Lab 9 for Math 17: Hypothesis Testing III

1 Hypothesis Testing for μ_d or $\mu_1 - \mu_2$ in Rcmdr

1.1 Data Formats

Certain functions will only work when data is supplied in a certain form. For example, when comparing two populations via two samples, you may need to have the data in one of these 2 common formats: two columns - one containing the data from sample 1 and the second with the data from sample 2 OR two columns - one containing the variable of interest and the second containing an indicator variable for whether that value is from sample 1 or sample 2. From here on, we will call having group data in separate columns Format 1 and the second format will be denoted Format 2. "Stacking" is when you start from Format 1 and convert to Format 2, since you basically "stack" the group data and create a new variable to denote which group it came from originally. To stack group data, follow these steps:
1. Make sure the variables you want to stack are in Format 1.
2. Under the Data menu, select Active data set.
3. In the submenu, select Stack variables in active data set.
4. A new window will open. In that window, select the variables you want to stack (select more than one by holding down ctrl as you click).
5. Now, set the name for a new data set where the stacked variable will be created, as well as a name for the variable and for the groups (this is the factor name: Type, Group, or just leaving it as factor are acceptable, but you might have something more specific, like gender or species).
6. Click Ok to perform the stacking.
7. The stacked variable data set will now be your active data set. You can swap back to the original data set by using select Active data set from Active data set in the Data menu if you need it.

1.2 Two Means

For the 2-sample t-test (test for 2 means), the data needs to be in Format 2. Note that you probably start off with the data in Format 1 for checking assumptions then stack it into Format 2.
To perform a two-sample t-test, under the Statistics menu, select Means and then select Independent samples t-test. In the window that opens, select the variable of interest as the response variable and the group variable (first box). Then, you need to set several options: whether your test is 2-sided (first option) or one-sided to a certain direction (second and third options), what confidence level you want (for CIs), and whether or not to assume equal variances (usually this is not assumed). When you have selected all options and both variables, click Ok.

1.3 Paired Mean

For the paired t-test, the data needs to be in Format 1 (example: one column for the "before", one column for the "after" values), or you need to have already computed the differences and can run a one-sample t-test on them. You will need to compute the differences anyway in order to check some assumptions.
To perform a paired t-test, under the Statistics menu, select Means and then select Paired t-test. In the window that opens, you will need to select the two variables. The test is done on the first variable minus the second (so you may prefer a certain order of subtraction). After setting that

up, you need to set several options: whether your test is 2-sided (first option) or one-sided to a certain direction (second and third options), and what confidence level you want (for CIs). When you have selected all options and both variables in the order you want, click Ok.

2 Advice to the Elves

Recall back during our first labs when we were trying to help the Keebler elves determine the number of chocolate chips in their cookies? Well, now we want to see if there is significant evidence that the competition (Chips Ahoy) has more chips than Keebler brand on average.

a. Discuss an appropriate design for this comparison. Will we be able to implement the optimal design?

b. What hypotheses do the elves want to test? Be sure to define your parameters.

c. Collect data with the class. (I will post it quickly).

d. List appropriate assumptions for the selected test. Do the assumptions appear valid?

e. Complete the test (even if assumptions don't look valid). Draw a picture of the p-value and label the distribution used to compute it.

f. What conclusion do you reach? What is your advice to the Keebler elves? Are assumption violations a concern? Any other concerns?

g. If the elves were interested in the magnitude of the difference (or lack thereof) in average chip numbers between the two brands, what estimates would you give them?

3 NO_x emissions

The NO_x emissions (in grams per kilowatt-hour) from 10 different engines were collected when both a baseline and reformulated gasoline were used (data from Journal of the Air and Water Management Association). The age of the engine was also recorded. The association in charge of the study is interested in whether or not there is a difference between the baseline and reformulated mean gasoline emissions as well as the magnitude of the difference (if one exists). The data are online in noxemissions.txt for your use in R/Rcmdr.

a. Is this data set an example of a paired data set or 2 independent samples?

b. To determine whether or not there is a difference between the emissions for the two types of gasoline is it more appropriate to use a hypothesis test or a confidence interval?

c. Perform a preliminary analysis of the baseline and reformulated gasoline data. What do you see?

d. Compute the differences (refer to the online hypothesis testing help if needed). Perform a preliminary analysis of the differences. What do you see?

e. Perform a hypothesis test to determine if there is a difference between the mean gasoline emissions for the two types of gasoline. Be sure to follow all steps and report your conclusion.

f. Provide an interpretation of your test statistic.

g. The association was also interested in the magnitude of the difference (if one exists). Is it more appropriate to address this question with a confidence interval or hypothesis test?

h. Obtain a 98 percent confidence interval for the difference between the mean NO_x emissions. (You do not need to comment on assumptions again).

i. Interpret your confidence interval and confidence level in context of the problem.

j. What is the margin of error for your confidence interval?

4 To Turn In

Thinking about your favorite sport, brainstorm research questions that if you had appropriately collected data would lead you to perform:
a. a confidence interval for a population mean or paired mean
b. a t-test for a difference in two population means
c. a confidence interval for proportions (state if your question results in a one or two proportion CI)