

**WINTER 2019/20 TERM 2 EPSE 592: ASSIGNMENT 4**

**Due: Apr. 23rd**

- Please make sure you write your answers to these questions in your own words. Even if you work with a group to formulate your responses, do not just copy someone else's sentences/words.
- There is no need to record more than 3 decimal places for any of these problems.
- All problem data are available online in .csv format.
- All R code is available online in plain text format. Copy and paste *\*exactly\** as recorded into the Rj editor of Jamovi. If Jamovi does not work on your computer: install base R from

<https://www.r-project.org/>

Then open R, copy and paste code *\*exactly\** as recorded, and hit 'enter'. **In either platform, the code may take 10-30 seconds to compile.**

**Question 1: (Analysis of categorical data with omnibus and post hoc chi-squared statistics)**

In the following sequence of questions, you will explore different applications of chi-squared tests for categorical data, and also investigate some of the problems with their *misapplication*.

- (a) Suppose you conduct an informal poll among undergraduate students on campus about their degree goals. You are interested in determining if students from different faculties exhibit different desires to pursue graduate degrees. For each of 406 students polled, you ask them (1) what the faculty of their current degree program is, and (2) how strongly they currently want to pursue graduate studies within one year of completing their undergraduate degree. Respondents indicate the strength of their desire to pursue graduate work on a Likert scale of 1 (little or no desire) to 5 (absolutely want to pursue graduate work). You collect the following responses (data available on webpage):

|                    |   | Faculty |         |           |                 |          |
|--------------------|---|---------|---------|-----------|-----------------|----------|
|                    |   | Arts    | Science | Education | Applied Science | Business |
| Likert<br>response | 1 | 9       | 18      | 22        | 25              | 27       |
|                    | 2 | 13      | 14      | 21        | 33              | 17       |
|                    | 3 | 10      | 15      | 19        | 13              | 18       |
|                    | 4 | 20      | 11      | 11        | 17              | 7        |
|                    | 5 | 17      | 13      | 9         | 15              | 12       |

Conduct a  $\chi^2$  test to see if there is evidence that students from different faculties express different desires to pursue graduate work. Report the p-value and interpret.

- (b) The  $\chi^2$  test above will indicate that evidence exists for a difference in desire to go to graduate school across faculties. But which faculties are different and which are alike? Before addressing this question, just how many different pairwise comparisons (across faculties) are possible?
- (c) Rather than running that many  $\chi^2$  tests pairwise between faculties, examine the *proportions* of the 5 different Likert responses for each faculty. In Jamovi, you can do this quickly by clicking on the 'Cells' dialogue box within your  $\chi^2$  analysis, then selecting the 'Column'

option under ‘Percentages’. This will give you proportions within each cell; i.e., the *sample estimate* of the *probability distribution* of Likert responses for each faculty. To help visualize, you can also produce barplots and view the literal *sample frequency distributions* for interest in pursuing graduate studies (quantified by the Likert response) for each faculty. Run the following code in Jamovi’s Rj-Editor to produce these plots (code available on the webpage - copy and paste from the .txt file into Jamovi):

```
mat <- matrix(c(9,22,25,27,18, 13,21,33,17,14, 10,19,13,18,15,
               20,11,17,7,11, 17,9,15,12,13),nrow=5,byrow=T)
par(mfrow=c(2,3))
barplot(mat[,3],names.arg=c("1","2","3","4","5"),
        main="Likert resp for Ap. Science",ylim=c(0,30))
barplot(mat[,1],names.arg=c("1","2","3","4","5"),
        main="Likert resp for Arts",ylim=c(0,30))
barplot(mat[,4],names.arg=c("1","2","3","4","5"),
        main="Likert resp for Business",ylim=c(0,30))
barplot(mat[,2],names.arg=c("1","2","3","4","5"),
        main="Likert resp for Education",ylim=c(0,30))
barplot(mat[,5],names.arg=c("1","2","3","4","5"),
        main="Likert resp for Science",ylim=c(0,30))
```

Examining these sample distributions, which faculty seems to be the *most* different from the rest, in terms of overall distribution of responses?

- (d) For the faculty you identified as, informally, most dislike the others, now formally compare this faculty to the others by using pairwise  $\chi^2$  tests (there will be 4 pairwise tests in total). [The easiest way to do this will be to copy and paste the relevant data cells into new columns for each chi-squared test.] Report the p-values and interpret for each of these tests.
- (e) Recall that when we perform multiple hypothesis tests, we need to (potentially) worry about inflating our type I error (false positive) rates. The analytical procedure you just followed should be reminiscent of an omnibus ANOVA, followed by a sequence of post hoc pairwise comparisons. Recall that we adjust for the multiple testing problem in an ANOVA framework by reporting *adjusted p-values*, usually using either Tukey’s or Scheffé’s method. Here, there are similar adjustments that can be made, but the only way we have learned how to adjust for multiple comparisons by hand is to perform a Bonferonni-type adjustment. Do this here as follows: divide your significance threshold by the number of post hoc tests performed (4 in total), and then reinterpret your p-values from part (d) with respect to this adjusted threshold. Note that this adjustment can change your interpretation *regardless* of if you use the traditional threshold or the orders of magnitude approach to interpreting p-values.

## Question 2: (Where chi-squared tests fail, Fisher’s exact tests prevail)

The research question in Q1 is a natural one, and we had lots of data available to address it. However, in practice, it is *quite* challenging, in terms of money and time spent, to survey 406 people. Instead, it is far more likely that you will find yourself in a data situation as below (data on webpage). Here, we have polled 54 undergraduate students about their desire to pursue graduate studies, using a 5-point Likert survey item.

- (a) Plot the sample frequency distributions for these responses per faculty (use the code on the webpage). Which faculty seems to be the *most* different from the rest, in terms of overall distribution of responses?

|                    |   | Faculty |         |           |                 |          |
|--------------------|---|---------|---------|-----------|-----------------|----------|
|                    |   | Arts    | Science | Education | Applied Science | Business |
| Likert<br>response | 1 | 1       | 4       | 3         | 4               | 4        |
|                    | 2 | 2       | 4       | 3         | 4               | 4        |
|                    | 3 | 0       | 3       | 4         | 2               | 4        |
|                    | 4 | 4       | 0       | 0         | 0               | 0        |
|                    | 5 | 3       | 0       | 0         | 1               | 0        |

- (b) Conduct a  $\chi^2$  test to see if there is evidence that students from different faculties express different desires to pursue graduate work. Report the p-value and interpret.
- (c) Recall one of the critical assumptions of the  $\chi^2$  test: cell counts should be at least 5 in most of the table's cells, and no cells should have a count of 0. Clearly, these data violate this assumption quite severely. Instead, we should be applying a *Fisher's exact test*, which does not rely on any assumptions about sample size (total or cell-wise). Use the code on the webpage to perform a Fisher's exact test on these data. [Note: the p-value for this test will change slightly if you run the code multiple times, since Jamovi only approximates via simulation; like approximating the probability of flipping heads by counting how many heads appear on 20 tosses of a coin.] Report the p-value and interpret. Compare with the inference you made in part (b) from (inappropriately) applying a  $\chi^2$  test.
- (d) Now use the code on the webpage to perform a Fisher's exact test on the large dataset (406 respondents) from Q1. Report the p-value and interpret. Compare with the inference you made in Q1(a) from (appropriately) applying a  $\chi^2$  test.