and standard deviation).

When $n = m$, we can visually explore this question by plotting $y_{(i)}$ vs. $x_{(i)}$, for $i = 1, 2, \dots, n$. The values in \underline{y} are spread out in a certain way depending on the true distribution: they may be more or less symmetric about some value (the population mean or median) or they may be skewed to the right or left; they may be concentrated close to the mean or median (platykurtic) or there may be several observations “far away” from the mean or median on either side (leptokurtic). Similarly, the values in \underline{x} are spread out in a certain way. If the values in \underline{x} and \underline{y} are spread out in the same way, then the plot of $y_{(i)}$ vs. $x_{(i)}$ will be approximately a straight line. If the cdf F is exactly the

same as the cdf G , then the plot of $y_{(i)}$ vs. $x_{(i)}$ will fall roughly on the straight line $y = x$. If F and G differ by a shift in location and scale, that is, if $F[(x - \mu)/\sigma] = G(x)$, then the plot will fall roughly on the line $y = \sigma x + \mu$.

When $n > m$, a slight adjustment has to be made to produce the plot. Let $\hat{p}_1, \hat{p}_2, \dots, \hat{p}_m$ denote the plotting positions corresponding to the m empirical quantiles for the y 's and let $\hat{p}_1^*, \hat{p}_2^*, \dots, \hat{p}_n^*$ denote the plotting positions corresponding to the n empirical quantiles for the x 's. Then we plot $y_{(j)}$ vs. $x_{(j)}^*$ for $j = 1, 2, \dots, m$ where

$$x_{(j)}^* = (1 - r)x_{(i)} + rx_{(i+1)} \quad (7)$$

$$r = \frac{\hat{p}_j - \hat{p}_i^*}{\hat{p}_{i+1}^* - \hat{p}_i^*} \quad (8)$$

$$\hat{p}_i^* \leq \hat{p}_j \leq \hat{p}_{i+1}^* \quad (9)$$

That is, the values for the $x_{(j)}^*$'s are determined by linear interpolation based on the values of the plotting positions for \underline{x} and \underline{y} .

A similar adjustment is made when $n < m$.

Note that the R function `qqplot` uses a different method than the one in Equation (7) above; it uses linear interpolation based on `1:n` and `m` by calling the `approx` function.

Value

`qqPlot` returns a list with components `x` and `y`, giving the (x, y) coordinates of the points that have been or would have been plotted. There are four cases to consider:

1. The argument `y` is not supplied and `plot.type="Q-Q"`.

<code>x</code>	the quantiles from the theoretical distribution.
<code>y</code>	the observed quantiles (order statistics) based on the data in the argument <code>x</code> .

2. The argument `y` is not supplied and `plot.type="Tukey Mean-Difference Q-Q"`.

<code>x</code>	the averages of the observed and theoretical quantiles.
<code>y</code>	the differences between the observed quantiles (order statistics) and the theoretical quantiles.

3. The argument `y` is supplied and `plot.type="Q-Q"`.

<code>x</code>	the observed quantiles based on the data in the argument <code>x</code> . Note that these are adjusted quantiles if the number of observations in the argument <code>x</code> is greater than the number of observations in the argument <code>y</code> .
<code>y</code>	the observed quantiles based on the data in the argument <code>y</code> . Note that these are adjusted quantiles if the number of observations in the argument <code>y</code> is greater than the number of observations in the argument <code>x</code> .

4. The argument `y` is supplied and `plot.type="Tukey Mean-Difference Q-Q"`.

<code>x</code>	the averages of the quantiles based on the argument <code>x</code> and the quantiles based on the argument <code>y</code> .
<code>y</code>	the differences between the quantiles based on the argument <code>x</code> and the quantiles based on the argument <code>y</code> .

Note

A *quantile-quantile (Q-Q) plot*, also called a *probability plot*, is a plot of the observed order statistics from a random sample (the empirical quantiles) against their (estimated) mean or median values based on an assumed distribution, or against the empirical quantiles of another set of data (Wilk and Gnanadesikan, 1968). Q-Q plots are used to assess whether data come from a particular distribution, or whether two datasets have the same parent distribution. If the distributions have the same shape (but not necessarily the same location or scale parameters), then the plot will fall roughly on a straight line. If the distributions are exactly the same, then the plot will fall roughly on the straight line $y = x$.

A *Tukey mean-difference Q-Q plot*, also called an *m-d plot*, is a modification of a Q-Q plot. Rather than plotting observed quantiles vs. theoretical quantiles or observed y -quantiles vs. observed x -quantiles, a Tukey mean-difference Q-Q plot plots the difference between the quantiles on the y -axis vs. the average of the quantiles on the x -axis (Cleveland, 1993, pp.22-23). If the two sets of quantiles come from the same parent distribution, then the points in this plot should fall roughly along the horizontal line $y = 0$. If one set of quantiles come from the same distribution with a shift in median, then the points in this plot should fall along a horizontal line above or below the line $y = 0$. A Tukey mean-difference Q-Q plot enhances our perception of how the points in the Q-Q plot deviate from a straight line, because it is easier to judge deviations from a horizontal line than from a line with a non-zero slope.

In a Q-Q plot, the extreme points have more variability than points toward the center. A U-shaped Q-Q plot indicates that the underlying distribution for the observations on the y -axis is skewed to the right relative to the underlying distribution for the observations on the x -axis. An upside-down-U-shaped Q-Q plot indicates the y -axis distribution is skewed left relative to the x -axis distribution. An S-shaped Q-Q plot indicates the y -axis distribution has shorter tails than the x -axis distribution. Conversely, a plot that is bent down on the left and bent up on the right indicates that the y -axis distribution has longer tails than the x -axis distribution.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Chambers, J.M., W.S. Cleveland, B. Kleiner, and P.A. Tukey. (1983). *Graphical Methods for Data Analysis*. Duxbury Press, Boston, MA, pp.11-16.
- Cleveland, W.S. (1993). *Visualizing Data*. Hobart Press, Summit, New Jersey, 360pp.
- D'Agostino, R.B. (1986a). Graphical Analysis. In: D'Agostino, R.B., and M.A. Stephens, eds. *Goodness-of-Fit Techniques*. Marcel Dekker, New York, Chapter 2, pp.7-62.

See Also

[ppoints](#), [ecdfPlot](#), [Distribution.df](#), [qqPlotGestalt](#), [qqPlotCensored](#), [qqnorm](#).

Examples

```
# The guidance document USEPA (1994b, pp. 6.22--6.25)
# contains measures of 1,2,3,4-Tetrachlorobenzene (TcCB)
# concentrations (in parts per billion) from soil samples
# at a Reference area and a Cleanup area. These data are stored
# in the data frame EPA.94b.tccb.df.
#
# Create an Q-Q plot for the reference area data first assuming a
```

```

# normal distribution, then a lognormal distribution, then a
# gamma distribution.

# Assume a normal distribution
#-----

windows()
with(EPA.94b.tccb.df, qqPlot(TcCB[Area == "Reference"]))

windows()
with(EPA.94b.tccb.df, qqPlot(TcCB[Area == "Reference"], add.line = TRUE))

windows()
with(EPA.94b.tccb.df, qqPlot(TcCB[Area == "Reference"],
  plot.type = "Tukey", add.line = TRUE))

# The Q-Q plot based on assuming a normal distribution shows a U-shape,
# indicating the Reference area TcCB data are skewed to the right
# compared to a normal distribuiton.

# Assume a lognormal distribution
#-----

windows()
with(EPA.94b.tccb.df,
  qqPlot(TcCB[Area == "Reference"], dist = "lnorm",
    digits = 2, points.col = "blue", add.line = TRUE))

windows()
with(EPA.94b.tccb.df,
  qqPlot(TcCB[Area == "Reference"], dist = "lnorm",
    digits = 2, plot.type = "Tukey", points.col = "blue",
    add.line = TRUE))

# Alternative parameterization

windows()
with(EPA.94b.tccb.df,
  qqPlot(TcCB[Area == "Reference"], dist = "lnormAlt",
    estimate.params = TRUE, digits = 2, points.col = "blue",
    add.line = TRUE))

windows()
with(EPA.94b.tccb.df,
  qqPlot(TcCB[Area == "Reference"], dist = "lnormAlt",
    digits = 2, plot.type = "Tukey", points.col = "blue",
    add.line = TRUE))

# The lognormal distribution appears to be an adequate fit.
# Now look at a Q-Q plot assuming a gamma distribution.
#-----

windows()
with(EPA.94b.tccb.df,
  qqPlot(TcCB[Area == "Reference"], dist = "gamma",

```

```

        estimate.params = TRUE, digits = 2, points.col = "blue",
        add.line = TRUE))

windows()
with(EPA.94b.tccb.df,
     qqPlot(TcCB[Area == "Reference"], dist = "gamma",
            digits = 2, plot.type = "Tukey", points.col = "blue",
            add.line = TRUE))

# Alternative Parameterization

windows()
with(EPA.94b.tccb.df,
     qqPlot(TcCB[Area == "Reference"], dist = "gammaAlt",
            estimate.params = TRUE, digits = 2, points.col = "blue",
            add.line = TRUE))

windows()
with(EPA.94b.tccb.df,
     qqPlot(TcCB[Area == "Reference"], dist = "gammaAlt",
            digits = 2, plot.type = "Tukey", points.col = "blue",
            add.line = TRUE))

#-----

# Generate 20 observations from a gamma distribution with parameters
# shape=2 and scale=2, then create a normal (Gaussian) Q-Q plot for these data.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(357)
dat <- rgamma(20, shape=2, scale=2)
windows()
qqPlot(dat, add.line = TRUE)

# Now assume a gamma distribution and estimate the parameters
#-----

windows()
qqPlot(dat, dist = "gamma", estimate.params = TRUE, add.line = TRUE)

# Clean up
#-----
rm(dat)
graphics.off()

```

Description

Produces a quantile-quantile (Q-Q) plot, also called a probability plot, for Type I censored data.

Usage

```
qqPlotCensored(x, censored, censoring.side = "left",
  prob.method = "michael-schucany", plot.pos.con = NULL,
  distribution = "norm", param.list = list(mean = 0, sd = 1),
  estimate.params = plot.type == "Tukey Mean-Difference Q-Q",
  est.arg.list = NULL, plot.type = "Q-Q", plot.it = TRUE,
  equal.axes = qq.line.type == "0-1" || estimate.params,
  add.line = FALSE, qq.line.type = "least squares",
  duplicate.points.method = "standard", points.col = 1, line.col = 1,
  line.lwd = par("cex"), line.lty = 1, digits = .Options$digits,
  include.cen = FALSE, cen.pch = ifelse(censoring.side == "left", 6, 2),
  cen.cex = par("cex"), cen.col = 4, ..., main = NULL, xlab = NULL,
  ylab = NULL, xlim = NULL, ylim = NULL)
```

Arguments

- | | |
|----------------|---|
| x | numeric vector of observations that is assumed to represent a sample from the hypothesized distribution specified by <code>distribution</code> . Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed. |
| censored | numeric or logical vector indicating which values of x are censored. This must be the same length as x. If the mode of censored is "logical", TRUE values correspond to elements of x that are censored, and FALSE values correspond to elements of x that are not censored. If the mode of censored is "numeric", it must contain only 1's and 0's; 1 corresponds to TRUE and 0 corresponds to FALSE. Missing (NA) values are allowed but will be removed. |
| censoring.side | character string indicating on which side the censoring occurs. The possible values are "left" (the default) and "right". |
| prob.method | character string indicating what method to use to compute the plotting positions (empirical probabilities). Possible values are "kaplan-meier" (product-limit method of Kaplan and Meier (1958)), "kaplan-meier with max" (same as "kaplan-meier" except the maximum value is plotted too), "nelson" (hazard plotting method of Nelson (1972)), "michael-schucany" (generalization of the product-limit method due to Michael and Schucany (1986)), and "hirsch-stedinger" (generalization of the product-limit method due to Hirsch and Stedinger (1987)). The default value is <code>prob.method="michael-schucany"</code> .

The "nelson" method is only available for <code>censoring.side="right"</code> , and the "kaplan-meier with max" method is only available for <code>censoring.side="left"</code> . See the DETAILS section for more explanation. |
| plot.pos.con | numeric scalar between 0 and 1 containing the value of the plotting position constant. The default value is <code>plot.pos.con=0.375</code> . See the DETAILS section for more information. This argument is used only if <code>prob.method</code> is equal to "michael-schucany" or "hirsch-stedinger". |
| distribution | a character string denoting the distribution abbreviation. The default value is <code>distribution="norm"</code> . See the help file for Distribution.df for a list of possible distribution abbreviations. |
| param.list | a list with values for the parameters of the distribution. The default value is <code>param.list=list(mean=0, sd=1)</code> . See the help file for Distribution.df for the names and possible values of the parameters associated with each distribution. This argument is ignored if <code>estimate.params=TRUE</code> . |

`estimate.params`

a logical scalar indicating whether to compute quantiles based on estimating the distribution parameters (`estimate.params=TRUE`) or using the known distribution parameters specified in `param.list` (`estimate.params=FALSE`, the default). The default value of `estimate.params` is `FALSE` if `plot.type="Q-Q"` because the default configuration is a standard normal (`mean=0`, `sd=1`) Q-Q plot, which will yield roughly a straight line if the observations in `x` are from any normal distribution. The default value of `estimate.params` is `TRUE` if `plot.type="Tukey Mean-Difference Q-Q"`.

You can set `estimate.params=TRUE` only when the argument `distribution` specifies a distribution that has an associated function for estimating distribution parameters in the case of Type I censored data. Currently this includes the normal (`dist="norm"`), lognormal (`dist="lnorm"` or `dist="lnormAlt"`), and Poisson (`dist="pois"`) distributions (see the section *Estimating Distribution Parameters* in the help file [EnvStats Functions for Censored Data](#)).

`est.arg.list`

a list whose components are optional arguments associated with the function used to estimate the parameters of the assumed distribution (see the section *Estimating Distribution Parameters* in the help file [EnvStats Functions for Censored Data](#)). For example, the function `enormCensored` has an optional argument called `method` that specifies the method to use to estimate the parameters. To override the default estimation method, supply the argument `est.arg.list` with a component called `method`; for example `est.arg.list=list(method="impute.w.qq.reg")`. The default value is `est.arg.list=NULL` so that all default values for the estimating function are used. This argument is ignored if `estimate.params=FALSE`.

`plot.type`

a character string denoting the kind of plot. Possible values are `"Q-Q"` (Quantile-Quantile plot, the default) and `"Tukey Mean-Difference Q-Q"` (Tukey mean-difference Q-Q plot). This argument may be abbreviated (e.g., `plot.type="T"`) to indicate a Tukey mean-difference Q-Q plot).

`plot.it`

a logical scalar indicating whether to create a plot on the current graphics device. The default value is `plot.it=TRUE`.

`equal.axes`

a logical scalar indicating whether to use the same range on the x - and y -axes when `plot.type="Q-Q"`. The default value is `TRUE` if `qq.line.type="0-1"` or `estimate.params=TRUE`, otherwise it is `FALSE`. This argument is ignored if `plot.type="Tukey Mean-Difference Q-Q"`.

`add.line`

a logical scalar indicating whether to add a line to the plot. If `add.line=TRUE` and `plot.type="Q-Q"`, a line determined by the value of `qq.line.type` is added to the plot. If `add.line=TRUE` and `plot.type="Tukey Mean-Difference Q-Q"`, a horizontal line at $y = 0$ is added to the plot. The default value is `add.line=FALSE`.

`qq.line.type`

character string determining what kind of line to add to the Q-Q plot. Possible values are `"least squares"` (the default), `"0-1"` and `"robust"`. For the value `"least squares"`, a least squares line is fit and added. For the value `"0-1"`, a line with intercept 0 and slope 1 is added. For the value `"robust"`, a line is fit through the first and third quartiles of the x and y data. This argument is ignored if `add.line=FALSE` or `plot.type="Tukey Mean-Difference Q-Q"`.

`duplicate.points.method`

a character string denoting how to plot points with duplicate (x, y) values. Possible values are `"standard"` (the default), `"jitter"`, and `"number"`. For the value `"standard"`, a single plotting symbol is plotted (this is the default behavior of R). For the value `"jitter"`, a separate plotting symbol is plotted for each duplicate point, where the plotting symbols cluster around the true value

	of x and y . For the value "number", a single number is plotted at (x, y) that represents how many duplicate points are at that (x, y) coordinate.
<code>points.col</code>	a numeric scalar or character string determining the color of the points in the plot. The default value is <code>points.col=1</code> . See the entry for <code>col</code> in the help file for par for more information.
<code>line.col</code>	a numeric scalar or character string determining the color of the line in the plot. The default value is <code>points.col=1</code> . See the entry for <code>col</code> in the help file for par for more information. This argument is ignored if <code>add.line=FALSE</code> .
<code>line.lwd</code>	a numeric scalar determining the width of the line in the plot. The default value is <code>line.lwd=par("cex")</code> . See the entry for <code>lwd</code> in the help file for par for more information. This argument is ignored if <code>add.line=FALSE</code> .
<code>line.lty</code>	a numeric scalar determining the line type of the line in the plot. The default value is <code>line.lty=1</code> . See the entry for <code>lty</code> in the help file for par for more information. This argument is ignored if <code>add.line=FALSE</code> .
<code>digits</code>	a scalar indicating how many significant digits to print for the distribution parameters. The default value is <code>digits=.Options\$digits</code> .
<code>include.cen</code>	logical scalar indicating whether to include censored values in the plot. The default value is <code>include.cen=FALSE</code> . If <code>include.cen=TRUE</code> , censored values are plotted using the plotting character indicated by the argument <code>cen.pch</code> (see below).
<code>cen.pch</code>	numeric scalar or character string indicating the plotting character to use to plot censored values. The default value is <code>cen.pch=2</code> (hollow triangle pointing up) when <code>censoring.side="right"</code> , and <code>cen.pch=6</code> (hollow triangle pointing down) when <code>censoring.side="left"</code> . See the help file for points for a list of other possible plotting characters. This argument is ignored if <code>include.cen=FALSE</code> .
<code>cen.cex</code>	numeric scalar that determines the size of the plotting character used to plot censored values. The default value is the current value of the <code>cex</code> graphics parameter. See the entry for <code>cex</code> in the help file for par for more information. This argument is ignored if <code>include.cen=FALSE</code> .
<code>cen.col</code>	numeric scalar or character string that determines the color of the plotting character used to plot censored values. The default value is <code>cen.col=4</code> . See the entry for <code>col</code> in the help file for par for more information. This argument is ignored if <code>include.cen=FALSE</code> .
<code>main, xlab, ylab, xlim, ylim, ...</code>	additional graphical parameters (see par).

Details

The function `qqPlotCensored` does exactly the same thing as `qqPlot` (when the argument `y` is not supplied to `qqPlot`), except `qqPlotCensored` calls the function `ppointsCensored` to compute the plotting positions (estimated cumulative probabilities).

The vector `x` is assumed to be a sample from the probability distribution specified by the argument `distribution` (and `param.list` if `estimate.params=FALSE`). When `plot.type="Q-Q"`, the quantiles of `x` are plotted on the y -axis against the quantiles of the assumed distribution on the x -axis.

When `plot.type="Tukey Mean-Difference Q-Q"`, the difference of the quantiles is plotted on the y -axis against the mean of the quantiles on the x -axis.

When `prob.method="kaplan-meier"` and `censoring.side="left"` and the assumed distribution has a maximum support of infinity (Inf; e.g., the normal or lognormal distribution), the point involving the largest value of x is not plotted because it corresponds to an estimated cumulative probability of 1 which corresponds to an infinite plotting position.

When `prob.method="kaplan-meier with max"` and `censoring.side="left"`, the estimated cumulative probability associated with the maximum value is modified from 1 to be $(n - .375)/(n + .25)$ where n denotes the sample size (i.e., the Blom plotting position) so that the point associated with the maximum value can be displayed.

Value

`qqPlotCensored` returns a list with the following components:

<code>x</code>	numeric vector of x -coordinates for the plot. When <code>plot.type="Q-Q"</code> these are the quantiles from the theoretical distribution. When <code>plot.type="Tukey Mean-Difference Q-Q"</code> these are the averages of the observed and theoretical quantiles.
<code>y</code>	numeric vector of y -coordinates for the plot. When <code>plot.type="Q-Q"</code> these are the observed quantiles (order statistics). When <code>plot.type="Tukey Mean-Difference Q-Q"</code> these are the differences between the observed quantiles (order statistics) and the theoretical quantiles.
<code>Order.Statistics</code>	numeric vector of the “ordered” observations. When <code>plot.type="Q-Q"</code> this component is exactly the same as the component <code>y</code> .
<code>Cumulative.Probabilities</code>	numeric vector of the plotting positions associated with the order statistics.
<code>Censored</code>	logical vector indicating which of the ordered observations are censored.
<code>Censoring.Side</code>	character string indicating whether the data are left- or right-censored. This is same value as the argument <code>censoring.side</code> .
<code>Prob.Method</code>	character string indicating what method was used to compute the plotting positions. This is the same value as the argument <code>prob.method</code> .
Optional Component (only present when <code>prob.method="michael-schucany"</code> or <code>prob.method="hirsch-stedinger"</code>):	
<code>Plot.Pos.Con</code>	numeric scalar containing the value of the plotting position constant that was used. This is the same as the argument <code>plot.pos.con</code> .

Note

A *quantile-quantile (Q-Q) plot*, also called a *probability plot*, is a plot of the observed order statistics from a random sample (the empirical quantiles) against their (estimated) mean or median values based on an assumed distribution, or against the empirical quantiles of another set of data (Wilk and Gnanadesikan, 1968). Q-Q plots are used to assess whether data come from a particular distribution, or whether two datasets have the same parent distribution. If the distributions have the same shape (but not necessarily the same location or scale parameters), then the plot will fall roughly on a straight line. If the distributions are exactly the same, then the plot will fall roughly on the straight line $y = x$.

A *Tukey mean-difference Q-Q plot*, also called an *m-d plot*, is a modification of a Q-Q plot. Rather than plotting observed quantiles vs. theoretical quantiles or observed y -quantiles vs. observed x -quantiles, a Tukey mean-difference Q-Q plot plots the difference between the quantiles on the y -axis vs. the average of the quantiles on the x -axis (Cleveland, 1993, pp.22-23). If the two sets

of quantiles come from the same parent distribution, then the points in this plot should fall roughly along the horizontal line $y = 0$. If one set of quantiles come from the same distribution with a shift in median, then the points in this plot should fall along a horizontal line above or below the line $y = 0$. A Tukey mean-difference Q-Q plot enhances our perception of how the points in the Q-Q plot deviate from a straight line, because it is easier to judge deviations from a horizontal line than from a line with a non-zero slope.

In a Q-Q plot, the extreme points have more variability than points toward the center. A U-shaped Q-Q plot indicates that the underlying distribution for the observations on the y -axis is skewed to the right relative to the underlying distribution for the observations on the x -axis. An upside-down-U-shaped Q-Q plot indicates the y -axis distribution is skewed left relative to the x -axis distribution. An S-shaped Q-Q plot indicates the y -axis distribution has shorter tails than the x -axis distribution. Conversely, a plot that is bent down on the left and bent up on the right indicates that the y -axis distribution has longer tails than the x -axis distribution.

Censored observations complicate the procedures used to graphically explore data. Techniques from survival analysis and life testing have been developed to generalize the procedures for constructing plotting positions, empirical cdf plots, and Q-Q plots to data sets with censored observations (see [ppointsCensored](#)).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Chambers, J.M., W.S. Cleveland, B. Kleiner, and P.A. Tukey. (1983). *Graphical Methods for Data Analysis*. Duxbury Press, Boston, MA, pp.11-16.
- Cleveland, W.S. (1993). *Visualizing Data*. Hobart Press, Summit, New Jersey, 360pp.
- D'Agostino, R.B. (1986a). Graphical Analysis. In: D'Agostino, R.B., and M.A. Stephens, eds. *Goodness-of-Fit Techniques*. Marcel Dekker, New York, Chapter 2, pp.7-62.
- Gillespie, B.W., Q. Chen, H. Reichert, A. Franzblau, E. Hedgeman, J. Lepkowski, P. Adriaens, A. Demond, W. Luksemburg, and D.H. Garabrant. (2010). Estimating Population Distributions When Some Data Are Below a Limit of Detection by Using a Reverse Kaplan-Meier Estimator. *Epidemiology* **21**(4), S64-S70.
- Helsel, D.R. (2012). *Statistics for Censored Environmental Data Using Minitab and R, Second Edition*. John Wiley & Sons, Hoboken, New Jersey.
- Helsel, D.R., and T.A. Cohn. (1988). Estimation of Descriptive Statistics for Multiply Censored Water Quality Data. *Water Resources Research* **24**(12), 1997-2004.
- Hirsch, R.M., and J.R. Stedinger. (1987). Plotting Positions for Historical Floods and Their Precision. *Water Resources Research* **23**(4), 715-727.
- Kaplan, E.L., and P. Meier. (1958). Nonparametric Estimation From Incomplete Observations. *Journal of the American Statistical Association* **53**, 457-481.
- Lee, E.T., and J. Wang. (2003). *Statistical Methods for Survival Data Analysis, Third Edition*. John Wiley and Sons, New York.
- Michael, J.R., and W.R. Schucany. (1986). Analysis of Data from Censored Samples. In D'Agostino, R.B., and M.A. Stephens, eds. *Goodness-of-Fit Techniques*. Marcel Dekker, New York, 560pp, Chapter 11, 461-496.
- Nelson, W. (1972). Theory and Applications of Hazard Plotting for Censored Failure Data. *Technometrics* **14**, 945-966.

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C. Chapter 15.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

See Also

[ppointsCensored](#), [EnvStats Functions for Censored Data](#), [qqPlot](#), [ecdfPlotCensored](#), [qqPlotGestalt](#).

Examples

```
# Generate 20 observations from a normal distribution with mean=20 and sd=5,
# censor all observations less than 18, then generate a Q-Q plot assuming
# a normal distribution for the complete data set and the censored data set.
# Note that the Q-Q plot for the censored data set starts at the first ordered
# uncensored observation, and that for values of x > 18 the two Q-Q plots are
# exactly the same. This is because there is only one censoring level and
# no uncensored observations fall below the censored observations.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(333)
x <- rnorm(20, mean=20, sd=5)
censored <- x < 18

sum(censored)
#[1] 7

new.x <- x
new.x[censored] <- 18

windows()
qqPlot(x, ylim = range(pretty(x)),
  main = "Q-Q Plot for\nComplete Data Set")

windows()
qqPlotCensored(new.x, censored, ylim = range(pretty(x)),
  main="Q-Q Plot for\nCensored Data Set")

# Clean up
#-----
rm(x, censored, new.x)

#-----

# Example 15-1 of USEPA (2009, page 15-10) gives an example of
# computing plotting positions based on censored manganese
# concentrations (ppb) in groundwater collected at 5 monitoring
# wells. The data for this example are stored in
# EPA.09.Ex.15.1.manganese.df. Here we will create a Q-Q
# plot based on the Kaplan-Meier method. First we'll assume
# a normal distribution, then a lognormal distribution, then a
# gamma distribution.
```

```

EPA.09.Ex.15.1.manganese.df
# Sample Well Manganese.Orig.ppb Manganese.ppb Censored
#1 1 Well.1 <5 5.0 TRUE
#2 2 Well.1 12.1 12.1 FALSE
#3 3 Well.1 16.9 16.9 FALSE
#4 4 Well.1 21.6 21.6 FALSE
#5 5 Well.1 <2 2.0 TRUE
#...
#21 1 Well.5 17.9 17.9 FALSE
#22 2 Well.5 22.7 22.7 FALSE
#23 3 Well.5 3.3 3.3 FALSE
#24 4 Well.5 8.4 8.4 FALSE
#25 5 Well.5 <2 2.0 TRUE

# Assume normal distribution
#-----

windows()
with(EPA.09.Ex.15.1.manganese.df,
  qqPlotCensored(Manganese.ppb, Censored,
    prob.method = "kaplan-meier", points.col = "blue", add.line = TRUE,
    main = paste("Normal Q-Q Plot of Manganese Data",
      "Based on Kaplan-Meier Plotting Positions", sep = "\n")))

# Include max value in the plot
#-----

windows()
with(EPA.09.Ex.15.1.manganese.df,
  qqPlotCensored(Manganese.ppb, Censored,
    prob.method = "kaplan-meier with max", points.col = "blue",
    add.line = TRUE,
    main = paste("Normal Q-Q Plot of Manganese Data",
      "Based on Kaplan-Meier Plotting Positions",
      "(Max Included)", sep = "\n")))

# Assume lognormal distribution
#-----

windows()
with(EPA.09.Ex.15.1.manganese.df,
  qqPlotCensored(Manganese.ppb, Censored, dist = "lnorm",
    prob.method = "kaplan-meier", points.col = "blue", add.line = TRUE,
    main = paste("Lognormal Q-Q Plot of Manganese Data",
      "Based on Kaplan-Meier Plotting Positions", sep = "\n")))

# Include max value in the plot
#-----

windows()
with(EPA.09.Ex.15.1.manganese.df,
  qqPlotCensored(Manganese.ppb, Censored, dist = "lnorm",
    prob.method = "kaplan-meier with max", points.col = "blue",
    add.line = TRUE,
    main = paste("Lognormal Q-Q Plot of Manganese Data",
      "Based on Kaplan-Meier Plotting Positions",

```

```

"(Max Included)", sep = "\n"))

# The lognormal distribution appears to be a better fit.
# Now create a Q-Q plot assuming a gamma distribution. Here we'll
# need to set estimate.params=TRUE.

windows()
with(EPA.09.Ex.15.1.manganese.df,
  qqPlotCensored(Manganese.ppb, Censored, dist = "gamma",
    estimate.params = TRUE, prob.method = "kaplan-meier",
    points.col = "blue", add.line = TRUE,
    main = paste("Gamma Q-Q Plot of Manganese Data",
      "Based on Kaplan-Meier Plotting Positions", sep = "\n"))

# Include max value in the plot
#-----

windows()
with(EPA.09.Ex.15.1.manganese.df,
  qqPlotCensored(Manganese.ppb, Censored, dist = "gamma",
    estimate.params = TRUE, prob.method = "kaplan-meier with max",
    points.col = "blue", add.line = TRUE,
    main = paste("Gamma Q-Q Plot of Manganese Data",
      "Based on Kaplan-Meier Plotting Positions",
      "(Max Included)", sep = "\n"))

#=====

# Clean up
#-----
graphics.off()

```

qqPlotGestalt

Develop Gestalt of Q-Q Plots for Specific Distributions

Description

Produce a series of quantile-quantile (Q-Q) plots (also called probability plots) or Tukey mean-difference Q-Q plots for a user-specified distribution.

Usage

```

qqPlotGestalt(distribution = "norm", param.list = list(mean = 0, sd = 1),
  estimate.params = FALSE, est.arg.list = NULL, sample.size = 10, num.pages = 2,
  num.plots.per.page = 4, nrow = ceiling(num.plots.per.page/2), plot.type = "Q-Q",
  plot.pos.con = switch(dist.abb, norm = , lnorm = , lnormAlt = , lnorm3 = 0.375,
    evd = 0.44, 0.4), equal.axes = (qq.line.type == "0-1" || estimate.params),
  margin.title = NULL, add.line = FALSE, qq.line.type = "least squares",
  duplicate.points.method = "standard", points.col = 1, line.col = 1,
  line.lwd = par("cex"), line.lty = 1, digits = .Options$digits,
  same.window = TRUE, ask = same.window & num.pages > 1,
  mfrow = c(nrow, num.plots.per.page/nrow),
  mar = c(4, 4, 1, 1) + 0.1, oma = c(0, 0, 7, 0), mgp = c(2, 0.5, 0), ...,
  main = NULL, xlab = NULL, ylab = NULL, xlim = NULL, ylim = NULL)

```

Arguments

<code>distribution</code>	a character string denoting the distribution abbreviation. The default value is <code>distribution="norm"</code> . See the help file for Distribution.df for a list of possible distribution abbreviations. This argument is ignored if <code>y</code> is supplied.
<code>param.list</code>	a list with values for the parameters of the distribution. The default value is <code>param.list=list(mean=0, sd=1)</code> . See the help file for Distribution.df for the names and possible values of the parameters associated with each distribution. This argument is ignored if <code>estimate.params=TRUE</code> .
<code>estimate.params</code>	a logical scalar indicating whether to compute quantiles based on estimating the distribution parameters (<code>estimate.params=TRUE</code>) or using the known distribution parameters specified in <code>param.list</code> (<code>estimate.params=FALSE</code> , the default). The default value of <code>estimate.params</code> is <code>FALSE</code> because the default configuration is to generate random numbers from a standard normal (<code>mean=0</code> , <code>sd=1</code>) distribution and produce a standard normal Q-Q plot.
<code>est.arg.list</code>	a list whose components are optional arguments associated with the function used to estimate the parameters of the assumed distribution (see the help file Estimating Distribution Parameters). For example, all functions used to estimate distribution parameters have an optional argument called <code>method</code> that specifies the method to use to estimate the parameters. (See the help file for Distribution.df for a list of available estimation methods for each distribution.) To override the default estimation method, supply the argument <code>est.arg.list</code> with a component called <code>method</code> ; for example <code>est.arg.list=list(method="mle")</code> . The default value is <code>est.arg.list=NULL</code> so that all default values for the estimating function are used. This argument is ignored if <code>estimate.params=FALSE</code> .
<code>sample.size</code>	numeric scalar indicating the number of observations to generate for each Q-Q plot. The default value is <code>sample.size=10</code> .
<code>num.pages</code>	numeric scalar indicating the number of pages of plots to generate. The default value is <code>num.pages=2</code> .
<code>num.plots.per.page</code>	numeric scalar indicating the number of plots per page. The default value is <code>num.pages=4</code> .
<code>nrow</code>	numeric scalar indicating the number of rows of plots on each page. The default value is the smallest integer greater than or equal to <code>num.plots.per.page/2</code> .
<code>plot.type</code>	a character string denoting the kind of plot. Possible values are "Q-Q" (Quantile-Quantile plot, the default) and "Tukey Mean-Difference Q-Q" (Tukey mean-difference Q-Q plot). This argument may be abbreviated (e.g., <code>plot.type="T"</code> to indicate a Tukey mean-difference Q-Q plot).
<code>plot.pos.con</code>	numeric scalar between 0 and 1 containing the value of the plotting position constant. The default value of <code>plot.pos.con</code> depends on the value of the argument <code>distribution</code> . For the normal, lognormal, three-parameter lognormal, zero-modified normal, and zero-modified lognormal distributions, the default value is <code>plot.pos.con=0.375</code> . For the Type I extreme value (Gumbel) distribution (<code>distribution="evd"</code>), the default value is <code>plot.pos.con=0.44</code> . For all other distributions, the default value is <code>plot.pos.con=0.4</code> . See the help file for qqPlot for the motivation behind these values for plotting positions.
<code>equal.axes</code>	logical scalar indicating whether to use the same range on the <i>x</i> - and <i>y</i> -axes when <code>plot.type="Q-Q"</code> . The default value is <code>TRUE</code> if <code>qq.line.type="0-1"</code> or <code>estimate.params=TRUE</code> , otherwise it is <code>FALSE</code> . This argument is ignored if <code>plot.type="Tukey Mean-Difference Q-Q"</code> .

<code>margin.title</code>	character string indicating the title printed in the top margin on each page of plots. The default value indicates the kind of Q-Q plot, the probability distribution, the sample size, and the estimation method used (if any).
<code>add.line</code>	logical scalar indicating whether to add a line to the plot. If <code>add.line=TRUE</code> and <code>plot.type="Q-Q"</code> , a line determined by the value of <code>qq.line.type</code> is added to the plot. If <code>add.line=TRUE</code> and <code>plot.type="Tukey Mean-Difference Q-Q"</code> , a horizontal line at $y = 0$ is added to the plot. The default value is <code>add.line=FALSE</code> .
<code>qq.line.type</code>	character string determining what kind of line to add to the Q-Q plot. Possible values are "least squares" (the default), "0-1" and "robust". For the value "least squares", a least squares line is fit and added. For the value "0-1", a line with intercept 0 and slope 1 is added. For the value "robust", a line is fit through the first and third quartiles of the x and y data. This argument is ignored if <code>add.line=FALSE</code> or <code>plot.type="Tukey Mean-Difference Q-Q"</code> .
<code>duplicate.points.method</code>	character string denoting how to plot points with duplicate (x, y) values. Possible values are "standard" (the default), "jitter", and "number". For the value "standard", a single plotting symbol is plotted (this is the default behavior of R). For the value "jitter", a separate plotting symbol is plotted for each duplicate point, where the plotting symbols cluster around the true value of x and y . For the value "number", a single number is plotted at (x, y) that represents how many duplicate points are at that (x, y) coordinate.
<code>points.col</code>	numeric scalar or character string determining the color of the points in the plot. The default value is <code>points.col=1</code> . See the entry for <code>col</code> in the help file for par for more information.
<code>line.col</code>	numeric scalar or character string determining the color of the line in the plot. The default value is <code>points.col=1</code> . See the entry for <code>col</code> in the help file for par for more information. This argument is ignored if <code>add.line=FALSE</code> .
<code>line.lwd</code>	numeric scalar determining the width of the line in the plot. The default value is <code>line.lwd=par("cex")</code> . See the entry for <code>lwd</code> in the help file for par for more information. This argument is ignored if <code>add.line=FALSE</code> .
<code>line.lty</code>	a numeric scalar determining the line type of the line in the plot. The default value is <code>line.lty=1</code> . See the entry for <code>lty</code> in the help file for par for more information. This argument is ignored if <code>add.line=FALSE</code> .
<code>digits</code>	a scalar indicating how many significant digits to print for the distribution parameters. The default value is <code>digits=.Options\$digits</code> .
<code>same.window</code>	logical scalar indicating whether to produce all plots in the same graphics window (<code>same.window=TRUE</code> ; the default), or to create a new graphics window for each separate plot (<code>same.window=FALSE</code>).
<code>ask</code>	logical scalar supplied to the function devAskNewPage , indicating whether to prompt the user before creating a new plot within a single graphics window. The default value is <code>TRUE</code> if <code>same.window=TRUE</code> and <code>num.pages > 1</code> , otherwise it is <code>FALSE</code> .
<code>mfrow, mar, oma, mgp, main, xlab, ylab, xlim, ylim, ...</code>	additional graphical parameters (see par).

Details

The function `qqPlotGestalt` allows the user to display several Q-Q plots or Tukey mean-difference Q-Q plots for a specified probability distribution. The distribution is specified with the arguments `distribution` and `param.list`. By default, [normal \(Gaussian\)](#) Q-Q plots are produced.

If `estimate.params=FALSE` (the default), the theoretical quantiles on the x -axis are computed using the known distribution parameters specified in `param.list`. If `estimate.params=TRUE`, the distribution parameters are estimated based on the sample, and these estimated parameters are then used to compute the theoretical quantiles. For distributions that can be specified by a location and scale parameter (e.g., Normal, Logistic, extreme value, etc.), the value of `estimate.params` will not affect the general shape of the plot, only the values recorded on the x -axis. For distributions that cannot be specified by a location and scale parameter (e.g., exponential, gamma, etc.), it is recommended that `estimate.params` be set to `TRUE` since in practice the values of the distribution parameters are not known but must be estimated from the sample.

The purpose of `qqPlotGestalt` is to allow the user to build-up a visual memory of “typical” Q-Q plots. A Q-Q plot is a graphical tool that allows you to assess how well a particular set of observations fit a particular probability distribution. The value of this tool depends on the user having an internal reference set of Q-Q plots with which to compare the current Q-Q plot.

See the help file for [qqPlot](#) for more information.

Value

The `NULL` value is returned.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

See the REFERENCES section for [qqPlot](#).

See Also

[qqPlot](#).

Examples

```
# Look at eight typical normal (Gaussian) Q-Q plots for random samples
# of size 10 from a N(0,1) distribution
# Are you surprised by the variability in the plots?
#
# (Note: you must use set.seed if you want to reproduce the exact
#       same plots more than once.)

set.seed(298)
qqPlotGestalt(same.window = FALSE)

# Add lines to these same Q-Q plots
#-----
set.seed(298)
qqPlotGestalt(same.window = FALSE, add.line = TRUE)

# Add lines to different Q-Q plots
#-----
qqPlotGestalt(same.window = FALSE, add.line = TRUE)

# Look at 4 sets of plots all in the same graphics window
#-----
# NOT DONE. Remove the hash sign (#) from the beginning
```

```

# of the next line to run this example.
# qqPlotGestalt(add.line = TRUE, num.pages = 4)

#=====

# Look at Q-Q plots for a gamma distribution
#-----

qqPlotGestalt(dist = "gammaAlt",
  param.list = list(mean = 10, cv = 1),
  estimate.params = TRUE, num.pages = 3,
  same.window = FALSE, add.line = TRUE)

# Look at Tukey Mean Difference Q-Q plots
# for a gamma distribution
#-----

qqPlotGestalt(dist = "gammaAlt",
  param.list = list(mean = 10, cv = 1),
  estimate.params = TRUE, num.pages = 3,
  plot.type = "Tukey", same.window = FALSE, add.line = TRUE)

#=====

# Clean up
#-----
graphics.off()

```

Refinery.CO.df

Carbon Monoxide Emissions from Oil Refinery.

Description

Carbon monoxide (CO) emissions (ppm) from an oil refinery near San Francisco. The refinery submitted 31 daily measurements from its stack for the period April 16, 1993 through May 16, 1993 to the Bay Area Air Quality Management District (BAAQMD). The BAAQMD made nine of its own indepent measurements for the period September 11, 1990 through March 30, 1993.

Usage

```
data(Refinery.CO.df)
```

Format

A data frame with 40 observations on the following 3 variables.

CO.ppm a numeric vector of CO emissions (ppm)

Source a factor indicating the source of the measurment (BAAQMD or refinery)

Date a Date object indicating the date the measurement was taken

Source

Data and Story Library, <http://lib.stat.cmu.edu/DASL/Datafiles/Refinery.html>.

References

Zou, G.Y., C.Y. Huo, and J. Taleban. (2009). Simple Confidence Intervals for Lognormal Means and their Differences with Environmental Applications. *Environmetrics*, **20**, 172–180.

simulateMvMatrix	<i>Simulate a Multivariate Matrix Based on a Specified Rank Correlation Mat</i>
------------------	---

Description

Simulate a multivariate matrix of random numbers from specified theoretical probability distributions and/or empirical probability distributions based on a specified rank correlation matrix, using either Latin Hypercube sampling or simple random sampling.

Usage

```
simulateMvMatrix(n, distributions = c(Var.1 = "norm", Var.2 = "norm"),
  param.list = list(Var.1 = list(mean = 0, sd = 1), Var.2 = list(mean = 0, sd = 1)),
  cor.mat = diag(length(distributions)), sample.method = "SRS", seed = NULL,
  left.tail.cutoff = ifelse(is.finite(supp.min), 0, .Machine$double.eps),
  right.tail.cutoff = ifelse(is.finite(supp.max), 0, .Machine$double.eps),
  tol.1 = .Machine$double.eps, tol.symmetry = .Machine$double.eps,
  tol.recip.cond.num = .Machine$double.eps, max.iter = 10)
```

Arguments

- | | |
|---------------|--|
| n | a positive integer indicating the number of random vectors (i.e., the number of rows of the matrix) to generate. |
| distributions | <p>a character vector of length k denoting the distribution abbreviations for each of the k distributions. If there is a names attribute associated with this character vector, these names will be the column names of the resulting matrix. The default value of distributions is <code>c(Var.1="norm", Var.2="norm")</code>, indicating that $k = 2$, both distributions are the normal distribution, and the column names of the resulting $n \times k$ matrix will be "Var.1" and "Var.2". See the help file for Distribution.df for a list of possible distribution abbreviations.</p> <p>Alternatively, the character string "emp" may be used to denote sampling from an empirical distribution based on a set of observations. The vector containing the observations is specified in the argument param.list.</p> |
| param.list | <p>a list containing k lists that specify the values for the parameters of the k distributions. If param.list has a names attribute (not necessary), the names attribute should be exactly the same as the names attribute of the argument distributions. The default value of param.list is <code>list(Var.1=list(mean=0, sd=1), Var.2=list(mean=0, sd=1))</code>. See the help file for Distribution.df for the names and possible values of the parameters associated with each distribution.</p> <p>Alternatively, if you specify an empirical distribution for the j'th distribution by setting the j'th element of distribution to "emp", then the j'th component of param.list must be a list of the form <code>list(obs=name)</code>, where name denotes the name of the vector containing the observations to use for the empirical distribution. In this case, you may also supply arguments to the qemp function</p> |

	through the j 'th component of <code>param.list</code> . For example, you may set this component to <code>list(obs=name, discrete=T)</code> to specify an empirical distribution based on a discrete random variable.
<code>cor.mat</code>	a $k \times k$ matrix specifying the rank correlations between the k distributions. This argument must be a positive definite symmetric matrix, with all 1's on the diagonal. All elements on the off-diagonal must be between -1 and 1. The default value is the $k \times k$ identity matrix, specifying no rank correlation between any of the variables.
<code>sample.method</code>	a character vector of length 1 or k indicating, for each distribution, whether to use Latin Hypercube sampling or simple random sampling. If <code>sample.method</code> is of length 1, it is replicated to length k . Each element of <code>sample.method</code> must be the character string "LHS" (Latin Hypercube sampling) or "SRS" (simple random sampling), or an abbreviation of one of these strings. The default value is "SRS", indicating simple random sampling for each distribution. Note that by specifying <code>sample.method</code> as a vector of length k , you may use different sampling methods for different distributions.
<code>seed</code>	integer to supply to the R function set.seed . The default value is <code>seed=NULL</code> , in which case the random seed is not set but instead based on the current value of <code>.Random.seed</code> .
<code>left.tail.cutoff</code>	a numeric vector of length k indicating, for each distribution, what proportion of the left-tail of the probability distribution to omit for Latin Hypercube sampling. All elements of <code>left.tail.cutoff</code> must be between 0 and 1. For densities with a finite support minimum (e.g., Lognormal or Empirical) the default value is <code>left.tail.cutoff=0</code> ; for densities with a support minimum of $-\infty$, the default value is <code>left.tail.cutoff=.Machine\$double.eps</code> . The j 'th element of this argument is ignored if the j 'th element of <code>sample.method</code> is equal to "SRS".
<code>right.tail.cutoff</code>	a numeric vector of length k indicating, for each distribution, what proportion of the right-tail of the probability distribution to omit for Latin Hypercube sampling. All elements of <code>right.tail.cutoff</code> must be between 0 and 1. For densities with a finite support maximum (e.g., Beta or Empirical) the default value is <code>right.tail.cutoff=0</code> ; for densities with a support maximum of ∞ , the default value is <code>right.tail.cutoff=.Machine\$double.eps</code> . The j 'th element of this argument is ignored if the j 'th element of <code>sample.method</code> is equal to "SRS".
<code>tol.1</code>	a positive numeric scalar indicating the allowable absolute deviation from 1 for the diagonal elements of <code>cor.mat</code> . The default value is <code>.Machine\$double.eps</code> .
<code>tol.symmetry</code>	a positive numeric scalar indicating the allowable absolute deviation from 0 for the difference between symmetric elements of <code>cor.mat</code> (e.g., <code>abs(cor.mat[3,2]-cor.mat[2,3])</code>). The default value is <code>.Machine\$double.eps</code> .
<code>tol.recip.cond.num</code>	a positive numeric scalar indicating the allowable minimum value of the reciprocal of the condition number for <code>cor.mat</code> . The condition number is defined to be the largest eigen value divided by the smallest eigen value. The reciprocal of the condition number is some number between 0 and 1. This value must be sufficiently large for <code>cor.mat</code> to be of full rank (i.e., to not be singular). The default value of <code>tol.recip.cond.num</code> is <code>.Machine\$double.eps</code> .
<code>max.iter</code>	a positive integer indicating the maximum number of iterations to use to produce the R matrix in the algorithm to create the output matrix. The sample correlation

matrix of R must be positive definite. The number of iterations will rarely be more than 2 for moderate to large sample sizes (e.g., $n > 2k$). The default value is `max.iter=10`. See the DETAILS section below for more information on the R matrix.

Details

Motivation

In risk assessment and Monte Carlo simulation, the outcome variable of interest, say Y , is usually some function of one or more other random variables:

$$Y = h(\underline{X}) = h(X_1, X_2, \dots, X_k) \quad (1)$$

For example, Y may be the incremental lifetime cancer risk due to ingestion of soil contaminated with benzene (Thompson et al., 1992; Hamed and Bedient, 1997). In this case the random vector \underline{X} may represent observations from several kinds of distributions that characterize exposure and dose-response, such as benzene concentration in the soil, soil ingestion rate, average body weight, the cancer potency factor for benzene, etc. These distributions may or may not be assumed to be independent of one another (Smith et al., 1992; Bukowski et al., 1995). Often, input variables in a Monte Carlo simulation are in fact known to be correlated, such as body weight and dermal area.

Characterizing the joint distribution of a random vector \underline{X} , where different elements of \underline{X} come from different distributions, is usually mathematically complex or impossible unless the elements (random variables) of \underline{X} are independent. Iman and Conover (1982) present an algorithm for creating a set of n multivariate observations with a rank correlation matrix that is approximately equal to a specified rank correlation matrix. This method allows for different probability distributions for each element of the multivariate vector. The details of this algorithm are as follows.

Algorithm

1. Specify n , the desired number of random vectors (i.e., number of rows of the $n \times k$ output matrix). This is specified by the argument `n` for the function `simulateMvMatrix`.
2. Create C , the desired $k \times k$ correlation matrix. This is specified by the argument `cor.mat`.
3. Compute P , where P is a lower triangular $k \times k$ matrix and

$$PP' = C \quad (2)$$

where P' denotes the transpose of P . The function `simulateMvMatrix` uses the Cholesky decomposition to compute P (see the R help file for [chol](#)).

4. Create R , an $n \times k$ matrix, whose columns represent k independent permutations of van der Waerden scores. That is, each column of R is a random permutation of the scores

$$\Phi^{-1}\left(\frac{i}{n+1}\right), \quad i = 1, 2, \dots, n \quad (3)$$

where Φ denotes the cumulative distribution function of the standard normal distribution.

5. Compute T , the $k \times k$ Pearson sample correlation matrix of R . Make sure T is positive definite; if it is not, then repeat step 4.
6. Compute Q , where Q is a lower triangular $k \times k$ matrix and

$$QQ' = T \quad (4)$$

The function `simulateMvMatrix` uses the Cholesky decomposition to compute Q (see the R help file for [chol](#)).

7. Compute the lower triangular $k \times k$ matrix S , where

$$S = PQ^{-1} \quad (5)$$

8. Compute the matrix R^* , where

$$R^* = RS' \quad (6)$$

9. Generate an $n \times k$ matrix of random numbers \underline{X} , where each column of \underline{X} comes from the distribution specified by the arguments `distributions` and `param.list`. Generate each column of random numbers independently of the other columns. If the j 'th element of `sample.method` equals "SRS", use simple random sampling to generate the random numbers for the j 'th column of \underline{X} . If the j 'th element of `sample.method` equals "LHS", use Latin Hypercube sampling to generate the random numbers for the j 'th column of \underline{X} . At this stage in the algorithm, the function `simulateMvMatrix` calls the function `simulateVector` to create each column of \underline{X} .
10. Order the observations within each column of \underline{X} so that the order of the ranks within each column of \underline{X} matches the order of the ranks within each column of R^* . This way, \underline{X} and R^* have exactly the same sample rank correlation matrix.

Explanation

Iman and Conover (1982) present two algorithms for computing an $n \times k$ output matrix with a specified rank correlation. The algorithm presented above is the second, more complicated one. In order to explain the reasoning behind this algorithm, we need to explain the simple algorithm first.

Simple Algorithm

Let R_i denote the i 'th row vector of the matrix R , the matrix of scores. This row vector has a population correlation matrix of I , where I denotes the $k \times k$ identity matrix. Thus, the $1 \times k$ vector $R_i P'$ has a population correlation matrix equal to C . Therefore, if we define R^* by

$$R^* = RP' \quad (7)$$

each row of R^* has the same multivariate distribution with population correlation matrix C . The rank correlation matrix of R^* should therefore be close to C . Ordering the columns of \underline{X} as described in Step 10 above will yield a matrix of observations with the specified distributions and the exact same rank correlation matrix as the rank correlation matrix of R^* .

Iman and Conover (1982) use van der Waerden scores instead of raw ranks to create R because van der Waerden scores yield more "natural-looking" pairwise scatterplots.

If the Pearson sample correlation matrix of R , denoted T in Step 5 above, is exactly equal to the true population correlation matrix I , then the sample correlation matrix of R^* is exactly equal to C , and the rank correlation matrix of R^* is approximately equal to C . The Pearson sample correlation matrix of R , however, is an estimate of the true population correlation matrix I , and is therefore "bouncing around" I . Likewise, the Pearson sample correlation matrix of R^* is an estimate of the true population correlation matrix C , and is therefore bouncing around C . Using this simple algorithm, the Pearson sample correlation matrix of R^* , as R^* is defined in Equation (7) above, may not be "close" enough to the desired rank correlation matrix C , and thus the rank correlation of R^* will not be close enough to C . Iman and Conover (1982), therefore present a more complicated algorithm.

More Complicated Algorithm

To get around the problem mentioned above, Iman and Conover (1982) find a $k \times k$ lower triangular matrix S such that the matrix R^* as defined in Equation (6) above has a correlation matrix exactly equal to C . The formula for S is given in Steps 6 and 7 of the algorithm above.

Iman and Conover (1982, p.330) note that even if the desired rank correlation matrix C is in fact the identity matrix I , this method of generating the matrix will produce a matrix with an associated

rank correlation that more closely resembles I than you would get by simply generating random numbers within each column of \underline{X} .

Value

A numeric matrix of dimension $n \times k$ of random numbers, where the j 'th column of numbers comes from the distribution specified by the j 'th elements of the arguments `distributions` and `param.list`, and the rank correlation of this matrix is approximately equal to the argument `cor.mat`. The value of n is determined by the argument `n`, and the value of k is determined by the length of the argument `distributions`.

Note

Monte Carlo simulation and risk assessment often involve looking at the distribution or characteristics of the distribution of some outcome variable that depends upon several input variables (see Equation (1) above). Usually these input variables can be considered random variables. An important part of both sensitivity analysis and uncertainty analysis involves looking at how the distribution of the outcome variable changes with changing assumptions on the input variables. One important assumption is the correlation between the input random variables.

Often, the input random variables are assumed to be independent when in fact they are known to be correlated (Smith et al., 1992; Bukowski et al., 1995). It is therefore important to assess the effect of the assumption of independence on the distribution of the outcome variable. One way to assess the effect of this assumption is to run the Monte Carlo simulation assuming independence and then also run it assuming certain forms of correlations among the input variables.

Iman and Davenport (1982) present a series of scatterplots showing “typical” scatterplots with various distributions on the x - and y -axes and various assumed rank correlations. These plots are meant to aid in developing reasonable estimates of rank correlation between input variables. These plots can easily be produced using the `simulateMvMatrix` and `plot` functions.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Bukowski, J., L. Korn, and D. Wartenberg. (1995). Correlated Inputs in Quantitative Risk Assessment: The Effects of Distributional Shape. *Risk Analysis* **15**(2), 215–219.
- Hamed, M., and P.B. Bedient. (1997). On the Effect of Probability Distributions of Input Variables in Public Health Risk Assessment. *Risk Analysis* **17**(1), 97–105.
- Iman, R.L., and W.J. Conover. (1980). Small Sample Sensitivity Analysis Techniques for Computer Models, With an Application to Risk Assessment (with Comments). *Communications in Statistics—Volume A, Theory and Methods*, **9**(17), 1749–1874.
- Iman, R.L., and W.J. Conover. (1982). A Distribution-Free Approach to Inducing Rank Correlation Among Input Variables. *Communications in Statistics—Volume B, Simulation and Computation*, **11**(3), 311–334.
- Iman, R.L., and J.M. Davenport. (1982). Rank Correlation Plots For Use With Correlated Input Variables. *Communications in Statistics—Volume B, Simulation and Computation*, **11**(3), 335–360.
- Iman, R.L., and J.C. Helton. (1988). An Investigation of Uncertainty and Sensitivity Analysis Techniques for Computer Models. *Risk Analysis* **8**(1), 71–90.
- Iman, R.L. and J.C. Helton. (1991). The Repeatability of Uncertainty and Sensitivity Analyses for Complex Probabilistic Risk Assessments. *Risk Analysis* **11**(4), 591–606.

McKay, M.D., R.J. Beckman., and W.J. Conover. (1979). A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output From a Computer Code. *Technometrics* **21**(2), 239–245.

Millard, S.P. (2013). *EnvStats: an R Package for Environmental Statistics*. Second Edition. Springer, New York.

Smith, A.E., P.B. Ryan, and J.S. Evans. (1992). The Effect of Neglecting Correlations When Propagating Uncertainty and Estimating the Population Distribution of Risk. *Risk Analysis* **12**(4), 467–474.

Thompson, K.M., D.E. Burmaster, and E.A.C. Crouch. (1992). Monte Carlo Techniques for Quantitative Uncertainty Analysis in Public Health Risk Assessments. *Risk Analysis* **12**(1), 53–63.

Vose, D. (2008). *Risk Analysis: A Quantitative Guide*. Third Edition. John Wiley & Sons, West Sussex, UK, 752 pp.

See Also

[Probability Distributions and Random Numbers](#), [Empirical](#), [simulateVector](#), [cor](#), [set.seed](#).

Examples

```
# Generate 5 observations from a standard bivariate normal distribution
# with a rank correlation matrix (approximately) equal to the 2 x 2
# identity matrix, using simple random sampling for each
# marginal distribution.

simulateMvMatrix(5, seed = 47)
#           Var.1      Var.2
#[1,]  0.01513086  0.03960243
#[2,] -1.08573747  0.09147291
#[3,] -0.98548216  0.49382018
#[4,] -0.25204590 -0.92245624
#[5,] -1.46575030 -1.82822917

#=====

# Look at the observed rank correlation matrix for 100 observations
# from a standard bivariate normal distribution with a rank correlation matrix
# (approximately) equal to the 2 x 2 identity matrix. Compare this observed
# rank correlation matrix with the observed rank correlation matrix based on
# generating two independent sets of standard normal random numbers.
# Note that the cross-correlation is closer to 0 for the matrix created with
# simulateMvMatrix.

cor(simulateMvMatrix(100, seed = 47), method = "spearman")
#           Var.1      Var.2
#Var.1  1.000000000 -0.005976598
#Var.2 -0.005976598  1.000000000

cor(matrix(simulateVector(200, seed = 47), 100 , 2), method = "spearman")
#           [,1]      [,2]
#[1,]  1.00000000 -0.05374137
#[2,] -0.05374137  1.00000000

#=====

# Generate 1000 observations from a bivariate distribution, where the first
```

```
# distribution is a normal distribution with parameters mean=10 and sd=2,
# the second distribution is a lognormal distribution with parameters
# mean=10 and cv=1, and the desired rank correlation between the two
# distributions is 0.8. Look at the observed rank correlation matrix, and
# plot the results.
```

```
mat <- simulateMvMatrix(1000,
  distributions = c(N.10.2 = "norm", LN.10.1 = "lnormAlt"),
  param.list = list(N.10.2 = list(mean=10, sd=2),
    LN.10.1 = list(mean=10, cv=1)),
  cor.mat = matrix(c(1, .8, .8, 1), 2, 2), seed = 47)

round(cor(mat, method = "spearman"), 2)
#      N.10.2 LN.10.1
#N.10.2    1.00    0.78
#LN.10.1    0.78    1.00

windows()
plot(mat, xlab = "Observations from N(10, 2)",
  ylab = "Observations from LN(mean=10, cv=1)",
  main = "Lognormal vs. Normal Deviates with Rank Correlation 0.8")
```

```
#-----
```

```
# Repeat the last example, but use Latin Hypercube sampling for both
# distributions. Note the wider range on the y-axis.
```

```
mat.LHS <- simulateMvMatrix(1000,
  distributions = c(N.10.2 = "norm", LN.10.1 = "lnormAlt"),
  param.list = list(N.10.2 = list(mean=10, sd=2),
    LN.10.1 = list(mean=10, cv=1)),
  cor.mat = matrix(c(1, .8, .8, 1), 2, 2),
  sample.method = "LHS", seed = 298)

round(cor(mat.LHS, method = "spearman"), 2)
#      N.10.2 LN.10.1
#N.10.2    1.00    0.79
#LN.10.1    0.79    1.00

windows()
plot(mat.LHS, xlab = "Observations from N(10, 2)",
  ylab = "Observations from LN(mean=10, cv=1)",
  main = paste("Lognormal vs. Normal Deviates with Rank Correlation 0.8",
    "(Latin Hypercube Sampling)", sep = "\n"))
```

```
#=====
```

```
# Generate 1000 observations from a multivariate distribution, where the
# first distribution is a normal distribution with parameters
# mean=10 and sd=2, the second distribution is a lognormal distribution
# with parameters mean=10 and cv=1, the third distribution is a beta
# distribution with parameters shape1=2 and shape2=3, and the fourth
# distribution is an empirical distribution of 100 observations that
# we'll generate from a Pareto distribution with parameters
# location=10 and shape=2. Set the desired rank correlation matrix to:
```

```
cor.mat <- matrix(c(1, .8, 0, .5, .8, 1, 0, .7,
```

```

0, 0, 1, .2, .5, .7, .2, 1), 4, 4)

cor.mat
#      [,1] [,2] [,3] [,4]
# [1,] 1.0  0.8  0.0  0.5
# [2,] 0.8  1.0  0.0  0.7
# [3,] 0.0  0.0  1.0  0.2
# [4,] 0.5  0.7  0.2  1.0

# Use Latin Hypercube sampling for each variable, look at the observed
# rank correlation matrix, and plot the results.

pareto.rns <- simulateVector(100, "pareto",
  list(location = 10, shape = 2), sample.method = "LHS",
  seed = 56)

mat <- simulateMvMatrix(1000,
  distributions = c(Normal = "norm", Lognormal = "lnormAlt",
    Beta = "beta", Empirical = "emp"),
  param.list = list(Normal = list(mean=10, sd=2),
    Lognormal = list(mean=10, cv=1),
    Beta = list(shape1 = 2, shape2 = 3),
    Empirical = list(obs = pareto.rns)),
  cor.mat = cor.mat, seed = 47, sample.method = "LHS")

round(cor(mat, method = "spearman"), 2)
#      Normal Lognormal  Beta Empirical
# Normal      1.00      0.78 -0.01      0.47
# Lognormal    0.78      1.00 -0.01      0.67
# Beta        -0.01     -0.01  1.00      0.19
# Empirical    0.47      0.67  0.19      1.00

windows()
pairs(mat)

#=====

# Clean up
#-----
rm(mat, mat.LHS, pareto.rns)
graphics.off()

```

simulateVector

Simulate a Vector of Random Numbers From a Specified Theoretical or Empirical Probability Distribution

Description

Simulate a vector of random numbers from a specified theoretical probability distribution or empirical probability distribution, using either Latin Hypercube sampling or simple random sampling.

Usage

```
simulateVector(n, distribution = "norm", param.list = list(mean = 0, sd = 1),
```

```
sample.method = "SRS", seed = NULL, sorted = FALSE,
left.tail.cutoff = ifelse(is.finite(supp.min), 0, .Machine$double.eps),
right.tail.cutoff = ifelse(is.finite(supp.max), 0, .Machine$double.eps))
```

Arguments

<code>n</code>	a positive integer indicating the number of random numbers to generate.
<code>distribution</code>	<p>a character string denoting the distribution abbreviation. The default value is <code>distribution="norm"</code>. See the help file for Distribution.df for a list of possible distribution abbreviations.</p> <p>Alternatively, the character string "emp" may be used to denote sampling from an empirical distribution based on a set of observations. The vector containing the observations is specified in the argument <code>param.list</code>.</p>
<code>param.list</code>	<p>a list with values for the parameters of the distribution. The default value is <code>param.list=list(mean=0, sd=1)</code>. See the help file for Distribution.df for the names and possible values of the parameters associated with each distribution.</p> <p>Alternatively, if you specify an empirical distribution by setting <code>distribution="emp"</code>, then <code>param.list</code> must be a list of the form <code>list(obs=name)</code>, where <i>name</i> denotes the name of the vector containing the observations to use for the empirical distribution. In this case, you may also supply arguments to the qemp function through <code>param.list</code>. For example, you may set <code>param.list=list(obs=name, discrete=T)</code> to specify an empirical distribution based on a discrete random variable.</p>
<code>sample.method</code>	a character string indicating whether to use simple random sampling (<code>sample.method="SRS"</code> , the default) or Latin Hypercube sampling (<code>sample.method="LHS"</code>).
<code>seed</code>	integer to supply to the R function set.seed . The default value is <code>seed=NULL</code> , in which case the random seed is not set but instead based on the current value of <code>.Random.seed</code> .
<code>sorted</code>	logical scalar indicating whether to return the random numbers in sorted (ascending) order. The default value is <code>sorted=FALSE</code> .
<code>left.tail.cutoff</code>	a scalar between 0 and 1 indicating what proportion of the left-tail of the probability distribution to omit for Latin Hypercube sampling. For densities with a finite support minimum (e.g., Lognormal or Empirical) the default value is <code>left.tail.cutoff=0</code> ; for densities with a support minimum of $-\infty$, the default value is <code>left.tail.cutoff=.Machine\$double.eps</code> . This argument is ignored if <code>sample.method="SRS"</code> .
<code>right.tail.cutoff</code>	a scalar between 0 and 1 indicating what proportion of the right-tail of the probability distribution to omit for Latin Hypercube sampling. For densities with a finite support maximum (e.g., Beta or Empirical) the default value is <code>right.tail.cutoff=0</code> ; for densities with a support maximum of ∞ , the default value is <code>right.tail.cutoff=.Machine\$double.eps</code> . This argument is ignored if <code>sample.method="SRS"</code> .

Details

Simple Random Sampling (`sample.method="SRS"`)

When `sample.method="SRS"`, the function `simulateVector` simply calls the function `rabb`, where *abb* denotes the abbreviation of the specified distribution (e.g., [rlnorm](#), [rempr](#), etc.).

Latin Hypercube Sampling (`sample.method="LHS"`)

When `sample.method="LHS"`, the function `simulateVector` generates n random numbers using Latin Hypercube sampling. The distribution is divided into n intervals of equal probability $1/n$ and simple random sampling is performed once within each interval; i.e., Latin Hypercube sampling is simply stratified sampling without replacement, where the strata are defined by the 0'th, $100(1/n)$ 'th, $100(2/n)$ 'th, ..., and 100'th percentiles of the distribution.

Latin Hypercube sampling, sometimes abbreviated **LHS**, is a method of sampling from a probability distribution that ensures all portions of the probability distribution are represented in the sample. It was introduced in the published literature by McKay et al. (1979) to overcome the following problem in Monte Carlo simulation based on simple random sampling (SRS). Suppose we want to generate random numbers from a specified distribution. If we use simple random sampling, there is a low probability of getting very many observations in an area of low probability of the distribution. For example, if we generate n observations from the distribution, the probability that none of these observations falls into the upper 98'th percentile of the distribution is 0.98^n . So, for example, there is a 13% chance that out of 100 random numbers, none will fall at or above the 98'th percentile. If we are interested in reproducing the shape of the distribution, we will need a very large number of observations to ensure that we can adequately characterize the tails of the distribution (Vose, 2008, pp. 59–62).

See Millard (2013) for a visual explanation of Latin Hypercube sampling.

Value

a numeric vector of random numbers from the specified distribution.

Note

Latin Hypercube sampling, sometimes abbreviated **LHS**, is a method of sampling from a probability distribution that ensures all portions of the probability distribution are represented in the sample. It was introduced in the published literature by McKay et al. (1979). Latin Hypercube sampling is often used in probabilistic risk assessment, specifically for sensitivity and uncertainty analysis (e.g., Iman and Conover, 1980; Iman and Helton, 1988; Iman and Helton, 1991; Vose, 1996).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Iman, R.L., and W.J. Conover. (1980). Small Sample Sensitivity Analysis Techniques for Computer Models, With an Application to Risk Assessment (with Comments). *Communications in Statistics—Volume A, Theory and Methods*, **9**(17), 1749–1874.
- Iman, R.L., and J.C. Helton. (1988). An Investigation of Uncertainty and Sensitivity Analysis Techniques for Computer Models. *Risk Analysis* **8**(1), 71–90.
- Iman, R.L. and J.C. Helton. (1991). The Repeatability of Uncertainty and Sensitivity Analyses for Complex Probabilistic Risk Assessments. *Risk Analysis* **11**(4), 591–606.
- McKay, M.D., R.J. Beckman., and W.J. Conover. (1979). A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output From a Computer Code. *Technometrics* **21**(2), 239–245.
- Millard, S.P. (2013). *EnvStats: an R Package for Environmental Statistics*. Second Edition. Springer, New York.
- Vose, D. (2008). *Risk Analysis: A Quantitative Guide*. Third Edition. John Wiley & Sons, West Sussex, UK, 752 pp.

See Also

[Probability Distributions and Random Numbers](#), [Empirical](#), [simulateMvMatrix](#), [set.seed](#).

Examples

```
# Generate 10 observations from a lognormal distribution with
# parameters mean=10 and cv=1 using simple random sampling:

simulateVector(10, distribution = "lnormAlt",
  param.list = list(mean = 10, cv = 1), seed = 47,
  sort = TRUE)
# [1] 2.086931 2.863589 3.112866 5.592502 5.732602 7.160707
# [7] 7.741327 8.251306 12.782493 37.214748

#-----

# Repeat the above example by calling rlnormAlt directly:

set.seed(47)
sort(rlnormAlt(10, mean = 10, cv = 1))
# [1] 2.086931 2.863589 3.112866 5.592502 5.732602 7.160707
# [7] 7.741327 8.251306 12.782493 37.214748

#-----

# Now generate 10 observations from the same lognormal distribution
# but use Latin Hypercube sampling. Note that the largest value
# is larger than for simple random sampling:

simulateVector(10, distribution = "lnormAlt",
  param.list = list(mean = 10, cv = 1), seed = 47,
  sample.method = "LHS", sort = TRUE)
# [1] 2.406149 2.848428 4.311175 5.510171 6.467852 8.174608
# [7] 9.506874 12.298185 17.022151 53.552699

#=====

# Generate 50 observations from a Pareto distribution with parameters
# location=10 and shape=2, then use this resulting vector of
# observations as the basis for generating 3 observations from an
# empirical distribution using Latin Hypercube sampling:

set.seed(321)
pareto.rns <- rpareto(50, location = 10, shape = 2)

simulateVector(3, distribution = "emp",
  param.list = list(obs = pareto.rns), sample.method = "LHS")
#[1] 11.50685 13.50962 17.47335

#=====

# Clean up
#-----
rm(pareto.rns)
```

Skagit.NH3_N.df	<i>Ammonia Nitrogen Concentrations in the Skagit River, Marblemount, Washington</i>
-----------------	---

Description

Ammonia nitrogen (NH₃—N) concentration (mg/L) in the Skagit River measured monthly from January 1978 through December 2010 at the Marblemount, Washington monitoring station.

Usage

Skagit.NH3_N.df

Format

A data frame with 396 observations on the following 6 variables.

Date Date of collection.

NH3_N.Orig.mg.per.L a character vector of the ammonia nitrogen concentrations where values for non-detects are preceded with the less-than sign (<).

NH3_N.mg.per.L a numeric vector of ammonia nitrogen concentrations; non-detects have been coded to their detection limit.

DQ1 factor of data qualifier values.

- U = The analyte was not detected at or above the reported result.
- J = The analyte was positively identified. The associated numerical result is an estimate.
- UJ = The analyte was not detected at or above the reported estimated result.

DQ2 factor of data qualifier values. An asterisk (*) indicates a possible quality problem for the result.

Censored a logical vector indicating which observations are censored.

Details

Station 04A100 - Skagit R \@ Marblemount. Located at the bridge on the Cascade River Road where Highway 20 (North Cascades Highway) turns 90 degrees in Marblemount.

Source

Washington State Department of Ecology.

<http://www.ecy.wa.gov/apps/watersheds/riv/station.asp?sta=04A100>

http://www.ecy.wa.gov/apps/watersheds/riv/parameters_ref.html

skewness	<i>Coefficient of Skewness</i>
----------	--------------------------------

Description

Compute the sample coefficient of skewness.

Usage

```
skewness(x, na.rm = FALSE, method = "fisher", l.moment.method = "unbiased",
  plot.pos.cons = c(a = 0.35, b = 0))
```

Arguments

<code>x</code>	numeric vector of observations.
<code>na.rm</code>	logical scalar indicating whether to remove missing values from <code>x</code> . If <code>na.rm=FALSE</code> (the default) and <code>x</code> contains missing values, then a missing value (NA) is returned. If <code>na.rm=TRUE</code> , missing values are removed from <code>x</code> prior to computing the coefficient of variation.
<code>method</code>	character string specifying what method to use to compute the sample coefficient of skewness. The possible values are "fisher" (ratio of unbiased moment estimators; the default), "moments" (ratio of product moment estimators), or "l.moments" (ratio of <i>L</i> -moment estimators).
<code>l.moment.method</code>	character string specifying what method to use to compute the <i>L</i> -moments when <code>method="l.moments"</code> . The possible values are "unbiased" (method based on the <i>U</i> -statistic; the default), or "plotting.position" (method based on the plotting position formula).
<code>plot.pos.cons</code>	numeric vector of length 2 specifying the constants used in the formula for the plotting positions when <code>method="l.moments"</code> and <code>l.moment.method="plotting.position"</code> . The default value is <code>plot.pos.cons=c(a=0.35, b=0)</code> . If this vector has a names attribute with the value <code>c("a", "b")</code> or <code>c("b", "a")</code> , then the elements will be matched by name in the formula for computing the plotting positions. Otherwise, the first element is mapped to the name "a" and the second element to the name "b".

Details

Let \underline{x} denote a random sample of n observations from some distribution with mean μ and standard deviation σ .

Product Moment Coefficient of Skewness (`method="moment"` or `method="fisher"`)

The **coefficient of skewness** of a distribution is the third standardized moment about the mean:

$$\eta_3 = \sqrt{\beta_1} = \frac{\mu_3}{\sigma^3} \quad (1)$$

where

$$\eta_r = E\left[\left(\frac{X - \mu}{\sigma}\right)^r\right] = \frac{1}{\sigma^r} E[(X - \mu)^r] = \frac{\mu_r}{\sigma^r} \quad (2)$$

and

$$\mu_r = E[(X - \mu)^r] \quad (3)$$

denotes the r 'th moment about the mean (central moment). That is, the coefficient of skewness is the third central moment divided by the cube of the standard deviation. The coefficient of skewness is 0 for a symmetric distribution. Distributions with positive skew have heavy right-hand tails, and distributions with negative skew have heavy left-hand tails.

When method="moment", the coefficient of skewness is estimated using the method of moments estimator for the third central moment and the method of moments estimator for the variance:

$$\hat{\eta}_3 = \frac{\hat{\mu}_3}{\sigma^3} = \frac{\frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^3}{[\frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2]^{3/2}} \quad (5)$$

where

$$\hat{\sigma}_m^2 = s_m^2 = \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2 \quad (6)$$

This form of estimation should be used when resampling (bootstrap or jackknife).

When method="fisher", the coefficient of kurtosis is estimated using the unbiased estimator for the fourth central moment (Serfling, 1980, p.73; Chen, 1995, p.769) and the unbiased estimator for the variance.

$$\hat{\eta}_3 = \frac{\frac{n}{(n-1)(n-2)} \sum_{i=1}^n (x_i - \bar{x})^3}{s^3} \quad (7)$$

where

$$\hat{\sigma}^2 = s^2 = \frac{1}{n-1} \sum_{i=1}^n (x_i - \bar{x})^2 \quad (8)$$

(Note that Serfling, 1980, p.73 contains a typographical error in the numerator for the unbiased estimator of the third central moment.)

L-Moment Coefficient of Kurtosis (method="l.moments")

Hosking (1990) defines the L -moment analog of the coefficient of kurtosis as:

$$\tau_3 = \frac{\lambda_3}{\lambda_2} \quad (9)$$

that is, the third L -moment divided by the second L -moment. He shows that this quantity lies in the interval $(-1, 1)$.

When l.moment.method="unbiased", the L -skewness is estimated by:

$$t_3 = \frac{l_3}{l_2} \quad (10)$$

that is, the unbiased estimator of the third L -moment divided by the unbiased estimator of the second L -moment.

When l.moment.method="plotting.position", the L -skewness is estimated by:

$$\tilde{\tau}_3 = \frac{\tilde{\lambda}_3}{\tilde{\lambda}_2} \quad (11)$$

that is, the plotting-position estimator of the third L -moment divided by the plotting-position estimator of the second L -moment.

See the help file for [lMoment](#) for more information on estimating L -moments.

Value

A numeric scalar – the sample coefficient of skewness.

Note

Traditionally, the coefficient of skewness has been estimated using product moment estimators. Sometimes an estimate of skewness is used in a goodness-of-fit test for normality (e.g., set `test="skew"` in the call to [gofTest](#)).

Hosking (1990) introduced the idea of *L*-moments and *L*-kurtosis.

Vogel and Fennessey (1993) argue that *L*-moment ratios should replace product moment ratios because of their superior performance (they are nearly unbiased and better for discriminating between distributions). They compare product moment diagrams with *L*-moment diagrams.

Hosking and Wallis (1995) recommend using unbiased estimators of *L*-moments (vs. plotting-position estimators) for almost all applications.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Berthouex, P.M., and L.C. Brown. (2002). *Statistics for Environmental Engineers, Second Edition*. Lewis Publishers, Boca Raton, FL.
- Chen, L. (1995). Testing the Mean of Skewed Distributions. *Journal of the American Statistical Association* **90**(430), 767–772.
- Helsel, D.R., and R.M. Hirsch. (1992). *Statistical Methods in Water Resources Research*. Elsevier, New York, NY.
- Ott, W.R. (1995). *Environmental Statistics and Data Analysis*. Lewis Publishers, Boca Raton, FL.
- Serfling, R.J. (1980). *Approximation Theorems of Mathematical Statistics*. John Wiley and Sons, New York, p.73.
- Taylor, J.K. (1990). *Statistical Techniques for Data Analysis*. Lewis Publishers, Boca Raton, FL.
- Vogel, R.M., and N.M. Fennessey. (1993). *L* Moment Diagrams Should Replace Product Moment Diagrams. *Water Resources Research* **29**(6), 1745–1752.
- Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ.

See Also

[var](#), [sd](#), [cv](#), [kurtosis](#), [summaryFull](#), [Summary Statistics](#).

Examples

```
# Generate 20 observations from a lognormal distribution with parameters
# mean=10 and cv=1, and estimate the coefficient of skewness.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)

dat <- rlnormAlt(20, mean = 10, cv = 1)

skewness(dat)
#[1] 0.9876632

skewness(dat, method = "moment")
#[1] 0.9119889
```

```

skewness(dat, meth = "l.moment")
#[1] 0.2656674

#-----
# Clean up
rm(dat)

```

stripChart

1-D Scatter Plots with Confidence Intervals

Description

stripChart is a modification of the R function [stripchart](#). It is a generic function used to produce one dimensional scatter plots (or dot plots) of the given data, along with text indicating sample size and estimates of location (mean or median) and scale (standard deviation or interquartile range), as well as confidence intervals for the population location parameter. One dimensional scatterplots are a good alternative to [boxplots](#) when sample sizes are small or moderate. The function invokes particular [methods](#) which depend on the [class](#) of the first argument.

Usage

```

stripChart(x, ...)

## S3 method for class 'formula'
stripChart(x, data = NULL, dlab = NULL,
  subset, na.action = NULL, ...)

## Default S3 method:
stripChart(x, method = "stack", seed = 47,
  jitter = 0.1 * cex, offset = 1/2, vertical = TRUE, group.names,
  drop.unused.levels = TRUE, add = FALSE, at = NULL,
  xlim = NULL, ylim = NULL, ylab = NULL, xlab = NULL,
  dlab = "", glab = "", log = "", pch = 1, col = par("fg"),
  cex = par("cex"), points.cex = cex, axes = TRUE, frame.plot = axes,
  show.ci = TRUE, location.pch = 16, location.cex = cex,
  conf.level = 0.95, min.n.for.ci = 2, ci.offset = 3,
  ci.bar.ends = TRUE, ci.bar.ends.size = 0.5 * cex,
  ci.bar.gap = FALSE, n.text = "bottom",
  n.text.line = ifelse(n.text == "bottom", 2, 0),
  n.text.cex = cex, location.scale.text = "top",
  location.scale.digits = 1, location.scale.text.line =
    ifelse(location.scale.text == "top", 0, 3.5),
  location.scale.text.cex = 0.8 * cex, p.value = FALSE,
  p.value.digits = 3, p.value.line = 2, p.value.cex = cex,
  group.difference.ci = p.value, group.difference.conf.level = 0.95,
  group.difference.digits = location.scale.digits,
  ci.and.test = "parametric", ci.arg.list = NULL,
  test.arg.list = NULL, alternative = "two.sided", ...)

```

Arguments

<code>x</code>	the data from which the plots are to be produced. In the default method the data can be specified as a single numeric vector, or as a list of numeric vectors, each corresponding to a component plot. In the formula method, a symbolic specification of the form <code>y ~ g</code> can be given, indicating the observations in the vector <code>y</code> are to be grouped according to the levels of the factor <code>g</code> (the form <code>y ~ 1</code> indicates no grouping). NAs are allowed in the data.
<code>data</code>	a data.frame (or list) from which the variables in <code>x</code> should be taken.
<code>subset</code>	an optional vector specifying a subset of observations to be used for plotting.
<code>na.action</code>	a function which indicates what should happen when the data contain NAs. The default is to ignore missing values in either the response or the group.
<code>...</code>	additional parameters passed to the default method, or by it to plot , points , axis , and title to control the appearance of the plot.
<code>method</code>	the method to be used to separate coincident points. The method "overplot" causes such points to be overplotted, but it is also possible to specify "jitter" to jitter the points, or "stack" to have coincident points stacked (the default). Note that the default value of method (method="stack") differs from the default value for the R function stripchart , which uses method="overplot" by default.
<code>seed</code>	when method="jitter" is used, the argument seed is passed to the R function set.seed . Since jittering depends on the R random number generator, using the same value of seed each time the same data are plotted with stripChart ensures that the resulting plot is the same.
<code>jitter</code>	when method="jitter" is used, jitter gives the amount of jittering applied.
<code>offset</code>	when stacking is used, points are stacked this many line-heights (symbol widths) apart.
<code>vertical</code>	when vertical=TRUE (the default), the plots are drawn vertically rather than horizontally.
<code>group.names</code>	group labels which will be printed alongside (or underneath) each plot.
<code>drop.unused.levels</code>	when drop.unused.levels=TRUE, groups with no observations are dropped.
<code>add</code>	logical, if true <i>add</i> the chart to the current plot.
<code>at</code>	numeric vector giving the locations where the charts should be drawn, particularly when add=TRUE; defaults to 1:n where n is the number of groups.
<code>xlim, ylim</code>	plot limits: see plot.window .
<code>ylab, xlab</code>	labels: see title .
<code>dlab, glab</code>	alternate way to specify axis labels. The dlab and glab labels may be used instead of xlab and ylab if those are not specified. dlab applies to the continuous data axis (the <i>y</i> -axis unless vertical=FALSE), and glab to the group axis.
<code>log</code>	on which axes to use a log scale: see plot.default .
<code>pch, col, cex</code>	Graphical parameters: see par .
<code>points.cex</code>	Sets the cex value for the points plotted.
<code>axes, frame.plot</code>	Axis control: see plot.default .
<code>show.ci</code>	logical scalar indicating whether to plot the confidence interval. The default is show.ci=TRUE.

location.pch	integer indicating which plotting character to use to indicate the estimate of location (mean or median) for each group (see the help file for plot.default). The default is location.pch=16, a filled circle.
location.cex	integer giving the amount by which the plotting characters indicating the estimate of location for each group should be scaled relative to the default (see the help file for plot.default). The default is the current value of the graphics parameter cex.
conf.level	numeric scalar between 0 and 1 indicating the confidence level associated with the confidence interval for the group location (population mean or median). The default value is conf.level=0.95.
min.n.for.ci	integer indicating the minimum sample size required in order to plot a confidence interval for the group location. The default value is min.n.for.ci=2.
ci.offset	numeric scalar in units of cex indicating the amount of space between the line showing the confidence interval and the center of the plotted points for a group. The default value is ci.offset=3.
ci.bar.ends	logical scalar indicating whether to add flat ends to the confidence interval bars. The default value is ci.bar.ends=TRUE.
ci.bar.ends.size	numeric scalar in units of cxy indicating the size of confidence interval bar ends. The default value is half of the current value of cex.
ci.bar.gap	logical scalar indicating with to add a gap between the estimate of group location and the confidence interval bar. The default value is ci.bar.gap=FALSE.
n.text	character string indicating whether and where to indicate the sample size for each group. Possible values are "bottom" (the default), "top", and "none".
n.text.line	integer indicating on which plot margin line to show the sample sizes for each group. The default value is n.text.line=2 when n.text="bottom" and 0 otherwise.
n.text.cex	integer giving the amount by which the text indicating the sample size for each group should be scaled relative to the default (see the help file for plot.default). The default is the current value of the graphics parameter cex.
location.scale.text	character string indicating whether and where to indicate the estimates of location (mean or median) and scale (standard deviation or interquartile range) for each group. Possible values are "top" (the default), "bottom", and "none".
location.scale.digits	integer indicating the number of digits to round the estimates of location and scale. The default value is location.scale.digits=1.
location.scale.text.line	integer indicating on which plot margin line to show the estimates of location and scale for each group. The default value is location.scale.text.line=0 when n.text="top" and 3.5 otherwise.
location.scale.text.cex	integer giving the amount by which the text indicating the estimates of location and scale for each group should be scaled relative to the default (see the help file for plot.default). The default is 80% of the current value of the graphics parameter cex.
p.value	logical scalar indicating whether to show the p-value associated with testing whether all groups have the same population location. The default value is p.value=TRUE. The p-value is displayed at the top of the graph.

<code>p.value.digits</code>	integer indicating the number of digits to round to when displaying the p-value associated with the test of equal group locations. The default value is <code>p.value.digits=3</code> .
<code>p.value.line</code>	integer indicating on which plot margin line to show the p-value associated with the test of equal group locations. The default value is <code>p.value.line=2</code> .
<code>p.value.cex</code>	integer giving the amount by which the text indicating the p-value associated with the test of equal group locations should be scaled relative to the default (see the help file for <code>plot.default</code>). The default is the current value of the graphics parameter <code>cex</code> .
<code>group.difference.ci</code>	for the case when there are just 2 groups, a logical scalar indicating whether to show the confidence interval for the difference between group locations. The default is the value of the <code>p.value</code> argument. The confidence interval is displayed at the top of the graph.
<code>group.difference.conf.level</code>	for the case when there are just 2 groups, a numeric scalar between 0 and 1 indicating the confidence level associated with the confidence interval for the difference between group locations. The default is <code>conf.level=0.95</code> .
<code>group.difference.digits</code>	for the case when there are just 2 groups, an integer indicating the number of digits to round to when displaying the confidence interval for the difference between group locations. The default value is <code>group.difference.digits=location.scale.digits</code> .
<code>ci.and.test</code>	character string indicating whether confidence intervals and tests should be based on parametric or nonparametric (<code>ci.and.test="nonparametric"</code>) methods. When <code>ci.and.test="parametric"</code> (the default), confidence intervals for the population mean are based on the one-sample t-test (see <code>t.test</code>), and the test of group differences is based on the two-sample t-test if there are two groups and the F-test (i.e., one-way analysis of variance, see <code>aoV</code>) if there are three or more groups. When <code>ci.and.test="nonparametric"</code> , confidence intervals for the population pseudo-median are based on the Wilcoxon signed rank test (see <code>wilcox.test</code> and page 56 of Hollander and Wolfe, 1999), and the test of group differences is based on the Wilcoxon rank sum test if there are two groups (see <code>kruskal.test</code>) and the Kruskal-Wallis test (see <code>kruskal.test</code>) if there are three or more groups.
<code>ci.arg.list</code>	an optional list of arguments to pass to the function used to compute confidence intervals. The default value is <code>ci.arg.list=NULL</code> .
<code>test.arg.list</code>	an optional list of arguments to pass to the function used to test for group differences in location. The default value is <code>test.arg.list=NULL</code> . In particular, in the case when there are two groups, <code>ci.and.test="parametric"</code> , and <code>ci.arg.list</code> is <code>NULL</code> or does not contain a component specifying the value for <code>var.equal</code> , this argument is updated to include the component <code>var.equal=TRUE</code> , which is not the default behavior of <code>t.test</code> .
<code>alternative</code>	character string describing the alternative hypothesis for the test of group differences in the case when there are two groups. Possible values are <code>"two.sided"</code> (the default), <code>"less"</code> , and <code>"greater"</code> .

Value

`stripChart` invisibly returns a list with the following components:

group.centers	numeric vector of values on the group axis (the <i>x</i> -axis unless vertical=FALSE) indicating the centers of the groups.
group.stats	a matrix with the number of rows equal to the number of groups and six columns indicating the sample size of the group (N), the estimate of the group location parameter (Mean or Median), the estimate of the group scale (SD or IQR), the lower confidence limit for the group location parameter (LCL), the upper confidence limit for the group location parameter (UCL), and the confidence level associated with the confidence interval (Conf.Level)
group.difference.p.value	numeric scalar indicating the p-value associated with the test of equal group locations.
group.difference.conf.int	numeric vector of two elements indicating the confidence interval for the difference between the group locations. Only present when there are two groups.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Hollander, M., and D.A. Wolfe. (1999). *Nonparametric Statistical Methods*. Second Edition. John Wiley and Sons, New York.
- Millard, S.P., and N.K. Neerchal. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, FL.
- Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ.

See Also

[stripchart](#), [t.test](#), [wilcox.test](#), [aov](#), [kruskal.test](#), [t.test](#).

Examples

```
# The guidance document USEPA (1994b, pp. 6.22--6.25)
# contains measures of 1,2,3,4-Tetrachlorobenzene (TcCB)
# concentrations (in parts per billion) from soil samples
# at a Reference area and a Cleanup area. These data are stored
# in the data frame EPA.94b.tccb.df.
#
# First create one-dimensional scatterplots to compare the
# TcCB concentrations between the areas and use a nonparametric
# test to test for a difference between areas.

windows()
stripChart(TcCB ~ Area, data = EPA.94b.tccb.df,
  p.value = TRUE, ci.and.test = "nonparametric",
  ylab = "TcCB (ppb)")

#-----

# Now log-transform the TcCB data and use a parametric test
# to compare the areas.

windows()
```

```

stripChart(log10(TcCB) ~ Area, data = EPA.94b.tccb.df,
  p.value = TRUE, ci.and.test = "parametric",
  ylab = "log10 [ TcCB (ppb) ]")

#-----

# Repeat the above procedure, but allow the variances to differ.

windows()
stripChart(log10(TcCB) ~ Area, data = EPA.94b.tccb.df,
  p.value = TRUE, ci.and.test = "parametric",
  ylab = "log10 [ TcCB (ppb) ]",
  test.arg.list = list(var.equal = FALSE))

#-----

# Repeat the above procedure, but jitter the points instead of
# stacking them.

windows()
stripChart(log10(TcCB) ~ Area, data = EPA.94b.tccb.df,
  p.value = TRUE, ci.and.test = "parametric",
  ylab = "log10 [ TcCB (ppb) ]",
  test.arg.list = list(var.equal = FALSE),
  method = "jitter", ci.offset = 4)

#=====

# Clean up
#-----
#graphics.off()

```

summaryFull

Full Complement of Summary Statistics

Description

summaryFull is a generic function used to produce a full complement of summary statistics. The function invokes particular [methods](#) which depend on the [class](#) of the first argument. The summary statistics include: sample size, number of missing values, mean, median, trimmed mean, geometric mean, skew, kurtosis, min, max, range, 1st quartile, 3rd quartile, standard deviation, geometric standard deviation, interquartile range, median absolute deviation, and coefficient of variation.

Usage

```

summaryFull(object, ...)

## S3 method for class 'formula'
summaryFull(object, data = NULL, subset,
  na.action = na.pass, ...)

## Default S3 method:
summaryFull(object, group = NULL,

```



```

combine.groups = FALSE, drop.unused.levels = TRUE,
rm.group.na = TRUE, stats = NULL, trim = 0.1,
sd.method = "sqrt.unbiased", geo.sd.method = "sqrt.unbiased",
skew.list = list(), kurtosis.list = list(),
cv.list = list(), digits = max(3, getOption("digits") - 3),
digit.type = "signif", stats.in.rows = TRUE,
drop0trailing = TRUE, data.name = deparse(substitute(object)),
...)

## S3 method for class 'data.frame'
summaryFull(object, ...)

## S3 method for class 'matrix'
summaryFull(object, ...)

## S3 method for class 'list'
summaryFull(object, ...)

```

Arguments

object	an object for which summary statistics are desired. In the default method, the argument object must be a numeric vector, a data frame, a matrix, or a list. When object is a data frame, all columns must be numeric. When object is a matrix, it must be a numeric matrix. When object is a list, all components must be numeric vectors. In the formula method, a symbolic specification of the form $y \sim g$ can be given, indicating the observations in the vector y are to be grouped according to the levels of the factor g (the form $y \sim 1$ indicates no grouping). NAs are allowed in the data.
data	when object is a formula, data specifies an optional data frame, list or environment (or object coercible by <code>as.data.frame</code> to a data frame) containing the variables in the model. If not found in data, the variables are taken from <code>environment(formula)</code> , typically the environment from which <code>summaryFull</code> is called.
subset	when object is a formula, subset specifies an optional vector specifying a subset of observations to be used.
na.action	when object is a formula, na.action specifies a function which indicates what should happen when the data contain NAs. The default is <code>na.pass</code> .
group	when object is a numeric vector, group is a factor or character vector indicating which group each observation belongs to. When object is a matrix or data frame this argument is ignored and the columns define the groups. When object is a formula, this argument is ignored and the right-hand side of the formula specifies the grouping variable.
combine.groups	logical scalar indicating whether to show summary statistics for all groups combined. The default value is FALSE.
drop.unused.levels	when <code>drop.unused.levels=TRUE</code> , groups with no observations are dropped.
rm.group.na	logical scalar indicating whether to remove missing values from the group argument. By default <code>rm.group.na=TRUE</code> .
stats	character vector indicating which statistics to compute. Possible elements of the character vector include: "all" (indicating to include all summary statistics), "for.non.pos" (only compute statistics that are meaningful for datasets with

	non-positive values), "n" (number of non-missing values), "n.miss" (number of missing values), "mean", "median", "trimmed.mean", "geo.mean", "skew", "kurtosis", "min", "max", "range", "1st.quart", "3rd.quart", "sd", "geo.sd", "iqr", "mad", "cv". The default value is stats="for.non.pos" when object contains non-positive values (i.e., values ≤ 0), and stats="all" when object contains only positive values.
trim	fraction (between 0 and 0.5 inclusive) of values to be trimmed from each end of the ordered data to compute the trimmed mean. The default value is trim=0.1. If trim=0.5, this yields the median.
sd.method	character string specifying what method to use to compute the sample standard deviation. The possible values are "sqrt.ubiased" (the square root of the unbiased estimate of variance; the default), or "moments" (the method of moments estimator).
geo.sd.method	character string specifying what method to use to compute the sample standard deviation of the log-transformed observations prior to exponentiating this quantity. The possible values are "sqrt.ubiased" (the square root of the unbiased estimate of variance; the default), or "moments" (the method of moments estimator). See the help file for geoSD for more information.
skew.list	list of arguments to supply to the skewness function. See the help file for skewness for more information. The default value is skew.list=list(), which results in using the default arguments to skewness .
kurtosis.list	list of arguments to supply to the kurtosis function. See the help file for kurtosis for more information. The default value is kurtosis.list=list(), which results in using the default arguments to kurtosis .
cv.list	list of arguments to supply to the cv function. See the help file for cv for more information. The default value is cv.list=list(), which results in using the default arguments to cv .
digits	integer indicating the number of digits to use for the summary statistics. When digit.type="signif", digits indicates the number of significant digits. When digit.type="round", digits indicates the number of decimal places to round to. The default value is max(3, getOption("digits") - 3), that is, the maximum of 3 versus the current setting of the "digits" component of .Options minus 3.
digit.type	character string indicating whether the digits argument refers to significant digits (digit.type="signif", the default), or how many decimal places to round to (digit.type="round").
stats.in.rows	logical scalar indicating whether to show the summary statistics in the rows or columns of the output. The default is stats.in.rows=TRUE.
drop0trailing	logical scalar indicating whether to drop trailing 0's when printing the summary statistics. The value of this argument is added as an attribute to the returned list and is used by the print.summaryStats function. The default value is TRUE.
data.name	character string indicating the name of the data used for the summary statistics.
...	additional arguments affecting the summary statistics produced.

Details

The function `summaryFull` returns summary statistics that are useful to describe various characteristics of one or more variables. It is an extended version of the built-in R function [summary](#)

specifically for non-factor numeric data. The table below shows what statistics are computed and what functions are called by `summaryFull` to compute these statistics.

The object returned by `summaryFull` is useful for printing or report purposes. You may also use the functions that `summaryFull` calls (see table below) to compute summary statistics to be used by other functions.

See the help files for the functions listed in the table below for more information on these summary statistics.

Summary Statistic	Function Used
Mean	mean
Median	median
Trimmed Mean	mean with <code>trim</code> argument
Geometric Mean	geoMean
Skew	skewness
Kurtosis	kurtosis
Min	min
Max	max
Range	range and diff
1st Quartile	quantile
3rd Quartile	quantile
Standard Deviation	sd
Geometric Standard Deviation	geoSD
Interquartile Range	iqr
Median Absolute Deviation	mad
Coefficient of Variation	cv

Value

an object of class "summaryStats" (see [summaryStats.object](#)). Objects of class "summaryStats" are numeric matrices that contain the summary statistics produced by a call to [summaryStats](#) or `summaryFull`. These objects have a special printing method that by default removes trailing zeros for sample size entries and prints blanks for statistics that are normally displayed as NA (see [print.summaryStats](#)).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Berthouex, P.M., and L.C. Brown. (2002). *Statistics for Environmental Engineers, Second Edition*. Lewis Publishers, Boca Raton, FL.
- Gilbert, R.O. (1987). *Statistical Methods for Environmental Pollution Monitoring*. Van Nostrand Reinhold, NY.
- Helsel, D.R., and R.M. Hirsch. (1992). *Statistical Methods in Water Resources Research*. Elsevier, New York, NY.
- Leidel, N.A., K.A. Busch, and J.R. Lynch. (1977). *Occupational Exposure Sampling Strategy Manual*. U.S. Department of Health, Education, and Welfare, Public Health Service, Center for Disease Control, National Institute for Occupational Safety and Health, Cincinnati, Ohio 45226, January, 1977, pp.102-103.

Millard, S.P., and N.K. Neerchal. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, FL.

Ott, W.R. (1995). *Environmental Statistics and Data Analysis*. Lewis Publishers, Boca Raton, FL.

Zar, J.H. (2010). *Biostatistical Analysis, Fifth Edition*. Prentice-Hall, Upper Saddle River, NJ.

See Also

[summary](#), [summaryStats](#).

Examples

```
# Generate 20 observations from a lognormal distribution with
# parameters mean=10 and cv=1, and compute the summary statistics.
# (Note: the call to set.seed simply allows you to reproduce this
# example.)
```

```
set.seed(250)
```

```
dat <- rlnormAlt(20, mean=10, cv=1)
```

```
summary(dat)
# Min. 1st Qu. Median Mean 3rd Qu. Max.
#2.608 4.995 6.235 7.490 9.295 15.440
```

```
summaryFull(dat)
#
#N 20
#Mean 7.49
#Median 6.235
#10% Trimmed Mean 7.125
#Geometric Mean 6.674
#Skew 0.9877
#Kurtosis -0.03539
#Min 2.608
#Max 15.44
#Range 12.83
#1st Quartile 4.995
#3rd Quartile 9.295
#Standard Deviation 3.803
#Geometric Standard Deviation 1.634
#Interquartile Range 4.3
#Median Absolute Deviation 2.607
#Coefficient of Variation 0.5078
```

```
#-----
```

```
# Compare summary statistics for normal and lognormal data:
log.dat <- log(dat)
```

```
summaryFull(list(dat = dat, log.dat = log.dat))
#
#N dat log.dat
#N 20 20
#Mean 7.49 1.898
#Median 6.235 1.83
#10% Trimmed Mean 7.125 1.902
#Geometric Mean 6.674 1.835
```

```

#Skew          0.9877  0.1319
#Kurtosis      -0.03539 -0.4288
#Min           2.608   0.9587
#Max           15.44   2.737
#Range         12.83   1.778
#1st Quartile  4.995   1.607
#3rd Quartile  9.295   2.227
#Standard Deviation 3.803 0.4913
#Geometric Standard Deviation 1.634 1.315
#Interquartile Range 4.3 0.62
#Median Absolute Deviation 2.607 0.4915
#Coefficient of Variation 0.5078 0.2588

```

```

# Clean up
rm(dat, log.dat)

```

```

#-----

```

```

# Compute summary statistics for 10 observations from a normal
# distribution with parameters mean=0 and sd=1. Note that the
# geometric mean and geometric standard deviation are not computed
# since some of the observations are non-positive.

```

```

set.seed(287)

```

```

dat <- rnorm(10)

```

```

summaryFull(dat)
#          dat
#N          10
#Mean       0.07406
#Median     0.1095
#10% Trimmed Mean 0.1051
#Skew      -0.1646
#Kurtosis  -0.7135
#Min      -1.549
#Max       1.449
#Range     2.998
#1st Quartile -0.5834
#3rd Quartile 0.6966
#Standard Deviation 0.9412
#Interquartile Range 1.28
#Median Absolute Deviation 1.05

```

```

# Clean up
rm(dat)

```

```

#-----

```

```

# Compute summary statistics for the TcCB data given in USEPA (1994b)
# (the data are stored in EPA.94b.tccb.df). Arbitrarily set the one
# censored observation to the censoring level. Group by the variable
# Area.

```

```

summaryFull(TcCB ~ Area, data = EPA.94b.tccb.df)
#          Cleanup Reference
#N           77         47

```

#Mean	3.915	0.5985
#Median	0.43	0.54
#10% Trimmed Mean	0.6846	0.5728
#Geometric Mean	0.5784	0.5382
#Skew	7.717	0.9019
#Kurtosis	62.67	0.132
#Min	0.09	0.22
#Max	168.6	1.33
#Range	168.5	1.11
#1st Quartile	0.23	0.39
#3rd Quartile	1.1	0.75
#Standard Deviation	20.02	0.2836
#Geometric Standard Deviation	3.898	1.597
#Interquartile Range	0.87	0.36
#Median Absolute Deviation	0.3558	0.2669
#Coefficient of Variation	5.112	0.4739

summaryStats

Summary Statistics

Description

summaryStats is a generic function used to produce summary statistics, confidence intervals, and results of hypothesis tests. The function invokes particular [methods](#) which depend on the [class](#) of the first argument.

The summary statistics include: sample size, number of missing values, mean, standard deviation, median, min, and max. Optional additional summary statistics include 1st quartile, 3rd quartile, and standard error.

Usage

```
summaryStats(object, ...)
```

```
## S3 method for class 'formula'
```

```
summaryStats(object, data = NULL, subset,
  na.action = na.pass, ...)
```

```
## Default S3 method:
```

```
summaryStats(object, group = NULL,
  drop.unused.levels = TRUE, se = FALSE, quartiles = FALSE,
  digits = max(3, getOption("digits") - 3),
  digit.type = "round", drop0trailing = TRUE,
  show.na = TRUE, show.0.na = FALSE, p.value = FALSE,
  p.value.digits = 2, p.value.digit.type = "signif",
  test = "parametric", test.arg.list = NULL,
  combine.groups = p.value, rm.group.na = TRUE,
  group.p.value.type = NULL, alternative = "two.sided",
  ci = NULL, ci.between = NULL, conf.level = 0.95,
  stats.in.rows = FALSE,
  data.name = deparse(substitute(object)), ...)
```

```
## S3 method for class 'factor'
```

```
summaryStats(object, group = NULL,
  drop.unused.levels = TRUE,
  digits = max(3, getOption("digits") - 3),
  digit.type = "round", drop0trailing = TRUE,
  show.na = TRUE, show.0.na = FALSE, p.value = FALSE,
  p.value.digits = 2, p.value.digit.type = "signif",
  test = "chisq", test.arg.list = NULL, combine.levels = TRUE,
  combine.groups = FALSE, rm.group.na = TRUE,
  ci = p.value & test != "chisq", conf.level = 0.95,
  stats.in.rows = FALSE, ...)

## S3 method for class 'data.frame'
summaryStats(object, ...)

## S3 method for class 'matrix'
summaryStats(object, ...)

## S3 method for class 'list'
summaryStats(object, ...)
```

Arguments

object	an object for which summary statistics are desired. In the default method, the argument object must be a numeric vector, a factor, a data frame, a matrix, or a list. When object is a data frame, all columns must be numeric or all columns must be factors. When object is a matrix, it must be a numeric or character matrix. When object is a list, all components must be numeric vectors. In the formula method, a symbolic specification of the form $y \sim g$ can be given, indicating the observations in the vector y are to be grouped according to the levels of the factor g (the form $y \sim 1$ indicates no grouping). NAs are allowed in the data.
data	when object is a formula, data specifies an optional data frame, list or environment (or object coercible by <code>as.data.frame</code> to a data frame) containing the variables in the model. If not found in data, the variables are taken from <code>environment(formula)</code> , typically the environment from which <code>summaryStats</code> is called.
subset	when object is a formula, subset specifies an optional vector specifying a subset of observations to be used.
na.action	when object is a formula, na.action specifies a function which indicates what should happen when the data contain NAs. The default is <code>na.pass</code> .
group	when object is a numeric vector or factor, group is a factor or character vector indicating which group each observation belongs to. When object is a matrix or data frame this argument is ignored and the columns define the groups. When object is a formula, this argument is ignored and the right-hand side of the formula specifies the grouping variable.
drop.unused.levels	when <code>drop.unused.levels=TRUE</code> , groups with no observations are dropped.
se	for numeric data, logical scalar indicating whether to include the standard error of the mean in the summary statistics. The default value is <code>se=FALSE</code> .
quartiles	for numeric data, logical scalar indicating whether to include the estimated 25th and 75th percentiles in the summary statistics. The default value is <code>quartiles=FALSE</code> .

<code>digits</code>	integer indicating the number of digits to use for the summary statistics. When <code>digit.type="signif"</code> , <code>digits</code> indicates the number of significant digits. When <code>digit.type="round"</code> , <code>digits</code> indicates the number of decimal places to round to. The default value is <code>max(3, getOption("digits") - 3)</code> , that is, the maximum of 3 versus the current setting of the "digits" component of <code>.Options</code> minus 3.
<code>digit.type</code>	character string indicating whether the <code>digits</code> argument refers to significant digits (<code>digit.type="signif"</code>), or how many decimal places to round to (<code>digit.type="round"</code> , the default).
<code>drop0trailing</code>	logical scalar indicating whether to drop trailing 0's when printing the summary statistics. The value of this argument is added as an attribute to the returned list and is used by the <code>print.summaryStats</code> function. The default value is <code>TRUE</code> .
<code>show.na</code>	logical scalar indicating whether to return the number of missing values. The default value is <code>show.na=TRUE</code> .
<code>show.0.na</code>	logical scalar indicating whether to display the number of missing values in the case when there are no missing values. The default value is <code>show.0.na=FALSE</code> .
<code>p.value</code>	logical scalar indicating whether to return the p-value associated with a test of hypothesis. The default value is <code>p.value=FALSE</code> . Numeric data: if there are no groups the p-value is associated with the t-test to test whether the mean is different from 0; if there are groups see the explanation for the argument <code>group.p.value.type</code> below. Factors: the p-value is associated with the test specified by the argument <code>test</code> (see below).
<code>p.value.digits</code>	integer indicating the number of digits to use for the p-value. When <code>p.value.digit.type="signif"</code> , <code>p.value.digits</code> indicates the number of significant digits. When <code>p.value.digit.type="round"</code> , <code>p.value.digits</code> indicates the number of decimal places to round to. The default value is <code>p.value.digits=2</code> .
<code>p.value.digit.type</code>	character string indicating whether the <code>p.value.digits</code> argument refers to significant digits (<code>p.value.digit.type="signif"</code> , the default), or how many decimal places to round to (<code>p.value.digit.type="round"</code>).
<code>test</code>	Numeric data: character string indicating whether to compute p-values and confidence intervals based on parametric (<code>test="parametric"</code> ; the default) or nonparametric (<code>test="nonparametric"</code>) tests when <code>p.value=TRUE</code> and/or <code>ci=TRUE</code> . When <code>test="parametric"</code> , confidence intervals are based on the t-test (see <code>t.test</code>) and p-values are based on the t-test or F-test (see <code>anova.lm</code>). When <code>test="nonparametric"</code> , confidence intervals are based on the Wilcoxon rank sum test (see <code>wilcox.test</code>) and p-values are based on the Wilcoxon rank sum test or the Kruskal-Wallis rank sum test (see <code>kruskal.test</code>). Factors: character string indicating which test to perform when <code>p.value=TRUE</code> . Possible values are <code>test="chisq"</code> for the chi-squared test as performed by the function <code>chisq.test</code> (the default), <code>test="prop"</code> for the chi-squared test as performed by the function <code>prop.test</code> , <code>test="fisher"</code> for Fisher's exact test as performed by the function <code>fisher.test</code> , and <code>test="binom"</code> for the one-sample exact binomial test as performed by <code>binom.test</code> . The chi-squared test as performed by <code>prop.test</code> is only available when the number of levels in object is 2 and either group is not supplied or the number of levels in group is 2. Fisher's exact test is only available when the number of levels in group is ≥ 2 . The exact binomial test is only available when group is not supplied and the number of levels in object is 2.

test.arg.list	a list with additional arguments to pass to the test used to compute p-values and confidence intervals. For numeric data, when test="parametric", p.value=TRUE, group.p.value.type="between" and there are two groups, if this argument is NULL or does not contain a component named var.equal, it will be modified to contain the component var.equal=TRUE. Note that this overrides the default behavior of <code>t.test</code> when there are two groups.
combine.groups	logical scalar indicating whether to show summary statistics for all groups combined. Numeric data: the default value is TRUE if p.value=TRUE, otherwise FALSE. Factors: the default value is FALSE.
rm.group.na	logical scalar indicating whether to remove missing values from the group argument. If rm.group.na=FALSE and group contains missing values then an error is returned. If rm.group.na=TRUE and group contains missing values then a warning is issued. By default rm.group.na=TRUE.
group.p.value.type	for numeric data, character string indicating which p-value(s) to compute when there is more than one group. When group.p.value.type="between" (the default when combine.groups=TRUE), the p-value is associated with the two-sample t-test (or the Wilcoxon rank sum test) in the case of two groups, and the one-way analysis of variance F-test (or Kruskal-Wallis test) in the case of three or more groups to test whether the group means (locations) are different from each other. When group.p.value.type="within" (the default when combine.groups=FALSE), the computed p-values for each group are associated with the one-sample t-test (or Wilcoxon signed rank test) to test whether the group mean (location) is different from 0.
alternative	for numeric data, character string indicating which alternative to assume for p-values and confidence intervals. Possible values are "two.sided" (the default), "less", and "greater". This argument is ignored for p-values in the case of three or more groups when group.p.value.type="between", and is ignored for confidence intervals in the case of three or more groups when ci.between=TRUE.
ci	Numeric data: logical scalar indicating whether to compute a confidence interval for the mean or each group mean. The default value is FALSE unless p.value=TRUE and there are no groups, or when p.value=TRUE and there are groups and group.p.value.type="within". Factors: logical scalar indicating whether to compute a confidence interval. A confidence interval is computed only if the number of levels in object is 2. When group is not supplied, if ci=TRUE and test="prop" or test="binom", a confidence interval for the <i>percent</i> (not probability) of the first level of object is computed. When group is supplied and the number of levels in group is 2, if ci=TRUE and test="prop", a confidence interval for the difference between <i>percents</i> (not proportions) is computed, and if test="fisher" a confidence interval for the odds ratio is computed.
ci.between	for numeric data, logical scalar indicating whether to compute a confidence interval for the difference between group means when there are two groups. The default value is ci.between=TRUE when p.value=TRUE and group.p.value.type="between", otherwise this argument is ignored.
conf.level	numeric scalar between 0 and 1 indicating the confidence level associated with the confidence intervals. The default value is conf.level=0.95.
stats.in.rows	logical scalar indicating whether to show the summary statistics in the rows or columns of the output. The default is stats.in.rows=FALSE.

data.name character string indicating the name of the data used for the summary statistics.
 combine.levels for factors, a logical scalar indicating whether to compute summary statistics
 based on combining all levels of a factor.
 ... additional arguments affecting the summary statistics produced.

Value

an object of class "summaryStats" (see [summaryStats.object](#)). Objects of class "summaryStats" are numeric matrices that contain the summary statistics produced by a call to `summaryStats` or `summaryFull`. These objects have a special printing method that by default removes trailing zeros for sample size entries and prints blanks for statistics that are normally displayed as NA (see [print.summaryStats](#)).

Summary statistics for numeric data include sample size, mean, standard deviation, median, min, and max. Options include the standard error of the mean (when `se=TRUE`), the estimated quartiles (when `quartiles=TRUE`), p-values (when `p.value=TRUE`), and/or confidence intervals (when `ci=TRUE` and/or `ci.between=TRUE`).

Summary statistics for factors include the sample size for each level of the factor and the percent of the total for that level. Options include a p-value (when `p.value=TRUE`).

Note that unlike the R function `summary` and the **EnvStats** function `summaryFull`, by default the `digits` argument for the **EnvStats** function `summaryStats` refers to how many decimal places to round to, not how many significant digits to use (see the explanation of the argument `digit.type` above).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Berthouex, P.M., and L.C. Brown. (2002). *Statistics for Environmental Engineers, Second Edition*. Lewis Publishers, Boca Raton, FL.
 Millard, S.P., and N.K. Neerchal. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, FL.
 Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ, Chapter 24.

See Also

[summary](#), [summaryFull](#), [t.test](#), [anova.lm](#), [wilcox.test](#), [kruskal.test](#), [chisq.test](#), [fisher.test](#), [binom.test](#).

Examples

```
# Page 9-3 of USEPA (2009) lists trichloroethene
# concentrations (TCE; mg/L) collected from groundwater at two wells.
# Here, the seven non-detects have been set to their detection limit.

#-----
# First, compute summary statistics for all TCE observations.

summaryStats(TCE.mg.per.L ~ 1, data = EPA.09.Table.9.1.TCE.df,
  digits = 3, data.name = "TCE")
```

```

#      N Mean    SD Median   Min  Max NA's N.Total
#TCE 27 0.09 0.064    0.1 0.004 0.25    3    30

summaryStats(TCE.mg.per.L ~ 1, data = EPA.09.Table.9.1.TCE.df,
  se = TRUE, quartiles = TRUE, digits = 3, data.name = "TCE")
#      N Mean    SD    SE Median   Min  Max 1st Qu. 3rd Qu. NA's N.Total
#TCE 27 0.09 0.064 0.012    0.1 0.004 0.25  0.031  0.12    3    30

#-----
# Now compute summary statistics by well.

summaryStats(TCE.mg.per.L ~ Well, data = EPA.09.Table.9.1.TCE.df,
  digits = 3)
#      N Mean    SD Median   Min  Max NA's N.Total
#Well.1 14 0.063 0.079  0.031 0.004 0.25    1    15
#Well.2 13 0.118 0.020  0.110 0.099 0.17    2    15

summaryStats(TCE.mg.per.L ~ Well, data = EPA.09.Table.9.1.TCE.df,
  digits = 3, stats.in.rows = TRUE)
#      Well.1 Well.2
#N      14      13
#Mean    0.063  0.118
#SD      0.079  0.02
#Median  0.031  0.11
#Min     0.004  0.099
#Max     0.25   0.17
#NA's    1      2
#N.Total 15     15

# If you want to keep trailing 0's, use the drop0trailing argument:
summaryStats(TCE.mg.per.L ~ Well, data = EPA.09.Table.9.1.TCE.df,
  digits = 3, stats.in.rows = TRUE, drop0trailing = FALSE)
#      Well.1 Well.2
#N      14.000 13.000
#Mean    0.063  0.118
#SD      0.079  0.020
#Median  0.031  0.110
#Min     0.004  0.099
#Max     0.250  0.170
#NA's    1.000  2.000
#N.Total 15.000 15.000

#-----

# Page 13-3 of USEPA (2009) lists iron concentrations (ppm) in
# groundwater collected from 6 wells.

#-----
# First, compute summary statistics for each well.

summaryStats(iron.ppm ~ Well, data = EPA.09.Ex.13.1.iron.df,
  combine.groups = FALSE, digits = 2, stats.in.rows = TRUE)
#      Well.1 Well.2 Well.3 Well.4 Well.5 Well.6
#N         4      4      4      4      4      4
#Mean     47.01  55.73  90.86  70.43 145.24 156.32
#SD       12.4   20.34  59.35  25.95  92.16  51.2

```

```

#Median 50.05 57.05 76.73 76.95 137.66 171.93
#Min    29.96 32.14 39.25 34.12 60.95 83.1
#Max    57.97 76.71 170.72 93.69 244.69 198.34

#-----
# Note the large differences in standard deviations between wells.
# Compute summary statistics for log(Iron), by Well.

summaryStats(log(Iron.ppm) ~ Well, data = EPA.09.Ex.13.1.iron.df,
  combine.groups = FALSE, digits = 2, stats.in.rows = TRUE)
#      Well.1 Well.2 Well.3 Well.4 Well.5 Well.6
#N      4      4      4      4      4      4
#Mean   3.82   3.97   4.35   4.19   4.8    5
#SD     0.3    0.4    0.66   0.45   0.7    0.4
#Median 3.91   4.02   4.29   4.34   4.8    5.14
#Min    3.4    3.47   3.67   3.53   4.11   4.42
#Max    4.06   4.34   5.14   4.54   5.5    5.29

#-----
# Include confidence intervals for the mean log(Fe) concentration
# at each well, and also the p-value from the one-way
# analysis of variance to test for a difference in well means.

summaryStats(log(Iron.ppm) ~ Well, data = EPA.09.Ex.13.1.iron.df,
  digits = 1, ci = TRUE, p.value = TRUE, stats.in.rows = TRUE)
#      Well.1 Well.2 Well.3 Well.4 Well.5 Well.6 Combined
#N      4      4      4      4      4      4      24
#Mean   3.8    4      4.3    4.2    4.8    5      4.4
#SD     0.3    0.4    0.7    0.5    0.7    0.4    0.6
#Median 3.9    4      4.3    4.3    4.8    5.1    4.3
#Min    3.4    3.5    3.7    3.5    4.1    4.4    3.4
#Max    4.1    4.3    5.1    4.5    5.5    5.3    5.5
#95%.LCL 3.3    3.3    3.3    3.5    3.7    4.4    4.1
#95%.UCL 4.3    4.6    5.4    4.9    5.9    5.6    4.6
#p.value.between                                0.025

#-----

# Using the built-in dataset HairEyeColor, summarize the frequencies
# of hair color and test whether there is a difference in proportions.
# NOTE: The data that was originally factor data has already been
# collapsed into frequency counts by category in the object
# HairEyeColor. In the examples in this section, we recreate
# the factor objects in order to show how summaryStats works
# for factor objects.

Hair <- apply(HairEyeColor, 1, sum)
Hair
#Black Brown Red Blond
# 108 286 71 127

Hair.color <- names(Hair)
Hair.fac <- factor(rep(Hair.color, times = Hair),
  levels = Hair.color)

#-----
# Compute summary statistics and perform the chi-square test

```

```

# for equal proportions of hair color

summaryStats(Hair.fac, digits = 1, p.value = TRUE)
#           N   Pct ChiSq_p
#Black      108  18.2
#Brown     286  48.3
#Red        71  12.0
#Blond     127  21.5
#Combined  592 100.0 2.5e-39

#-----
# Now test the hypothesis that 10% of the population from which
# this sample was drawn has Red hair, and compute a 95% confidence
# interval for the percent of subjects with red hair.

Red.Hair.fac <- factor(Hair.fac == "Red", levels = c(TRUE, FALSE),
  labels = c("Red", "Not Red"))

summaryStats(Red.Hair.fac, digits = 1, p.value = TRUE,
  ci = TRUE, test = "binom", test.arg.list = list(p = 0.1))
#           N Pct Exact_p 95%.LCL 95%.UCL
#Red         71  12          9.5    14.9
#Not Red    521  88
#Combined  592 100    0.11

#-----
# Now test whether the percent of people with Green eyes is the
# same for people with and without Red hair.

HairEye <- apply(HairEyeColor, 1:2, sum)
Hair.color <- rownames(HairEye)
Eye.color <- colnames(HairEye)

n11 <- HairEye[Hair.color == "Red", Eye.color == "Green"]
n12 <- sum(HairEye[Hair.color == "Red", Eye.color != "Green"])
n21 <- sum(HairEye[Hair.color != "Red", Eye.color == "Green"])
n22 <- sum(HairEye[Hair.color != "Red", Eye.color != "Green"])

Hair.fac <- factor(rep(c("Red", "Not Red"), c(n11+n12, n21+n22)),
  levels = c("Red", "Not Red"))
Eye.fac <- factor(c(rep("Green", n11), rep("Not Green", n12),
  rep("Green", n21), rep("Not Green", n22)),
  levels = c("Green", "Not Green"))

#-----
# Here are the results using the chi-square test and computing
# confidence limits for the difference between the two percentages

summaryStats(Eye.fac, group = Hair.fac, digits = 1,
  p.value = TRUE, ci = TRUE, test = "prop",
  stats.in.rows = TRUE, test.arg.list = list(correct = FALSE))
#           Green Not Green Combined
#Red(N)         14    57         71
#Red(Pct)       19.7  80.3        100
#Not Red(N)      50  471        521
#Not Red(Pct)    9.6  90.4        100

```

```

#ChiSq_p                0.01
#95%.LCL.between         0.5
#95%.UCL.between         19.7

#-----
# Here are the results using Fisher's exact test and computing
# confidence limits for the odds ratio

summaryStats(Eye.fac, group = Hair.fac, digits = 1,
  p.value = TRUE, ci = TRUE, test = "fisher",
  stats.in.rows = TRUE)
#           Green Not Green Combined
#Red(N)      14    57      71
#Red(Pct)    19.7  80.3    100
#Not Red(N)   50   471    521
#Not Red(Pct) 9.6  90.4    100
#Fisher_p                0.015
#95%.LCL.OR             1.1
#95%.UCL.OR             4.6

rm(Hair, Hair.color, Hair.fac, Red.Hair.fac, HairEye, Eye.color,
  n11, n12, n21, n22, Eye.fac)

#-----

# The data set EPA.89b.cadmium.df contains information on
# cadmium concentrations in groundwater collected from a
# background and compliance well. Compare detection frequencies
# between the well types and test for a difference using
# Fisher's exact test.

summaryStats(factor(Censored) ~ Well.type,
  data = EPA.89b.cadmium.df, digits = 1, p.value = TRUE,
  test = "fisher", stats.in.rows = TRUE)
#           FALSE TRUE Combined
#Background(N)    8    16    24
#Background(Pct)  33.3  66.7  100
#Compliance(N)   24    40    64
#Compliance(Pct) 37.5  62.5  100
#Fisher_p                0.81
#95%.LCL.OR             0.3
#95%.UCL.OR             2.5

```

tolIntGamma

Tolerance Interval for a Gamma Distribution

Description

Construct a β -content or β -expectation tolerance interval for a [gamma distribution](#).

Usage

```

tolIntGamma(x, coverage = 0.95, cov.type = "content",
  ti.type = "two-sided", conf.level = 0.95, method = "exact",

```

```

est.method = "mle", normal.approx.transform = "kulkarni.powar")

tolIntGammaAlt(x, coverage = 0.95, cov.type = "content",
  ti.type = "two-sided", conf.level = 0.95, method = "exact",
  est.method = "mle", normal.approx.transform = "kulkarni.powar")

```

Arguments

<code>x</code>	numeric vector of non-negative observations. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
<code>coverage</code>	a scalar between 0 and 1 indicating the desired coverage of the tolerance interval. The default value is <code>coverage=0.95</code> . If <code>cov.type="expectation"</code> , this argument is ignored.
<code>cov.type</code>	character string specifying the coverage type for the tolerance interval. The possible values are "content" (β -content; the default), and "expectation" (β -expectation). See the DETAILS section for more information.
<code>ti.type</code>	character string indicating what kind of tolerance interval to compute. The possible values are "two-sided" (the default), "lower", and "upper".
<code>conf.level</code>	a scalar between 0 and 1 indicating the confidence level associated with the tolerance interval. The default value is <code>conf.level=0.95</code> .
<code>method</code>	for the case of a two-sided tolerance interval, a character string specifying the method for constructing the two-sided normal distribution tolerance interval using the transformed data. This argument is ignored if <code>ti.type="lower"</code> or <code>ti.type="upper"</code> . The possible values are "exact" (the default) and "wald.wolfowitz" (the Wald-Wolfowitz approximation). See the DETAILS section of the help file for tolIntNorm for more information.
<code>est.method</code>	character string specifying the method of estimation for the shape and scale distribution parameters. The possible values are "mle" (maximum likelihood; the default), "bcmle" (bias-corrected mle), "mme" (method of moments), and "mmue" (method of moments based on the unbiased estimator of variance). See the DETAILS section of the help file for egamma for more information.
<code>normal.approx.transform</code>	character string indicating which power transformation to use. Possible values are "kulkarni.powar" (the default), "cube.root", and "fourth.root". See the DETAILS section for more information.

Details

The function `tolIntGamma` returns a tolerance interval as well as estimates of the shape and scale parameters. The function `tolIntGammaAlt` returns a tolerance interval as well as estimates of the mean and coefficient of variation.

The tolerance interval is computed by 1) using a power transformation on the original data to induce approximate normality, 2) using [tolIntNorm](#) to compute the tolerance interval, and then 3) back-transforming the interval to create a tolerance interval on the original scale. (Krishnamoorthy et al., 2008). The value `normal.approx.transform="cube.root"` uses the cube root transformation suggested by Wilson and Hilferty (1931) and used by Krishnamoorthy et al. (2008) and Singh et al. (2010b), and the value `normal.approx.transform="fourth.root"` uses the fourth root transformation suggested by Hawkins and Wixley (1986) and used by Singh et al. (2010b). The default value `normal.approx.transform="kulkarni.powar"` uses the "Optimum Power Normal Approximation Method" of Kulkarni and Powar (2010). The "optimum" power p is determined by:

$$p = -0.0705 - 0.178 \text{ shape} + 0.475 \sqrt{\text{shape}} \quad \text{if } \text{shape} \leq 1.5$$

$$p = 0.246 \quad \text{if } \text{shape} > 1.5$$

where *shape* denotes the estimate of the shape parameter. Although Kulkarni and Powar (2010) use the maximum likelihood estimate of shape to determine the power *p*, for the functions `tolIntGamma` and `tolIntGammaAlt` the power *p* is based on whatever estimate of shape is used (e.g., `est.method="mle"`, `est.method="bcmle"`, etc.).

Value

A list of class "estimate" containing the estimated parameters, the tolerance interval, and other information. See [estimate.object](#) for details.

In addition to the usual components contained in an object of class "estimate", the returned value also includes an additional component within the "interval" component:

```
normal.transform.power
      the value of the power used to transform the original data to approximate normality.
```

Warning

It is possible for the lower tolerance limit based on the transformed data to be less than 0. In this case, the lower tolerance limit on the original scale is set to 0 and a warning is issued stating that the normal approximation is not accurate in this case.

Note

The gamma distribution takes values on the positive real line. Special cases of the gamma are the [exponential](#) distribution and the [chi-square](#) distributions. Applications of the gamma include life testing, statistical ecology, queuing theory, inventory control, and precipitation processes. A gamma distribution starts to resemble a normal distribution as the shape parameter *a* tends to infinity.

Some EPA guidance documents (e.g., Singh et al., 2002; Singh et al., 2010a,b) strongly recommend against using a lognormal model for environmental data and recommend trying a gamma distribution instead.

Tolerance intervals have long been applied to quality control and life testing problems (Hahn, 1970b,c; Hahn and Meeker, 1991). References that discuss tolerance intervals in the context of environmental monitoring include: Berthouex and Brown (2002, Chapter 21), Gibbons et al. (2009), Millard and Neerchal (2001, Chapter 6), Singh et al. (2010b), and USEPA (2009).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Berthouex, P.M., and L.C. Brown. (2002). *Statistics for Environmental Engineers*. Lewis Publishers, Boca Raton.
- Draper, N., and H. Smith. (1998). *Applied Regression Analysis*. Third Edition. John Wiley and Sons, New York.
- Ellison, B.E. (1964). On Two-Sided Tolerance Intervals for a Normal Distribution. *Annals of Mathematical Statistics* **35**, 762-772.

- Evans, M., N. Hastings, and B. Peacock. (1993). *Statistical Distributions*. Second Edition. John Wiley and Sons, New York, Chapter 18.
- Gibbons, R.D., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*, Second Edition. John Wiley & Sons, Hoboken.
- Guttman, I. (1970). *Statistical Tolerance Regions: Classical and Bayesian*. Hafner Publishing Co., Darien, CT.
- Hahn, G.J. (1970b). Statistical Intervals for a Normal Population, Part I: Tables, Examples and Applications. *Journal of Quality Technology* **2**(3), 115-125.
- Hahn, G.J. (1970c). Statistical Intervals for a Normal Population, Part II: Formulas, Assumptions, Some Derivations. *Journal of Quality Technology* **2**(4), 195-206.
- Hahn, G.J., and W.Q. Meeker. (1991). *Statistical Intervals: A Guide for Practitioners*. John Wiley and Sons, New York.
- Hawkins, D. M., and R.A.J. Wixley. (1986). A Note on the Transformation of Chi-Squared Variables to Normality. *The American Statistician*, **40**, 296–298.
- Johnson, N.L., S. Kotz, and N. Balakrishnan. (1994). *Continuous Univariate Distributions, Volume 1*. Second Edition. John Wiley and Sons, New York, Chapter 17.
- Krishnamoorthy K., T. Mathew, and S. Mukherjee. (2008). Normal-Based Methods for a Gamma Distribution: Prediction and Tolerance Intervals and Stress-Strength Reliability. *Technometrics*, **50**(1), 69–78.
- Krishnamoorthy K., and T. Mathew. (2009). *Statistical Tolerance Regions: Theory, Applications, and Computation*. John Wiley and Sons, Hoboken.
- Kulkarni, H.V., and S.K. Powar. (2010). A New Method for Interval Estimation of the Mean of the Gamma Distribution. *Lifetime Data Analysis*, **16**, 431–447.
- Millard, S.P., and N.K. Neerchal. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton.
- Singh, A., A.K. Singh, and R.J. Iaci. (2002). *Estimation of the Exposure Point Concentration Term Using a Gamma Distribution*. EPA/600/R-02/084. October 2002. Technology Support Center for Monitoring and Site Characterization, Office of Research and Development, Office of Solid Waste and Emergency Response, U.S. Environmental Protection Agency, Washington, D.C.
- Singh, A., R. Maichle, and N. Armbya. (2010a). *ProUCL Version 4.1.00 User Guide (Draft)*. EPA/600/R-07/041, May 2010. Office of Research and Development, U.S. Environmental Protection Agency, Washington, D.C.
- Singh, A., N. Armbya, and A. Singh. (2010b). *ProUCL Version 4.1.00 Technical Guide (Draft)*. EPA/600/R-07/041, May 2010. Office of Research and Development, U.S. Environmental Protection Agency, Washington, D.C.
- Wilson, E.B., and M.M. Hilferty. (1931). The Distribution of Chi-Squares. *Proceedings of the National Academy of Sciences*, **17**, 684–688.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.
- USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

See Also

[GammaDist](#), [estimate.object](#), [egamma](#), [tolIntNorm](#), [predIntGamma](#).

Examples

```

# Generate 20 observations from a gamma distribution with parameters
# shape=3 and scale=2, then create a tolerance interval.
# (Note: the call to set.seed simply allows you to reproduce this
# example.)

set.seed(250)
dat <- rgamma(20, shape = 3, scale = 2)
tolIntGamma(dat)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Gamma
#
#Estimated Parameter(s):      shape = 2.203862
#                               scale = 2.174928
#
#Estimation Method:           mle
#
#Data:                         dat
#
#Sample Size:                  20
#
#Tolerance Interval Coverage:  95%
#
#Coverage Type:                content
#
#Tolerance Interval Method:    Exact using
#                               Kulkarni & Powar (2010)
#                               transformation to Normality
#                               based on mle of 'shape'
#
#Tolerance Interval Type:      two-sided
#
#Confidence Level:             95%
#
#Number of Future Observations: 1
#
#Tolerance Interval:           LTL = 0.2340438
#                               UTL = 21.2996464
#-----

# Using the same data as in the previous example, create an upper
# one-sided tolerance interval and use the bias-corrected estimate of
# shape.

tolIntGamma(dat, ti.type = "upper", est.method = "bcmle")

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Gamma
#
#Estimated Parameter(s):      shape = 1.906616

```

```

#                               scale = 2.514005
#
#Estimation Method:             bcmle
#
#Data:                          dat
#
#Sample Size:                   20
#
#Tolerance Interval Coverage:    95%
#
#Coverage Type:                 content
#
#Tolerance Interval Method:      Exact using
#                               Kulkarni & Powar (2010)
#                               transformation to Normality
#                               based on bcmle of 'shape'
#
#Tolerance Interval Type:        upper
#
#Confidence Level:              95%
#
#Tolerance Interval:             LTL = 0.00000
#                               UTL = 17.72107
#
#-----

# Clean up
rm(dat)

#-----

# Example 17-3 of USEPA (2009, p. 17-17) shows how to construct a
# beta-content upper tolerance limit with 95% coverage and 95%
# confidence using chrysene data and assuming a lognormal
# distribution. Here we will use the same chrysene data but assume a
# gamma distribution.

attach(EPA.09.Ex.17.3.chrysene.df)
Chrysene <- Chrysene.ppb[Well.type == "Background"]

#-----
# First perform a goodness-of-fit test for a gamma distribution

gofTest(Chrysene, dist = "gamma")

#Results of Goodness-of-Fit Test
#-----
#
#Test Method:                   Shapiro-Wilk GOF Based on
#                               Chen & Balakrishnan (1995)
#
#Hypothesized Distribution:      Gamma
#
#Estimated Parameter(s):        shape = 2.806929
#                               scale = 5.286026
#
#Estimation Method:             mle

```

```

#
#Data:                      Chrysene
#
#Sample Size:               8
#
#Test Statistic:           W = 0.9156306
#
#Test Statistic Parameter:  n = 8
#
#P-value:                   0.3954223
#
#Alternative Hypothesis:    True cdf does not equal the
#                           Gamma Distribution.

#-----
# Now compute the upper tolerance limit

tolIntGamma(Chrysene, ti.type = "upper", coverage = 0.95,
  conf.level = 0.95)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:      Gamma
#
#Estimated Parameter(s):   shape = 2.806929
#                           scale = 5.286026
#
#Estimation Method:        mle
#
#Data:                     Chrysene
#
#Sample Size:              8
#
#Tolerance Interval Coverage: 95%
#
#Coverage Type:            content
#
#Tolerance Interval Method: Exact using
#                           Kulkarni & Powar (2010)
#                           transformation to Normality
#                           based on mle of 'shape'
#
#Tolerance Interval Type:  upper
#
#Confidence Level:         95%
#
#Tolerance Interval:       LTL = 0.00000
#                           UTL = 69.32425

#-----
# Compare this upper tolerance limit of 69 ppb to the upper tolerance limit
# assuming a lognormal distribution.

tolIntLnorm(Chrysene, ti.type = "upper", coverage = 0.95,
  conf.level = 0.95)$interval$limits["UTL"]

```

```

#      UTL
#90.9247

#-----
# Clean up

rm(Chrysene)
detach("EPA.09.Ex.17.3.chrysene.df")

#-----

# Reproduce some of the example on page 73 of
# Krishnamoorthy et al. (2008), which uses alkalinity concentrations
# reported in Gibbons (1994) and Gibbons et al. (2009) to construct
# two-sided and one-sided upper tolerance limits for various values
# of coverage using a 95% confidence level.

tolIntGamma(Gibbons.et.al.09.Alkilinity.vec, ti.type = "upper",
  coverage = 0.9, normal.approx.transform = "cube.root")

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Gamma
#
#Estimated Parameter(s):      shape = 9.375013
#                               scale = 6.202461
#
#Estimation Method:           mle
#
#Data:                        Gibbons.et.al.09.Alkilinity.vec
#
#Sample Size:                  27
#
#Tolerance Interval Coverage:  90%
#
#Coverage Type:                content
#
#Tolerance Interval Method:    Exact using
#                               Wilson & Hilferty (1931) cube-root
#                               transformation to Normality
#
#Tolerance Interval Type:      upper
#
#Confidence Level:             95%
#
#Tolerance Interval:           LTL = 0.00000
#                               UTL = 97.70502

tolIntGamma(Gibbons.et.al.09.Alkilinity.vec,
  coverage = 0.99, normal.approx.transform = "cube.root")

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Gamma
#

```

```
#Estimated Parameter(s):      shape = 9.375013
#                               scale = 6.202461
#
#Estimation Method:           mle
#
#Data:                        Gibbons.et.al.09.Alkilinity.vec
#
#Sample Size:                 27
#
#Tolerance Interval Coverage:  99%
#
#Coverage Type:               content
#
#Tolerance Interval Method:    Exact using
#                               Wilson & Hilferty (1931) cube-root
#                               transformation to Normality
#
#Tolerance Interval Type:      two-sided
#
#Confidence Level:            95%
#
#Tolerance Interval:          LTL = 13.14318
#                               UTL = 148.43876
```

tolIntLnorm	<i>Tolerance Interval for a Lognormal Distribution</i>
-------------	--

Description

Estimate the mean and standard deviation on the log-scale for a [lognormal distribution](#), or estimate the mean and coefficient of variation for a [lognormal distribution \(alternative parameterization\)](#), and construct a β -content or β -expectation tolerance interval.

Usage

```
tolIntLnorm(x, coverage = 0.95, cov.type = "content", ti.type = "two-sided",
  conf.level = 0.95, method = "exact")

tolIntLnormAlt(x, coverage = 0.95, cov.type = "content", ti.type = "two-sided",
  conf.level = 0.95, method = "exact", est.method = "mvue")
```

Arguments

x For tolIntLnorm, x can be a numeric vector of positive observations, or an object resulting from a call to an estimating function that assumes a lognormal distribution (i.e., [elnorm](#) or [elnormCensored](#)). You *cannot* supply objects resulting from a call to estimating functions that use the alternative parameterization such as [elnormAlt](#) or [elnormAltCensored](#).
For tolIntLnormAlt, a numeric vector of positive observations.
If x is a numeric vector, missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.

coverage	a scalar between 0 and 1 indicating the desired coverage of the tolerance interval. The default value is coverage=0.95. If cov.type="expectation", this argument is ignored.
cov.type	character string specifying the coverage type for the tolerance interval. The possible values are "content" (β -content; the default), and "expectation" (β -expectation). See the DETAILS section for more information.
ti.type	character string indicating what kind of tolerance interval to compute. The possible values are "two-sided" (the default), "lower", and "upper".
conf.level	a scalar between 0 and 1 indicating the confidence level associated with the tolerance interval. The default value is conf.level=0.95.
method	for the case of a two-sided tolerance interval, a character string specifying the method for constructing the tolerance interval. This argument is ignored if ti.type="lower" or ti.type="upper". The possible values are "exact" (the default) and "wald.wolfowitz" (the Wald-Wolfowitz approximation). See the DETAILS section in this help file and the DETAILS section in the help file for tolIntNorm for more information.
est.method	for tolIntLnormAlt, a character string specifying the method of estimating the mean and coefficient of variation. <i>This argument has no effect on the method of constructing the tolerance interval.</i> Possible values are "mvue" (minimum variance unbiased; the default), "qmlle" (quasi maximum likelihood), "mle" (maximum likelihood), "mme" (method of moments), and "mmue" (method of moments based on the unbiased estimate of variance). See the DETAILS section of elnormAlt for more information on these estimation methods.

Details

The function `tolIntLnorm` returns a tolerance interval as well as estimates of the meanlog and sdlog parameters. The function `tolIntLnormAlt` returns a tolerance interval as well as estimates of the mean and coefficient of variation.

A tolerance interval for a lognormal distribution is constructed by taking the natural logarithm of the observations and constructing a tolerance interval based on the normal (Gaussian) distribution by calling [tolIntNorm](#). These tolerance limits are then exponentiated to produce a tolerance interval on the original scale of the data.

Value

If `x` is a numeric vector, a list of class "estimate" containing the estimated parameters, a component called `interval` containing the tolerance interval information, and other information. See [estimate.object](#) for details.

If `x` is the result of calling an estimation function, a list whose class is the same as `x`. The list contains the same components as `x`. If `x` already has a component called `interval`, this component is replaced with the tolerance interval information.

Note

Tolerance intervals have long been applied to quality control and life testing problems (Hahn, 1970b,c; Hahn and Meeker, 1991; Krishnamoorthy and Mathew, 2009). References that discuss tolerance intervals in the context of environmental monitoring include: Berthouex and Brown (2002, Chapter 21), Gibbons et al. (2009), Millard and Neerchal (2001, Chapter 6), Singh et al. (2010b), and USEPA (2009).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Berthouex, P.M., and L.C. Brown. (2002). *Statistics for Environmental Engineers*. Lewis Publishers, Boca Raton.
- Draper, N., and H. Smith. (1998). *Applied Regression Analysis*. Third Edition. John Wiley and Sons, New York.
- Ellison, B.E. (1964). On Two-Sided Tolerance Intervals for a Normal Distribution. *Annals of Mathematical Statistics* **35**, 762-772.
- Gibbons, R.D., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*, Second Edition. John Wiley & Sons, Hoboken.
- Guttman, I. (1970). *Statistical Tolerance Regions: Classical and Bayesian*. Hafner Publishing Co., Darien, CT.
- Hahn, G.J. (1970b). Statistical Intervals for a Normal Population, Part I: Tables, Examples and Applications. *Journal of Quality Technology* **2**(3), 115-125.
- Hahn, G.J. (1970c). Statistical Intervals for a Normal Population, Part II: Formulas, Assumptions, Some Derivations. *Journal of Quality Technology* **2**(4), 195-206.
- Hahn, G.J., and W.Q. Meeker. (1991). *Statistical Intervals: A Guide for Practitioners*. John Wiley and Sons, New York.
- Krishnamoorthy K., and T. Mathew. (2009). *Statistical Tolerance Regions: Theory, Applications, and Computation*. John Wiley and Sons, Hoboken.
- Millard, S.P., and N.K. Neerchal. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton.
- Odeh, R.E., and D.B. Owen. (1980). *Tables for Normal Tolerance Limits, Sampling Plans, and Screening*. Marcel Dekker, New York.
- Owen, D.B. (1962). *Handbook of Statistical Tables*. Addison-Wesley, Reading, MA.
- Singh, A., R. Maichle, and N. Armbya. (2010a). *ProUCL Version 4.1.00 User Guide (Draft)*. EPA/600/R-07/041, May 2010. Office of Research and Development, U.S. Environmental Protection Agency, Washington, D.C.
- Singh, A., N. Armbya, and A. Singh. (2010b). *ProUCL Version 4.1.00 Technical Guide (Draft)*. EPA/600/R-07/041, May 2010. Office of Research and Development, U.S. Environmental Protection Agency, Washington, D.C.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.
- USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.
- Wald, A., and J. Wolfowitz. (1946). Tolerance Limits for a Normal Distribution. *Annals of Mathematical Statistics* **17**, 208-215.

See Also

[tolIntNorm](#), [Lognormal](#), [LognormalAlt](#), [estimate.object](#), [elnorm](#), [elnormAlt](#), [eqlnorm](#), [predIntLnorm](#), [Tolerance Intervals](#), [Estimating Distribution Parameters](#), [Estimating Distribution Quantiles](#).

Examples

```

# Generate 20 observations from a lognormal distribution with parameters
# meanlog=0 and sdlog=1. Use tolIntLnorm to estimate
# the mean and standard deviation of the log of the true distribution, and
# construct a two-sided 90% beta-content tolerance interval with associated
# confidence level 95%.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
dat <- rlnorm(20)
tolIntLnorm(dat, coverage = 0.9)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Lognormal
#
#Estimated Parameter(s):      meanlog = -0.06941976
#                             sdlog   =  0.59011300
#
#Estimation Method:           mvue
#
#Data:                         dat
#
#Sample Size:                  20
#
#Tolerance Interval Coverage:  90%
#
#Coverage Type:                content
#
#Tolerance Interval Method:    Exact
#
#Tolerance Interval Type:      two-sided
#
#Confidence Level:             95%
#
#Tolerance Interval:           LTL = 0.237457
#                             UTL = 3.665369

# The exact two-sided interval that contains 90% of this distribution
# is given by: [0.193, 5.18].

qlnorm(p = c(0.05, 0.95))
#[1] 0.1930408 5.1802516

# Clean up
rm(dat)

#=====

# Example 17-3 of USEPA (2009, p. 17-17) shows how to construct a
# beta-content upper tolerance limit with 95% coverage and 95%
# confidence using chrysene data and assuming a lognormal distribution.
# The data for this example are stored in EPA.09.Ex.17.3.chrysene.df,
# which contains chrysene concentration data (ppb) found in water

```

```

# samples obtained from two background wells (Wells 1 and 2) and
# three compliance wells (Wells 3, 4, and 5). The tolerance limit
# is based on the data from the background wells.

head(EPA.09.Ex.17.3.chrysene.df)
# Month Well Well.type Chrysene.ppb
#1 1 Well.1 Background 19.7
#2 2 Well.1 Background 39.2
#3 3 Well.1 Background 7.8
#4 4 Well.1 Background 12.8
#5 1 Well.2 Background 10.2
#6 2 Well.2 Background 7.2

longToWide(EPA.09.Ex.17.3.chrysene.df, "Chrysene.ppb", "Month", "Well")
# Well.1 Well.2 Well.3 Well.4 Well.5
#1 19.7 10.2 68.0 26.8 47.0
#2 39.2 7.2 48.9 17.7 30.5
#3 7.8 16.1 30.1 31.9 15.0
#4 12.8 5.7 38.1 22.2 23.4

with(EPA.09.Ex.17.3.chrysene.df,
     tolIntLnorm(Chrysene.ppb[Well.type == "Background"],
                 ti.type = "upper", coverage = 0.95, conf.level = 0.95))

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution: Lognormal
#
#Estimated Parameter(s): meanlog = 2.5085773
#                        sdlog = 0.6279479
#
#Estimation Method: mvue
#
#Data: Chrysene.ppb[Well.type == "Background"]
#
#Sample Size: 8
#
#Tolerance Interval Coverage: 95%
#
#Coverage Type: content
#
#Tolerance Interval Method: Exact
#
#Tolerance Interval Type: upper
#
#Confidence Level: 95%
#
#Tolerance Interval: LTL = 0.0000
#                  UTL = 90.9247
#-----

# Repeat the above example, but estimate the mean and
# coefficient of variation on the original scale
#-----

```

```
with(EPA.09.Ex.17.3.chrysene.df,
  tolIntLnrmAlt(Chrysene.ppb[Well.type == "Background"],
    ti.type = "upper", coverage = 0.95, conf.level = 0.95))

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:      Lognormal
#
#Estimated Parameter(s):   mean = 14.5547353
#                           cv   = 0.6390825
#
#Estimation Method:        mvue
#
#Data:                     Chrysene.ppb[Well.type == "Background"]
#
#Sample Size:              8
#
#Tolerance Interval Coverage: 95%
#
#Coverage Type:            content
#
#Tolerance Interval Method: Exact
#
#Tolerance Interval Type:   upper
#
#Confidence Level:         95%
#
#Tolerance Interval:       LTL = 0.0000
#                           UTL = 90.9247
```

tolIntNorm	<i>Tolerance Interval for a Normal Distribution</i>
------------	---

Description

Construct a β -content or β -expectation tolerance interval for a [normal distribution](#).

Usage

```
tolIntNorm(x, coverage = 0.95, cov.type = "content",
  ti.type = "two-sided", conf.level = 0.95, method = "exact")
```

Arguments

- x numeric vector of observations, or an object resulting from a call to an estimating function that assumes a normal (Gaussian) distribution (i.e., [enorm](#) or [enormCensored](#)). If x is a numeric vector, missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
- coverage a scalar between 0 and 1 indicating the desired coverage of the tolerance interval. The default value is coverage=0.95. If cov.type="expectation", this argument is ignored.

<code>cov.type</code>	character string specifying the coverage type for the tolerance interval. The possible values are "content" (β -content; the default), and "expectation" (β -expectation). See the DETAILS section for more information.
<code>ti.type</code>	character string indicating what kind of tolerance interval to compute. The possible values are "two-sided" (the default), "lower", and "upper".
<code>conf.level</code>	a scalar between 0 and 1 indicating the confidence level associated with the tolerance interval. The default value is <code>conf.level=0.95</code> .
<code>method</code>	for the case of a two-sided tolerance interval, a character string specifying the method for constructing the tolerance interval. This argument is ignored if <code>ti.type="lower"</code> or <code>ti.type="upper"</code> . The possible values are "exact" (the default) and "wald.wolfowitz" (the Wald-Wolfowitz approximation). See the DETAILS section for more information.

Details

If `x` contains any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

A tolerance interval for some population is an interval on the real line constructed so as to contain $100\beta\%$ of the population (i.e., $100\beta\%$ of all future observations), where $0 < \beta < 1$. The quantity $100\beta\%$ is called the *coverage*.

There are two kinds of tolerance intervals (Guttman, 1970):

- A β -content tolerance interval with confidence level $100(1 - \alpha)\%$ is constructed so that it contains at least $100\beta\%$ of the population (i.e., the coverage is at least $100\beta\%$) with probability $100(1 - \alpha)\%$, where $0 < \alpha < 1$. The quantity $100(1 - \alpha)\%$ is called the confidence level or confidence coefficient associated with the tolerance interval.
- A β -expectation tolerance interval is constructed so that the *average* coverage of the interval is $100\beta\%$.

Note: A β -expectation tolerance interval with coverage $100\beta\%$ is equivalent to a prediction interval for one future observation with associated confidence level $100\beta\%$. Note that there is no explicit confidence level associated with a β -expectation tolerance interval. If a β -expectation tolerance interval is treated as a β -content tolerance interval, the confidence level associated with this tolerance interval is usually around 50% (e.g., Guttman, 1970, Table 4.2, p.76).

For a normal distribution, the form of a two-sided $100(1 - \alpha)\%$ tolerance interval is:

$$[\bar{x} - Ks, \bar{x} + Ks]$$

where \bar{x} denotes the sample mean, s denotes the sample standard deviation, and K denotes a constant that depends on the sample size n , the coverage, and, for a β -content tolerance interval (but not a β -expectation tolerance interval), the confidence level.

Similarly, the form of a one-sided lower tolerance interval is:

$$[\bar{x} - Ks, \infty]$$

and the form of a one-sided upper tolerance interval is:

$$[-\infty, \bar{x} + Ks]$$

but K differs for one-sided versus two-sided tolerance intervals. The derivation of the constant K is explained in the help file for [tolIntNormK](#).

Value

If x is a numeric vector, `tolIntNorm` returns a list of class "estimate" containing the estimated parameters, a component called `interval` containing the tolerance interval information, and other information. See [estimate.object](#) for details.

If x is the result of calling an estimation function, `tolIntNorm` returns a list whose class is the same as x . The list contains the same components as x . If x already has a component called `interval`, this component is replaced with the tolerance interval information.

Note

Tolerance intervals have long been applied to quality control and life testing problems (Hahn, 1970b,c; Hahn and Meeker, 1991; Krishnamoorthy and Mathew, 2009). References that discuss tolerance intervals in the context of environmental monitoring include: Berthouex and Brown (2002, Chapter 21), Gibbons et al. (2009), Millard and Neerchal (2001, Chapter 6), Singh et al. (2010b), and USEPA (2009).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Berthouex, P.M., and L.C. Brown. (2002). *Statistics for Environmental Engineers*. Lewis Publishers, Boca Raton.
- Draper, N., and H. Smith. (1998). *Applied Regression Analysis*. Third Edition. John Wiley and Sons, New York.
- Ellison, B.E. (1964). On Two-Sided Tolerance Intervals for a Normal Distribution. *Annals of Mathematical Statistics* **35**, 762-772.
- Gibbons, R.D., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*, Second Edition. John Wiley & Sons, Hoboken.
- Guttman, I. (1970). *Statistical Tolerance Regions: Classical and Bayesian*. Hafner Publishing Co., Darien, CT.
- Hahn, G.J. (1970b). Statistical Intervals for a Normal Population, Part I: Tables, Examples and Applications. *Journal of Quality Technology* **2**(3), 115-125.
- Hahn, G.J. (1970c). Statistical Intervals for a Normal Population, Part II: Formulas, Assumptions, Some Derivations. *Journal of Quality Technology* **2**(4), 195-206.
- Hahn, G.J., and W.Q. Meeker. (1991). *Statistical Intervals: A Guide for Practitioners*. John Wiley and Sons, New York.
- Krishnamoorthy K., and T. Mathew. (2009). *Statistical Tolerance Regions: Theory, Applications, and Computation*. John Wiley and Sons, Hoboken.
- Millard, S.P., and N.K. Neerchal. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton.
- Odeh, R.E., and D.B. Owen. (1980). *Tables for Normal Tolerance Limits, Sampling Plans, and Screening*. Marcel Dekker, New York.
- Owen, D.B. (1962). *Handbook of Statistical Tables*. Addison-Wesley, Reading, MA.
- Singh, A., R. Maichle, and N. Armbya. (2010a). *ProUCL Version 4.1.00 User Guide (Draft)*. EPA/600/R-07/041, May 2010. Office of Research and Development, U.S. Environmental Protection Agency, Washington, D.C.

Singh, A., N. Armbya, and A. Singh. (2010b). *ProUCL Version 4.1.00 Technical Guide (Draft)*. EPA/600/R-07/041, May 2010. Office of Research and Development, U.S. Environmental Protection Agency, Washington, D.C.

USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.

USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.

Wald, A., and J. Wolfowitz. (1946). Tolerance Limits for a Normal Distribution. *Annals of Mathematical Statistics* **17**, 208-215.

See Also

[tolIntNormK](#), [tolIntLnorm](#), [Normal](#), [estimate.object](#), [enorm](#), [eqnorm](#), [predIntNorm](#), [Tolerance Intervals](#), [Estimating Distribution Parameters](#), [Estimating Distribution Quantiles](#).

Examples

```
# Generate 20 observations from a normal distribution with parameters
# mean=10 and sd=2, then create a tolerance interval.
# (Note: the call to set.seed simply allows you to reproduce this
# example.)
```

```
set.seed(250)
dat <- rnorm(20, mean = 10, sd = 2)
tolIntNorm(dat)
```

```
#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:      Normal
#
#Estimated Parameter(s):   mean = 9.861160
#                           sd   = 1.180226
#
#Estimation Method:        mvue
#
#Data:                     dat
#
#Sample Size:              20
#
#Tolerance Interval Coverage: 95%
#
#Coverage Type:            content
#
#Tolerance Interval Method: Exact
#
#Tolerance Interval Type:   two-sided
#
#Confidence Level:         95%
#
#Tolerance Interval:       LTL = 6.603328
#                           UTL = 13.118993
```

```

#-----

# Clean up
rm(dat)

#-----

# Example 17-3 of USEPA (2009, p. 17-17) shows how to construct a
# beta-content upper tolerance limit with 95% coverage and 95%
# confidence using chrysene data and assuming a lognormal distribution.
# The data for this example are stored in EPA.09.Ex.17.3.chrysene.df,
# which contains chrysene concentration data (ppb) found in water
# samples obtained from two background wells (Wells 1 and 2) and
# three compliance wells (Wells 3, 4, and 5). The tolerance limit
# is based on the data from the background wells.

# Here we will first take the log of the data and
# then construct the tolerance interval; note however that it is
# easier to call the function tolIntNorm instead using the
# original data.

head(EPA.09.Ex.17.3.chrysene.df)
# Month Well Well.type Chrysene.ppb
#1 1 Well.1 Background 19.7
#2 2 Well.1 Background 39.2
#3 3 Well.1 Background 7.8
#4 4 Well.1 Background 12.8
#5 1 Well.2 Background 10.2
#6 2 Well.2 Background 7.2

longToWide(EPA.09.Ex.17.3.chrysene.df, "Chrysene.ppb", "Month", "Well")
# Well.1 Well.2 Well.3 Well.4 Well.5
#1 19.7 10.2 68.0 26.8 47.0
#2 39.2 7.2 48.9 17.7 30.5
#3 7.8 16.1 30.1 31.9 15.0
#4 12.8 5.7 38.1 22.2 23.4

tol.int.list <- with(EPA.09.Ex.17.3.chrysene.df,
  tolIntNorm(log(Chrysene.ppb[Well.type == "Background"]),
    ti.type = "upper", coverage = 0.95, conf.level = 0.95))

tol.int.list

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution: Normal
#
#Estimated Parameter(s): mean = 2.5085773
# sd = 0.6279479
#
#Estimation Method: mvue
#
#Data: log(Chrysene.ppb[Well.type == "Background"])
#
#Sample Size: 8

```

```

#
#Tolerance Interval Coverage:    95%
#
#Coverage Type:                  content
#
#Tolerance Interval Method:      Exact
#
#Tolerance Interval Type:        upper
#
#Confidence Level:               95%
#
#Tolerance Interval:             LTL =      -Inf
#                               UTL = 4.510032

# Compute the upper tolerance interval on the original scale
# by exponentiating the upper tolerance limit:

exp(tol.int.list$interval$limits["UTL"])
#    UTL
#90.9247

#-----

# Clean up

rm(tol.int.list)

```

tolIntNormCensored	<i>Tolerance Interval for a Normal Distribution Based on Censored Data</i>
--------------------	--

Description

Construct a β -content or β -expectation tolerance interval for a normal distribution based on Type I or Type II censored data.

Usage

```
tolIntNormCensored(x, censored, censoring.side = "left", coverage = 0.95,
  cov.type = "content", ti.type = "two-sided", conf.level = 0.95,
  method = "mle", ti.method = "exact.for.complete", seed = NULL,
  nmc = 1000)
```

Arguments

x	numeric vector of observations. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
censored	numeric or logical vector indicating which values of x are censored. This must be the same length as x. If the mode of censored is "logical", TRUE values correspond to elements of x that are censored, and FALSE values correspond to elements of x that are not censored. If the mode of censored is "numeric", it must contain only 1's and 0's; 1 corresponds to TRUE and 0 corresponds to FALSE. Missing (NA) values are allowed but will be removed.

censoring.side	character string indicating on which side the censoring occurs. The possible values are "left" (the default) and "right".
coverage	a scalar between 0 and 1 indicating the desired coverage of the tolerance interval. The default value is coverage=0.95.
cov.type	character string specifying the coverage type for the tolerance interval. The possible values are "content" (β -content; the default), and "expectation" (β -expectation). See the DETAILS section for more information.
ti.type	character string indicating what kind of tolerance interval to compute. The possible values are "two-sided" (the default), "lower", and "upper".
conf.level	a scalar between 0 and 1 indicating the confidence level associated with the tolerance interval. The default value is conf.level=0.95.
method	character string indicating the method to use for parameter estimation. For singly censored data, possible values are "mle" (the default), "bcmle", "qq.reg", "qq.reg.w.cen.level", "impute.w.qq.reg", "impute.w.qq.reg.w.cen.level", "impute.w.mle", "iterative.impute.w.qq.reg", "m.est", and "half.cen.level". See the help file for enormCensored for details. For multiply censored data, possible values are "mle" (the default), "qq.reg", "impute.w.qq.reg", and "half.cen.level". See the help file for enormCensored for details.
ti.method	character string specifying the method for constructing the tolerance interval. The default method is "gpq" (Generalized Pivotal Quantity). Other possible values are "exact.for.complete", and, for a two-sided tolerance interval (i.e., when ti.type="two-sided"), "wald.wolfowitz.for.complete". See the DETAILS section for more information.
seed	for the case when ti.method="gpq", a positive integer to pass to the function gpqTolIntNormSinglyCensored or gpqTolIntNormMultiplyCensored . This argument is ignored if seed=NULL (the default). Using the seed argument lets you reproduce the exact same result if all other arguments stay the same.
nmc	for the case when ti.method="gpq", a positive integer ≥ 10 indicating the number of Monte Carlo trials to run in order to compute the GPQ(s).

Details

See the help file for [tolIntNorm](#) for an explanation of tolerance intervals. When ti.method="gpq", the tolerance interval is constructed using the method of Generalized Pivotal Quantities as explained in Krishnamoorthy and Mathew (2009, p. 327). When ti.method="exact.for.complete" or ti.method="wald.wolfowitz.for.complete", the tolerance interval is constructed by first computing the maximum likelihood estimates of the mean and standard deviation by calling [enormCensored](#), then passing these values to the function [tolIntNorm](#) to produce the tolerance interval as if the estimates were based on complete rather than censored data. These last two methods are purely ad-hoc and their properties need to be studied.

Value

A list of class "estimateCensored" containing the estimated parameters, the tolerance interval, and other information. See [estimateCensored.object](#) for details.

Note

Tolerance intervals have long been applied to quality control and life testing problems (Hahn, 1970b,c; Hahn and Meeker, 1991; Krishnamoorthy and Mathew, 2009). References that discuss tolerance intervals in the context of environmental monitoring include: Berthouex and Brown (2002, Chapter 21), Gibbons et al. (2009), Millard and Neerchal (2001, Chapter 6), Singh et al. (2010b), and USEPA (2009).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

Krishnamoorthy K., and T. Mathew. (2009). *Statistical Tolerance Regions: Theory, Applications, and Computation*. John Wiley and Sons, Hoboken.

[tolIntNorm](#)

See Also

[gpqTolIntNormSinglyCensored](#), [eqnormCensored](#), [enormCensored](#), [estimateCensored.object](#).

Examples

```
# Generate 20 observations from a normal distribution with parameters
# mean=10 and sd=3, censor the observations less than 9,
# then create a one-sided upper tolerance interval with 90%
# coverage and 95% confidence based on these Type I left, singly
# censored data.
# (Note: the call to set.seed and supplying the seed argument to
# tolIntNormCensored simply allows you to reproduce this example.)

set.seed(250)
dat <- rnorm(20, mean = 10, sd = 3)
censored <- dat < 9
dat[censored] <- 9

tolIntNormCensored(dat, censored, coverage = 0.9, ti.type="upper",
  seed = 432)

# Results of Distribution Parameter Estimation
# Based on Type I Censored Data
# -----

# Assumed Distribution:          Normal

# Censoring Side:               left

# Censoring Level(s):          9

# Estimated Parameter(s):      mean = 9.700962
#                               sd   = 1.845067

# Estimation Method:           MLE

# Data:                         dat
```

```

# Censoring Variable:          censored

# Sample Size:                20

# Percent Censored:           35%

# Tolerance Interval Coverage: 90%

# Coverage Type:              content

# Tolerance Interval Method:   Generalized Pivotal Quantity

# Number of Monte Carlos:     1000

# Tolerance Interval Type:    upper

# Confidence Level:            95%

# Tolerance Interval:         LTL =      -Inf
#                             UTL = 13.56826

# Note: The true 90'th percentile is 13.84465
#-----
qnorm(0.9, mean = 10, sd = 3)
# [1] 13.84465

# Compare the result using the method "exact.for.complete"
tolIntNormCensored(dat, censored, coverage = 0.9, ti.type="upper",
  ti.method = "exact.for.complete")$interval$limits
#      LTL      UTL
#      -Inf 13.25454

# Clean Up
#-----
rm(dat, censored)

#-----

# Example 15-1 of USEPA (2009, p. 15-10) shows how to estimate
# the mean and standard deviation using log-transformed multiply
# left-censored manganese concentration data. Here we'll construct a
# 95% upper tolerance limit with 90% coverage using these data.

EPA.09.Ex.15.1.manganese.df
#   Sample  Well Manganese.Orig.ppb Manganese.ppb Censored
# 1      1 Well.1          <5          5.0      TRUE
# 2      2 Well.1         12.1         12.1     FALSE
# 3      3 Well.1         16.9         16.9     FALSE
# ...
# 23     3 Well.5          3.3          3.3     FALSE
# 24     4 Well.5          8.4          8.4     FALSE
# 25     5 Well.5          <2          2.0      TRUE

with(EPA.09.Ex.15.1.manganese.df,
  tolIntNormCensored(log(Manganese.ppb), Censored, coverage = 0.9,

```

```

    ti.type = "upper", seed = 47))

# Results of Distribution Parameter Estimation
# Based on Type I Censored Data
# -----

# Assumed Distribution:      Normal

# Censoring Side:           left

# Censoring Level(s):       0.6931472 1.6094379

# Estimated Parameter(s):   mean = 2.215905
#                           sd   = 1.356291

# Estimation Method:        MLE

# Data:                     log(Manganese.ppb)

# Censoring Variable:       censored

# Sample Size:              25

# Percent Censored:         24%

# Tolerance Interval Coverage: 90%

# Coverage Type:            content

# Tolerance Interval Method: Generalized Pivotal Quantity

# Number of Monte Carlos:   1000

# Tolerance Interval Type:  upper

# Confidence Level:         95%

# Tolerance Interval:       LTL =      -Inf
#                           UTL = 4.657662

```

tolIntNormK

Compute the Value of K for a Tolerance Interval for a Normal Distribution

Description

Compute the value of K (the multiplier of estimated standard deviation) used to construct a tolerance interval based on data from a normal distribution.

Usage

```

tolIntNormK(n, df = n - 1, coverage = 0.95, cov.type = "content",
  ti.type = "two-sided", conf.level = 0.95, method = "exact",
  rel.tol = 1e-07, abs.tol = rel.tol)

```

Arguments

n	a positive integer greater than 2 indicating the sample size upon which the tolerance interval is based.
df	the degrees of freedom associated with the tolerance interval. The default is $df=n-1$.
coverage	a scalar between 0 and 1 indicating the desired coverage of the tolerance interval. The default value is $coverage=0.95$.
cov.type	character string specifying the coverage type for the tolerance interval. The possible values are "content" (β -content; the default), and "expectation" (β -expectation). See the help file for tolIntNorm for more information on the difference between β -content and β -expectation tolerance intervals.
ti.type	character string indicating what kind of tolerance interval to compute. The possible values are "two-sided" (the default), "lower", and "upper".
conf.level	a scalar between 0 and 1 indicating the confidence level associated with the tolerance interval. The default value is $conf.level=0.95$.
method	for the case of a two-sided tolerance interval, a character string specifying the method for constructing the tolerance interval. This argument is ignored if $ti.type="lower"$ or $ti.type="upper"$. The possible values are "exact" (the default) and "wald.wolfowitz" (the Wald-Wolfowitz approximation). See the DETAILS section for more information.
rel.tol	in the case when $ti.type="two-sided"$ and $method="exact"$, the argument $rel.tol$ is passed to the function integrate . The default value is $rel.tol=1e-07$.
abs.tol	in the case when $ti.type="two-sided"$ and $method="exact"$, the argument $abs.tol$ is passed to the function integrate . The default value is the value of $rel.tol$.

Details

A tolerance interval for some population is an interval on the real line constructed so as to contain $100\beta\%$ of the population (i.e., $100\beta\%$ of all future observations), where $0 < \beta < 1$. The quantity $100\beta\%$ is called the coverage.

There are two kinds of tolerance intervals (Guttman, 1970):

- A β -content tolerance interval with confidence level $100(1 - \alpha)\%$ is constructed so that it contains at least $100\beta\%$ of the population (i.e., the coverage is at least $100\beta\%$) with probability $100(1 - \alpha)\%$, where $0 < \alpha < 1$. The quantity $100(1 - \alpha)\%$ is called the confidence level or confidence coefficient associated with the tolerance interval.
- A β -expectation tolerance interval is constructed so that the *average* coverage of the interval is $100\beta\%$.

Note: A β -expectation tolerance interval with coverage $100\beta\%$ is equivalent to a prediction interval for one future observation with associated confidence level $100\beta\%$. Note that there is no explicit confidence level associated with a β -expectation tolerance interval. If a β -expectation tolerance interval is treated as a β -content tolerance interval, the confidence level associated with this tolerance interval is usually around 50% (e.g., Guttman, 1970, Table 4.2, p.76).

For a normal distribution, the form of a two-sided $100(1 - \alpha)\%$ tolerance interval is:

$$[\bar{x} - Ks, \bar{x} + Ks]$$

where \bar{x} denotes the sample mean, s denotes the sample standard deviation, and K denotes a constant that depends on the sample size n , the coverage, and, for a β -content tolerance interval (but not a β -expectation tolerance interval), the confidence level.

Similarly, the form of a one-sided lower tolerance interval is:

$$[\bar{x} - Ks, \infty]$$

and the form of a one-sided upper tolerance interval is:

$$[-\infty, \bar{x} + Ks]$$

but K differs for one-sided versus two-sided tolerance intervals.

The Derivation of K for a β -Content Tolerance Interval

One-Sided Case

When `ti.type="upper"` or `ti.type="lower"`, the constant K for a $100\beta\%$ β -content tolerance interval with associated confidence level $100(1 - \alpha)\%$ is given by:

$$K = t(n - 1, 1 - \alpha, z_\beta \sqrt{n})$$

where $t(\nu, p, \delta)$ denotes the p 'th quantile of a non-central t-distribution with ν degrees of freedom and noncentrality parameter δ (see the help file for [TDist](#)), and z_p denotes the p 'th quantile of a standard normal distribution.

Two-Sided Case

When `ti.type="two-sided"` and `method="exact"`, the exact formula for the constant K for a $100\beta\%$ β -content tolerance interval with associated confidence level $100(1 - \alpha)\%$ requires numerical integration and has been derived by several different authors, including Odeh (1978), Eberhardt et al. (1989), Jilek (1988), Fujino (1989), and Janiga and Miklos (2001). Specifically, for given values of the sample size n , degrees of freedom ν , confidence level $(1 - \alpha)$, and coverage β , the constant K is the solution to the equation:

$$\sqrt{\frac{n}{2\pi}} \int_{-\infty}^{\infty} F(x, K, \nu, R) e^{(-nx^2)/2} dx = 1 - \alpha$$

where $F(x, K, \nu, R)$ denotes the upper-tail area from $(\nu R^2)/K^2$ to ∞ of the chi-squared distribution with ν degrees of freedom, and R is the solution to the equation:

$$\Phi(x + R) - \Phi(x - R) = \beta$$

where $\Phi()$ denotes the standard normal cumulative distribution function.

When `ti.type="two-sided"` and `method="wald.wolfowitz"`, the approximate formula due to Wald and Wolfowitz (1946) for the constant K for a $100\beta\%$ β -content tolerance interval with associated confidence level $100(1 - \alpha)\%$ is given by:

$$K \approx r u$$

where r is the solution to the equation:

$$\Phi\left(\frac{1}{\sqrt{n}} + r\right) - \Phi\left(\frac{1}{\sqrt{n}} - r\right) = \beta$$

$\Phi()$ denotes the standard normal cumulative distribuiton function, and u is given by:

$$u = \sqrt{\frac{n-1}{\chi^2(n-1, \alpha)}}$$

where $\chi^2(\nu, p)$ denotes the p 'th quantile of the chi-squared distribution with ν degrees of freedom.

The Derivation of K for a β -Expectation Tolerance Interval

As stated above, a β -expectation tolerance interval with coverage $100\beta\%$ is equivalent to a prediction interval for one future observation with associated confidence level $100\beta\%$. This is because the probability that any single future observation will fall into this interval is $100\beta\%$, so the distribution of the number of N future observations that will fall into this interval is binomial with parameters size = N and prob = β (see the help file for [Binomial](#)). Hence the expected proportion of future observations that will fall into this interval is $100\beta\%$ and is independent of the value of N . See the help file for [predIntNormK](#) for information on how to derive K for these intervals.

Value

The value of K , a numeric scalar used to construct tolerance intervals for a normal (Gaussian) distribution.

Note

Tabled values of K are given in Gibbons et al. (2009), Gilbert (1987), Guttman (1970), Krishnamoorthy and Mathew (2009), Owen (1962), Odeh and Owen (1980), and USEPA (2009).

Tolerance intervals have long been applied to quality control and life testing problems (Hahn, 1970b,c; Hahn and Meeker, 1991; Krishnamoorthy and Mathew, 2009). References that discuss tolerance intervals in the context of environmental monitoring include: Berthouex and Brown (2002, Chapter 21), Gibbons et al. (2009), Millard and Neerchal (2001, Chapter 6), Singh et al. (2010b), and USEPA (2009).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Berthouex, P.M., and L.C. Brown. (2002). *Statistics for Environmental Engineers*. Lewis Publishers, Boca Raton.
- Draper, N., and H. Smith. (1998). *Applied Regression Analysis*. Third Edition. John Wiley and Sons, New York.
- Eberhardt, K.R., R.W. Mee, and C.P. Reeve. (1989). Computing Factors for Exact Two-Sided Tolerance Limits for a Normal Distribution. *Communications in Statistics, Part B-Simulation and Computation* **18**, 397-413.
- Ellison, B.E. (1964). On Two-Sided Tolerance Intervals for a Normal Distribution. *Annals of Mathematical Statistics* **35**, 762-772.
- Fujino, T. (1989). Exact Two-Sided Tolerance Limits for a Normal Distribution. *Japanese Journal of Applied Statistics* **18**, 29-36.
- Gibbons, R.D., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*, Second Edition. John Wiley & Sons, Hoboken.

- Gilbert, R.O. (1987). *Statistical Methods for Environmental Pollution Monitoring*. Van Nostrand Reinhold, New York.
- Guttman, I. (1970). *Statistical Tolerance Regions: Classical and Bayesian*. Hafner Publishing Co., Darien, CT.
- Hahn, G.J. (1970b). Statistical Intervals for a Normal Population, Part I: Tables, Examples and Applications. *Journal of Quality Technology* **2**(3), 115-125.
- Hahn, G.J. (1970c). Statistical Intervals for a Normal Population, Part II: Formulas, Assumptions, Some Derivations. *Journal of Quality Technology* **2**(4), 195-206.
- Hahn, G.J., and W.Q. Meeker. (1991). *Statistical Intervals: A Guide for Practitioners*. John Wiley and Sons, New York.
- Jilek, M. (1988). *Statistické Toleranční Meze*. SNTL, Praha.
- Krishnamoorthy K., and T. Mathew. (2009). *Statistical Tolerance Regions: Theory, Applications, and Computation*. John Wiley and Sons, Hoboken.
- Janiga, I., and R. Miklos. (2001). Statistical Tolerance Intervals for a Normal Distribution. *Measurement Science Review* **11**, 29-32.
- Millard, S.P., and N.K. Neerchal. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton.
- Odeh, R.E. (1978). Tables of Two-Sided Tolerance Factors for a Normal Distribution. *Communications in Statistics, Part B-Simulation and Computation* **7**, 183-201.
- Odeh, R.E., and D.B. Owen. (1980). *Tables for Normal Tolerance Limits, Sampling Plans, and Screening*. Marcel Dekker, New York.
- Owen, D.B. (1962). *Handbook of Statistical Tables*. Addison-Wesley, Reading, MA.
- Singh, A., R. Maichle, and N. Armbya. (2010a). *ProUCL Version 4.1.00 User Guide (Draft)*. EPA/600/R-07/041, May 2010. Office of Research and Development, U.S. Environmental Protection Agency, Washington, D.C.
- Singh, A., N. Armbya, and A. Singh. (2010b). *ProUCL Version 4.1.00 Technical Guide (Draft)*. EPA/600/R-07/041, May 2010. Office of Research and Development, U.S. Environmental Protection Agency, Washington, D.C.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.
- USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.
- Wald, A., and J. Wolfowitz. (1946). Tolerance Limits for a Normal Distribution. *Annals of Mathematical Statistics* **17**, 208-215.

See Also

[tolIntNorm](#), [predIntNorm](#), [Normal](#), [estimate.object](#), [enorm](#), [eqnorm](#), [Tolerance Intervals](#), [Prediction Intervals](#), [Estimating Distribution Parameters](#), [Estimating Distribution Quantiles](#).

Examples

```
# Compute the value of K for a two-sided 95% beta-content
# tolerance interval with associated confidence level 95%
# given a sample size of n=20.
```



```

#-----
# Exact method

tolIntNormK(n = 20)
#[1] 2.760346

#-----
# Approximate method due to Wald and Wolfowitz (1946)

tolIntNormK(n = 20, method = "wald")
# [1] 2.751789

#-----

# Compute the value of K for a one-sided upper tolerance limit
# with 99% coverage and associated confidence level 90%
# given a sample size of n=20.

tolIntNormK(n = 20, ti.type = "upper", coverage = 0.99,
  conf.level = 0.9)
#[1] 3.051543

#-----

# Example 17-3 of USEPA (2009, p. 17-17) shows how to construct a
# beta-content upper tolerance limit with 95% coverage and 95%
# confidence using chrysene data and assuming a lognormal
# distribution. The sample size is n = 8 observations from
# the two compliance wells. Here we will compute the
# multiplier for the log-transformed data.

tolIntNormK(n = 8, ti.type = "upper")
#[1] 3.187294

```

tolIntNpar

Nonparametric Tolerance Interval for a Continuous Distribution

Description

Construct a β -content or β -expectation tolerance interval nonparametrically without making any assumptions about the form of the distribution except that it is continuous.

Usage

```

tolIntNpar(x, coverage, conf.level, cov.type = "content",
  ltl.rank = ifelse(ti.type == "upper", 0, 1),
  n.plus.one.minus.utl.rank = ifelse(ti.type == "lower", 0, 1),
  lb = -Inf, ub = Inf, ti.type = "two-sided")

```

Arguments

<code>x</code>	numeric vector of observations. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
<code>coverage</code>	a scalar between 0 and 1 indicating the desired coverage of the β -content tolerance interval. The default value is <code>coverage=0.95</code> . If <code>cov.type="content"</code> , you must supply a value for <code>coverage</code> or a value for <code>conf.level</code> , but not both. If <code>cov.type="expectation"</code> , this argument is ignored.
<code>conf.level</code>	a scalar between 0 and 1 indicating the confidence level associated with the β -content tolerance interval. The default value is <code>conf.level=0.95</code> . If <code>cov.type="content"</code> , you must supply a value for <code>coverage</code> or a value for <code>conf.level</code> , but not both. If <code>cov.type="expectation"</code> , this argument is ignored.
<code>cov.type</code>	character string specifying the coverage type for the tolerance interval. The possible values are "content" (β -content; the default), and "expectation" (β -expectation). See the DETAILS section for more information.
<code>ltl.rank</code>	positive integer indicating the rank of the order statistic to use for the lower bound of the tolerance interval. If <code>ti.type="two-sided"</code> or <code>ti.type="lower"</code> , the default value is <code>ltl.rank=1</code> (implying the minimum value of <code>x</code> is used as the lower bound of the tolerance interval). If <code>ti.type="upper"</code> , this argument is set equal to 0 and the value of <code>lb</code> is used as the lower bound of the tolerance interval.
<code>n.plus.one.minus.utl.rank</code>	positive integer related to the rank of the order statistic to use for the upper bound of the tolerance interval. Set $v = n.plus.one.minus.utl.rank$, let w denote the rank of the order statistic used as the upper bound (<code>utl.rank</code>), and let n denote the sample size. The relationship between v and w is given by: $v = n + 1 - w$. If <code>ti.type="two-sided"</code> or <code>ti.type="upper"</code> , the default value is <code>n.plus.one.minus.utl.rank=1</code> (so $v = n$, implying the maximum value of <code>x</code> is used as the upper bound of the tolerance interval). If <code>ti.type="lower"</code> , this argument is set equal to 0 and the value of <code>ub</code> is used as the upper bound of the tolerance interval.
<code>lb, ub</code>	scalars indicating lower and upper bounds on the distribution. By default, <code>lb=-Inf</code> and <code>ub=Inf</code> . If you are constructing a tolerance interval for a distribution that you know has a lower bound other than <code>-Inf</code> (e.g., 0), set <code>lb</code> to this value. Similarly, if you know the distribution has an upper bound other than <code>Inf</code> , set <code>ub</code> to this value. The argument <code>lb</code> is ignored if <code>ti.type="two-sided"</code> or <code>ti.type="lower"</code> . The argument <code>ub</code> is ignored if <code>ti.type="two-sided"</code> or <code>ti.type="upper"</code> .
<code>ti.type</code>	character string indicating what kind of tolerance interval to compute. The possible values are "two-sided" (the default), "lower", and "upper".

Details

A tolerance interval for some population is an interval on the real line constructed so as to contain $100\beta\%$ of the population (i.e., $100\beta\%$ of all future observations), where $0 < \beta < 1$. The quantity $100\beta\%$ is called the *coverage*.

There are two kinds of tolerance intervals (Guttman, 1970):

- A β -content tolerance interval with confidence level $100(1 - \alpha)\%$ is constructed so that it contains at least $100\beta\%$ of the population (i.e., the coverage is at least $100\beta\%$) with probability $100(1 - \alpha)\%$, where $0 < \alpha < 1$. The quantity $100(1 - \alpha)\%$ is called the confidence level or confidence coefficient associated with the tolerance interval.

- A β -expectation tolerance interval is constructed so that the *average* coverage of the interval is $100\beta\%$.

Note: A β -expectation tolerance interval with coverage $100\beta\%$ is equivalent to a prediction interval for one future observation with associated confidence level $100\beta\%$. Note that there is no explicit confidence level associated with a β -expectation tolerance interval. If a β -expectation tolerance interval is treated as a β -content tolerance interval, the confidence level associated with this tolerance interval is usually around 50% (e.g., Guttman, 1970, Table 4.2, p.76).

The Form of a Nonparametric Tolerance Interval

Let \underline{x} denote a random sample of n independent observations from some continuous distribution and let $x_{(i)}$ denote the i 'th order statistic in \underline{x} . A two-sided nonparametric tolerance interval is constructed as:

$$[x_{(u)}, x_{(v)}] \quad (1)$$

where u and v are positive integers between 1 and n , and $u < v$. That is, u denotes the rank of the lower tolerance limit, and v denotes the rank of the upper tolerance limit. To make it easier to write some equations later on, we can also write the tolerance interval (1) in a slightly different way as:

$$[x_{(u)}, x_{(n+1-w)}] \quad (2)$$

where

$$w = n + 1 - v \quad (3)$$

so that w is a positive integer between 1 and $n - 1$, and $u < n + 1 - w$. In terms of the arguments to the function `tolIntNpar`, the argument `ltl.rank` corresponds to u , and the argument `n.plus.one.minus.utl.rank` corresponds to w .

If we allow $u = 0$ and $w = 0$ and define lower and upper bounds as:

$$x_{(0)} = lb \quad (4)$$

$$x_{(n+1)} = ub \quad (5)$$

then equation (2) above can also represent a one-sided lower or one-sided upper tolerance interval as well. That is, a one-sided lower nonparametric tolerance interval is constructed as:

$$[x_{(u)}, x_{(n+1)}] = [x_{(u)}, ub] \quad (6)$$

and a one-sided upper nonparametric tolerance interval is constructed as:

$$[x_{(0)}, x_{(v)}] = [lb, x_{(v)}] \quad (7)$$

Usually, $lb = -\infty$ or $lb = 0$ and $ub = \infty$.

Let C be a random variable denoting the coverage of the above nonparametric tolerance intervals. Wilks (1941) showed that the distribution of C follows a [beta distribution](#) with parameters $\text{shape1} = v - u$ and $\text{shape2} = w + u$ when the unknown distribution is continuous.

Computations for a β -Content Tolerance Interval

For a β -content tolerance interval, if the coverage $C = \beta$ is specified, then the associated confidence level $(1 - \alpha)100\%$ is computed as:

$$1 - \alpha = 1 - F(\beta, v - u, w + u) \quad (8)$$

where $F(y, \delta, \gamma)$ denotes the cumulative distribution function of a [beta random variable](#) with parameters $\text{shape1} = \delta$ and $\text{shape2} = \gamma$ evaluated at y .

Similarly, if the confidence level associated with the tolerance interval is specified as $(1 - \alpha)100\%$, then the coverage $C = \beta$ is computed as:

$$\beta = B(\alpha, v - u, w + u) \quad (9)$$

where $B(p, \delta, \gamma)$ denotes the p 'th quantile of a [beta distribution](#) with parameters $\text{shape1}=\delta$ and $\text{shape2}=\gamma$.

Computations for a β -Expectation Tolerance Interval

For a β -expectation tolerance interval, the expected coverage is simply the mean of a [beta random variable](#) with parameters $\text{shape1}=v - u$ and $\text{shape2}=w + u$, which is given by:

$$E(C) = \frac{v - u}{n + 1} \quad (10)$$

As stated above, a β -expectation tolerance interval with coverage $\beta 100\%$ is equivalent to a prediction interval for one future observation with associated confidence level $\beta 100\%$. This is because the probability that any single future observation will fall into this interval is $\beta 100\%$, so the distribution of the number of N future observations that will fall into this interval is [binomial](#) with parameters $\text{size}=N$ and $\text{prob}=\beta$. Hence the expected proportion of future observations that fall into this interval is $\beta 100\%$ and is independent of the value of N . See the help file for [predIntNpar](#) for more information on constructing a nonparametric prediction interval.

Value

A list of class "estimate" containing the estimated parameters, the tolerance interval, and other information. See [estimate.object](#) for details.

Note

Tolerance intervals have long been applied to quality control and life testing problems (Hahn, 1970b,c; Hahn and Meeker, 1991; Krishnamoorthy and Mathew, 2009). References that discuss tolerance intervals in the context of environmental monitoring include: Berthouex and Brown (2002, Chapter 21), Gibbons et al. (2009), Millard and Neerchal (2001, Chapter 6), Singh et al. (2010b), and USEPA (2009).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Conover, W.J. (1980). *Practical Nonparametric Statistics*. Second Edition. John Wiley and Sons, New York.
- Danziger, L., and S. Davis. (1964). Tables of Distribution-Free Tolerance Limits. *Annals of Mathematical Statistics* **35**(5), 1361–1365.
- Davis, C.B. (1994). Environmental Regulatory Statistics. In Patil, G.P., and C.R. Rao, eds., *Handbook of Statistics, Vol. 12: Environmental Statistics*. North-Holland, Amsterdam, a division of Elsevier, New York, NY, Chapter 26, 817–865.
- Davis, C.B., and R.J. McNichols. (1994a). Ground Water Monitoring Statistics Update: Part I: Progress Since 1988. *Ground Water Monitoring and Remediation* **14**(4), 148–158.
- Gibbons, R.D. (1991b). Statistical Tolerance Limits for Ground-Water Monitoring. *Ground Water* **29**, 563–570.

- Gibbons, R.D., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*, Second Edition. John Wiley & Sons, Hoboken.
- Guttman, I. (1970). *Statistical Tolerance Regions: Classical and Bayesian*. Hafner Publishing Co., Darien, CT, Chapter 2.
- Hahn, G.J., and W.Q. Meeker. (1991). *Statistical Intervals: A Guide for Practitioners*. John Wiley and Sons, New York, 392pp.
- Helsel, D.R., and R.M. Hirsch. (1992). *Statistical Methods in Water Resources Research*. Elsevier, New York, NY, pp.88-90.
- Krishnamoorthy K., and T. Mathew. (2009). *Statistical Tolerance Regions: Theory, Applications, and Computation*. John Wiley and Sons, Hoboken.
- Millard, S.P., and N.K. Neerchal. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.
- Wilks, S.S. (1941). Determination of Sample Sizes for Setting Tolerance Limits. *Annals of Mathematical Statistics* **12**, 91–96.

See Also

[eqnpar](#), [estimate.object](#), [tolIntNparN](#), [Tolerance Intervals](#), [Estimating Distribution Parameters](#), [Estimating Distribution Quantiles](#).

Examples

```
# Generate 20 observations from a \link[=LognormalMixAlt]{lognormal mixture distribution}
# with parameters mean1=1, cv1=0.5, mean2=5, cv2=1, and p.mix=0.1.
# The exact two-sided interval that contains 90% of this distribution is given by:
# [0.682312, 13.32052]. Use tolIntNpar to construct a two-sided 90%
# \eqn{\beta}-content tolerance interval. Note that the associated confidence level
# is only 61%. A larger sample size is required to obtain a larger confidence
# level (see the help file for \link{tolIntNparN}).
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(23)
dat <- rlnormMixAlt(20, 1, 0.5, 5, 1, 0.1)
tolIntNpar(dat, coverage = 0.9)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:      None
#
#Data:                     dat
#
#Sample Size:              20
#
#Tolerance Interval Coverage: 90%
#
#Coverage Type:            content
#
#Tolerance Interval Method: Exact
```

```

#
#Tolerance Interval Type:      two-sided
#
#Confidence Level:           60.8253%
#
#Tolerance Limit Rank(s):     1 20
#
#Tolerance Interval:          LTL = 0.5035035
#                             UTL = 9.9504662

#-----

# Clean up
rm(dat)

#-----

# Reproduce Example 17-4 on page 17-21 of USEPA (2009). This example uses
# copper concentrations (ppb) from 3 background wells to set an upper
# limit for 2 compliance wells. The maximum value from the 3 wells is set
# to the 95% confidence upper tolerance limit, and we need to determine the
# coverage of this tolerance interval. The data are stored in EPA.92c.copper2.df.
# Note that even though these data are Type I left singly censored, it is still
# possible to compute an upper tolerance interval using any of the uncensored
# observations as the upper limit.

EPA.92c.copper2.df
#   Copper.orig Copper Censored Month Well Well.type
#1          <5    5.0     TRUE     1    1 Background
#2          <5    5.0     TRUE     2    1 Background
#3          7.5    7.5    FALSE     3    1 Background
#...
#9          9.2    9.2    FALSE     1    2 Background
#10         <5    5.0     TRUE     2    2 Background
#11         <5    5.0     TRUE     3    2 Background
#...
#17         <5    5.0     TRUE     1    3 Background
#18         5.4    5.4    FALSE     2    3 Background
#19         6.7    6.7    FALSE     3    3 Background
#...
#29         6.2    6.2    FALSE     5    4 Compliance
#30         <5    5.0     TRUE     6    4 Compliance
#31         7.8    7.8    FALSE     7    4 Compliance
#...
#38         <5    5.0     TRUE     6    5 Compliance
#39         5.6    5.6    FALSE     7    5 Compliance
#40         <5    5.0     TRUE     8    5 Compliance

with(EPA.92c.copper2.df,
     tolIntNpar(Copper[Well.type=="Background"],
               conf.level = 0.95, lb = 0, ti.type = "upper"))

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:      None
#

```

```

#Data:                                Copper[Well.type == "Background"]
#
#Sample Size:                          24
#
#Tolerance Interval Coverage:          88.26538%
#
#Coverage Type:                        content
#
#Tolerance Interval Method:            Exact
#
#Tolerance Interval Type:              upper
#
#Confidence Level:                     95%
#
#Tolerance Limit Rank(s):              24
#
#Tolerance Interval:                   LTL = 0.0
#                                       UTL = 9.2

#-----

# Repeat the last example, except compute an upper
# \eqn{\beta}-expectation tolerance interval:

with(EPA.92c.copper2.df,
      tolIntNpar(Copper[Well.type=="Background"],
                  cov.type = "expectation", lb = 0, ti.type = "upper"))

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:                  None
#
#Data:                                Copper[Well.type == "Background"]
#
#Sample Size:                          24
#
#Tolerance Interval Coverage:          96%
#
#Coverage Type:                        expectation
#
#Tolerance Interval Method:            Exact
#
#Tolerance Interval Type:              upper
#
#Tolerance Limit Rank(s):              24
#
#Tolerance Interval:                   LTL = 0.0
#                                       UTL = 9.2

```

tolIntPois

Tolerance Interval for a Poisson Distribution

Description

Construct a β -content or β -expectation tolerance interval for a [Poisson distribution](#).

Usage

```
tolIntPois(x, coverage = 0.95, cov.type = "content", ti.type = "two-sided",
  conf.level = 0.95)
```

Arguments

<code>x</code>	numeric vector of observations, or an object resulting from a call to an estimating function that assumes a Poisson distribution (i.e., <code>epois</code> or <code>epoisCensored</code>). If <code>cov.type="content"</code> then <code>x</code> must be a numeric vector. If <code>x</code> is a numeric vector, missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
<code>coverage</code>	a scalar between 0 and 1 indicating the desired coverage of the tolerance interval. The default value is <code>coverage=0.95</code> . If <code>cov.type="expectation"</code> , this argument is ignored.
<code>cov.type</code>	character string specifying the coverage type for the tolerance interval. The possible values are "content" (β -content; the default), and "expectation" (β -expectation). See the DETAILS section for more information.
<code>ti.type</code>	character string indicating what kind of tolerance interval to compute. The possible values are "two-sided" (the default), "lower", and "upper".
<code>conf.level</code>	a scalar between 0 and 1 indicating the confidence level associated with the tolerance interval. The default value is <code>conf.level=0.95</code> .

Details

If `x` contains any missing (NA), undefined (NaN) or infinite (Inf, -Inf) values, they will be removed prior to performing the estimation.

A tolerance interval for some population is an interval on the real line constructed so as to contain $100\beta\%$ of the population (i.e., $100\beta\%$ of all future observations), where $0 < \beta < 1$. The quantity $100\beta\%$ is called the *coverage*.

There are two kinds of tolerance intervals (Guttman, 1970):

- A β -content tolerance interval with confidence level $100(1 - \alpha)\%$ is constructed so that it contains at least $100\beta\%$ of the population (i.e., the coverage is at least $100\beta\%$) with probability $100(1 - \alpha)\%$, where $0 < \alpha < 1$. The quantity $100(1 - \alpha)\%$ is called the confidence level or confidence coefficient associated with the tolerance interval.
- A β -expectation tolerance interval is constructed so that the *average* coverage of the interval is $100\beta\%$.

Note: A β -expectation tolerance interval with coverage $100\beta\%$ is equivalent to a prediction interval for one future observation with associated confidence level $100\beta\%$. Note that there is no explicit confidence level associated with a β -expectation tolerance interval. If a β -expectation tolerance interval is treated as a β -content tolerance interval, the confidence level associated with this tolerance interval is usually around 50% (e.g., Guttman, 1970, Table 4.2, p.76).

Because of the discrete nature of the [Poisson distribution](#), even true tolerance intervals (tolerance intervals based on the true value of λ) will usually not contain exactly $\beta\%$ of the population. For example, for the Poisson distribution with parameter `lambda=2`, the interval `[0, 4]` contains 94.7% of this distribution and the interval `[0, 5]` contains 98.3% of this distribution. Thus, no interval can contain exactly 95% of this distribution.

β -Content Tolerance Intervals for a Poisson Distribution

Zacks (1970) showed that for monotone likelihood ratio (MLR) families of discrete distributions, a

uniformly most accurate upper $\beta 100\%$ β -content tolerance interval with associated confidence level $(1 - \alpha)100\%$ is constructed by finding the upper $(1 - \alpha)100\%$ confidence limit for the parameter associated with the distribution, and then computing the β 'th quantile of the distribution assuming the true value of the parameter is equal to the upper confidence limit. This idea can be extended to one-sided lower and two-sided tolerance limits.

It can be shown that all distributions that are one parameter exponential families have the MLR property, and the Poisson distribution is a one-parameter exponential family, so the method of Zacks (1970) can be applied to a Poisson distribution.

Let X denote a [Poisson random variable](#) with parameter λ . Let $x_{p|\lambda}$ denote the p 'th quantile of this distribution. That is,

$$Pr(X < x_{p|\lambda}) \leq p \leq Pr(X \leq x_{p|\lambda}) \quad (1)$$

Note that due to the discrete nature of the Poisson distribution, there will be several values of p associated with one value of X . For example, for $\lambda = 2$, the value 1 is the p 'th quantile for any value of p between 0.140 and 0.406.

Let \underline{x} denote a vector of n observations from a [Poisson distribution](#) with parameter λ . When `ti.type="upper"`, the first step is to compute the one-sided upper $(1 - \alpha)100\%$ confidence limit for λ based on the observations \underline{x} (see the help file for [epois](#)). Denote this upper confidence limit by UCL . The one-sided upper $\beta 100\%$ tolerance limit is then given by:

$$[0, x_{\beta|\lambda=UCL}] \quad (2)$$

Similarly, when `ti.type="lower"`, the first step is to compute the one-sided lower $(1 - \alpha)100\%$ confidence limit for λ based on the observations \underline{x} . Denote this lower confidence limit by LCL . The one-sided lower $\beta 100\%$ tolerance limit is then given by:

$$[x_{1-\beta|\lambda=LCL}, \infty] \quad (3)$$

Finally, when `ti.type="two-sided"`, the first step is to compute the two-sided $(1 - \alpha)100\%$ confidence limits for λ based on the observations \underline{x} . Denote these confidence limits by LCL and UCL . The two-sided $\beta 100\%$ tolerance limit is then given by:

$$[x_{\frac{1-\beta}{2}|\lambda=LCL}, x_{\frac{1+\beta}{2}|\lambda=UCL}] \quad (4)$$

Note that the function `tolIntPois` uses the exact confidence limits for λ when computing β -content tolerance limits (see [epois](#)).

β -Expectation Tolerance Intervals for a Poisson Distribution

As stated above, a β -expectation tolerance interval with coverage $\beta 100\%$ is equivalent to a prediction interval for one future observation with associated confidence level $\beta 100\%$. This is because the probability that any single future observation will fall into this interval is $\beta 100\%$, so the distribution of the number of N future observations that will fall into this interval is [binomial](#) with parameters `size=N` and `prob= β` . Hence the expected proportion of future observations that fall into this interval is $\beta 100\%$ and is independent of the value of N . See the help file for [predIntPois](#) for information on how these intervals are constructed.

Value

If \mathbf{x} is a numeric vector, `tolIntPois` returns a list of class "estimate" containing the estimated parameters, a component called `interval` containing the tolerance interval information, and other information. See [estimate.object](#) for details.

If \mathbf{x} is the result of calling an estimation function, `tolIntPois` returns a list whose class is the same as \mathbf{x} . The list contains the same components as \mathbf{x} . If \mathbf{x} already has a component called `interval`, this component is replaced with the tolerance interval information.

Note

Tolerance intervals have long been applied to quality control and life testing problems (Hahn, 1970b,c; Hahn and Meeker, 1991; Krishnamoorthy and Mathew, 2009). References that discuss tolerance intervals in the context of environmental monitoring include: Berthouex and Brown (2002, Chapter 21), Gibbons et al. (2009), Millard and Neerchal (2001, Chapter 6), Singh et al. (2010b), and USEPA (2009).

Gibbons (1987b) used the Poisson distribution to model the number of detected compounds per scan of the 32 volatile organic priority pollutants (VOC), and also to model the distribution of chemical concentration (in ppb). He explained the derivation of a one-sided upper β -content tolerance limit for a Poisson distribution based on the work of Zacks (1970) using the Pearson-Hartley approximation to the confidence limits for the mean parameter λ (see the help file for [epois](#)). Note that there are several typographical errors in the derivation and examples on page 575 of Gibbons (1987b) because there is confusion between where the value of β (the coverage) should be and where the value of $1 - \alpha$ (the confidence level) should be. Gibbons et al. (2009, pp.103-104) gives correct formulas.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Gibbons, R.D. (1987b). Statistical Models for the Analysis of Volatile Organic Compounds in Waste Disposal Sites. *Ground Water* **25**, 572–580.
- Gibbons, R.D., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*, Second Edition. John Wiley & Sons, Hoboken.
- Guttman, I. (1970). *Statistical Tolerance Regions: Classical and Bayesian*. Hafner Publishing Co., Darien, CT.
- Hahn, G.J., and W.Q. Meeker. (1991). *Statistical Intervals: A Guide for Practitioners*. John Wiley and Sons, New York.
- Johnson, N. L., S. Kotz, and A. Kemp. (1992). *Univariate Discrete Distributions*. Second Edition. John Wiley and Sons, New York, Chapter 4.
- Krishnamoorthy K., and T. Mathew. (2009). *Statistical Tolerance Regions: Theory, Applications, and Computation*. John Wiley and Sons, Hoboken.
- Millard, S.P., and N.K. Neerchal. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton.
- Zacks, S. (1970). Uniformly Most Accurate Upper Tolerance Limits for Monotone Likelihood Ratio Families of Discrete Distributions. *Journal of the American Statistical Association* **65**, 307–316.

See Also

[Poisson](#), [epois](#), [eqpois](#), [estimate.object](#), [Tolerance Intervals](#), [Estimating Distribution Parameters](#), [Estimating Distribution Quantiles](#).

Examples

```
# Generate 20 observations from a Poisson distribution with parameter
# lambda=2. The interval [0, 4] contains 94.7% of this distribution and
# the interval [0,5] contains 98.3% of this distribution. Thus, because
# of the discrete nature of the Poisson distribution, no interval contains
# exactly 95% of this distribution. Use tolIntPois to estimate the mean
```

```

# parameter of the true distribution, and construct a one-sided upper 95%
# beta-content tolerance interval with associated confidence level 90%.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(250)
dat <- rpois(20, 2)
tolIntPois(dat, conf.level = 0.9)

#Results of Distribution Parameter Estimation
#-----
#
#Assumed Distribution:          Poisson
#
#Estimated Parameter(s):      lambda = 1.8
#
#Estimation Method:           mle/mme/mvue
#
#Data:                         dat
#
#Sample Size:                  20
#
#Tolerance Interval Coverage:  95%
#
#Coverage Type:                content
#
#Tolerance Interval Method:    Zacks
#
#Tolerance Interval Type:      two-sided
#
#Confidence Level:             90%
#
#Tolerance Interval:           LTL = 0
#                               UTL = 6
#
#-----

# Clean up
rm(dat)

```

Total.P.df

*Total Phosphorus Data from Chesapeake Bay***Description**

Monthly estimated total phosphorus mass (mg) within a water column at two different stations for the 5-year time period October 1984 to September 1989 from a study on phosphorus concentration conducted in the Chesapeake Bay.

Usage

Total.P.df

Format

A data frame with 60 observations on the following 4 variables.

CB3.1 a numeric vector of phosphorus concentrations at station CB3.1

CB3.3e a numeric vector phosphorus concentrations at station CB3.3e

Month a factor indicating the month the observation was taken

Year a numeric vector indicating the year an observation was taken

Source

Neerchal, N. K., and S. L. Brunenmeister. (1993). Estimation of Trend in Chesapeake Bay Water Quality Data. In Patil, G.P., and C.R. Rao, eds., *Handbook of Statistics, Vol. 6: Multivariate Environmental Statistics*. North-Holland, Amsterdam, Chapter 19, 407-422.

Triangular	<i>The Triangular Distribution</i>
------------	------------------------------------

Description

Density, distribution function, quantile function, and random generation for the triangular distribution with parameters min, max, and mode.

Usage

```
dtri(x, min = 0, max = 1, mode = 1/2)
ptri(q, min = 0, max = 1, mode = 1/2)
qtri(p, min = 0, max = 1, mode = 1/2)
rtri(n, min = 0, max = 1, mode = 1/2)
```

Arguments

- x vector of quantiles. Missing values (NAs) are allowed.
- q vector of quantiles. Missing values (NAs) are allowed.
- p vector of probabilities between 0 and 1. Missing values (NAs) are allowed.
- n sample size. If length(n) is larger than 1, then length(n) random values are returned.
- min vector of minimum values of the distribution of the random variable. The default value is min=0.
- max vector of maximum values of the random variable. The default value is max=1.
- mode vector of modes of the random variable. The default value is mode=1/2.

Details

Let X be a triangular random variable with parameters $\min=a$, $\max=b$, and $\text{mode}=c$.

Probability Density and Cumulative Distribution Function

The density function of X is given by:

$$f(x; a, b, c) = \begin{cases} \frac{2(x-a)}{(b-a)(c-a)} & \text{for } a \leq x \leq c \\ \frac{2(b-x)}{(b-a)(b-c)} & \text{for } c \leq x \leq b \end{cases}$$

where $a < c < b$.

The cumulative distribution function of X is given by:

$$F(x; a, b, c) = \begin{cases} \frac{(x-a)^2}{(b-a)(c-a)} & \text{for } a \leq x \leq c \\ 1 - \frac{(b-x)^2}{(b-a)(b-c)} & \text{for } c \leq x \leq b \end{cases}$$

where $a < c < b$.

Quantiles

The p^{th} quantile of X is given by:

$$x_p = \begin{cases} a + \sqrt{(b-a)(c-a)p} & \text{for } 0 \leq p \leq F(c) \\ b - \sqrt{(b-a)(b-c)(1-p)} & \text{for } F(c) \leq p \leq 1 \end{cases}$$

where $0 \leq p \leq 1$.

Random Numbers

Random numbers are generated using the inverse transformation method:

$$x = F^{-1}(u)$$

where u is a random deviate from a uniform $[0, 1]$ distribution.

Mean and Variance

The mean and variance of X are given by:

$$E(X) = \frac{a + b + c}{3}$$

$$Var(X) = \frac{a^2 + b^2 + c^2 - ab - ac - bc}{18}$$

Value

`dtri` gives the density, `ptri` gives the distribution function, `qtri` gives the quantile function, and `rtri` generates random deviates.

Note

The triangular distribution is so named because of the shape of its probability density function. The average of two independent identically distributed uniform random variables with parameters $\min=\alpha$ and $\max=\beta$ has a triangular distribution with parameters $\min=\alpha$, $\max=\beta$, and $\text{mode}=(\beta-\alpha)/2$.

The triangular distribution is sometimes used as an input distribution in probability risk assessment.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Forbes, C., M. Evans, N. Hastings, and B. Peacock. (2011). *Statistical Distributions*. Fourth Edition. John Wiley and Sons, Hoboken, NJ.
- Johnson, N. L., S. Kotz, and N. Balakrishnan. (1995). *Continuous Univariate Distributions, Volume 2*. Second Edition. John Wiley and Sons, New York.

See Also

[Uniform, Probability Distributions and Random Numbers.](#)

Examples

```
# Density of a triangular distribution with parameters
# min=10, max=15, and mode=12, evaluated at 12, 13 and 14:

dtri(12:14, 10, 15, 12)
#[1] 0.4000000 0.2666667 0.1333333

#-----

# The cdf of a triangular distribution with parameters
# min=2, max=7, and mode=5, evaluated at 3, 4, and 5:

ptri(3:5, 2, 7, 5)
#[1] 0.06666667 0.26666667 0.60000000

#-----

# The 25'th percentile of a triangular distribution with parameters
# min=1, max=4, and mode=3:

qtri(0.25, 1, 4, 3)
#[1] 2.224745

#-----

# A random sample of 4 numbers from a triangular distribution with
# parameters min=3 , max=20, and mode=12.
# (Note: the call to set.seed simply allows you to reproduce this example.)

set.seed(10)
rtri(4, 3, 20, 12)
#[1] 11.811593  9.850955 11.081885 13.539496
```

varGroupTest

Test for Homogeneity of Variance Among Two or More Groups

Description

Test the null hypothesis that the variances of two or more normal distributions are the same using Levene's or Bartlett's test.

Usage

```
varGroupTest(object, ...)

## S3 method for class 'formula'
varGroupTest(object, data = NULL, subset,
  na.action = na.pass, ...)
```

```
## Default S3 method:
varGroupTest(object, group, test = "Levene",
  correct = TRUE, data.name = NULL, group.name = NULL,
  parent.of.data = NULL, subset.expression = NULL, ...)

## S3 method for class 'data.frame'
varGroupTest(object, ...)

## S3 method for class 'matrix'
varGroupTest(object, ...)

## S3 method for class 'list'
varGroupTest(object, ...)
```

Arguments

object	an object containing data for 2 or more groups whose variances are to be compared. In the default method, the argument <code>object</code> must be a numeric vector. When <code>object</code> is a data frame, all columns must be numeric. When <code>object</code> is a matrix, it must be a numeric matrix. When <code>object</code> is a list, all components must be numeric vectors. In the formula method, a symbolic specification of the form <code>y ~ g</code> can be given, indicating the observations in the vector <code>y</code> are to be grouped according to the levels of the factor <code>g</code> . Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
data	when <code>object</code> is a formula, <code>data</code> specifies an optional data frame, list or environment (or object coercible by <code>as.data.frame</code> to a data frame) containing the variables in the model. If not found in <code>data</code> , the variables are taken from <code>environment(formula)</code> , typically the environment from which <code>summaryStats</code> is called.
subset	when <code>object</code> is a formula, <code>subset</code> specifies an optional vector specifying a subset of observations to be used.
na.action	when <code>object</code> is a formula, <code>na.action</code> specifies a function which indicates what should happen when the data contain NAs. The default is na.pass .
group	when <code>object</code> is a numeric vector, <code>group</code> is a factor or character vector indicating which group each observation belongs to. When <code>object</code> is a matrix or data frame this argument is ignored and the columns define the groups. When <code>object</code> is a list this argument is ignored and the components define the groups. When <code>object</code> is a formula, this argument is ignored and the right-hand side of the formula specifies the grouping variable.
test	character string indicating which test to use. The possible values are "Levene" (Levene's test; the default) and "Bartlett" (Bartlett's test).
correct	logical scalar indicating whether to use the correction factor for Bartlett's test. The default value is <code>correct=TRUE</code> . This argument is ignored if <code>test="Levene"</code> .
data.name	character string indicating the name of the data used for the group variance test. The default value is <code>data.name=deparse(substitute(object))</code> .
group.name	character string indicating the name of the data used to create the groups. The default value is <code>group.name=deparse(substitute(group))</code> .
parent.of.data	character string indicating the source of the data used for the group variance test.
subset.expression	character string indicating the expression used to subset the data.
...	additional arguments affecting the group variance test.

Details

The function `varGroupTest` performs Levene's or Bartlett's test for homogeneity of variance among two or more groups. The R function `var.test` compares two variances.

Bartlett's test is very sensitive to the assumption of normality and will tend to give significant results even when the null hypothesis is true if the underlying distributions have long tails (e.g., are leptokurtic). Levene's test is almost as powerful as Bartlett's test when the underlying distributions are normal, and unlike Bartlett's test it tends to maintain the assumed *alpha*-level when the underlying distributions are not normal (Snedecor and Cochran, 1989, p.252; Milliken and Johnson, 1992, p.22; Conover et al., 1981). Thus, Levene's test is generally recommended over Bartlett's test.

Value

a list of class "htest" containing the results of the group variance test. Objects of class "htest" have special printing and plotting methods. See the help file for `htest.object` for details.

Note

Chapter 11 of USEPA (2009) discusses using Levene's test to test the assumption of equal variances between monitoring wells or to test that the variance is stable over time when performing intrawell tests.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Conover, W.J., M.E. Johnson, and M.M. Johnson. (1981). A Comparative Study of Tests for Homogeneity of Variances, with Applications to the Outer Continental Shelf Bidding Data. *Technometrics* **23**(4), 351-361.
- Davis, C.B. (1994). Environmental Regulatory Statistics. In Patil, G.P., and C.R. Rao, eds., *Handbook of Statistics, Vol. 12: Environmental Statistics*. North-Holland, Amsterdam, a division of Elsevier, New York, NY, Chapter 26, 817-865.
- Milliken, G.A., and D.E. Johnson. (1992). *Analysis of Messy Data, Volume I: Designed Experiments*. Chapman & Hall, New York.
- Snedecor, G.W., and W.G. Cochran. (1989). *Statistical Methods, Eighth Edition*. Iowa State University Press, Ames Iowa.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.
- USEPA. (2010). *Errata Sheet - March 2009 Unified Guidance*. EPA 530/R-09-007a, August 9, 2010. Office of Resource Conservation and Recovery, Program Information and Implementation Division. U.S. Environmental Protection Agency, Washington, D.C.
- Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ.

See Also

`var.test`, `varTest`.

Examples

```
# Example 11-2 of USEPA (2009, page 11-7) gives an example of
# testing the assumption of equal variances across wells for arsenic
# concentrations (ppb) in groundwater collected at 6 monitoring
# wells over 4 months. The data for this example are stored in
# EPA.09.Ex.11.1.arsenic.df.
```

```
head(EPA.09.Ex.11.1.arsenic.df)
# Arsenic.ppb Month Well
#1      22.9      1      1
#2       3.1      2      1
#3      35.7      3      1
#4       4.2      4      1
#5       2.0      1      2
#6       1.2      2      2
```

```
longToWide(EPA.09.Ex.11.1.arsenic.df, "Arsenic.ppb", "Month", "Well",
  paste.row.name = TRUE, paste.col.name = TRUE)
#      Well.1 Well.2 Well.3 Well.4 Well.5 Well.6
#Month.1  22.9   2.0   2.0   7.8  24.9   0.3
#Month.2   3.1   1.2 109.4   9.3   1.3   4.8
#Month.3  35.7   7.8   4.5  25.9   0.8   2.8
#Month.4   4.2  52.0   2.5   2.0  27.0   1.2
```

```
varGroupTest(Arsenic.ppb ~ Well, data = EPA.09.Ex.11.1.arsenic.df)
```

```
#Results of Hypothesis Test
#-----
#
#Null Hypothesis:          Ratio of each pair of variances = 1
#
#Alternative Hypothesis:   At least one variance differs
#
#Test Name:               Levene's Test for
#                          Homogeneity of Variance
#
#Estimated Parameter(s):  Well.1 = 246.8158
#                          Well.2 = 592.6767
#                          Well.3 = 2831.4067
#                          Well.4 = 105.2967
#                          Well.5 = 207.4467
#                          Well.6 = 3.9025
#
#Data:                    Arsenic.ppb
#
#Grouping Variable:       Well
#
#Data Source:             EPA.09.Ex.11.1.arsenic.df
#
#Sample Sizes:            Well.1 = 4
#                          Well.2 = 4
#                          Well.3 = 4
#                          Well.4 = 4
#                          Well.5 = 4
#                          Well.6 = 4
#
```

```
#Test Statistic:          F = 4.564176
#
#Test Statistic Parameters:  num df   = 5
#                           denom df = 18
#
#P-value:                  0.007294084
```

varTest

One-Sample Chi-Squared Test on Variance

Description

Estimate the variance, test the null hypothesis using the chi-squared test that the variance is equal to a user-specified value, and create a confidence interval for the variance.

Usage

```
varTest(x, alternative = "two.sided", conf.level = 0.95,
        sigma.squared = 1, data.name = NULL)
```

Arguments

x	numeric vector of observations. Missing (NA), undefined (NaN), and infinite (Inf, -Inf) values are allowed but will be removed.
alternative	character string indicating the kind of alternative hypothesis. The possible values are "two.sided" (the default), "greater", and "less".
conf.level	numeric scalar between 0 and 1 indicating the confidence level associated with the confidence interval for the population variance. The default value is conf.level=0.95.
sigma.squared	a numeric scalar indicating the hypothesized value of the variance. The default value is sigma.squared=1.
data.name	character string indicating the name of the data used for the test of variance.

Details

The function `varTest` performs the one-sample chi-squared test of the hypothesis that the population variance is equal to the user specified value given by the argument `sigma.squared`, and it also returns a confidence interval for the population variance. The R function `var.test` performs the F-test for comparing two variances.

Value

A list of class "htest" containing the results of the hypothesis test. See the help file for `htest.object` for details.

Note

Just as you can perform tests of hypothesis on measures of location (mean, median, percentile, etc.), you can do the same thing for measures of spread or variability. Usually, we are interested in estimating variability only because we want to quantify the uncertainty of our estimated location or percentile. Sometimes, however, we are interested in estimating variability and quantifying the uncertainty in our estimate of variability (for example, for performing a sensitivity analysis for power or sample size calculations), or testing whether the population variability is equal to a certain value. There are at least two possible methods of performing a one-sample hypothesis test on variability:

- Perform a hypothesis test for the population variance based on the chi-squared statistic, assuming the underlying population is normal.
- Perform a hypothesis test for any kind of measure of spread assuming any kind of underlying distribution based on a bootstrap confidence interval (using, for example, the package **boot**).

You can use `varTest` for the first method.

Note: For a one-sample test of location, Student's t-test is fairly robust to departures from normality (i.e., the Type I error rate is maintained), as long as the sample size is reasonably "large." The chi-squared test on the population variance, however, is extremely sensitive to departures from normality. For example, if the underlying population is skewed, the actual Type I error rate will be larger than assumed.

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- van Belle, G., L.D. Fisher, Heagerty, P.J., and Lumley, T. (2004). *Biostatistics: A Methodology for the Health Sciences, 2nd Edition*. John Wiley & Sons, New York.
- Millard, S.P., and N.K. Neerchal. (2001). *Environmental Statistics with S-PLUS*. CRC Press, Boca Raton, FL.
- Zar, J.H. (2010). *Biostatistical Analysis*. Fifth Edition. Prentice-Hall, Upper Saddle River, NJ.

See Also

[var.test](#), [varGroupTest](#).

Examples

```
# Generate 20 observations from a normal distribution with parameters
# mean=2 and sd=1. Test the null hypothesis that the true variance is
# equal to 0.5 against the alternative that the true variance is not
# equal to 0.5.
# (Note: the call to set.seed allows you to reproduce this example).

set.seed(23)
dat <- rnorm(20, mean = 2, sd = 1)
varTest(dat, sigma.squared = 0.5)

#Results of Hypothesis Test
#-----
#
```

```

#Null Hypothesis:          variance = 0.5
#
#Alternative Hypothesis:    True variance is not equal to 0.5
#
#Test Name:                Chi-Squared Test on Variance
#
#Estimated Parameter(s):   variance = 0.753708
#
#Data:                     dat
#
#Test Statistic:           Chi-Squared = 28.64090
#
#Test Statistic Parameter: df = 19
#
#P-value:                  0.1436947
#
#95% Confidence Interval:   LCL = 0.4359037
#                           UCL = 1.6078623

# Note that in this case we would not reject the
# null hypothesis at the 5% or even the 10% level.

# Clean up
rm(dat)

```

ZeroModifiedLognormal *The Zero-Modified Lognormal (Delta) Distribution*

Description

Density, distribution function, quantile function, and random generation for the zero-modified log-normal distribution with parameters meanlog, sdlog, and p.zero.

The zero-modified lognormal (delta) distribution is the mixture of a lognormal distribution with a positive probability mass at 0.

Usage

```

dzmlnorm(x, meanlog = 0, sdlog = 1, p.zero = 0.5)
pzmlnorm(q, meanlog = 0, sdlog = 1, p.zero = 0.5)
qzmlnorm(p, meanlog = 0, sdlog = 1, p.zero = 0.5)
rzmlnorm(n, meanlog = 0, sdlog = 1, p.zero = 0.5)

```

Arguments

x	vector of quantiles.
q	vector of quantiles.
p	vector of probabilities between 0 and 1.
n	sample size. If length(n) is larger than 1, then length(n) random values are returned.
meanlog	vector of means of the normal (Gaussian) part of the distribution on the log scale. The default is meanlog=0.

sdlog	vector of (positive) standard deviations of the normal (Gaussian) part of the distribution on the log scale. The default is sdlog=1.
p.zero	vector of probabilities between 0 and 1 indicating the probability the random variable equals 0. For rzm1norm this must be a single, non-missing number.

Details

The zero-modified lognormal (delta) distribution is the mixture of a lognormal distribution with a positive probability mass at 0. This distribution was introduced without a name by Aitchison (1955), and the name Δ -distribution was coined by Aitchison and Brown (1957, p.95). It is a special case of a “zero-modified” distribution (see Johnson et al., 1992, p. 312).

Let $f(x; \mu, \sigma)$ denote the density of a [lognormal random variable](#) X with parameters meanlog= μ and sdlog= σ . The density function of a zero-modified lognormal (delta) random variable Y with parameters meanlog= μ , sdlog= σ , and p.zero= p , denoted $h(y; \mu, \sigma, p)$, is given by:

$$h(y; \mu, \sigma, p) = \begin{cases} p & \text{for } y = 0 \\ (1 - p)f(y; \mu, \sigma) & \text{for } y > 0 \end{cases}$$

Note that μ is *not* the mean of the zero-modified lognormal distribution on the log scale; it is the mean of the lognormal part of the distribution on the log scale. Similarly, σ is *not* the standard deviation of the zero-modified lognormal distribution on the log scale; it is the standard deviation of the lognormal part of the distribution on the log scale.

Let γ and δ denote the mean and standard deviation of the overall zero-modified lognormal distribution on the log scale. Aitchison (1955) shows that:

$$E[\log(Y)] = \gamma = (1 - p)\mu$$

$$Var[\log(Y)] = \delta^2 = (1 - p)\sigma^2 + p(1 - p)\mu^2$$

Note that when p.zero= $p=0$, the zero-modified lognormal distribution simplifies to the lognormal distribution.

Value

dzm1norm gives the density, pzm1norm gives the distribution function, qzm1norm gives the quantile function, and rzm1norm generates random deviates.

Note

The zero-modified lognormal (delta) distribution is sometimes used to model chemical concentrations for which some observations are reported as “Below Detection Limit” (the nondetects are assumed equal to 0). See, for example, Gilliom and Helsel (1986), Owen and DeRouen (1980), and Gibbons et al. (2009, Chapter 12). USEPA (2009, Chapter 15) recommends this strategy only in specific situations, and Helsel (2012, Chapter 1) strongly discourages this approach to dealing with non-detects.

A variation of the zero-modified lognormal (delta) distribution is the [zero-modified normal distribution](#), in which a normal distribution is mixed with a positive probability mass at 0.

One way to try to assess whether a zero-modified lognormal (delta), zero-modified normal, censored normal, or censored lognormal is the best model for the data is to construct both censored and detects-only probability plots (see [qqPlotCensored](#)).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Aitchison, J. (1955). On the Distribution of a Positive Random Variable Having a Discrete Probability Mass at the Origin. *Journal of the American Statistical Association* **50**, 901-908.
- Aitchison, J., and J.A.C. Brown (1957). *The Lognormal Distribution (with special reference to its uses in economics)*. Cambridge University Press, London. pp.94-99.
- Crow, E.L., and K. Shimizu. (1988). *Lognormal Distributions: Theory and Applications*. Marcel Dekker, New York, pp.47-51.
- Gibbons, RD., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*. Second Edition. John Wiley and Sons, Hoboken, NJ.
- Gilliom, R.J., and D.R. Helsel. (1986). Estimation of Distributional Parameters for Censored Trace Level Water Quality Data: 1. Estimation Techniques. *Water Resources Research* **22**, 135-146.
- Helsel, D.R. (2012). *Statistics for Censored Environmental Data Using Minitab and R*. Second Edition. John Wiley and Sons, Hoboken, NJ, Chapter 1.
- Johnson, N. L., S. Kotz, and A.W. Kemp. (1992). *Univariate Discrete Distributions*. Second Edition. John Wiley and Sons, New York, p.312.
- Owen, W., and T. DeRouen. (1980). Estimation of the Mean for Lognormal Data Containing Zeros and Left-Censored Values, with Applications to the Measurement of Worker Exposure to Air Contaminants. *Biometrics* **36**, 707-719.
- USEPA (1992c). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities: Addendum to Interim Final Guidance*. Office of Solid Waste, Permits and State Programs Division, US Environmental Protection Agency, Washington, D.C.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.

See Also

[Zero-Modified Lognormal \(Alternative Parameterization\)](#), [Lognormal](#), [LognormalAlt](#), [Zero-Modified Normal](#), [ezm1norm](#), [Probability Distributions and Random Numbers](#).

Examples

```
# Density of the zero-modified lognormal (delta) distribution with
# parameters meanlog=0, sdlog=1, and p.zero=0.5, evaluated at
# 0, 0.5, 1, 1.5, and 2:

dzm1norm(seq(0, 2, by = 0.5))
#[1] 0.50000000 0.31374804 0.19947114 0.12248683
#[5] 0.07843701

#-----

# The cdf of the zero-modified lognormal (delta) distribution with
# parameters meanlog=1, sdlog=2, and p.zero=0.1, evaluated at 4:

pzm1norm(4, 1, 2, .1)
```

```

#[1] 0.6189203

#-----

# The median of the zero-modified lognormal (delta) distribution with
# parameters meanlog=2, sdlog=3, and p.zero=0.1:

qzmlnorm(0.5, 2, 3, 0.1)
#[1] 4.859177

#-----

# Random sample of 3 observations from the zero-modified lognormal
# (delta) distribution with parameters meanlog=1, sdlog=2, and p.zero=0.4.
# (Note: The call to set.seed simply allows you to reproduce this example.)

set.seed(20)
rzmlnorm(3, 1, 2, 0.4)
#[1] 0.000000 0.000000 3.146641

```

ZeroModifiedLognormalAlt

The Zero-Modified Lognormal (Delta) Distribution (Alternative Parameterization)

Description

Density, distribution function, quantile function, and random generation for the zero-modified lognormal distribution with parameters mean, cv, and p.zero.

The zero-modified lognormal (delta) distribution is the mixture of a lognormal distribution with a positive probability mass at 0.

Usage

```

dzmlnormAlt(x, mean = exp(1/2), cv = sqrt(exp(1) - 1), p.zero = 0.5)
pzmlnormAlt(q, mean = exp(1/2), cv = sqrt(exp(1) - 1), p.zero = 0.5)
qzmlnormAlt(p, mean = exp(1/2), cv = sqrt(exp(1) - 1), p.zero = 0.5)
rzmlnormAlt(n, mean = exp(1/2), cv = sqrt(exp(1) - 1), p.zero = 0.5)

```

Arguments

x	vector of quantiles.
q	vector of quantiles.
p	vector of probabilities between 0 and 1.
n	sample size. If length(n) is larger than 1, then length(n) random values are returned.
mean	vector of means of the lognormal part of the distribution on the. The default is mean=exp(1/2).
cv	vector of (positive) coefficients of variation of the lognormal part of the distribution. The default is cv=sqrt(exp(1) - 1).
p.zero	vector of probabilities between 0 and 1 indicating the probability the random variable equals 0. For rzmlnormAlt this must be a single, non-missing number.

Details

The zero-modified lognormal (delta) distribution is the mixture of a lognormal distribution with a positive probability mass at 0. This distribution was introduced without a name by Aitchison (1955), and the name Δ -distribution was coined by Aitchison and Brown (1957, p.95). It is a special case of a “zero-modified” distribution (see Johnson et al., 1992, p. 312).

Let $f(x; \theta, \tau)$ denote the density of a [lognormal random variable](#) X with parameters $\text{mean}=\theta$ and $\text{cv}=\tau$. The density function of a zero-modified lognormal (delta) random variable Y with parameters $\text{mean}=\theta$, $\text{cv}=\tau$, and $\text{p.zero}=p$, denoted $h(y; \theta, \tau, p)$, is given by:

$$h(y; \theta, \tau, p) = \begin{cases} p & \text{for } y = 0 \\ (1 - p)f(y; \theta, \tau) & \text{for } y > 0 \end{cases}$$

Note that θ is *not* the mean of the zero-modified lognormal distribution; it is the mean of the lognormal part of the distribution. Similarly, τ is *not* the coefficient of variation of the zero-modified lognormal distribution; it is the coefficient of variation of the lognormal part of the distribution.

Let γ , δ , and ω denote the mean, standard deviation, and coefficient of variation of the overall zero-modified lognormal distribution. Let η denote the standard deviation of the lognormal part of the distribution, so that $\eta = \theta\tau$. Aitchison (1955) shows that:

$$E(Y) = \gamma = (1 - p)\theta$$

$$\text{Var}(Y) = \delta^2 = (1 - p)\eta^2 + p(1 - p)\theta^2$$

so that

$$\omega = \sqrt{(\tau^2 + p)/(1 - p)}$$

Note that when $\text{p.zero}=p=0$, the zero-modified lognormal distribution simplifies to the lognormal distribution.

Value

`dzmlnormAlt` gives the density, `pzmlnormAlt` gives the distribution function, `qzmlnormAlt` gives the quantile function, and `rzmlnormAlt` generates random deviates.

Note

The zero-modified lognormal (delta) distribution is sometimes used to model chemical concentrations for which some observations are reported as “Below Detection Limit” (the nondetects are assumed equal to 0). See, for example, Gilliom and Helsel (1986), Owen and DeRouen (1980), and Gibbons et al. (2009, Chapter 12). USEPA (2009, Chapter 15) recommends this strategy only in specific situations, and Helsel (2012, Chapter 1) strongly discourages this approach to dealing with non-detects.

A variation of the zero-modified lognormal (delta) distribution is the [zero-modified normal distribution](#), in which a normal distribution is mixed with a positive probability mass at 0.

One way to try to assess whether a zero-modified lognormal (delta), zero-modified normal, censored normal, or censored lognormal is the best model for the data is to construct both censored and detects-only probability plots (see [qqPlotCensored](#)).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Aitchison, J. (1955). On the Distribution of a Positive Random Variable Having a Discrete Probability Mass at the Origin. *Journal of the American Statistical Association* **50**, 901-908.
- Aitchison, J., and J.A.C. Brown (1957). *The Lognormal Distribution (with special reference to its uses in economics)*. Cambridge University Press, London. pp.94-99.
- Crow, E.L., and K. Shimizu. (1988). *Lognormal Distributions: Theory and Applications*. Marcel Dekker, New York, pp.47-51.
- Gibbons, RD., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*. Second Edition. John Wiley and Sons, Hoboken, NJ.
- Gilliom, R.J., and D.R. Helsel. (1986). Estimation of Distributional Parameters for Censored Trace Level Water Quality Data: 1. Estimation Techniques. *Water Resources Research* **22**, 135-146.
- Helsel, D.R. (2012). *Statistics for Censored Environmental Data Using Minitab and R*. Second Edition. John Wiley and Sons, Hoboken, NJ, Chapter 1.
- Johnson, N. L., S. Kotz, and A.W. Kemp. (1992). *Univariate Discrete Distributions*. Second Edition. John Wiley and Sons, New York, p.312.
- Owen, W., and T. DeRouen. (1980). Estimation of the Mean for Lognormal Data Containing Zeros and Left-Censored Values, with Applications to the Measurement of Worker Exposure to Air Contaminants. *Biometrics* **36**, 707-719.
- USEPA (1992c). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities: Addendum to Interim Final Guidance*. Office of Solid Waste, Permits and State Programs Division, US Environmental Protection Agency, Washington, D.C.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.

See Also

[Zero-Modified Lognormal](#), [LognormalAlt](#), [ezmlnormAlt](#), [Probability Distributions and Random Numbers](#).

Examples

```
# Density of the zero-modified lognormal (delta) distribution with
# parameters mean=10, cv=1, and p.zero=0.5, evaluated at
# 9, 10, and 11:

dzmlnormAlt(9:11, mean = 10, cv = 1, p.zero = 0.5)
#[1] 0.02552685 0.02197043 0.01891924

#-----

# The cdf of the zero-modified lognormal (delta) distribution with
# parameters mean=10, cv=2, and p.zero=0.1, evaluated at 8:

pzmlnormAlt(8, 10, 2, .1)
#[1] 0.709009

#-----

# The median of the zero-modified lognormal (delta) distribution with
```

```
# parameters mean=10, cv=2, and p.zero=0.1:

qzmlnormAlt(0.5, 10, 2, 0.1)
#[1] 3.74576

#-----

# Random sample of 3 observations from the zero-modified lognormal
# (delta) distribution with parameters mean=10, cv=2, and p.zero=0.4.
# (Note: The call to set.seed simply allows you to reproduce this example.)

set.seed(20)
rzmlnormAlt(3, 10, 2, 0.4)
#[1] 0.000000 0.000000 4.907131
```

ZeroModifiedNormal	<i>The Zero-Modified Normal Distribution</i>
--------------------	--

Description

Density, distribution function, quantile function, and random generation for the zero-modified normal distribution with parameters mean, sd, and p.zero.

The zero-modified normal distribution is the mixture of a normal distribution with a positive probability mass at 0.

Usage

```
dzmnorm(x, mean = 0, sd = 1, p.zero = 0.5)
pzmnorm(q, mean = 0, sd = 1, p.zero = 0.5)
qzmnorm(p, mean = 0, sd = 1, p.zero = 0.5)
rzmnorm(n, mean = 0, sd = 1, p.zero = 0.5)
```

Arguments

x	vector of quantiles.
q	vector of quantiles.
p	vector of probabilities between 0 and 1.
n	sample size. If length(n) is larger than 1, then length(n) random values are returned.
mean	vector of means of the normal (Gaussian) part of the distribution. The default is mean=0.
sd	vector of (positive) standard deviations of the normal (Gaussian) part of the distribution. The default is sd=1.
p.zero	vector of probabilities between 0 and 1 indicating the probability the random variable equals 0. For rzmnorm this must be a single, non-missing number.

Details

The zero-modified normal distribution is the mixture of a normal distribution with a positive probability mass at 0.

Let $f(x; \mu, \sigma)$ denote the density of a [normal \(Gaussian\) random variable](#) X with parameters $\text{mean}=\mu$ and $\text{sd}=\sigma$. The density function of a zero-modified normal random variable Y with parameters $\text{mean}=\mu$, $\text{sd}=\sigma$, and $\text{p.zero}=p$, denoted $h(y; \mu, \sigma, p)$, is given by:

$$h(y; \mu, \sigma, p) = \begin{cases} p & \text{for } y = 0 \\ (1 - p)f(y; \mu, \sigma) & \text{for } y \neq 0 \end{cases}$$

Note that μ is *not* the mean of the zero-modified normal distribution; it is the mean of the normal part of the distribution. Similarly, σ is *not* the standard deviation of the zero-modified normal distribution; it is the standard deviation of the normal part of the distribution.

Let γ and δ denote the mean and standard deviation of the overall zero-modified normal distribution. Aitchison (1955) shows that:

$$E(Y) = \gamma = (1 - p)\mu$$

$$\text{Var}(Y) = \delta^2 = (1 - p)\sigma^2 + p(1 - p)\mu^2$$

Note that when $\text{p.zero}=p=0$, the zero-modified normal distribution simplifies to the normal distribution.

Value

`dzmnorm` gives the density, `pzmnorm` gives the distribution function, `qzmnorm` gives the quantile function, and `rzmnorm` generates random deviates.

Note

The zero-modified normal distribution is sometimes used to model chemical concentrations for which some observations are reported as “Below Detection Limit”. See, for example USEPA (1992c, pp.27-34) and Gibbons et al. (2009, Chapter 12). Note, however, that USEPA (1992c) has been superseded by USEPA (2009) which recommends this strategy only in specific situations (see Chapter 15 of the document). This strategy is strongly discouraged by Helsel (2012, Chapter 1).

In cases where you want to model chemical concentrations for which some observations are reported as “Below Detection Limit” and you want to treat the non-detects as equal to 0, it will usually be more appropriate to model the data with a [zero-modified lognormal \(delta\) distribution](#) since chemical concentrations are bounded below at 0 (e.g., Gilliom and Helsel, 1986; Owen and DeRouen, 1980).

One way to try to assess whether a zero-modified lognormal (delta), zero-modified normal, censored normal, or censored lognormal is the best model for the data is to construct both censored and detects-only probability plots (see [qqPlotCensored](#)).

Author(s)

Steven P. Millard (<EnvStats@ProbStatInfo.com>)

References

- Aitchison, J. (1955). On the Distribution of a Positive Random Variable Having a Discrete Probability Mass at the Origin. *Journal of the American Statistical Association* **50**, 901-908.
- Gilliom, R.J., and D.R. Helsel. (1986). Estimation of Distributional Parameters for Censored Trace Level Water Quality Data: 1. Estimation Techniques. *Water Resources Research* **22**, 135-146.
- Gibbons, RD., D.K. Bhaumik, and S. Aryal. (2009). *Statistical Methods for Groundwater Monitoring*. Second Edition. John Wiley and Sons, Hoboken, NJ.
- Helsel, D.R. (2012). *Statistics for Censored Environmental Data Using Minitab and R*. Second Edition. John Wiley and Sons, Hoboken, NJ, Chapter 1.
- Johnson, N. L., S. Kotz, and A.W. Kemp. (1992). *Univariate Discrete Distributions*. Second Edition. John Wiley and Sons, New York, p.312.
- Owen, W., and T. DeRouen. (1980). Estimation of the Mean for Lognormal Data Containing Zeros and Left-Censored Values, with Applications to the Measurement of Worker Exposure to Air Contaminants. *Biometrics* **36**, 707-719.
- USEPA (1992c). *Statistical Analysis of Ground-Water Monitoring Data at RCRA Facilities: Addendum to Interim Final Guidance*. Office of Solid Waste, Permits and State Programs Division, US Environmental Protection Agency, Washington, D.C.
- USEPA. (2009). *Statistical Analysis of Groundwater Monitoring Data at RCRA Facilities, Unified Guidance*. EPA 530/R-09-007, March 2009. Office of Resource Conservation and Recovery Program Implementation and Information Division. U.S. Environmental Protection Agency, Washington, D.C.

See Also

[Zero-Modified Lognormal](#), [Normal](#), [ezmnorm](#), [Probability Distributions and Random Numbers](#).

Examples

```
# Density of the zero-modified normal distribution with parameters
# mean=2, sd=1, and p.zero=0.5, evaluated at 0, 0.5, 1, 1.5, and 2:

dzmnorm(seq(0, 2, by = 0.5), mean = 2)
#[1] 0.5000000 0.0647588 0.1209854 0.1760327 0.1994711

#-----

# The cdf of the zero-modified normal distribution with parameters
# mean=3, sd=2, and p.zero=0.1, evaluated at 4:

pzmnorm(4, 3, 2, .1)
#[1] 0.7223162

#-----

# The median of the zero-modified normal distribution with parameters
# mean=3, sd=1, and p.zero=0.1:

qzmnorm(0.5, 3, 1, 0.1)
#[1] 2.86029

#-----
```

```
# Random sample of 3 observations from the zero-modified normal distribution
# with parameters mean=3, sd=1, and p.zero=0.4.
# (Note: The call to set.seed simply allows you to reproduce this example.)

set.seed(20)
rzmnorm(3, 3, 1, 0.4)
#[1] 0.000000 0.000000 3.073168
```

Index

*Topic **classes**

boxcox.object, [31](#)
boxcoxCensored.object, [42](#)
boxcoxLm.object, [45](#)
estimate.object, [376](#)
estimateCensored.object, [380](#)
gof.object, [429](#)
gofGroup.object, [432](#)
gofTwoSample.object, [455](#)

*Topic **datagen**

Chi, [63](#)
Empirical, [212](#)
EVD, [388](#)
GammaAlt, [419](#)
GEVD, [426](#)
Lognormal3, [494](#)
LognormalAlt, [498](#)
LognormalMix, [500](#)
LognormalMixAlt, [502](#)
LognormalTrunc, [504](#)
LognormalTruncAlt, [507](#)
NormalMix, [514](#)
NormalTrunc, [516](#)
Pareto, [521](#)
simulateMvMatrix, [716](#)
simulateVector, [723](#)
Triangular, [788](#)
ZeroModifiedLognormal, [796](#)
ZeroModifiedLognormalAlt, [799](#)
ZeroModifiedNormal, [802](#)

*Topic **datasets**

Benthic.df, [20](#)
Distribution.df, [91](#)
Environmental, [249](#)
EPA.02d.Ex.2.ug.per.L.vec, [251](#)
EPA.02d.Ex.4.mg.per.kg.vec, [251](#)
EPA.02d.Ex.6.mg.per.kg.vec, [252](#)
EPA.02d.Ex.9.mg.per.L.vec, [252](#)
EPA.09.Ex.10.1.nickel.df, [253](#)
EPA.09.Ex.11.1.arsenic.df, [253](#)
EPA.09.Ex.12.1.cc14.df, [254](#)
EPA.09.Ex.12.4.naphthalene.df, [255](#)
EPA.09.Ex.13.1.iron.df, [255](#)

EPA.09.Ex.14.1.manganese.df, [256](#)
EPA.09.Ex.14.3.alkalinity.df, [257](#)
EPA.09.Ex.14.4.arsenic.df, [257](#)
EPA.09.Ex.14.8.df, [258](#)
EPA.09.Ex.15.1.manganese.df, [259](#)
EPA.09.Ex.16.1.sulfate.df, [259](#)
EPA.09.Ex.16.2.benzene.df, [260](#)
EPA.09.Ex.16.4.copper.df, [261](#)
EPA.09.Ex.16.5.PCE.df, [261](#)
EPA.09.Ex.17.1.loglead.df, [262](#)
EPA.09.Ex.17.2.toluene.df, [263](#)
EPA.09.Ex.17.3.chrysene.df, [263](#)
EPA.09.Ex.17.3.log.chrysene.df, [264](#)
EPA.09.Ex.17.4.copper.df, [265](#)
EPA.09.Ex.17.5.chloride.df, [265](#)
EPA.09.Ex.17.6.sulfate.df, [266](#)
EPA.09.Ex.17.7.sodium.df, [267](#)
EPA.09.Ex.18.1.arsenic.df, [268](#)
EPA.09.Ex.18.2.chrysene.df, [268](#)
EPA.09.Ex.18.3.TCE.df, [269](#)
EPA.09.Ex.18.4.xylene.df, [270](#)
EPA.09.Ex.19.1.sulfate.df, [270](#)
EPA.09.Ex.19.2.chloride.df, [271](#)
EPA.09.Ex.19.5.mercury.df, [272](#)
EPA.09.Ex.20.1.nickel.df, [272](#)
EPA.09.Ex.21.1.aldicarb.df, [273](#)
EPA.09.Ex.21.2.benzene.df, [274](#)
EPA.09.Ex.21.5.beryllium.df, [274](#)
EPA.09.Ex.21.6.nitrate.df, [275](#)
EPA.09.Ex.21.7.TCE.df, [276](#)
EPA.09.Ex.22.1.VC.df, [276](#)
EPA.09.Ex.22.2.Specific.Conductance.df, [277](#)
EPA.09.Ex.6.3.sulfate.df, [278](#)
EPA.09.Ex.7.1.arsenic.df, [278](#)
EPA.09.Table.9.1.TCE.df, [279](#)
EPA.09.Table.9.3.df, [280](#)
EPA.09.Table.9.4.nickel.vec, [280](#)
EPA.89b.aldicarb1.df, [281](#)
EPA.89b.aldicarb2.df, [282](#)
EPA.89b.benzene.df, [282](#)
EPA.89b.cadmium.df, [283](#)

EPA.89b.chlordane1.df, [283](#)
 EPA.89b.chlordane2.df, [284](#)
 EPA.89b.edb.df, [284](#)
 EPA.89b.lead.df, [285](#)
 EPA.89b.loglead.df, [286](#)
 EPA.89b.manganese.df, [286](#)
 EPA.89b.sulfate.df, [287](#)
 EPA.89b.t29.df, [287](#)
 EPA.89b.toc.vec, [288](#)
 EPA.92c.arsenic1.df, [288](#)
 EPA.92c.arsenic2.df, [289](#)
 EPA.92c.arsenic3.df, [289](#)
 EPA.92c.benzene1.df, [290](#)
 EPA.92c.benzene2.df, [290](#)
 EPA.92c.cc14.df, [291](#)
 EPA.92c.chrysene.df, [292](#)
 EPA.92c.copper1.df, [292](#)
 EPA.92c.copper2.df, [293](#)
 EPA.92c.lognickel1.df, [293](#)
 EPA.92c.nickel1.df, [294](#)
 EPA.92c.nickel2.df, [294](#)
 EPA.92c.toluene.df, [295](#)
 EPA.92c.zinc.df, [295](#)
 EPA.92d.chromium.df, [296](#)
 EPA.92d.chromium.vec, [297](#)
 EPA.94b.lead.df, [297](#)
 EPA.94b.tccb.df, [298](#)
 EPA.97.cadmium.111.df, [298](#)
 Gibbons.et.al.09.Alkalinity.vec,
 [428](#)
 Gibbons.et.al.09.Vinyl.Chloride.vec,
 [428](#)
 Graham.et.al.75.etu.df, [463](#)
 Helsel.Cohn.88.app.b.df, [464](#)
 Helsel.Cohn.88.silver.df, [464](#)
 Lin.Evans.80.df, [489](#)
 Millard.Deverel.88.df, [512](#)
 Modified.TcCB.df, [513](#)
 NIOSH.89.air.lead.vec, [514](#)
 Olympic.NH4.df, [519](#)
 Ozone.NE.df, [520](#)
 Refinery.CO.df, [715](#)
 Skagit.NH3_N.df, [727](#)
 Total.P.df, [787](#)

***Topic design**

aovN, [13](#)
 aovPower, [16](#)
 ciBinomHalfWidth, [64](#)
 ciBinomN, [70](#)
 ciNormHalfWidth, [77](#)
 ciNormN, [80](#)
 plotAovDesign, [550](#)

plotCiBinomDesign, [553](#)
 plotCiNormDesign, [559](#)
 propTestN, [679](#)
 propTestPower, [685](#)
***Topic distribution**
 CastilloAndHadi1994, [54](#)
 cdfCompare, [56](#)
 cdfPlot, [60](#)
 Chi, [63](#)
 ciBinomHalfWidth, [64](#)
 ciBinomN, [70](#)
 ciNormHalfWidth, [77](#)
 ciNormN, [80](#)
 ebeta, [98](#)
 ebinom, [101](#)
 ecdfPlot, [107](#)
 ecdfPlotCensored, [111](#)
 eevd, [116](#)
 eexp, [122](#)
 egamma, [124](#)
 egammaAltCensored, [132](#)
 egammaCensored, [140](#)
 egeom, [148](#)
 egevd, [150](#)
 ehyper, [156](#)
 elnorm, [159](#)
 elnorm3, [161](#)
 elnormAlt, [173](#)
 elnormAltCensored, [185](#)
 elnormCensored, [200](#)
 elogis, [209](#)
 Empirical, [212](#)
 enbinom, [216](#)
 enorm, [219](#)
 enormCensored, [223](#)
 enparCensored, [242](#)
 epareto, [299](#)
 epdfPlot, [301](#)
 epois, [304](#)
 epoisCensored, [307](#)
 eqbeta, [315](#)
 eqbinom, [317](#)
 eqevd, [320](#)
 eqexp, [323](#)
 eqgamma, [325](#)
 eqgeom, [330](#)
 eqgevd, [332](#)
 eqhyper, [335](#)
 eqlnorm, [338](#)
 eqlnorm3, [342](#)
 eqlogis, [345](#)
 eqnbinom, [347](#)

- eqnorm, 349
- eqnpar, 353
- eqpareto, 361
- eqpois, 363
- equnif, 367
- eqweibull, 369
- eqzmlnorm, 371
- eqzmnorm, 374
- EulersConstant, 384
- eunif, 385
- EVD, 388
- eweibull, 390
- ezmlnorm, 393
- ezmnorm, 398
- GammaAlt, 419
- GEVD, 426
- gpqCiNormCensored, 457
- gpqTolIntNormCensored, 460
- HoskingEtAl1985, 465
- lMoment, 489
- Lognormal3, 494
- LognormalAlt, 498
- LognormalMix, 500
- LognormalMixAlt, 502
- LognormalTrunc, 504
- LognormalTruncAlt, 507
- NormalMix, 514
- NormalTrunc, 516
- Pareto, 521
- pdfPlot, 523
- plotCiBinomDesign, 553
- plotCiNormDesign, 559
- ppointsCensored, 567
- predIntGamma, 583
- predIntGammaSimultaneous, 589
- predIntLnorm, 598
- predIntLnormSimultaneous, 606
- predIntNorm, 614
- predIntNormK, 620
- predIntNormSimultaneous, 626
- predIntNormSimultaneousK, 636
- predIntNpar, 644
- predIntNparSimultaneous, 653
- predIntPois, 664
- pwMoment, 691
- qqPlot, 695
- qqPlotCensored, 703
- qqPlotGestalt, 711
- simulateMvMatrix, 716
- simulateVector, 723
- tolIntGamma, 750
- tolIntLnorm, 758
- tolIntNorm, 763
- tolIntNormCensored, 768
- tolIntNormK, 772
- tolIntNpar, 777
- tolIntPois, 783
- Triangular, 788
- ZeroModifiedLognormal, 796
- ZeroModifiedLognormalAlt, 799
- ZeroModifiedNormal, 802
- *Topic **dplot**
 - ppointsCensored, 567
- *Topic **hplot**
 - cdfCompare, 56
 - cdfPlot, 60
 - ecdfPlot, 107
 - ecdfPlotCensored, 111
 - epdfPlot, 301
 - pdfPlot, 523
 - plotAovDesign, 550
 - plotCiBinomDesign, 553
 - qqPlot, 695
 - qqPlotCensored, 703
 - qqPlotGestalt, 711
 - stripChart, 731
- *Topic **htest**
 - aovN, 13
 - aovPower, 16
 - ciBinomHalfWidth, 64
 - ciBinomN, 70
 - ciNormHalfWidth, 77
 - ciNormN, 80
 - ebeta, 98
 - ebinom, 101
 - eevd, 116
 - eexp, 122
 - egamma, 124
 - egammaAltCensored, 132
 - egammaCensored, 140
 - egeom, 148
 - egevd, 150
 - ehyper, 156
 - elnorm, 159
 - elnorm3, 161
 - elnormAlt, 173
 - elnormAltCensored, 185
 - elnormCensored, 200
 - elogis, 209
 - enbinom, 216
 - enorm, 219
 - enormCensored, 223
 - enparCensored, 242
 - epareto, 299

- epois, 304
- epoisCensored, 307
- eqbeta, 315
- eqbinom, 317
- eqevd, 320
- eqexp, 323
- eqgamma, 325
- eqgeom, 330
- eqgevd, 332
- eqhyper, 335
- eqlnorm, 338
- eqlnorm3, 342
- eqlogis, 345
- eqnbinom, 347
- eqnorm, 349
- eqnpar, 353
- eqpareto, 361
- eqpois, 363
- equnif, 367
- eqweibull, 369
- eqzmlnorm, 371
- eqzmnorm, 374
- eunif, 385
- eweibull, 390
- ezmlnorm, 393
- ezmnorm, 398
- gofGroupTest, 435
- gofTest, 441
- gpqCiNormCensored, 457
- gpqTolIntNormCensored, 460
- kendallSeasonalTrendTest, 474
- kendallTrendTest, 481
- lMoment, 489
- plotAovDesign, 550
- plotCiBinomDesign, 553
- plotCiNormDesign, 559
- predIntGamma, 583
- predIntGammaSimultaneous, 589
- predIntLnorm, 598
- predIntLnormSimultaneous, 606
- predIntNorm, 614
- predIntNormK, 620
- predIntNormSimultaneous, 626
- predIntNormSimultaneousK, 636
- predIntNpar, 644
- predIntNparSimultaneous, 653
- predIntPois, 664
- propTestN, 679
- propTestPower, 685
- pwMoment, 691
- stripChart, 731
- summaryStats, 742
- tolIntGamma, 750
- tolIntLnorm, 758
- tolIntNorm, 763
- tolIntNormCensored, 768
- tolIntNormK, 772
- tolIntNpar, 777
- tolIntPois, 783
- varGroupTest, 790
- varTest, 794
- *Topic **manip**
 - longToWide, 510
- *Topic **math**
 - base, 19
- *Topic **models**
 - anovaPE, 12
 - aovN, 13
 - aovPower, 16
 - boxcox, 22
 - boxcoxCensored, 33
 - boxcoxTransform, 47
 - calibrate, 51
 - detectionLimitCalibrate, 86
 - gofGroupTest, 435
 - gofTest, 441
 - inversePredictCalibrate, 469
 - plotAovDesign, 550
 - pointwise, 563
 - predict.lm, 578
 - propTestN, 679
 - propTestPower, 685
 - varGroupTest, 790
 - varTest, 794
- *Topic **nonparametric**
 - kendallSeasonalTrendTest, 474
 - kendallTrendTest, 481
- *Topic **package**
 - EnvStats-package, 7
 - FcnsByCat, 402
 - FcnsByCatCalibration, 403
 - FcnsByCatCensoredData, 403
 - FcnsByCatDataTrans, 406
 - FcnsByCatEstDistParams, 407
 - FcnsByCatEstDistQuants, 407
 - FcnsByCatGOFTests, 408
 - FcnsByCatHypothTests, 409
 - FcnsByCatMCandRisk, 410
 - FcnsByCatPlotProbDists, 410
 - FcnsByCatPower, 411
 - FcnsByCatPredInts, 414
 - FcnsByCatPrintPlot, 415
 - FcnsByCatProbDists, 417
 - FcnsByCatSumStats, 418

- FcnsByCatToInts, 418
- FcnsByCatTrend, 419
- *Topic **plot**
 - plot.boxcox, 526
 - plot.boxcoxCensored, 529
 - plot.boxcoxLm, 533
 - plot.gof, 536
 - plot.gofGroup, 541
 - plot.gofTwoSample, 545
- *Topic **print**
 - print.boxcox, 673
 - print.boxcoxCensored, 674
 - print.boxcoxLm, 675
 - print.gof, 676
 - print.gofGroup, 677
 - print.gofTwoSample, 678
- *Topic **regression**
 - anovaPE, 12
 - aovN, 13
 - aovPower, 16
 - calibrate, 51
 - detectionLimitCalibrate, 86
 - inversePredictCalibrate, 469
 - kendallSeasonalTrendTest, 474
 - kendallTrendTest, 481
 - plotAovDesign, 550
 - pointwise, 563
 - predict.lm, 578
- *Topic **univar**
 - boxcox, 22
 - boxcoxCensored, 33
 - boxcoxTransform, 47
 - cv, 83
 - geoMean, 422
 - geoSD, 424
 - iqr, 472
 - kurtosis, 485
 - skewness, 728
 - summaryFull, 736
 - summaryStats, 742
- (Alternative), 94–98
- (Delta), 95, 98
- abstract for Castillo and Hadi (1994), 152
- abstract for Hosking et al. (1985), 152, 153
- Air.df (Environmental), 249
- Aldicarb (EPA.89b.aldicarb2.df), 282
- anova, 12
- anova.lm, 12, 744, 746
- anovaPE, 12, 403
- aov, 15, 17, 552, 734, 735
- aovN, 13, 17, 411, 551, 552
- aovPower, 14, 15, 16, 412, 551, 552
- approx, 215, 700
- as.data.frame, 51
- axis, 732
- barplot, 418
- base, 19
- Benthic.df, 20
- Benzene (EPA.89b.benzene.df), 282
- Beta, 93, 95, 100, 316, 407, 408, 717, 724
- beta, 417
- beta distribution, 98, 99, 315, 639, 656, 779, 780
- beta function, 639, 692
- beta random variable, 779, 780
- binom, 417
- binom.test, 66, 68, 74, 102, 104, 558, 681, 682, 687, 688, 744, 746
- Binomial, 60, 93–95, 97, 104, 319, 407, 408, 524, 666, 775
- binomial, 667, 780, 785
- binomial distribution, 157, 306, 317, 336, 365
- binomial random variable, 355, 616
- Box-Cox (boxcox), 22
- Box-Cox Censored (boxcoxCensored), 33
- Box-Cox data transformations, 8
- Box-Cox Transformation (boxcox), 22
- Box-Cox Transformation for Censored Data (boxcoxCensored), 33
- BoxCox (boxcox), 22
- boxcox, 22, 31, 32, 39, 45, 46, 48, 50, 406, 526, 528, 529, 533, 535, 673–676
- BoxCox Censored (boxcoxCensored), 33
- BoxCox Transformation (boxcox), 22
- BoxCox Transformation for Censored Data (boxcoxCensored), 33
- boxcox.default (boxcox), 22
- boxcox.lm (boxcox), 22
- boxcox.object, 26, 27, 31, 46, 527–529, 673, 674
- boxcoxCensored, 33, 42, 43, 403, 529, 530, 532, 674, 675
- boxcoxCensored.object, 37, 39, 42, 530–532, 674, 675
- boxcoxLm.object, 26, 27, 32, 45, 533, 535, 675, 676
- boxcoxTransform, 24, 27, 35, 47, 406
- boxplot, 418, 473
- boxplots, 303, 731
- Cadmium (EPA.89b.cadmium.df), 283

- calibrate, [12](#), [51](#), [87](#), [90](#), [403](#), [470](#), [471](#), [565](#), [580](#), [581](#)
- Calibration, [402](#)
- Calibration (FcnsByCatCalibration), [403](#)
- calibration, [8](#)
- Castillo and Hadi (1994), [153](#)
- Castillo and Hadi 1994
(CastilloAndHadi1994), [54](#)
- CastilloAndHadi1994, [54](#)
- Cauchy, [93](#), [95](#)
- cauchy, [417](#)
- cdfCompare, [56](#), [58](#), [62](#), [110](#), [215](#), [303](#), [411](#), [540](#), [549](#)
- cdfCompareCensored, [114](#), [115](#), [405](#)
- cdfPlot, [58](#), [59](#), [60](#), [110](#), [114](#), [115](#), [303](#), [411](#), [526](#)
- censored (less-than-detection-limit)
data, [8](#)
- Censored Data, [92](#), [93](#), [402](#)
- Censored Data (FcnsByCatCensoredData),
[403](#)
- chenTTest, [409](#)
- Chi, [63](#), [93](#), [96](#)
- chi, [417](#)
- chi random variable, [623](#)
- Chi-square, [93](#), [96](#)
- chi-square, [128](#), [585](#), [592](#), [752](#)
- chi-square distribution, [123](#), [167](#), [177](#),
[305](#), [420](#)
- chi-square distributions, [326](#)
- chi-squared, [63](#)
- chi-squared distribution, [128](#), [135](#), [143](#),
[192](#), [232](#), [311](#)
- chisq, [417](#)
- chisq.test, [445](#), [449](#), [744](#), [746](#)
- Chisquare, [64](#)
- Chisquare distribution, [543](#)
- Chlordane (EPA.89b.chlordane1.df), [283](#)
- chol, [718](#)
- ciBinomHalfWidth, [64](#), [72](#), [74](#), [104](#), [411](#), [556](#),
[558](#)
- ciBinomN, [67](#), [68](#), [70](#), [73](#), [104](#), [411](#), [556](#)–[558](#)
- ciNormHalfWidth, [77](#), [81](#), [82](#), [411](#), [561](#)
- ciNormN, [78](#), [80](#), [81](#), [411](#), [561](#), [562](#)
- class, [12](#), [23](#), [52](#), [731](#), [736](#), [742](#)
- Coefficient of Variation (cv), [83](#)
- coefficient of variation (cv), [83](#)
- Compare CDFs (cdfCompare), [56](#)
- cor, [721](#)
- cor.test, [478](#), [482](#)–[484](#)
- Cumulative Distribution (cdfPlot), [60](#)
- CV (cv), [83](#)
- cv, [83](#), [418](#), [488](#), [494](#), [730](#), [738](#), [739](#)
- Data Transformations, [27](#), [39](#), [50](#), [402](#), [526](#),
[529](#), [532](#), [533](#), [535](#), [673](#)–[676](#)
- Data Transformations
(FcnsByCatDataTrans), [406](#)
- data.frame, [12](#), [511](#)
- dbeta, [93](#)
- dchi (Chi), [63](#)
- Delta Distribution
(ZeroModifiedLognormal), [796](#)
- DeltaDist (ZeroModifiedLognormal), [796](#)
- DeltaDistAlt
(ZeroModifiedLognormalAlt), [799](#)
- demp (Empirical), [212](#)
- density, [212](#), [213](#), [215](#), [302](#), [303](#)
- detectionLimitCalibrate, [53](#), [86](#), [403](#), [471](#),
[565](#), [581](#)
- devAskNewPage, [527](#), [530](#), [534](#), [537](#), [542](#), [546](#),
[713](#)
- devd (EVD), [388](#)
- dgamma, [420](#)
- dgammaAlt (GammaAlt), [419](#)
- dgevd (GEVD), [426](#)
- diff, [739](#)
- digamma, [388](#)
- digamma function, [99](#), [385](#)
- distribution quantiles, [8](#)
- Distribution.df, [57](#), [60](#), [62](#), [91](#), [376](#), [379](#),
[380](#), [383](#), [417](#), [429](#), [430](#), [432](#), [433](#),
[436](#), [437](#), [442](#)–[445](#), [524](#), [526](#), [695](#),
[696](#), [701](#), [704](#), [712](#), [716](#), [724](#)
- dlnorm, [495](#), [498](#)
- dlnorm3 (Lognormal3), [494](#)
- dlnormAlt (LognormalAlt), [498](#)
- dlnormMix (LognormalMix), [500](#)
- dlnormMixAlt (LognormalMixAlt), [502](#)
- dlnormTrunc (LognormalTrunc), [504](#)
- dlnormTruncAlt (LognormalTruncAlt), [507](#)
- dnormMix (NormalMix), [514](#)
- dnormTrunc (NormalTrunc), [516](#)
- dotchart, [418](#)
- dpareto (Pareto), [521](#)
- dtri (Triangular), [788](#)
- dzmlnorm (ZeroModifiedLognormal), [796](#)
- dzmlnormAlt (ZeroModifiedLognormalAlt),
[799](#)
- dzmnorm (ZeroModifiedNormal), [802](#)
- ebeta, [93](#), [98](#), [315](#), [316](#), [407](#)
- ebinom, [68](#), [74](#), [101](#), [317](#), [319](#), [407](#), [558](#), [667](#)
- ecdfPlot, [57](#)–[59](#), [62](#), [107](#), [113](#), [115](#), [136](#), [144](#),
[194](#), [215](#), [234](#), [245](#), [303](#), [312](#), [411](#),

- [538, 547, 569, 571, 575, 698, 701](#)
 ecdfPlotCensored, [110, 111, 243, 247, 405, 574, 575, 709](#)
 eevd, [116, 155, 320–322, 335, 385, 389, 407, 694](#)
 eexp, [122, 323, 324, 407](#)
 egamma, [124, 138, 146, 325, 326, 328, 407, 437, 443, 584, 586, 590, 595, 751, 753](#)
 egammaAlt, [325, 407, 421, 595](#)
 egammaAlt(egamma), [124](#)
 egammaAltCensored, [132, 146, 404](#)
 egammaCensored, [138, 140, 404](#)
 egeom, [148, 218, 331, 332, 348, 407](#)
 egevd, [54, 56, 150, 333, 335, 407, 427, 469, 694](#)
 ehyper, [156, 336, 337, 407](#)
 elnorm, [159, 206, 338, 340, 407, 425, 449, 599, 601, 607, 610, 758, 760](#)
 elnorm3, [161, 342–344, 407, 449, 497](#)
 elnormAlt, [26, 49, 173, 188, 189, 193, 198, 407, 423, 449, 500, 599, 601, 607, 610, 758–760](#)
 elnormAltCensored, [38, 185, 404, 599, 607, 758](#)
 elnormCensored, [189, 198, 200, 404, 599, 607, 758](#)
 elnormMultiplyCensored, [338](#)
 elnormSinglyCensored, [338](#)
 elogis, [209, 345, 346, 407](#)
 Empirical, [212, 303, 410, 717, 721, 724, 726](#)
 empirical cdf plots, [61](#)
 empirical PDF (epdfPlot), [301](#)
 empirical pdf plots, [525](#)
 enbinom, [149, 216, 332, 347, 348, 407](#)
 enorm, [78, 82, 160, 219, 229, 231, 239, 349–351, 377, 407, 449, 562, 614, 617, 625, 626, 631, 642, 763, 766, 776](#)
 enormCensored, [37, 188, 198, 203, 206, 223, 381, 404, 458, 459, 461, 462, 614, 626, 705, 763, 769, 770](#)
 enormMultiplyCensored, [349](#)
 enormSinglyCensored, [349](#)
 enparCensored, [241, 404](#)
 Environmental, [249](#)
 EnvStats (EnvStats-package), [7](#)
 EnvStats Functions for Censored Data, [380, 383, 705, 709](#)
 EnvStats Functions for Censored Data (FcnsByCatCensoredData), [403](#)
 EnvStats Functions for
 Goodness-of-Fit Tests
 (FcnsByCatGOFTests), [408](#)
 EnvStats-package, [7](#)
 EPA.02d.Ex.2.ug.per.L.vec, [251](#)
 EPA.02d.Ex.4.mg.per.kg.vec, [251](#)
 EPA.02d.Ex.6.mg.per.kg.vec, [252](#)
 EPA.02d.Ex.9.mg.per.L.vec, [252](#)
 EPA.09.Ex.10.1.nickel.df, [253](#)
 EPA.09.Ex.11.1.arsenic.df, [253](#)
 EPA.09.Ex.12.1.cc14.df, [254](#)
 EPA.09.Ex.12.4.naphthalene.df, [255](#)
 EPA.09.Ex.13.1.iron.df, [255](#)
 EPA.09.Ex.14.1.manganese.df, [256](#)
 EPA.09.Ex.14.3.alkalinity.df, [257](#)
 EPA.09.Ex.14.4.arsenic.df, [257](#)
 EPA.09.Ex.14.8.df, [258](#)
 EPA.09.Ex.15.1.manganese.df, [259](#)
 EPA.09.Ex.16.1.sulfate.df, [259](#)
 EPA.09.Ex.16.2.benzene.df, [260](#)
 EPA.09.Ex.16.4.copper.df, [261](#)
 EPA.09.Ex.16.5.PCE.df, [261](#)
 EPA.09.Ex.17.1.loglead.df, [262](#)
 EPA.09.Ex.17.2.toluene.df, [263](#)
 EPA.09.Ex.17.3.chrysene.df, [263](#)
 EPA.09.Ex.17.3.log.chrysene.df, [264](#)
 EPA.09.Ex.17.4.copper.df, [265](#)
 EPA.09.Ex.17.5.chloride.df, [265](#)
 EPA.09.Ex.17.6.sulfate.df, [266](#)
 EPA.09.Ex.17.7.sodium.df, [267](#)
 EPA.09.Ex.18.1.arsenic.df, [268](#)
 EPA.09.Ex.18.2.chrysene.df, [268](#)
 EPA.09.Ex.18.3.TCE.df, [269](#)
 EPA.09.Ex.18.4.xylene.df, [270](#)
 EPA.09.Ex.19.1.sulfate.df, [270](#)
 EPA.09.Ex.19.2.chloride.df, [271](#)
 EPA.09.Ex.19.5.mercury.df, [272](#)
 EPA.09.Ex.20.1.nickel.df, [272](#)
 EPA.09.Ex.21.1.aldicarb.df, [273](#)
 EPA.09.Ex.21.2.benzene.df, [274](#)
 EPA.09.Ex.21.5.beryllium.df, [274](#)
 EPA.09.Ex.21.6.nitrate.df, [275](#)
 EPA.09.Ex.21.7.TCE.df, [276](#)
 EPA.09.Ex.22.1.VC.df, [276](#)
 EPA.09.Ex.22.2.Specific.Conductance.df, [277](#)
 EPA.09.Ex.6.3.sulfate.df, [278](#)
 EPA.09.Ex.7.1.arsenic.df, [278](#)
 EPA.09.Table.9.1.TCE.df, [279](#)
 EPA.09.Table.9.3.df, [280](#)
 EPA.09.Table.9.4.nickel.vec, [280](#)
 EPA.89b.aldicarb1.df, [281](#)
 EPA.89b.aldicarb2.df, [282](#)

- EPA.89b.benzene.df, [282](#)
- EPA.89b.cadmium.df, [283](#)
- EPA.89b.chlordane1.df, [283](#)
- EPA.89b.chlordane2.df, [284](#)
- EPA.89b.edb.df, [284](#)
- EPA.89b.lead.df, [285](#)
- EPA.89b.loglead.df, [286](#)
- EPA.89b.manganese.df, [286](#)
- EPA.89b.sulfate.df, [287](#)
- EPA.89b.t29.df, [287](#)
- EPA.89b.toc.vec, [288](#)
- EPA.92c.arsenic1.df, [288](#)
- EPA.92c.arsenic2.df, [289](#)
- EPA.92c.arsenic3.df, [289](#)
- EPA.92c.benzene1.df, [290](#)
- EPA.92c.benzene2.df, [290](#)
- EPA.92c.ccl4.df, [291](#)
- EPA.92c.chrysene.df, [292](#)
- EPA.92c.copper1.df, [292](#)
- EPA.92c.copper2.df, [293](#)
- EPA.92c.lognickel1.df, [293](#)
- EPA.92c.nickel1.df, [294](#)
- EPA.92c.nickel2.df, [294](#)
- EPA.92c.toluene.df, [295](#)
- EPA.92c.zinc.df, [295](#)
- EPA.92d.chromium.df, [296](#)
- EPA.92d.chromium.vec, [297](#)
- EPA.94b.lead.df, [297](#)
- EPA.94b.tccb.df, [298](#), [513](#)
- EPA.97.cadmium.111.df, [298](#)
- epareto, [299](#), [361](#), [362](#), [407](#), [523](#)
- epdfPlot, [215](#), [301](#), [411](#), [526](#)
- epois, [304](#), [310](#), [314](#), [363–366](#), [407](#), [665](#), [670](#), [784–786](#)
- epoisCensored, [307](#), [404](#), [665](#), [784](#)
- eqbeta, [315](#), [408](#)
- eqbinom, [317](#), [408](#)
- eqevd, [320](#), [408](#)
- eqexp, [323](#), [408](#)
- eqgamma, [129](#), [325](#), [408](#)
- eqgammaAlt, [408](#)
- eqgammaAlt (eqgamma), [325](#)
- eqgeom, [330](#), [408](#)
- eqgevd, [332](#), [408](#)
- eqhyper, [335](#), [408](#)
- eqlnorm, [179](#), [338](#), [408](#), [760](#)
- eqlnorm3, [342](#), [408](#)
- eqlnormCensored, [404](#)
- eqlogis, [345](#), [408](#)
- eqnbinom, [347](#), [408](#)
- eqnorm, [326](#), [328](#), [339](#), [340](#), [349](#), [377](#), [408](#), [614](#), [617](#), [625](#), [626](#), [766](#), [776](#)
- eqnormCensored, [381](#), [404](#), [461](#), [462](#), [770](#)
- eqnpar, [173](#), [177](#), [215](#), [353](#), [404](#), [405](#), [781](#)
- eqpareto, [361](#), [408](#), [523](#)
- eqpois, [363](#), [408](#), [786](#)
- equnif, [367](#), [408](#)
- eqweibull, [369](#), [408](#)
- eqzmlnorm, [371](#), [408](#)
- eqzmlnormAlt, [408](#)
- eqzmlnormAlt (eqzmlnorm), [371](#)
- eqzmnorm, [374](#), [408](#)
- estimate (estimate.object), [376](#)
- Estimate distribution parameters, [8](#)
- estimate.object, [99](#), [102](#), [119](#), [123](#), [128](#), [129](#), [149](#), [153](#), [157](#), [160](#), [168](#), [178](#), [211](#), [217](#), [221](#), [300](#), [305](#), [316](#), [318](#), [319](#), [321–324](#), [326](#), [328](#), [331–333](#), [335–337](#), [339](#), [340](#), [343](#), [346](#), [348](#), [350](#), [351](#), [357](#), [358](#), [362](#), [365](#), [366](#), [368](#), [370](#), [372](#), [374](#), [375](#), [376](#), [383](#), [386](#), [387](#), [391](#), [392](#), [395](#), [399](#), [400](#), [585](#), [586](#), [591](#), [595](#), [599](#), [601](#), [608](#), [610](#), [616](#), [617](#), [625](#), [628](#), [631](#), [642](#), [646](#), [648](#), [657](#), [658](#), [669](#), [670](#), [752](#), [753](#), [759](#), [760](#), [765](#), [766](#), [776](#), [780](#), [781](#), [785](#), [786](#)
- estimateCensored
(estimateCensored.object), [380](#)
- estimateCensored.object, [137](#), [138](#), [145](#), [146](#), [195](#), [198](#), [203](#), [206](#), [236](#), [239](#), [246](#), [247](#), [313](#), [314](#), [379](#), [380](#), [459](#), [462](#), [769](#), [770](#)
- Estimating Distribution Parameters, [57](#), [78](#), [82](#), [92](#), [93](#), [377](#), [379](#), [402](#), [437](#), [443](#), [562](#), [670](#), [696](#), [712](#), [760](#), [766](#), [776](#), [781](#), [786](#)
- Estimating Distribution Parameters
(FcnsByCatEstDistParams), [407](#)
- Estimating Distribution Quantiles, [92](#), [358](#), [377](#), [379](#), [402](#), [404](#), [405](#), [760](#), [766](#), [776](#), [781](#), [786](#)
- Estimating Distribution Quantiles
(FcnsByCatEstDistQuants), [407](#)
- Euler's Constant, [120](#)
- Euler's constant, [118](#), [119](#), [388](#)
- Eulers Constant (EulersConstant), [384](#)
- EulersConstant, [384](#)
- eunif, [367](#), [368](#), [385](#), [407](#)
- EVD, [370](#), [388](#), [392](#), [427](#)
- evd, [417](#)
- evNormOrdStats, [165](#)
- eweibull, [369](#), [370](#), [390](#), [407](#)
- exp, [417](#)

- Exponential, [93](#), [96](#), [124](#), [324](#), [370](#), [392](#), [407](#), [408](#), [523](#)
- exponential, [128](#), [326](#), [389](#), [585](#), [592](#), [752](#)
- exponential distribution, [120](#), [122](#), [123](#), [300](#), [321](#), [323](#), [362](#), [370](#), [389](#), [391](#), [420](#), [522](#)
- exponential random variable, [120](#), [321](#)
- Extreme, [94](#), [96](#)
- Extreme Value, [407](#), [408](#)
- Extreme Value Distribution, [120](#), [155](#), [322](#), [335](#), [385](#)
- Extreme Value Distribution (EVD), [388](#)
- extreme value distribution, [116](#), [117](#), [123](#), [320](#), [323](#), [426](#)
- extreme value distributions, [154](#), [334](#)
- extreme value distributions (EVD), [427](#)
- ezmlnorm, [371–373](#), [375](#), [393](#), [400](#), [407](#), [449](#), [798](#)
- ezmlnormAlt, [371–373](#), [407](#), [449](#), [801](#)
- ezmlnormAlt (ezmlnorm), [393](#)
- ezmnorm, [374](#), [398](#), [407](#), [449](#), [804](#)
- F, [94](#), [96](#)
- f, [417](#)
- F-distribution, [667](#)
- FcnsByCat, [402](#)
- FcnsByCatCalibration, [403](#)
- FcnsByCatCensoredData, [403](#), [574](#)
- FcnsByCatDataTrans, [406](#)
- FcnsByCatEstDistParams, [407](#)
- FcnsByCatEstDistQuants, [407](#)
- FcnsByCatGOFTests, [408](#)
- FcnsByCatHypothTests, [409](#)
- FcnsByCatMCandRisk, [410](#)
- FcnsByCatPlotProbDists, [410](#)
- FcnsByCatPower, [411](#)
- FcnsByCatPredInts, [414](#)
- FcnsByCatPrintPlot, [415](#)
- FcnsByCatProbDists, [417](#)
- FcnsByCatSumStats, [418](#)
- FcnsByCatTolInts, [418](#)
- FcnsByCatTrend, [419](#)
- fisher.test, [744](#), [746](#)
- Fishers's exact test, [158](#), [337](#)
- floor, [157](#)
- format, [673–678](#)
- formula, [51](#)
- Functions By Category, [7](#)
- Functions By Category (FcnsByCat), [402](#)
- Functions for Censored Data (FcnsByCatCensoredData), [403](#)
- Gamma, [94](#), [96](#), [124](#), [407](#), [408](#)
- gamma, [48](#), [417](#)
- Gamma Distribution (GammaAlt), [419](#)
- gamma distribution, [123](#), [125](#), [132](#), [133](#), [140](#), [141](#), [323](#), [325](#), [585](#), [589](#), [592](#), [750](#)
- gamma distribution (alternative parameterization), [589](#)
- Gamma function, [127](#)
- gamma function, [391](#), [466](#)
- GammaAlt, [419](#), [586](#), [595](#)
- GammaDist, [129](#), [138](#), [146](#), [328](#), [421](#), [586](#), [595](#), [753](#)
- gammaAlt, [417](#)
- Generalized, [94](#), [96](#)
- Generalized Extreme Value, [407](#), [408](#)
- Generalized Extreme Value Distribution, [56](#), [155](#), [335](#), [469](#)
- Generalized Extreme Value Distribution (GEVD), [426](#)
- generalized extreme value distribution, [54](#), [150](#), [152](#), [154](#), [332](#), [334](#), [465](#), [693](#)
- geom, [417](#)
- geoMean, [418](#), [422](#), [425](#), [739](#)
- Geometric, [94](#), [96](#), [149](#), [218](#), [332](#), [348](#), [407](#), [408](#)
- geometric distribution, [148](#), [149](#), [217](#), [330](#), [331](#), [348](#)
- Geometric Mean (geoMean), [422](#)
- geometric mean (geoMean), [422](#)
- Geometric SD (geoSD), [424](#)
- geometric SD (geoSD), [424](#)
- Geometric Standard Deviation (geoSD), [424](#)
- geometric standard deviation (geoSD), [424](#)
- geoSD, [418](#), [423](#), [424](#), [738](#), [739](#)
- GEVD, [152](#), [389](#), [426](#)
- gevd, [417](#)
- Gibbons.et.al.09.Alkilinity.vec, [428](#)
- Gibbons.et.al.09.Vinyl.Chloride.vec, [428](#)
- GOF (FcnsByCatGOFTests), [408](#)
- gof (gof.object), [429](#)
- gof.object, [429](#), [447](#), [449](#), [537](#), [539](#), [540](#), [546](#), [676](#)
- gofGroup (gofGroup.object), [432](#)
- gofGroup.object, [432](#), [438](#), [439](#), [541](#), [543](#), [544](#), [677](#)
- gofGroupTest, [409](#), [432](#), [433](#), [435](#), [447](#), [541](#), [544](#), [677](#)
- gofTest, [24](#), [25](#), [37](#), [59](#), [110](#), [166](#), [168](#), [408](#),

- [429, 430, 438, 439, 441, 455, 536, 540, 545, 547, 549, 678, 730](#)
- [gofTestCensored, 36, 405](#)
- [gofTwoSample \(gofTwoSample.object\), 455](#)
- [gofTwoSample.object, 447, 455, 548, 549, 678](#)
- [Goodness-of-Fit Tests, 27, 39, 50, 402, 409, 430, 433, 456, 536, 540, 541, 544, 545, 549, 676–678](#)
- [Goodness-of-Fit Tests \(FcnsByCatGOFTests\), 408](#)
- [goodness-of-fit tests, 8](#)
- [gpqCiNormCensored, 457](#)
- [gpqCiNormMultiplyCensored, 236, 404](#)
- [gpqCiNormMultiplyCensored \(gpqCiNormCensored\), 457](#)
- [gpqCiNormSinglyCensored, 235, 404](#)
- [gpqCiNormSinglyCensored \(gpqCiNormCensored\), 457](#)
- [gpqTolIntNormCensored, 460](#)
- [gpqTolIntNormMultiplyCensored, 405, 769](#)
- [gpqTolIntNormMultiplyCensored \(gpqTolIntNormCensored\), 460](#)
- [gpqTolIntNormSinglyCensored, 405, 769, 770](#)
- [gpqTolIntNormSinglyCensored \(gpqTolIntNormCensored\), 460](#)
- [Graham.et.al.75.etu.df, 463](#)
- [Gumbel Distribution \(EVD\), 388](#)
- [Helsel.Cohn.88.app.b.df, 464](#)
- [Helsel.Cohn.88.silver.df, 464](#)
- [hist, 418, 540](#)
- [histograms, 303](#)
- [Hosking et al 1985 \(HoskingEtAl1985\), 465](#)
- [Hosking et al., 1985\), 54, 56](#)
- [HoskingEtAl1985, 465](#)
- [htest.object, 476, 478, 482, 484, 792, 794](#)
- [hyper, 417](#)
- [Hypergeometric, 94, 96, 158, 337, 407, 408](#)
- [hypergeometric distribution, 156, 157, 335, 336](#)
- [Hypothesis Tests, 402, 419](#)
- [Hypothesis Tests \(FcnsByCatHypothTests\), 409](#)
- [hypothesis tests, 8](#)
- [integrate, 637, 654, 773](#)
- [Interquartile Range \(iqr\), 472](#)
- [inversePredictCalibrate, 52, 53, 90, 403, 469, 565, 581](#)
- [iqr, 418, 472, 739](#)
- [kendallSeasonalTrendTest, 409, 474, 476, 482, 484](#)
- [kendallTrendTest, 409, 476, 478, 481](#)
- [kruskal.test, 734, 735, 744, 746](#)
- [ks.test, 438, 443–445, 449](#)
- [Kurtosis \(kurtosis\), 485](#)
- [kurtosis, 86, 418, 485, 494, 730, 738, 739](#)
- [L-moments, 693](#)
- [Lin.Evans.80.df, 489](#)
- [linearTrendTestN, 412](#)
- [linearTrendTestPower, 412](#)
- [linearTrendTestScaledMds, 412](#)
- [lines, 58, 61, 108, 113, 302](#)
- [lm, 12, 23, 52, 53, 471, 565, 581](#)
- [lmle, mme,, 97](#)
- [lMoment, 85, 418, 487, 489, 694, 729](#)
- [lmsreg, 55, 153](#)
- [lnorm, 417](#)
- [lnorm3, 417](#)
- [lnormAlt, 417](#)
- [lnormMix, 417](#)
- [lnormMixAlt, 417](#)
- [lnormTrunc, 417](#)
- [lnormTruncAlt, 417](#)
- [logis, 417](#)
- [Logistic, 94, 96, 211, 346, 407, 408](#)
- [logistic distribution, 209, 211, 300, 345, 346, 362, 522](#)
- [Lognormal, 60, 94–98, 161, 170, 181, 206, 340, 344, 373, 397, 404, 407, 408, 437, 443–445, 449, 497, 500, 502, 504, 506, 524, 601, 610, 717, 724, 760, 798](#)
- [lognormal, 538](#)
- [lognormal distribution, 48, 159, 160, 173, 174, 185, 187, 200, 203, 338, 420, 425, 495, 498, 499, 503, 598, 606, 758](#)
- [lognormal distribution \(alternative parameterization\), 598, 606, 758](#)
- [lognormal random variable, 501, 503, 797, 800](#)
- [Lognormal3, 170, 344, 449, 494](#)
- [LognormalAlt, 161, 170, 174, 176, 181, 187, 198, 344, 498, 504, 507, 509, 610, 760, 798, 801](#)
- [LognormalMix, 500, 504, 516](#)
- [LognormalMixAlt, 502](#)
- [LognormalTrunc, 504](#)
- [LognormalTruncAlt, 507](#)
- [longToWide, 510](#)
- [lse, mle, 97](#)

- ltsreg, [153](#)
- mad, [739](#)
- matrix, [511](#)
- max, [739](#)
- mean, [418](#), [423](#), [739](#)
- median, [153](#), [418](#), [423](#), [739](#)
- methods, [23](#), [731](#), [736](#), [742](#)
- Millard.Deverel.88.df, [512](#)
- min, [739](#)
- Mixture, [94–97](#)
- mle, bcmle, mme, mmue, [96](#)
- mle, mme, mmue, [95–97](#)
- mle, mme, mmue,, [96](#)
- mle, mme, mmue, pwme, [96](#)
- mle, mvue, [96](#)
- mle, pwme, tsoe, [96](#)
- mle/mme, [96](#)
- mle/mme, mvue, [96](#), [97](#)
- mle/mme/mvue, [95](#), [97](#)
- mmue, mmme,, [97](#)
- model.matrix, [52](#)
- Modified.TcCB.df, [513](#)
- Monte Carlo simulation and
probabilistic risk
assessment, [8](#)
- Monte Carlo Simulation and Risk
Assessment, [402](#)
- Monte Carlo Simulation and Risk
Assessment
(FcnsByCatMCandRisk), [410](#)
- mvue, [98](#)
- mvue, qmle, [96](#)
- na.exclude, [51](#)
- na.fail, [51](#)
- na.omit, [51](#)
- na.pass, [436](#), [442](#), [475](#), [482](#), [737](#), [743](#), [791](#)
- napredict, [579](#)
- nbinom, [417](#)
- Negative, [94](#), [97](#)
- Negative Binomial, [407](#), [408](#)
- Negative Binomial distribution, [656](#)
- negative binomial distribution, [149](#),
[216](#), [217](#), [331](#), [347](#), [348](#)
- negative binomial distributions, [217](#)
- NegBinomial, [149](#), [218](#), [332](#), [348](#)
- NIOSH.89.air.lead.vec, [514](#)
- nlnmb, [24](#), [35](#), [152](#), [164](#), [639](#)
- non-central Student's t-distribution,
[639](#), [640](#)
- Nonparametric Prediction Interval
(predIntNpar), [644](#)
- Nonparametric Simultaneous Prediction
Interval
(predIntNparSimultaneous), [653](#)
- norm, [417](#)
- Normal, [15](#), [17](#), [64](#), [78](#), [82](#), [92](#), [93](#), [95](#), [97](#), [98](#),
[161](#), [170](#), [181](#), [222](#), [239](#), [344](#), [351](#),
[375](#), [400](#), [404](#), [407](#), [408](#), [436](#), [437](#),
[442–445](#), [449](#), [516](#), [518](#), [552](#), [562](#),
[617](#), [625](#), [631](#), [642](#), [766](#), [776](#), [804](#)
- normal, [63](#), [538](#)
- normal (Gaussian), [713](#)
- normal (Gaussian) distribution, [160](#),
[203](#), [219](#), [220](#), [224](#)
- normal (Gaussian) random variable, [803](#)
- normal distribuiton, [177](#)
- Normal distribution, [457](#)
- normal distribution, [177](#), [226](#), [349](#), [614](#),
[615](#), [620–622](#), [626](#), [627](#), [636–639](#),
[763](#)
- normal random variable, [515](#)
- Normal(0,1) distribution, [543](#)
- NormalMix, [502](#), [504](#), [514](#)
- NormalTrunc, [516](#)
- normMix, [417](#)
- normTrunc, [417](#)
- Olympic.NH4.df, [519](#)
- oneSamplePermutationTest, [20](#), [409](#)
- options, [51](#), [551](#), [556](#), [561](#)
- Ozone.NE.df, [520](#)
- par, [57](#), [58](#), [61](#), [108](#), [112](#), [113](#), [302](#), [524](#), [525](#),
[527](#), [528](#), [530](#), [531](#), [534](#), [535](#), [538](#),
[539](#), [543](#), [547](#), [548](#), [551](#), [556](#), [558](#),
[561](#), [697](#), [706](#), [713](#), [732](#)
- Parameter, [94](#), [97](#)
- Pareto, [95](#), [97](#), [299](#), [300](#), [362](#), [407](#), [408](#), [521](#)
- pareto, [417](#)
- Pareto distribution, [299](#), [361](#)
- pbeta, [93](#)
- pchi (Chi), [63](#)
- pdfPlot, [62](#), [303](#), [411](#), [523](#), [540](#)
- pemp (Empirical), [212](#)
- pevd (EVD), [388](#)
- pgamma, [420](#)
- pgammaAlt (GammaAlt), [419](#)
- pgev (GEVD), [426](#)
- plnorm, [495](#), [498](#)
- plnorm3 (Lognormal3), [494](#)
- plnormAlt (LognormalAlt), [498](#)
- plnormMix (LognormalMix), [500](#)
- plnormMixAlt (LognormalMixAlt), [502](#)
- plnormTrunc (LognormalTrunc), [504](#)

- plnormTruncAlt (LognormalTruncAlt), 507
- plot, 430, 433, 456, 526, 528, 529, 531–533, 535, 536, 539–541, 543–545, 548, 549, 720, 732
- Plot CDF (cdfPlot), 60
- Plot Cumulative Distribution (cdfPlot), 60
- Plot PDF (pdfPlot), 523
- Plot Probability Density (pdfPlot), 523
- Plot Probability Distributions (FcnsByCatPlotProbDists), 410
- Plot probability distributions, 8
- plot.bboxcox, 27, 32, 406, 416, 526, 674
- plot.bboxcoxCensored, 39, 43, 403, 416, 529, 675
- plot.bboxcoxLm, 27, 46, 406, 416, 533, 676
- plot.default, 525, 732–734
- plot.gof, 409, 416, 430, 449, 536
- plot.gofCensored, 405, 416
- plot.gofGroup, 409, 416, 433, 439, 541
- plot.gofTwoSample, 409, 416, 456, 545
- plot.permutationTest, 416
- plot.window, 732
- plotAovDesign, 14, 15, 17, 412, 550
- plotCiBinomDesign, 67, 68, 73, 74, 104, 411, 553, 557
- plotCiNormDesign, 78, 81, 82, 411, 559
- plotCiNormHalfWidth, 562
- plotCiNparDesign, 411
- plotLinearTrendTestDesign, 412
- plotPredIntLnormAltSimultaneousTestPowerCurve, 413
- plotPredIntLnormAltTestPowerCurve, 413
- plotPredIntNormDesign, 413
- plotPredIntNormSimultaneousTestPowerCurve, 413
- plotPredIntNormTestPowerCurve, 413
- plotPredIntNparDesign, 413, 648
- plotPredIntNparSimultaneousDesign, 414
- plotPredIntNparSimultaneousTestPowerCurve, 414
- plotPropTestDesign, 412, 682, 688
- Plotting Probability Distributions, 402, 418
- Plotting Probability Distributions (FcnsByCatPlotProbDists), 410
- plotTolIntNormDesign, 414
- plotTolIntNparDesign, 414
- plotTTestDesign, 412
- plotTTestLnormAltDesign, 412
- pnormMix (NormalMix), 514
- pnormTrunc (NormalTrunc), 516
- points, 113, 706, 732
- pointwise, 90, 403, 471, 563, 578
- pois, 417
- Poisson, 95, 97, 306, 314, 366, 407, 408, 670, 786
- Poisson distribution, 48, 304, 306, 307, 309, 363, 365, 665, 666, 783–785
- Poisson random variable, 364, 785
- Power and Sample Size, 415, 418
- Power and Sample Size (FcnsByCatPower), 411
- power and sample size, 8
- Power and Sample Size Calculations, 402, 409
- Power and Sample Size Calculations (FcnsByCatPower), 411
- ppareto (Pareto), 521
- ppoints, 110, 115, 575, 701
- ppointsCensored, 113–115, 187, 189, 190, 202, 225, 229, 405, 567, 706, 708, 709
- predict, 563–565, 581
- predict.lm, 403, 578
- Prediction Intervals, 92, 377, 379, 402, 406, 670, 776
- Prediction Intervals (FcnsByCatPredInts), 414
- prediction intervals, simultaneous prediction intervals, 8
- predIntGamma, 129, 415, 583, 753
- predIntGammaAlt, 415
- predIntGammaAlt (predIntGamma), 583
- predIntGammaAltSimultaneous, 415
- predIntGammaAltSimultaneous (predIntGammaSimultaneous), 589
- predIntGammaSimultaneous, 415, 589
- predIntLnorm, 415, 598, 617, 625, 760
- predIntLnormAlt, 415
- predIntLnormAlt (predIntLnorm), 598
- predIntLnormAltSimultaneous, 415, 601
- predIntLnormAltSimultaneous (predIntLnormSimultaneous), 606
- predIntLnormAltSimultaneousTestPower, 413, 610
- predIntLnormAltTestPower, 413
- predIntLnormSimultaneous, 415, 601, 606, 631, 642
- predIntNorm, 63, 64, 377, 415, 584, 586, 595, 599, 601, 610, 614, 620, 622, 625, 627, 628, 631, 637, 639, 642, 667, 668, 766, 776
- predIntNormHalfWidth, 412

- predIntNormK, [412](#), [415](#), [599](#), [601](#), [614](#), [615](#),
[617](#), [620](#), [642](#), [775](#)
- predIntNormN, [412](#)
- predIntNormSimultaneous, [415](#), [591](#), [595](#),
[608](#), [610](#), [617](#), [625](#), [626](#), [636–638](#),
[642](#)
- predIntNormSimultaneousK, [415](#), [628](#), [631](#),
[636](#)
- predIntNormSimultaneousTestPower, [413](#),
[595](#), [610](#), [631](#), [639](#), [642](#)
- predIntNormTestPower, [413](#)
- predIntNpar, [406](#), [415](#), [644](#), [654](#), [655](#), [658](#),
[780](#)
- predIntNparConfLevel, [413](#), [648](#)
- predIntNparN, [413](#), [648](#)
- predIntNparSimultaneous, [406](#), [415](#), [652](#)
- predIntNparSimultaneousConfLevel, [414](#),
[658](#)
- predIntNparSimultaneousN, [414](#), [658](#)
- predIntNparSimultaneousTestPower, [414](#),
[658](#)
- predIntPois, [415](#), [664](#), [785](#)
- print, [32](#), [43](#), [46](#), [378](#), [382](#), [430](#), [433](#), [456](#),
[673–678](#)
- print.boxcox, [27](#), [32](#), [406](#), [416](#), [529](#), [673](#)
- print.boxcoxCensored, [39](#), [43](#), [403](#), [416](#),
[532](#), [674](#)
- print.boxcoxLm, [27](#), [46](#), [406](#), [416](#), [535](#), [675](#)
- print.estimate, [416](#)
- print.estimateCensored, [404](#), [416](#)
- print.gof, [409](#), [416](#), [430](#), [449](#), [540](#), [676](#)
- print.gofCensored, [405](#), [416](#)
- print.gofGroup, [409](#), [416](#), [433](#), [439](#), [544](#), [677](#)
- print.gofTwoSample, [409](#), [416](#), [456](#), [549](#),
[678](#)
- print.htest, [416](#)
- print.htestCensored, [405](#), [416](#)
- print.intervalEstimate, [416](#)
- print.intervalEstimateCensored, [416](#)
- print.permutationTest, [416](#)
- print.summaryFull, [416](#)
- print.summaryStats, [416](#), [738](#), [739](#), [744](#),
[746](#)
- Printing and Plotting Methods, [402](#)
- Printing and Plotting Methods
(FcnsByCatPrintPlot), [415](#)
- Probability Density (pdfPlot), [523](#)
- Probability Distributions
(FcnsByCatProbDists), [417](#)
- probability distributions, [8](#)
- Probability Distributions and Random
Numbers, [64](#), [389](#), [402](#), [421](#), [427](#),
[497](#), [500](#), [502](#), [504](#), [506](#), [509](#), [516](#),
[518](#), [523](#), [721](#), [726](#), [790](#), [798](#), [801](#),
[804](#)
- Probability Distributions and Random
Numbers (FcnsByCatProbDists),
[417](#)
- probability plots, [214](#)
- probability-weighted moment, [466](#)
- prop.test, [66](#), [68](#), [74](#), [102](#), [104](#), [558](#), [681](#),
[682](#), [687](#), [688](#), [744](#)
- propTestMdd, [412](#), [682](#), [688](#)
- propTestN, [412](#), [679](#), [688](#)
- propTestPower, [412](#), [682](#), [685](#)
- ptri (Triangular), [788](#)
- pwMoment, [117](#), [118](#), [151](#), [418](#), [491](#), [494](#), [691](#)
- pzmlnorm (ZeroModifiedLognormal), [796](#)
- pzmlnormAlt (ZeroModifiedLognormalAlt),
[799](#)
- pzmnorm (ZeroModifiedNormal), [802](#)
- qbeta, [93](#), [315](#)
- qbinom, [317](#)
- qchi (Chi), [63](#)
- qemp, [716](#), [724](#)
- qemp (Empirical), [212](#)
- qevd, [321](#)
- qevd (EVD), [388](#)
- qexp, [323](#)
- qgamma, [326](#), [420](#)
- qgammaAlt (GammaAlt), [419](#)
- qgeom, [331](#)
- qgev, [333](#)
- qgev (GEVD), [426](#)
- qhyper, [336](#)
- qlnorm, [495](#), [498](#)
- qlnorm3, [343](#)
- qlnorm3 (Lognormal3), [494](#)
- qlnormAlt (LognormalAlt), [498](#)
- qlnormMix (LognormalMix), [500](#)
- qlnormMixAlt (LognormalMixAlt), [502](#)
- qlnormTrunc (LognormalTrunc), [504](#)
- qlnormTruncAlt (LognormalTruncAlt), [507](#)
- qlogis, [345](#)
- qnb, [347](#)
- qnorm, [350](#)
- qnormMix (NormalMix), [514](#)
- qnormTrunc (NormalTrunc), [516](#)
- qpareto, [362](#)
- qpareto (Pareto), [521](#)
- qpois, [364](#)
- qqnorm, [695](#), [701](#)
- qqPlot, [57](#), [59](#), [109](#), [110](#), [114](#), [115](#), [228](#), [303](#),
[411](#), [439](#), [447](#), [449](#), [527–530](#), [534](#),

- [535, 538, 540, 543, 544, 547, 549, 575, 695, 706, 709, 712, 714](#)
- [qqplot, 215, 695, 700](#)
- [qqPlotCensored, 114, 115, 243, 247, 372, 375, 396, 400, 405, 530, 532, 574, 575, 701, 703, 797, 800, 803](#)
- [qqPlotGestalt, 411, 701, 709, 711](#)
- [qtri \(Triangular\), 788](#)
- [quantile, 136, 144, 194, 215, 234, 245, 312, 355, 358, 418, 458, 461, 473, 739](#)
- [quantile-quantile plot, 25, 36](#)
- [quantile-quantile plots, 48](#)
- [quantileTest, 409](#)
- [quantileTestPValue, 409](#)
- [qunif, 368](#)
- [qweibull, 369](#)
- [qzmlnorm, 372](#)
- [qzmlnorm \(ZeroModifiedLognormal\), 796](#)
- [qzmlnormAlt, 372](#)
- [qzmlnormAlt \(ZeroModifiedLognormalAlt\), 799](#)
- [qzmnorm, 374](#)
- [qzmnorm \(ZeroModifiedNormal\), 802](#)
-
- [range, 418, 739](#)
- [Rank Sum, 95, 98](#)
- [rbeta, 93](#)
- [rchi \(Chi\), 63](#)
- [Refinery.CO.df, 715](#)
- [remp, 724](#)
- [remp \(Empirical\), 212](#)
- [reshape, 510, 511](#)
- [revd \(EVD\), 388](#)
- [rgamma, 420](#)
- [rgammaAlt \(GammaAlt\), 419](#)
- [rgevd \(GEVD\), 426](#)
- [rlnorm, 495, 498, 724](#)
- [rlnorm3 \(Lognormal3\), 494](#)
- [rlnormAlt \(LognormalAlt\), 498](#)
- [rlnormMix \(LognormalMix\), 500](#)
- [rlnormMixAlt \(LognormalMixAlt\), 502](#)
- [rlnormTrunc \(LognormalTrunc\), 504](#)
- [rlnormTruncAlt \(LognormalTruncAlt\), 507](#)
- [rnormMix \(NormalMix\), 514](#)
- [rnormTrunc \(NormalTrunc\), 516](#)
- [royston.skew,, 97](#)
- [rpareto \(Pareto\), 521](#)
- [rtri \(Triangular\), 788](#)
- [rzmlnorm \(ZeroModifiedLognormal\), 796](#)
- [rzmlnormAlt \(ZeroModifiedLognormalAlt\), 799](#)
- [rzmnorm \(ZeroModifiedNormal\), 802](#)
-
- [SafePrediction, 581](#)
- [sample, 214, 215](#)
- [sample kurtosis \(kurtosis\), 485](#)
- [sd, 86, 474, 488, 730, 739](#)
- [serialCorrelationTest, 409](#)
- [set.seed, 202, 225, 458, 461, 717, 721, 724, 726, 732](#)
- [shapiro.test, 449](#)
- [signTest, 409](#)
- [simulateMvMatrix, 215, 410, 716, 726](#)
- [simulateVector, 215, 410, 719, 721, 723](#)
- [Skagit.NH3_N.df, 727](#)
- [Skew \(skewness\), 728](#)
- [skew \(skewness\), 728](#)
- [Skewness \(skewness\), 728](#)
- [skewness, 86, 418, 488, 494, 728, 738, 739](#)
- [standard normal distribution, 136, 144, 192, 233, 244, 311](#)
- [stripChart, 418, 731](#)
- [stripchart, 418, 731, 732, 735](#)
- [Student's t, 95, 97](#)
- [Student's t-distribution, 136, 144, 192, 193, 233, 244, 311](#)
- [Student's t-distribution, 118, 167, 177, 210, 220, 231, 356, 395, 399, 622, 623, 668](#)
- [summary, 418, 738, 740, 746](#)
- [Summary Plots \(FcnsByCatSumStats\), 418](#)
- [summary plots, 7](#)
- [Summary Statistics, 86, 402, 474, 488, 730](#)
- [Summary Statistics \(FcnsByCatSumStats\), 418](#)
- [summary statistics, 7](#)
- [summaryFull, 86, 418, 423, 425, 474, 488, 730, 736, 746](#)
- [summaryStats, 418, 739, 740, 742](#)
- [summaryStats.object, 739, 746](#)
- [survfit, 574, 575](#)
-
- [t, 417](#)
- [t.test, 78, 82, 562, 734, 735, 744–746](#)
- [TDist, 774](#)
- [the lognormal distribution, 505](#)
- [the normal distribution, 517](#)
- [Three Parameter Lognormal \(Lognormal3\), 494](#)
- [Three-, 94, 97](#)
- [Three-Parameter Lognormal, 407, 408, 437, 443, 444](#)
- [three-parameter lognormal distribution, 161, 162, 342](#)
- [three-parameter lognormal distributions, 538](#)

- title, [732](#)
- Tolerance Intervals, [358, 377, 379, 402, 406, 760, 766, 776, 781, 786](#)
- Tolerance Intervals (FcnsByCatTolInts), [418](#)
- tolerance intervals, [8](#)
- tolIntGamma, [129, 328, 419, 586, 595, 750](#)
- tolIntGammaAlt, [419](#)
- tolIntGammaAlt (tolIntGamma), [750](#)
- tolIntLnorm, [419, 601, 610, 758, 766](#)
- tolIntLnormAlt, [419, 601](#)
- tolIntLnormAlt (tolIntLnorm), [758](#)
- tolIntLnormCensored, [406](#)
- tolIntNorm, [350, 351, 377, 419, 615, 617, 622, 625, 628, 631, 638, 642, 751, 753, 759, 760, 763, 769, 770, 773, 776](#)
- tolIntNormCensored, [381, 406, 461, 462, 768](#)
- tolIntNormHalfWidth, [414](#)
- tolIntNormK, [414, 419, 764, 766, 772](#)
- tolIntNormN, [414](#)
- tolIntNpar, [357, 358, 406, 419, 646, 658, 777](#)
- tolIntNparConfLevel, [414](#)
- tolIntNparCoverage, [414](#)
- tolIntNparN, [414, 781](#)
- tolIntPois, [365, 419, 670, 783](#)
- Total.P.df, [787](#)
- Trend Analysis, [402](#)
- Trend Analysis (FcnsByCatTrend), [419](#)
- tri, [417](#)
- Triangular, [95, 97, 788](#)
- Truncated, [94, 95, 97](#)
- tTestLnormAltN, [412](#)
- tTestLnormAltPower, [412](#)
- tTestLnormAltRatioOfMeans, [412](#)
- tTestN, [412](#)
- tTestPower, [412](#)
- tTestScaledMdd, [412](#)
- two-parameter lognormal distribution, [495, 496](#)
- twoSampleLinearRankTest, [410](#)
- twoSampleLinearRankTestCensored, [405](#)
- twoSamplePermutationTestLocation, [410](#)
- twoSamplePermutationTestProportion, [410](#)
- Type I (Gumbel) extreme value distribution, [152](#)
- Type I Extreme Value (Gumbel) distribution, [693](#)
- Type I extreme value (Gumbel) distribution, [370, 385, 392, 427, 538](#)
- Type I, also called the Gumbel extreme value distribution or simply Gumbel distribution, [120, 321](#)
- unif, [417](#)
- Uniform, [91–93, 95, 97, 99, 316, 368, 387, 407, 408, 790](#)
- uniform (0,1), [698, 699](#)
- Uniform [0,1] distribution, [438, 543](#)
- uniform distribution, [367, 368, 385, 386](#)
- uniroot, [72, 74](#)
- Value, [94, 96](#)
- var, [86, 418, 474, 488, 730](#)
- var.test, [792, 794, 795](#)
- varGroupTest, [410, 790, 795](#)
- varTest, [410, 792, 794](#)
- Weibull, [95, 97, 370, 392, 407, 408](#)
- weibull, [417](#)
- Weibull distribution, [120, 321, 369, 370, 389–391](#)
- Weibull random variable, [120, 321, 389](#)
- wilcox, [417](#)
- wilcox.test, [67, 73, 557, 734, 735, 744, 746](#)
- Wilcoxon, [95, 97](#)
- Zero Modified Lognormal (ZeroModifiedLognormal), [796](#)
- Zero Modified Normal (ZeroModifiedNormal), [802](#)
- Zero-Modified, [95, 98](#)
- Zero-Modified Lognormal, [373, 397, 801, 804](#)
- Zero-Modified Lognormal (ZeroModifiedLognormal), [796](#)
- Zero-Modified Lognormal (Alternative Parameterization), [373, 798](#)
- Zero-Modified Lognormal (Delta), [407, 408, 437, 443–445, 449](#)
- zero-modified lognormal (delta), [375, 400](#)
- zero-modified lognormal (delta) distribution, [803](#)
- zero-modified lognormal distribution, [371, 393, 394](#)
- zero-modified lognormal distribution (alternative parameterization), [371, 393, 394](#)
- Zero-Modified Normal, [373, 397, 407, 408, 437, 443–445, 449, 798](#)

Zero-Modified Normal
 (ZeroModifiedNormal), 802
zero-modified normal, 375, 400
zero-modified normal distribution, 372,
 374, 375, 396, 398, 400, 797, 800
zero.skew, 97
ZeroModifiedLognormal, 375, 400, 796
ZeroModifiedLognormalAlt, 799
ZeroModifiedNormal, 375, 400, 802
zmlnorm, 417
zmlnormAlt, 417
zmnorm, 418
zTestGevdShape, 154, 155, 334, 410, 427, 466