1. 36 points. Polio vaccine effectiveness

   (a) 3 pts. **prospective binomial**

   (b) 3 pts. Treatments are assigned to kids, then the response is measured (1 year later)

   (c) 5 pts. **0.028% = 0.00028 se = 0.000037 = 3.7e-5**
   Notes: Calculated as 56/200000 and $\sqrt{0.00028 * (1 - 0.00028)/200000}$
   Some folks divided by 400,000 not 200,000. The only individuals informing the requested proportion are the 200,000 in the polio vaccine group.

   (d) 5 pts. **(0.000207, 0.000353)**, computed as $0.00028 \pm 1.960 \times 0.000037 =$
   Notes: The quantile you want is the 0.975 quantile of the normal (Z) distribution
   We were looking for the correct choice of quantile. Full credit if you used 1.96 even if the estimate or standard error were wrong.

   (e) 5 pts. **99**
   Note: Computed as 198*200000 / 400000

   (f) 5 pts. **2.54**
   Note: Computed from the odds of polio in the placebo group = 142/199858 = 0.000710 and the odds of polio in the vaccinated group = 56/199944 = 0.00028. The odds ratio = 0.000710 / 0.00028 = 2.54

   (g) 5 pts. **Z = 5.90 (5.8983)**
   Note: This is easiest done by inference on the log odds ratio.
   Computed as: log odds ratio $= \log 2.54 = 0.9309$, se $\log or = \sqrt{1/56 + 1/199944 + 1/142 + 1/199858} =$ 0.1578.
   $Z = (0.9398 - \log 1)/0.1578 = 5.8983$

   (h) 5 pts. **No, "natural control" is not randomly assigned**
   Note: Comparing vaccine to natural controls changes this to an observational study. The natural controls could differ in important ways from the students in the experimental group.

2. 31 pts. Leaf blight infection vs light levels

   (a) 3 pts. **light, infect. No evidence of lack of fit or unequal variance**
   Note: Plot is flat (no evidence of lack of fit) and residuals have equal variance (same vertical spread in each vertical group of observations). Other 3 models show lack of fit and those with log infect have unequal variance (right-most group is **less** variable than the others).

   (b) 5 pts.
   intercept: **Mean % infection at 0 light level**
   slope: **Change in average % infection when light level increased by 1 unit**
   Note: A causal claim about the slope is appropriate because this is an experimental study.

   (c) 5 pts. The regression model assumes that the means for each unique X value fall along the straight line. The ANOVA makes no such assumption - each group is allowed to have its own unique mean.

   (d) 6 pts.

| Source | d.f. | SS | MS |
|---|---|---|---|
| Difference | 3 | 1.17 | 0.39 |
| Full model | 70 | 59.28 | 0.847 |
| Reduced model | 73 | 60.45 | |

Notes: The reduced model is the regression. A number of folks used the intercept only (equal means) model. Points were given for using the correct two models and doing the rest of the calculations correctly.

(e) 3 pts. **0.50**
Note: Computed as 15.2 - 5.25*2.8

(f) 3 pts. **Yes, 2.8 is larger than any light level in the study**. Using the regression to predict at light = 2.8 assumes that the linear regression continues to values well beyond the range of the data.
Notes: A number of folks suggested there were no concerns because the relationship is a straight line. But, you only know that for light intensities in the range of the data. If I had predicted # flowers at light = 3, you would have calculated $-0.55$ flowers, which is clearly impossible.

(g) 3 pts. **0.98%**

(h) 3 pts. The sd for a predicted value is larger than the error sd (0.91%).
Note: Quite a few folks didn't notice that the question asked about an individual plant (or didn't appreciate the difference between predicting the mean and predicting an individual observation). If you calculate the sd for predicting an individual plant, it's $\sqrt{0.37^2 + 0.91^2} = 0.982$

3. 33 pts.

(a) 1 pt. Which packet are you using? Answers will vary - all are correct

(b) 5 pts. **log hardness** The residual plot is closer to a flat fat sausage
Note: The residual plot for hardness clearly shows evidence of increasing error variance

(c) 5 pts. **p = 0.055. Weak evidence that at least one mean is different.**
Note: Key part of the conclusion is “at least one mean”.

(d) 6 pts.

| Treatment | coefficients for contrast: | | |
|---|---|---|---|
| | # 1 | #2 | #3 |
| Control | 1 | -1/3 | -1 |
| Low dose | 0 | -1/3 | -0.5 |
| High dose | 0 | -1/3 | 1.5 |
| Heat | -1 | 1 | 0 |

Note: Contrast 3 is calculated from the radiation doses in the 3 radiation treatments (0, 0.5, 2.5 kGray). Their average is 1.0, so the coefficients, $X_i - \overline{X}$, are -1, -0.5, and 1.5.

(e) 3 pts. **Contrast # 2 will have a smaller se than # 1**. The average of 3 treatments (part of contrast # 2) is more precise than a single mean.

(f) 5 pts. Answer depends on which model was chosen. Best is the log hardness analysis
Heat multiplies the mean hardness by **1.49** units (95% CI: **0.98**, **2.26**).
If you chose the hardness analysis, your answer would be:
Heat increases the mean hardness by **8.79** units (95% CI: **0.675**, **16.90**).
Note: The computer code has the contrast coefficients as 1, 0, 0, -1. You want -1, 0, 0, 1 because the wording is in terms of the increase or multiplicative effect associated with heat.

(g) 5 pts. **3. Do the test using a multiple comparisons adjustment. Tukey adjustment**. The comparison was chosen because of the data; it was not an a-priori question.
Note: Some folks explained the choice of Tukey and didn't explain why some form on multiple comparisons adjustment is needed. That got partial credit.

(h) 3 pts. **3. something like 0.0030**. The p-value adjusted for multiple comparisons is larger (less significant) than the unadjusted value