Assignment #3: ANOVA

1. An ANOVA was conducted to compare multiple group means. Fill-in the 6 blanks in the following table:

Source of variation	**SS**	**df**	**MS**	**MSR (F)**
Treatments (among groups)	_____	3	158.96	_____
Error (within groups)	_____	32	62.81	(blank)
Total variation	_____	_____	(blank)	η^2 = _____

Based on the above ANOVA summary table:

a. There were _____ groups (treatments) in this study.

b. There were _____ total observations (subjects) in this study.

c. The total variance of all observations in this study was equal to _____

2. An ANOVA was conducted to compare multiple group means. Fill-in the 6 blanks in the following table:

Source of variation	**SS**	**df**	**MS**	**MSR (F)**
Treatments (among groups)	_____	7	_____	5.01
Error (within groups)	_____	_____	3.62	(blank)
Total variation	_____	29	(blank)	η^2 = _____

Based on the above ANOVA summary table:

a. There were _____ groups (treatments) in this study.

b. There were _____ total observations (subjects) in this study.

c. The total variance of all observations in this study was equal to _____

Scenario: A sample of 20 different types of cereals was taken from each of three grocery store shelves (the lowest, middle, and highest shelves). A summary of the sugar content (grams per serving) and dietary fiber (grams per serving) of the cereals is given below.

	Sugar content (g per serving)			**Fiber content (g per serving)**		
	n	$\bar{X}$	**s**	**n**	$\bar{X}$	**s**
Lowest shelf	n = 20	4.80	2.138	n = 20	1.68	1.166
Middle shelf	n = 20	9.85	1.985	n = 20	1.32	1.162
Highest shelf	n = 20	6.10	1.865	n = 20	2.11	1.277
All shelves combined	N = 60	M = 6.9167	s = 2.9194	N = 60	M=1.703	s = 1.2263

3. In order to conduct an ANOVA, we need to make some assumptions: independence, normality, & equal variances. Determine whether the equal variance assumption is satisfied for each study (sugar and fiber) by conducting an Fmax test.

   For the sugar study: Fmax = _____. The equal variance assumption **IS** **IS NOT** satisfied.

   For the fiber study: Fmax = _____. The equal variance assumption **IS** **IS NOT** satisfied.

4. We want to determine if the (population) mean sugar levels differ among the 3 shelves. What are the hypotheses?

   Null hypothesis: _____ Alternate hypothesis: _____

5. For the sugar data, fill-in the following ANOVA summary table by calculating SS, df, MS, MSR, and eta-squared. Just this once, try to calculate all of these things by hand (using a calculator; not a statistical program or applet).

Source of variation	**SS**	**df**	**MS**	**MSR (F)**
Shelves	_____	_____	_____	_____
Error	_____	_____	_____	(blank)
Total	_____	_____	MS_{total}	η^2 = _____

Once you have the table completed, go ahead and check your work with http://danielsoper.com/statcalc3/calc.aspx?id=43

6. Calculate MS_{total} for the sugar data (use the ANOVA summary table you just created). The MS_{total} was actually given to you in the scenario. What does MS_{total} represent?

   MS_{total} represents: _____

7. If the null hypothesis were true, how unlikely were we to get the sugar data that we observed in this study? Use an F-distribution table or calculator to estimate the p-value for your MSR.

   Calculator: http://lock5stat.com/statkey/theoretical_distribution/theoretical_distribution.html#F
   F-table: http://bradthiessen.com/html5/stats/m301/ftable.pdf

   Identify the degrees-of-freedom for your F-statistic (MSR): Numerator df = _____ Denominator df = _____

   Approximate p-value: $P(F \geq MSR)$ =_____

8. From this p-value write out any conclusion(s) you can make for this sugar study:

   Conclusions: _____

9. Look at the eta-squared you calculated in the ANOVA summary table. Interpret this eta-squared with regards to this study.

   Interpretation: _____

10. Conduct an ANOVA for the fiber data. Use an online calculator to complete the ANOVA summary table, estimate a p-value, calculate eta-squared, and write out any conclusions you can make.

Source of variation	**SS**	**df**	**MS**	**MSR (F)**
Shelves	_____	_____	_____	_____
Error	_____	_____	_____	p = _____
Total	_____	_____	MS_{total}	η^2 = _____

   Conclusions: _____

Scenario: Suppose we're interested in testing the critical thinking skills of seniors at St. Ambrose. We randomly sample seniors majoring in each of the following 6 groups of programs:

- STEM majors (e.g., math, chemistry, engineering, computer science)
- Health science majors (e.g., biology, nursing, exercise science)
- Social science majors (e.g., psychology, sociology, criminal justice, political science)
- Humanities majors (e.g., English, philosophy, theology, history)
- Creative arts majors
- Education majors

We then administer a test of critical thinking to these students.

11. We want to compare the means of 6 different groups. To do this, we could decide to conduct a series of t-tests to compare all possible pairs of group means. Answer the following:

    a. How many t-tests would we need to conduct to compare all possible pairs of our 6 group means? _____

    b. Suppose we set $\alpha = 0.05$ for each of those t-tests. What would be the probability of making at least one α-error across all our t-tests?

    Answer: _____

    c. Suppose you conduct an ANOVA to compare the 6 group means. What would happen to your F-statistic (MSR) and the p-value under each of these scenarios:

    i. If we sample a larger number of students in each group, the F-statistic will **INCREASE** **DECREASE**

    ii. If we sample a larger number of students in each group, the p-value will **INCREASE** **DECREASE**

    iii. If the population (group) means remain the same but the group standard deviations all increase, the F-statistic will: **INCREASE** **DECREASE**

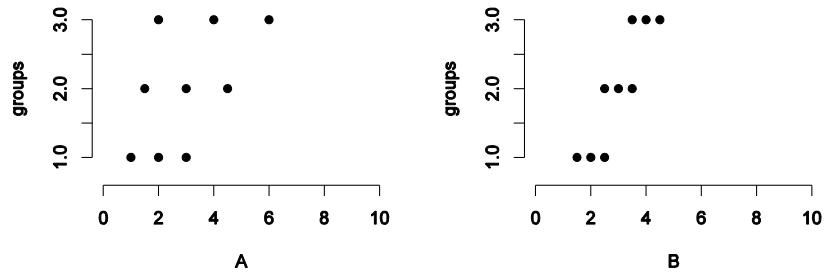

12. Datasets A and B are displayed above. Use that pair of plots to answer the following:

    a. If you conducted an ANOVA for each dataset, which one would have the larger MSA (between groups)?

    Circle your answer: **DATASET A** **DATASET B** **They are equal** **Unable to determine**

    b. If you conducted an ANOVA for each, which dataset would have the larger MSE (within groups)?

    Circle your answer: **DATASET A** **DATASET B** **They are equal** **Unable to determine**

    c. Based on the MSR (F-statistic) you would calculate for each, which would yield the smaller p-value?

    Circle your answer: **DATASET A** **DATASET B** **They are equal** **Unable to determine**

Scenario: Does the amount of time doctors spend with patients depend on whether the patient is obese?

A 2001 study examined physicians' behavioral intentions as well as their expressed attitudes towards average-weight and obese patients. 71 primary care physicians in Houston participated in this study. The doctors were sent a packet containing a medical chart similar to the one they view upon seeing a patient. This chart portrayed a patient who was displaying symptoms of a migraine headache but was otherwise healthy. The weight of the patient was manipulated so that:

- 33 doctors received a chart from a patient of average weight (body mass index = 23)
- 38 doctors received a chart from an obese patient (body mass index = 36)

The doctors were instructed to examine the charts and then asked, among other questions, how much time they believed they would spend with the patient. Here's a summary of that data:

weight	N	mean	p50	sd
average	33	31.36364	30	9.864134
obese	38	24.73684	25	9.652571
Total	71	27.8169	30	10.23762

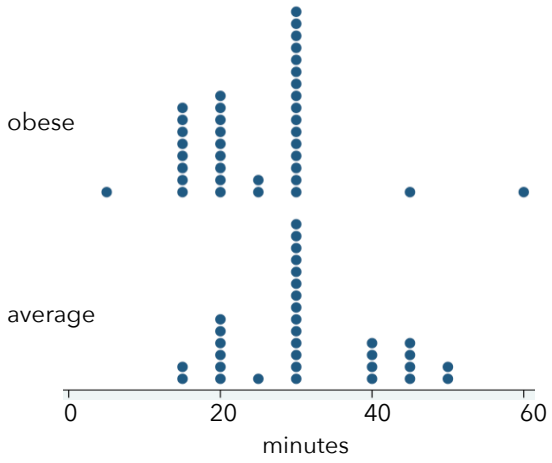

13. If we were to conduct an independent samples t-test (assuming the population variances are equal), we would calculate an observed t-statistic of:

$$t_{33+38-2} = \frac{31.36364 - 24.73684}{\sqrt{\frac{1}{33}+\frac{1}{38}}\sqrt{\frac{(33-1)(9.864134)^2+(38-1)(9.652571)^2}{33+38-2}}} = \frac{6.6268}{(0.23795)(9.75125)} = 2.856$$

If we were running a two-tailed test (i.e., our alternative hypothesis was $\mu_{average} = \mu_{obese}$), we would estimate the p-value as p= 0.0028 x 2 = 0.0057 (as displayed in the t-distribution to the right)

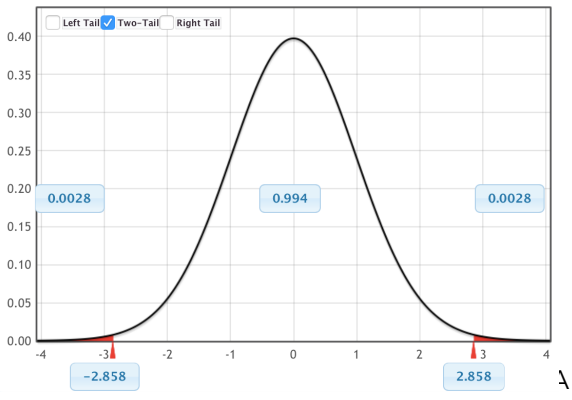

If we conducted an ANOVA on this same data (to compare the two group means), what would we calculate for our F-statistic (our MSR)? What would our p-value be?

If you don't know, you could always type the sample sizes, means, calculator: http://danielsoper.com/statcalc3/calc.aspx?id=43

The F-statistic (MSR) would equal: _____ (look at how that compares to the t-statistic calculated above).

The p-value would be: _____.

Scenario: A 2003 study investigated whether the type of background music played in a restaurant affected the amount of money diners spent on their meals. A restaurant alternated among 3 types of background music -- silence, popular music, and classical music -- over the course of 18 days. Each type of music was played for 6 nights (with the order of the music being randomly determined).

The data from this study are available at: http://www.bradthiessen.com/html5/data/music.csv

Copy that data and paste it into the StatKey applet:
http://lock5stat.com/statkey/advanced_1_quant_1_cat/advanced_1_quant_1_cat.html

You'll know you entered the data correctly if you get the following table at the top-right:

n = 393, F = 27.822

Statistics	Classical	Pop	None	Overall
Sample Size	120	142	131	393
Mean	24.1	21.9	21.7	22.5
Standard Deviation	2.3	2.7	3.4	3.0

14. Click the ANOVA Table button on the top-right of the screen and fill-in the ANOVA summary table below. Calculate eta-squared and estimate the p-value. It might be easiest to estimate the p-value by comparing the MSR to the F-distribution on this page: http://lock5stat.com/statkey/theoretical_distribution/theoretical_distribution.html#F

Source of variation	SS	df	MS	MSR (F)
Music	_____	_____	_____	_____
Error	_____	_____	_____	p-value= _____
Total	_____	_____	MS_{total}	η^2 = _____

15. Now that we've seen the results from our (theoretical) ANOVA, let's try a randomization-based approach. Go to the StatKey applet (where you've pasted your data) and generate at least 10,000 samples. Remember that the applet is randomly swapping the music types for each dollar amount in the dataset (because our null hypothesis is that the music type doesn't matter). Each dot on the distribution represents an F-statistic calculated from each of these samples.

    Look at the F-statistic (MSR) you recorded in the above ANOVA table. Assuming the music type doesn't influence the amount spent by diners, what's the likelihood of observing an F-statistic that or more extreme? Use the applet to estimate this p-value:

    p-value = _____

Scenario: Does having a disability influence how others perceive you perform on a job interview?

To study this, researchers prepared five videotaped job interviews using the same two male actors for each. A set script was designed to reflect an interview with an applicant of average qualifications. The tapes differed only in that the applicant appeared with a different handicap. In one, he appeared in a wheelchair; in a second, he appeared on crutches; in another, his hearing was impaired; in a fourth, he appeared to have one leg amputated; and in the final tape, he appeared to have no handicap.

70 undergraduate students from an American university were randomly assigned to view the tapes, 14 to each tape. After viewing the tape, each subject rated the qualifications of the applicant on a 0-10 point scale. Their ratings were as follows:

	No Disability	Amputee	Crutches	Hearing	Wheelchair
	1.90	1.90	3.70	1.40	1.70
	2.50	2.50	4.00	2.10	2.80
	3.00	2.60	4.30	2.40	3.50
	3.60	3.20	4.30	2.90	4.70
	4.10	3.60	5.10	3.40	4.80
	4.20	3.80	5.80	3.70	5.00
	4.90	4.00	6.00	3.90	5.30
	5.10	4.60	6.20	4.20	6.10
	5.40	5.30	6.30	4.30	6.10
	5.90	5.50	6.40	4.70	6.20
	6.10	5.80	7.40	5.50	6.40
	6.70	5.90	7.40	5.80	7.20
	7.40	6.10	7.50	5.90	7.40
	7.80	7.20	8.50	6.50	7.60
Mean	4.9000	4.4286	5.9124	4.0500	5.3429
StDev	1.7936	1.5857	1.4818	1.5325	1.7483

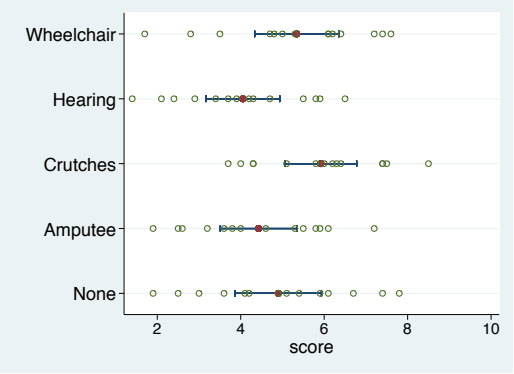

You can download and copy this data at: http://www.bradthiessen.com/html5/data/disability.csv

Source: Cesare, Tannenbaum, Dalessio (1990). *Interviewers' Decisions Related to Applicant Handicap Type and Rater Empathy*. Human Performance, 3(3)

16. Go to http://lock5stat.com/statkey/advanced_1_quant_1_cat/advanced_1_quant_1_cat.html and paste the data. Click the **ANOVA TABLE** button to conduct an ANOVA and record the summary table below. Calculate eta-squared. The p-value has been recorded for you:

Source of variation	**SS**	**df**	**MS**	**MSR (F)**
Music	_____	_____	_____	_____
Error	_____	_____	_____	p-value= 0.0301
Total	_____	_____	MS_{total}	η^2 = _____

17. Generate at least 10,000 samples and estimate the p-value using this randomization-based approach. Record the p-value here:

    p-value = _____

18. Copy the data once again and paste it into a different applet: http://www.rossmanchance.com/applets/AnovaShuffle.htm?hideExtras=2. Click **USE DATA** to get the data into this applet. You should see the distributions shown to the right.

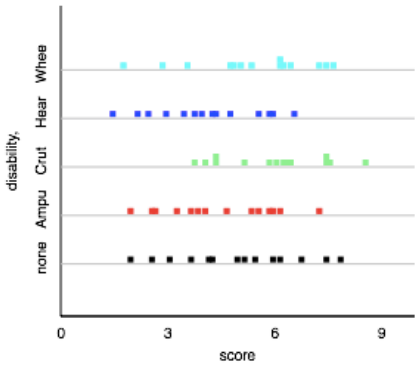

    At the bottom-left of the screen, make sure the **MAD** statistic is selected:

    We can see that for the data in this study, the observed MAD = 0.931.

    What does this value represent?

    The MAD represents: _____

    On the top-right, check the **SHOW SHUFFLE OPTIONS** box. Shuffle the data at least 10,000 times. Again, this is simply shuffling the disability associated with each interview score (because we're assuming disability status does not matter).

    To the right, you should now see a distribution of at least 10,000 MAD statistics. Estimate the p-value by finding the proportion of MAD values as or more extreme than 0.931.

    p-value = _____

    Finally, select a new statistic from the menu at the bottom-left of the screen. Select **MAX-MIN**:

    This statistic is simply the largest group mean (crutches) minus the smallest group mean (hearing impaired). Generate 10,000 shuffles of this statistic and estimate the p-value:

    p-value = _____