

PH142 Fall 2019 Final Examination

Academic Integrity Pledge

I, _____, affirm that I will not plagiarize answers from those around me or seek to communicate with others either verbally or non-verbally about the examination.

- Students continuing to write after the examination has finished will receive a grade of zero.
- Students must keep their student IDs on their desks and must show them to the professor or GSI when submitting their test.

1. [1 point] The only additional information needed to run a one sample z-test instead of a one sample t-test is the population standard deviation.

- True
- False

2. [1 point] Fill in the bubble(s) next to the appropriate null hypothesis for a chi-square test for goodness of fit. More than one may be applicable.

- The distribution of weights of the population is the same as in a 1970 study.
- Heart disease is independent of anger levels.
- The mean difference in a subject's weight under Diet 1 or Diet 2 is 0.
- $p_1 = p_2 \dots = p_k$

3. [2 points] Circle the TWO most helpful figures for determining if my sample is approximately normally distributed (no credit if more than two answers are selected):

- Histogram
- Pie Chart
- Boxplot
- Q-Q Plot

4. [1 point] Permutation test and bootstrap sampling require the data to come from a simple random sample with a large sample size.

- True
- False

5. [1 point] The statistic from the Chi-square test for independence follows the Chi-square distribution with $k - 1$ degrees of freedom where k is the number of categories.

- True
- False

6. [1 point] One condition for running a two sample t-test for the difference in population means is that the standard deviations of both populations are the same.

- True
- False

7. [1 point] To put bounds on our best guess of the mean response y for a given x value from a line of best fit we calculate a prediction interval.

- True
- False

8. [1 point] If a paired t-test has 30 observations in one group and 30 observations in the second group, then the number of degrees of freedom for running the paired t-test is 59.

- True
- False

9. [1 point] ANOVA's test statistic follows the F distribution with $N_{total} - 1$ in the numerator, and $N_{total} - k_{group}$ in the denominator.

- True
- False

10. [1 point] Using the same random sample from a population, you compute both a 95% confidence interval and a 99% confidence interval for a population proportion. The 99% confidence interval will always be wider than the 95% confidence interval.

- True
- False

11. [1 point] The underlying distribution of a population has to be Normal for the sampling distribution of the mean to become Normal as sample size becomes larger.

- True
- False

12. [1 point] We can get a p-value of 0 from a permutation test.

- True
- False

13. [1 point] This code will produce the correct p-value from the chi-square distribution for a specified `test_stat` and `degf`:

```
pchisq(q = test_stat, df = degf)
```

- True
- False

14. [1 point] You want to find the appropriate sample size to obtain a confidence interval with your desired confidence level and margin of error. To do this, you have to input a guessed value of \hat{p} , which we label p^* . If we have no best guess for our sample proportion, what is the most conservative guess you can make for p^* ?

- $p^* = 0$
- $p^* = 0.5$
- $p^* = 1$
- It depends on the desired confidence level, z^*
- It depends on the desired margin of error, m

15. [1 point] Which of the following are true about Tukey's HSD test?

- It maintains a 5% *experimentwise* or "*family*" error rate no matter how many tests you conduct.
- It overcomes the issue of multiple testing.
- It compares all possible pairs of means.
- All of the above.

Question 16 [3 points total]

It has been claimed that beetroot juice reduces blood pressure. Scientists investigate this question using three different experimental designs. For each design, fill in the circle next to the most appropriate test.

- a) [1 point] A group of randomly sampled individuals are randomly assigned to two groups. One group is given beetroot juice everyday for 4 weeks; the other group is given a red drink everyday for 4 weeks. After 4 weeks, the blood pressure for both groups is measured.

- one-sample proportion
- one-sample t-test
- ANOVA
- paired t-test
- independent two-sample t-test

- b) [1 point] The blood pressure of a group of randomly sampled individuals was measured before the study. Then every day for four weeks they were given beetroot juice and their blood pressure was measured again after the study, the difference was analyzed.

- one-sample proportion
- one-sample t-test
- ANOVA
- paired t-test
- independent two-sample t-test

- c) [1 point] A group of randomly sampled individuals with known high blood pressure was given beetroot juice everyday for 4 weeks. After 4 weeks their blood pressure was measured. The average blood pressure at the end of the study is compared with 140/90mmHg (which is the level where high blood pressure is diagnosed).

- one-sample proportion
- one-sample t-test
- ANOVA
- paired t-test
- independent two-sample t-test

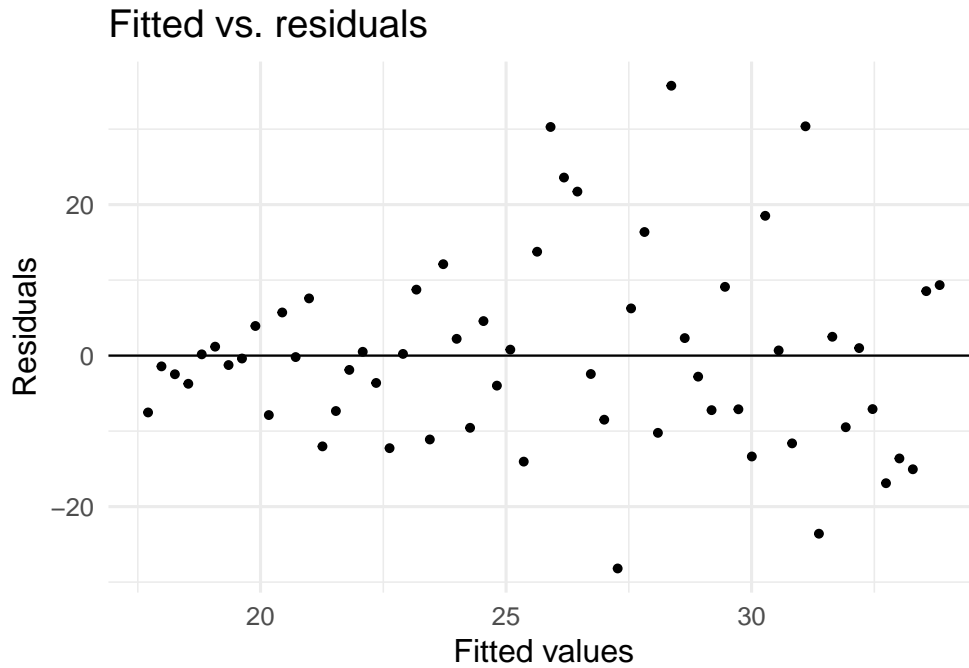
17. [1 point] Fill in the blanks with the type of variable. ANOVA test involves one _____ explanatory variable, and one _____ response variable.

Question 18 [5 points total]

Consider the line of best fit between x and y .

- a) What do a 95% confidence interval for the predicted value for a given x and the 95% prediction interval for an individual value given the same x have in common? Which interval is wider? Why? [3 points]

- b) [2 points] Suppose you were presented with the following plot based on a linear regression. What is the pattern in the plot known as? Which assumption of linear regression is most clearly violated from the plot?



Question 19 [5 points total]

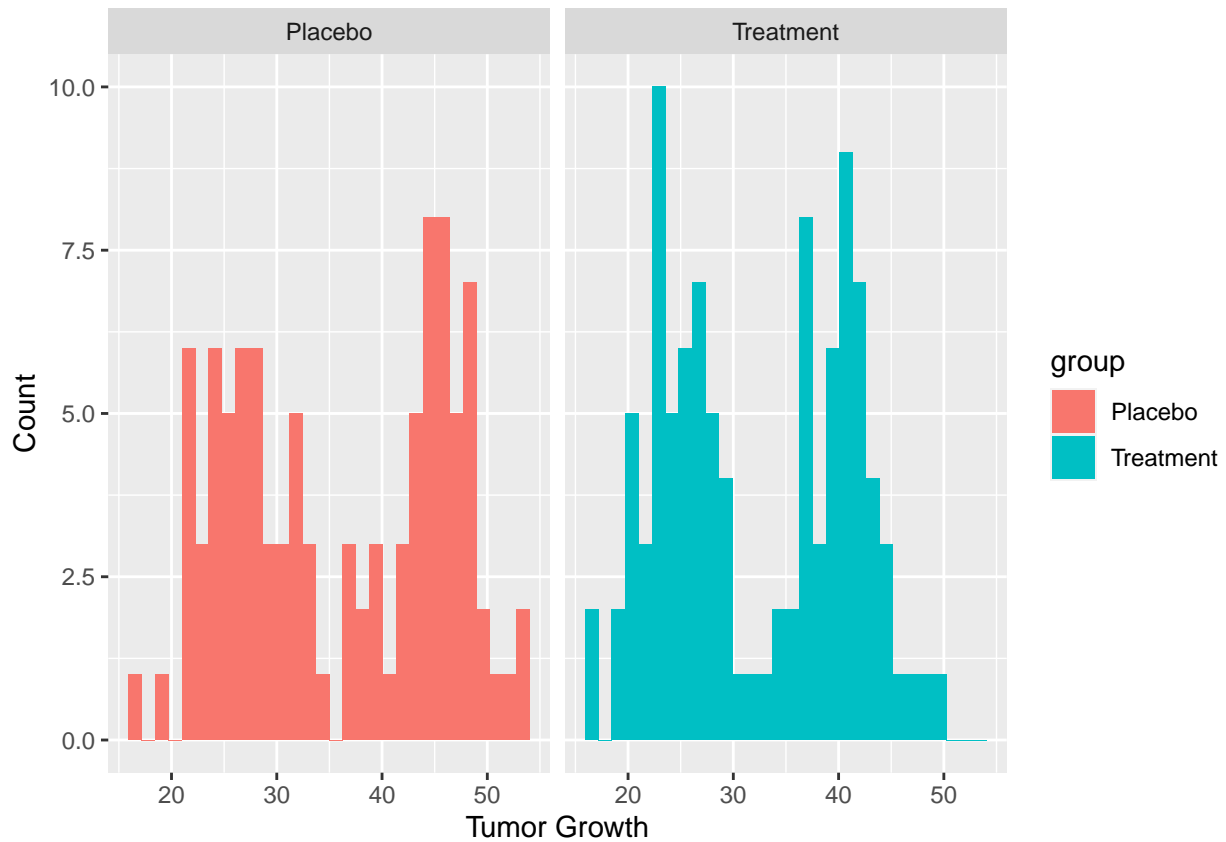
A researcher wants to assess the effects of watching news during mealtime on amount of food intake. A random sample of 50 individuals was assigned to the study: on day one, they had to eat a pasta meal while watching a 30-minute news segment; on day two, they had to eat the same pasta meal while watching nothing. Pasta consumption (in grams) during dinner on each day was collected.

- a) [1 point] The p-value that is obtained from running an independent, two-sample t-test on this data will most likely be _____ than that of the paired t-test p-value.
- equal
 - smaller than
 - greater than
 - cannot say without additional information
- b) [1 point] In one sentence, justify your above response.
- c) [1 point] This study design can be improved. Provide one change that can be made in order to improve the design. (Do not undo the pairing in the design.)
- d) [1 point] Suppose the mean difference in food consumption between days one and two was -30g and the standard deviation of the differences is 43 g. Calculate the appropriate test statistic. Round your answer to two decimal places.
- e) [1 point] Write the R code to calculate the two-sided p-value using the test statistic you calculated above.

Question 20 [9.5 points total]

Pharmaceutical company JoeSchmovartis has developed a drug to slow lung tumor growth. In order to demonstrate the drug's efficacy to the FDA, they run a randomized trial on 200 lung cancer patients, randomly assigning 100 to take a placebo and 100 to test the new drug. Growth in the tumor (in mms) is recorded after two months.

Below is a plot of tumor growth for the placebo group and the treatment group:



a) [1 point] The distribution of tumor growth is:

- Approximately normal
- Skewed Right
- Skewed Left
- Bimodal

- b) [1 point] Fill in the blank. Even though we have evidence that the population distribution of tumor growth is not normal, since our sample sizes are large, we apply the Central Limit Theorem to conclude the _____ for our two sample t-test is approximately normal.

We proceed with a two sample t-test for difference of group means to see if the mean tumor growth for the placebo group is less than the mean tumor growth for patients on the drug.

The sample means and standard deviations for both groups are reported below, as well as the degrees of freedom:

```
## 'summarise()' ungrouping output (override with '.groups' argument)
```

```
## # A tibble: 2 x 3
##   group      sample_mean sample_sd
##   <chr>          <dbl>     <dbl>
## 1 Placebo         36.0      10.0
## 2 Treatment       32.2       8.67
```

```
## There are 193.9652 degrees of freedom
```

- c) [1.5 points] Write the R code to calculate the critical value for a 99% confidence interval for the difference in group means.

The critical value you calculated in c) is:

```
## [1] 2.6
```

- d) [3 points] Compute a 99% confidence interval for the difference between the two means (treatment - placebo)

[Hint for this part and future parts: the confidence interval does not include 0]

e) [1 point] Would you reject a two-sided $\alpha = 1\%$ null hypothesis that the difference between the two means are the same? Why or why not?

f) [1 point] Based on your confidence interval, place an upper or lower bound on the p-value of the t-statistic you would calculate if you carried out the test suggested in 10(f).

BONUS [1 point] True or False. Based on these results, I know that I would reject a one-sided test that the treatment mean tumor growth is **less** than the placebo mean tumor growth.

- True
- False

Question 21 [13 points total]

Suppose that you're attending your maternal and child health seminar and learned about differences in birth weight among babies born to mothers of different race/ethnicities. In your internship, you have access to a random sample of births at the UCSF Benioff Children's Hospital in Oakland and want to know if babies born at this hospital also exhibit these differences by race/ethnicity.

You first run these commands on your random sample of births that you stored in the data frame called `birth_data`:

```
dim(birth_data)
head(birth_data)
str(birth_data)
```

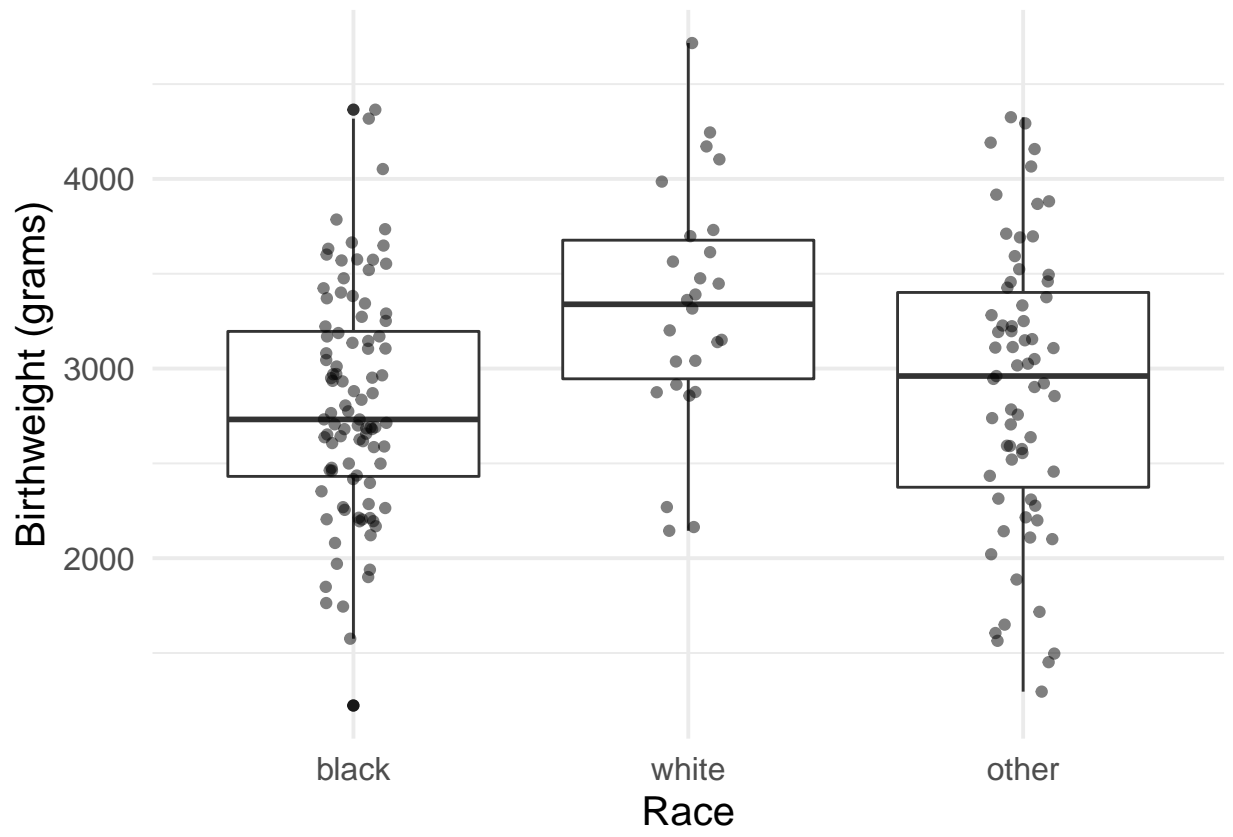
You then use `dplyr` to examine characteristics of a random sample of births at the hospital:

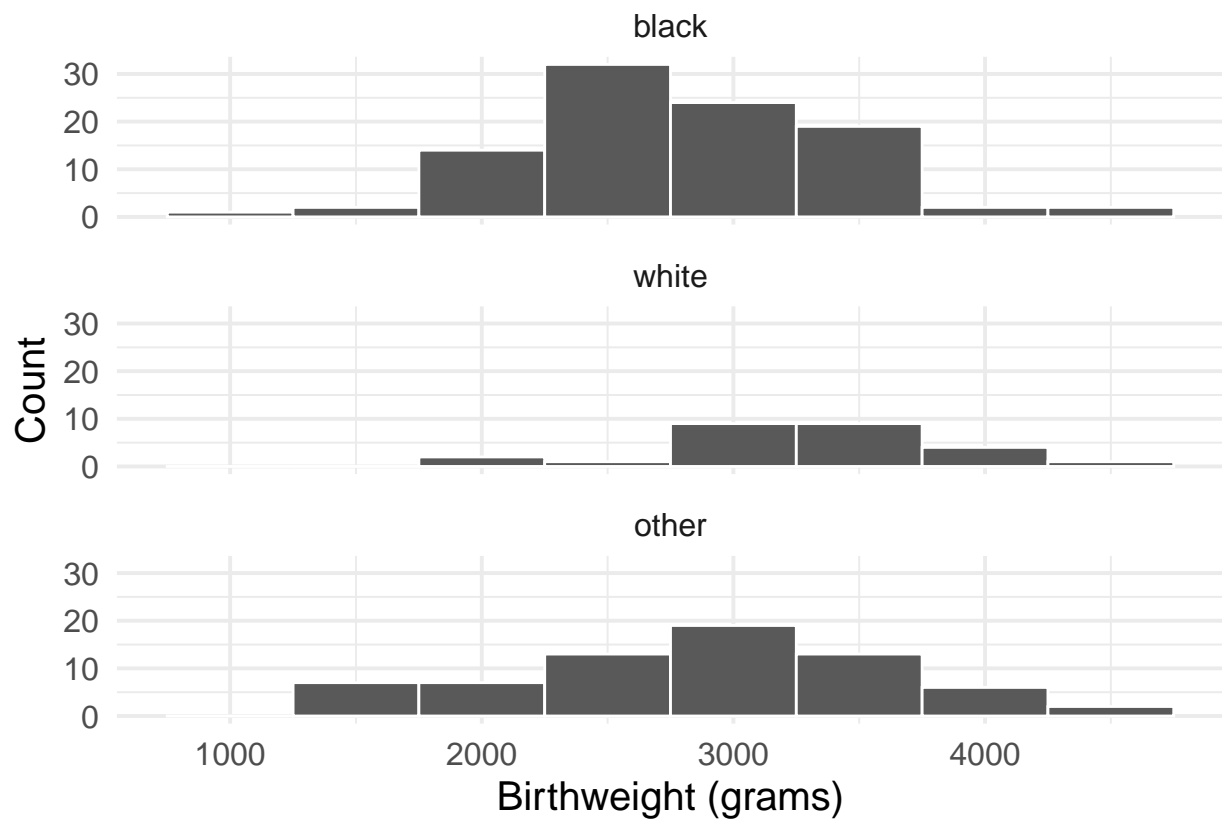
```
birth_data %>%
  group_by(race_order) %>%
  summarize(n = n(),
            mean = round(mean(birthweight), 1),
            sd = round(sd(birthweight), 1))
```

```
## 'summarise()' ungrouping output (override with '.groups' argument)
```

```
## # A tibble: 3 x 4
##   race_order     n mean   sd
##   <fct>         <int> <dbl> <dbl>
## 1 black           96 2813  594.
## 2 white           26 3327. 630.
## 3 other           67 2893. 747.
```

You also make the following descriptive plots to visually explore the relationship between race/ethnicity and birthweight.





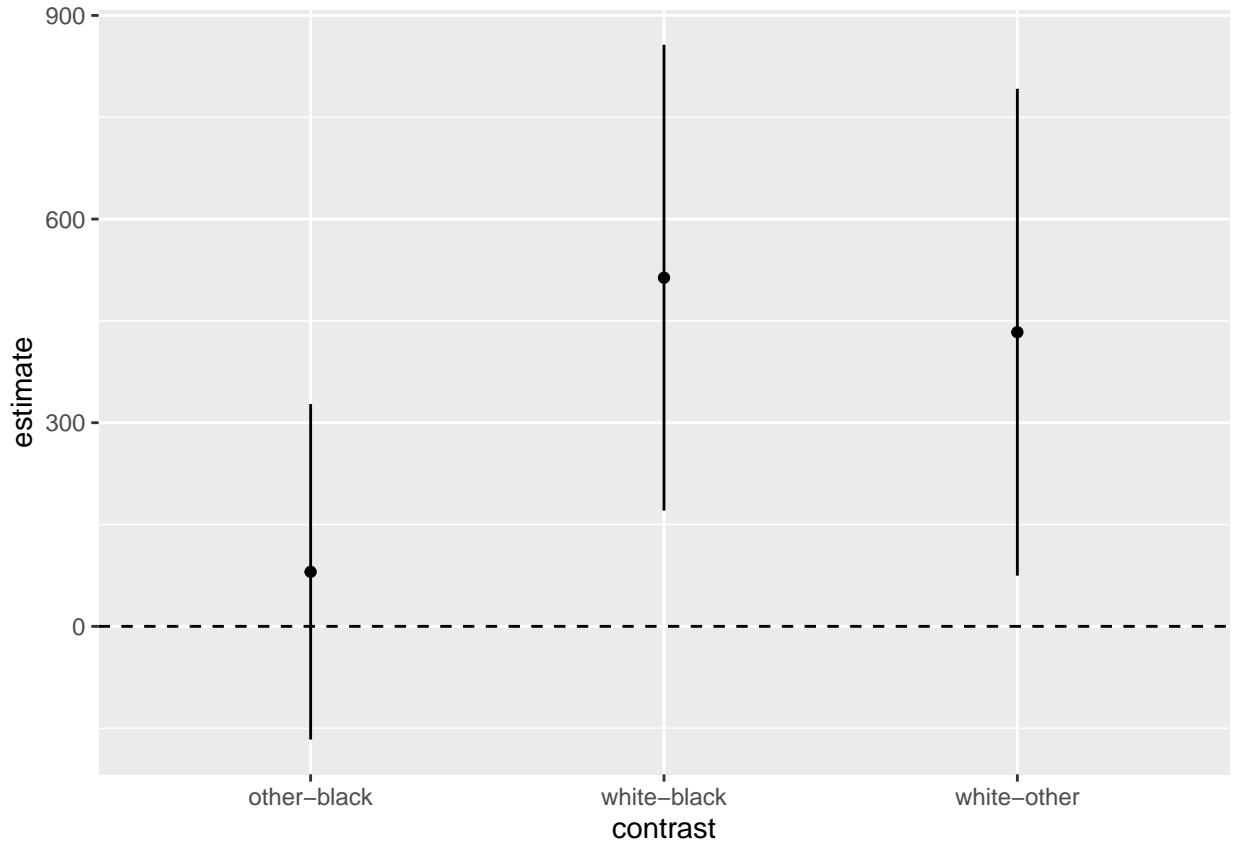
a) [3 points] Write the code required to produce the output below. Load any package that may be required.

```
## # A tibble: 2 x 6
##   term      df      sumsq  meansq statistic  p.value
##   <chr>    <dbl>    <dbl>    <dbl>    <dbl>    <dbl>
## 1 race      NA    5434359. 2717180.      6.30 0.00226
## 2 Residuals NA    80277279. 431598.      NA    NA
```

- b) [2 points] The above output corresponded to a hypothesis test with a specific null and alternative hypothesis. Write the null and alternative hypotheses.
- c) [3 points] Based on the table above, report the value of the test statistic and state its distribution. What are the degrees of freedom for this distribution? (Hint: some of the earlier report information will be helpful!)
- d) [1 point] Interpret the p-value corresponding to the test statistic.

e) [2 points] Because the p-value is so small, you decide to run the following code and make the following plot:

```
## # A tibble: 3 x 7
##   term contrast    null.value estimate conf.low conf.high adj.p.value
##   <chr> <chr>          <dbl>     <dbl>   <dbl>   <dbl>   <dbl>
## 1 race  other-black      0      80.3   -167.   327.    0.723
## 2 race  white-black      0     514.    170.   857.    0.00148
## 3 race  white-other       0     433.    74.7   792.    0.0133
```



Based on the results, comment on which groups appear the most different.

f) [2 points] In the table produced from the TukeyHSD code, `adj.p.value` stands for adjusted p-value. What is it adjusted for? Will it be higher or lower than an unadjusted p-value?

Question 22 [10 points total]

In a 2015 study examining the use of tobacco among students in California, students were asked about their tobacco usage and the number of school days they missed in the past thirty days. Students were not asked why they missed school. We are interested in looking at the relationship between missing school and tobacco usage. The two-way table below summarizes the findings of the study.

	Currently Using	Not Currently Using	Total
0 days	1,421	14,364	15,785
1-5 days	2,863	17,584	20,447
6+ days	1,346	3,678	5,024
Total	5,630	35,626	41,256

- a) [2 points] State which test you would like to conduct followed by your null and alternative hypotheses (in words) in the context of this question.

Test:

Null hypothesis:

Alternative hypothesis:

- b) [1 point] Restate your hypotheses using probability notation.

Null:

Alternative:

- c) [3 points] Calculate the expected values for this test by creating a table with 6 cells and filling it in with your 6 calculated values.
- d) [2 points] Write the formula for calculating the test statistic. From your values in part c), write the first term (based on the upper left cell of your table) of the test statistic. You do not need to calculate the test statistic completely.
- e) [1 point] The test statistic equals 1027.769. Write a line of code that you could use to calculate your p-value. Make sure to specify your degrees of freedom.
- f) [1 point] Running your code, R returns a p-value for this test of 6.649627e-224. Interpret the p-value.

Question 23 [5 points total]

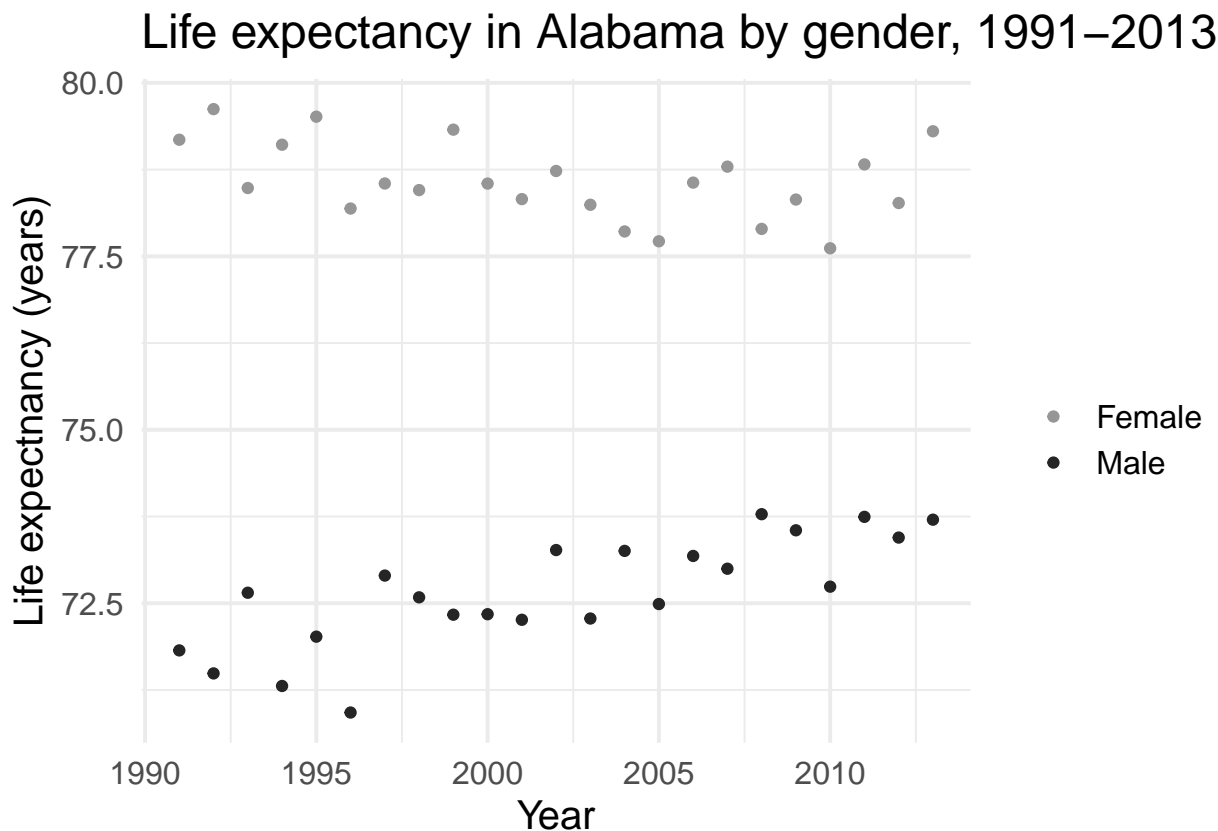
Two cities are concerned that wheezing may be related to pollution from traffic. The residents of each city had their wheezing symptoms compared after a traffic bypass was constructed in one city in 2018 to remove congestion. Data was collected after the bypass was constructed in order to assess the impact of reduction in air pollution on wheezing. Residents reported wheezing symptoms experienced in 2017 and 2018. In the city that had a bypass constructed, 45 out of 282 people reported an improvement in wheezing symptoms. In the city that did not have a bypass constructed, 21 out of 162 people reported an improvement. Assume these represent random samples.

- a) Perform a two-sided test (but not a chi-square test) to conclude whether the bypass city has a statistically different proportion of improvement in wheezing symptoms relative to the city that did not get a bypass. Assume the significance level is 5%. Make sure to:
- State your hypotheses,
 - Name and calculate your test statistic, and
 - Write the code to find the p-value (you do not have to calculate the actual p-value)

- b) Assume the p-value comes out to be 0.393. Interpret this value. Is this evidence consistent with the null hypothesis or in favor of the alternative?

Question 25 [7 points total]

Suppose that you had data on life expectancy for white men and women in Alabama between 1990 and 2013, where such data was based on a random sample of Alabama white residents across those years. You first decided to graph these data as shown below:



You decide to run two models, one for women and the other for men, looking at the relationship between year and life expectancy. Here are the model outputs:

Model for women:

```
## # A tibble: 2 x 5
##   term      estimate std.error statistic  p.value
##   <chr>      <dbl>    <dbl>    <dbl>    <dbl>
## 1 (Intercept) 149.      32.8      4.55 0.000174
## 2 year      -0.0354   0.0164   -2.16 0.0427
```

Model for men:

```
## # A tibble: 2 x 5
##   term      estimate std.error statistic  p.value
##   <chr>      <dbl>    <dbl>    <dbl>    <dbl>
## 1 (Intercept) -115.     31.1     -3.68 0.00138
## 2 year         0.0935   0.0155    6.02 0.00000565
```

- a) [3 points] You are interested in knowing whether life expectancy changed linearly over the time period examined for both men and women. Fill in the following table.

Group	Estimate of the slope	Test statistic	p-value
Women			
Men			

- b) [1 point] For which group (men, women, or both) do you have evidence against the null hypothesis in favor of the alternative? (assume a 5% significance level).
- c) [3 points] Interpret the slope coefficient for men in the context of the question. Create a 95% confidence interval for the slope coefficient for men (Hint: the output provided below will be helpful).

[1] 2.079614

- d) [BONUS point] The intercept estimate is -115 for men. What does this mean?