Activity 5: AxB (Two-Way) ANOVA

Scenario: To test the effectiveness of a new cholesterol drug, you gather 40 subjects and randomly assign half of them to take the drug and the other half to take a placebo. After 6 months, you measure the cholesterol level of each group. Here's the data you obtain:

Placebo	Drug	Total
n = 20	n = 20	N = 40
mean = 98.0	mean = 88.0	mean = 93.0
std. dev = 8.8	std. dev = 8.4	std. dev = 9.887

You calculate the following independent samples t-test:

$$t_{n_1+n_2-2} = \frac{\text{(Observed) - (Hypothesized)}}{\text{(Standard Error)}} = \frac{\left(\bar{X}_1 - \bar{X}_2\right) - 0}{\sqrt{s^2_{pooled}}\sqrt{\frac{1}{n_1} + \frac{1}{n_2}}} = \frac{(98-88)}{\sqrt{74}\sqrt{\frac{1}{20}+\frac{1}{20}}} = 3.676$$

where:

$$s^2_{pooled} = \frac{(n_1-1)s_1^2 + (n_2-1)s_2^2}{n_1+n_2-2} = \frac{(20-1)8.8^2 + (20-1)8.4^2}{20+20-2} = 74$$

and a critical value for t is found to be: $t^{\alpha=0.025}_{df=38} = 2.024$

1) Notice that we could have also chosen to conduct an ANOVA on this data. Based on the results from the t-test, fill-in the blanks in the following ANOVA summary table. What conclusions can we make from either the t-test or the ANOVA?

Source	SS	df	MS	MSR
Drug	1000	__________	__________	__________
Error	2812	__________	__________	F^{1}_{38} = ______
Total	3812	__________	__________	

2) You publish the results of this ground-breaking study in a prestigious journal and receive the following feedback:

I question the results of your study stating that this new drug lowers cholesterol. It is widely known that individuals who are overweight have significantly higher cholesterol levels than individuals who are not overweight. I believe your "placebo" group had more obese individuals than your "drug" group. The "placebo" group's cholesterol level was higher simply because the group contained many obese individuals.

Uh oh - you didn't weigh the subjects in this study. Is this a legitimate concern?

3) We're going to learn how to conduct an ANOVA when we have two independent variables of interest (drug and weight, in this example) and one dependent variable (cholesterol). This is called an AxB ANOVA (also called a two-way ANOVA or a factorial design). To see how this is going to work, it may help to visualize the data in a 2-dimensional table:

	Placebo (D_1)	**Drug (D_2)**	**Total**
Not obese (W_1)	n_{11} $\bar{X}_{11}$	n_{12} $\bar{X}_{12}$	$n_{1\bullet}$ $\bar{X}_{1\bullet}$
Obese (W_2)	n_{21} $\bar{X}_{21}$	n_{22} $\bar{X}_{22}$	$n_{2\bullet}$ $\bar{X}_{2\bullet}$
Total	$n_{\bullet 1}$ $\bar{X}_{\bullet 1}$	$n_{\bullet 2}$ $\bar{X}_{\bullet 2}$	N M

Before we begin worrying about the concepts and calculations behind an AxB ANOVA, let's consider some possible outcomes from our study. For each of the following scenarios, graph the mean cholesterol levels for the non-obese and obese groups. Also, describe the outcome(s) of the study in each scenario:

<u>Scenario A</u>

	Placebo (D_1)	**Drug (D_2)**	**Total**
Not obese (W_1)	$n_{11} = 10$ $\bar{X}_{11} = 93$	$n_{12} = 10$ $\bar{X}_{12} = 93$	$n_{1\bullet} = 20$ $\bar{X}_{1\bullet} = 93$
Obese (W_2)	$n_{21} = 10$ $\bar{X}_{21} = 93$	$n_{22} = 10$ $\bar{X}_{22} = 93$	$n_{2\bullet} = 20$ $\bar{X}_{2\bullet} = 93$
Total	$n_{\bullet 1} = 20$ $\bar{X}_{\bullet 1} = 93$	$n_{\bullet 2} = 20$ $\bar{X}_{\bullet 2} = 93$	$N = 40$ $M = 93$

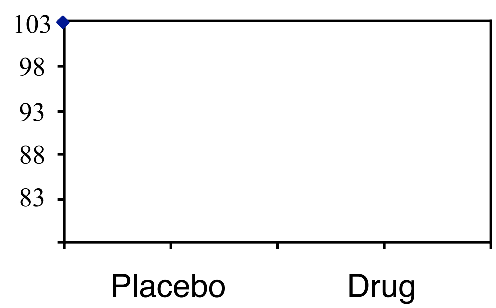

Does the drug impact cholesterol? ______________________________

Does obesity impact cholesterol? ______________________________

<u>Scenario B</u>

	Placebo (D_1)	**Drug (D_2)**	**Total**
Not obese (W_1)	Mean = 93	Mean = 83	Mean = 88
Obese (W_2)	Mean = 103	Mean = 93	Mean = 98
Total	Mean = 98	Mean = 88	Mean = 93

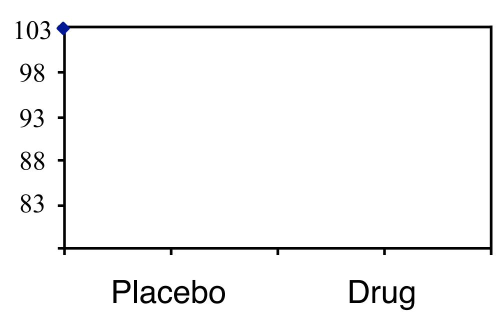

Does the drug impact cholesterol? ______________________________

Does obesity impact cholesterol? ______________________________

Scenario C

	Placebo (D_1)	Drug (D_2)	Total
Not obese (W_1)	Mean = 88	Mean = 83	Mean = 85.5
Obese (W_2)	Mean = 98	Mean = 88	Mean = 93
Total	Mean = 93	Mean = 85.5	Mean=89.25

Does the drug impact cholesterol? ____

Does obesity impact cholesterol? ____

Scenario D

	Placebo (D_1)	Drug (D_2)	Total
Not obese (W_1)	Mean = 88	Mean = 98	Mean = 93
Obese (W_2)	Mean = 98	Mean = 88	Mean = 93
Total	Mean = 93	Mean = 93	Mean = 93

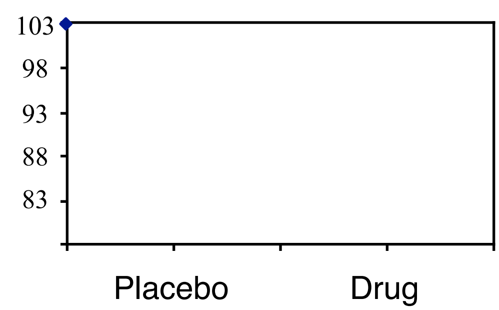

Does the drug impact cholesterol? ____

Does obesity impact cholesterol? ____

In an AxB ANOVA, we will analyze the differences (distances) among the means of the row variables and the differences among the means of the column variables separately. These two sets of differences are called *main effects*. Calculate the main effects (and the sum of those main effects) for each scenario:

Scenario A: Drug main effect = Obesity main effect = Sum =

Scenario B: Drug main effect = Obesity main effect = Sum =

Scenario C: Drug main effect = Obesity main effect = Sum =

Scenario D: Drug main effect = Obesity main effect = Sum =

4) In addition to main effects, an AxB ANOVA allows us to investigate the *interaction* between our two independent variables. Interaction is present when the two independent variables in combination produce effects that are different from the sum of their main effects.

Take Scenario B, for example. We found the drug main effect was equal to 5 points (that's the average distance from the drug and placebo means to the overall mean). We also found the obesity main effect to equal 5 points. The sum of the effects is, thus, 10. Now let's look at the combination of the drug and obesity effects in the study. Suppose we have an average person (with a cholesterol level of 93, the grand mean). If that person becomes non-obese **and** takes the drug, we expect that person to have a cholesterol of 83. So, in combination, the drug and obesity effects equal 10 (which is what we might expect). This is an example of a scenario with **no interaction** -- we can tell because the graph of means shows parallel lines.

Now look again at Scenario D. In this scenario, we found no drug or obesity main effects. However, when you take an average person (with a cholesterol of 93) and that person becomes non-obese **and** takes the drug, we expect that person to have a cholesterol of 98. So even though the drug and obesity effects, taken separately, should not change that person's cholesterol, the combination of those effects does have an impact. This is an example of interaction -- we can tell because the graph of means shows nonparallel lines.

A more subtle form of interaction is found in Scenario C. Here we find main effects of 3.75 (for Drug) and 3.75 (for Obesity). Summed, we would expect a total effect of 7.5 points. Our average person with a cholesterol of 89.25 drops down to 83 (only 6.25 points) with non-obesity and the drug. Once again, the drug and obesity effects in combination differ from their simple sum, and we can see this with the non-parallel lines in our graph.

What does the presence of interaction mean? We'll investigate this more in a bit. For now, we'll just mention:

Interaction means the effect of one independent variable (on the dependent var.) depends on the other independent var.

Procedurally, an AxB ANOVA is similar to the one-way ANOVA we've been studying. In a one-way ANOVA, we partitioned the total variation into two components: a between-groups component and a within-groups (error) component. We then took the ratio of those two estimates of variance and compared it to an F-distribution.

In an AxB ANOVA, we're going to partition the total variation into between-groups, within-groups, and interaction components. Before we do that, let's investigate the formal model for an AxB ANOVA:

Pretend you're an individual in this cholesterol study. What factors could influence your cholesterol score? Which of these factors are of interest to the researcher?

An individual's cholesterol = __________________ + __________________ + __________________ + __________________

We'll write the formal model as: $x_{ijk} = \mu + \alpha_j + \beta_k + \alpha\beta_{jk} + \varepsilon_{ijk}$ where

$$\alpha_j = \mu_{Aj} - \mu$$
$$\beta_k = \mu_{Bk} - \mu$$
$$\alpha\beta_{jk} = \mu_{AjBk} - \mu_{Aj} - \mu_{Bk} + \mu$$

5) The table below displays the actual data from this cholesterol study. As you can see, the first individual in the study had a cholesterol level of 85.08. According to our formal model, why did this individual have this cholesterol level? What are the values of α, β, αβ, and ε?

	Placebo (D_1)		Drug (D_2)		Total
Not obese (W_1)	**85.08**	92.68	80.11	73.14	Overall Non-Obese 20 subjects Avg. Cholesterol = 85.1 St. Deviation = 9.7
	90.60	89.86	95.13	75.59	
	76.85	77.63	**79.20**	67.83	
	90.24	107.21	70.05	91.61	
	91.38	90.24	89.25	87.96	
	Mean = 89.2 StDev = 8.5		Mean = 81.0 StDev = 9.5		
Obese (W_2)	103.13	95.17	92.97	101.78	Overall Obese 20 subjects Avg. Cholesterol = 97.8 St. Deviation = 10.0
	116.50	100.63	81.51	90.67	
	102.87	115.14	86.53	90.39	
	101.10	100.16	83.55	102.52	
	109.19	94.27	104.24	82.86	
	Mean = 103.8 StDev = 7.6		Mean = 91.7 StDev = 8.5		
Total	20 subjects Avg. Cholesterol = 96.5 St. Deviation = 10.8		20 subjects Avg. Cholesterol = 86.3 St. Deviation = 10.4		40 subjects Avg. Cholesterol = 91.4 St. Deviation = 11.7

Individual Score X_{ijk}	Overall Mean μ	Drug effect $\alpha_j = \mu_{Aj} - \mu$	Obesity effect $\beta_k = \mu_{Bk} - \mu$	Interaction $\alpha\beta_{jk} = \mu_{AjBk} - \mu_{Aj} - \mu_{Bk} + \mu$	Error ε
85.08					
79.20					
109.19					

So, all Greek letters aside, our formal model is simply stating our expectations for the factors that may attribute to variation in cholesterol levels.

6) Assuming the model makes sense, let's begin partitioning the variation in cholesterol levels. Explain each step of the derivation, paying particular attention to any assumptions we're making along the way:

We start with this tautology. What makes it true?

$$X_{ijk} = M + \left(\bar{X}_{Aj} - M\right) + \left(\bar{X}_{Bk} - M\right) + \left(\bar{X}_{AjBk} - \bar{X}_{Aj} - \bar{X}_{Bk} + M\right) + \left(X_{ijk} - \bar{X}_{AjBk}\right)$$

$$\sum_{j=1}^{a}\sum_{k=1}^{b}\sum_{i=1}^{n_{ajbk}}\left(X_{ijk} - M\right) = \sum_{j=1}^{a}\sum_{k=1}^{b}\sum_{i=1}^{n_{ajbk}}\left(\bar{X}_{Aj} - M\right) + \sum_{j=1}^{a}\sum_{k=1}^{b}\sum_{i=1}^{n_{ajbk}}\left(\bar{X}_{Bk} - M\right) + \sum_{j=1}^{a}\sum_{k=1}^{b}\sum_{i=1}^{n_{ajbk}}\left(\bar{X}_{AjBk} - \bar{X}_{Aj} - \bar{X}_{Bk} + M\right) + \sum_{j=1}^{a}\sum_{k=1}^{b}\sum_{i=1}^{n_{ajbk}}\left(X_{ijk} - \bar{X}_{AjBk}\right)$$

$$\sum_{j=1}^{a}\sum_{k=1}^{b}\sum_{i=1}^{n_{ajbk}}\left(X_{ijk} - M\right)^2 = \sum_{j=1}^{a}\sum_{k=1}^{b}\sum_{i=1}^{n_{ajbk}}\left(\bar{X}_{Aj} - M\right)^2 + \sum_{j=1}^{a}\sum_{k=1}^{b}\sum_{i=1}^{n_{ajbk}}\left(\bar{X}_{Bk} - M\right)^2 + \sum_{j=1}^{a}\sum_{k=1}^{b}\sum_{i=1}^{n_{ajbk}}\left(\bar{X}_{AjBk} - \bar{X}_{Aj} - \bar{X}_{Bk} + M\right)^2 + \sum_{j=1}^{a}\sum_{k=1}^{b}\sum_{i=1}^{n_{ajbk}}\left(X_{ijk} - \bar{X}_{AjBk}\right)^2$$

$$SS_{\text{Total}} = SS_{\text{A}} + SS_{\text{B}} + SS_{\text{AxB}} + SS_{\text{Error}}$$

Explain what these sums of squares represent.

7) We will convert these SS values into MS values by dividing each component by its degrees of freedom. We'll then take ratios of these MS values to compare them. What ratio of mean squares will we calculate to determine if the drug had a significant effect on cholesterol? What about to determine the significance of the obesity effect? How will we test for interaction? Which of these MS ratios will we want to calculate first? Why?

8) In a visual form, here's how a one-way and an AxB ANOVA partition variance:

one-way ANOVA

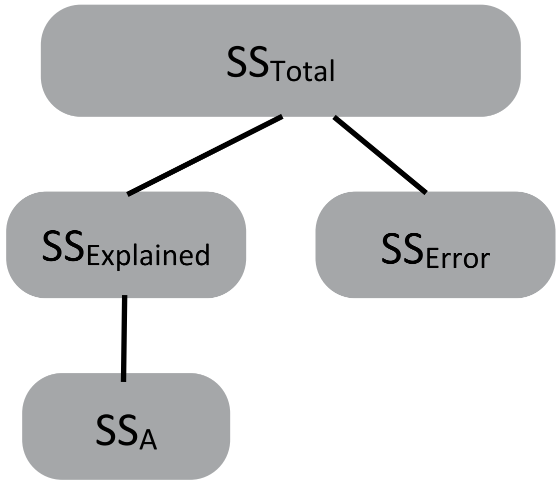

AxB ANOVA

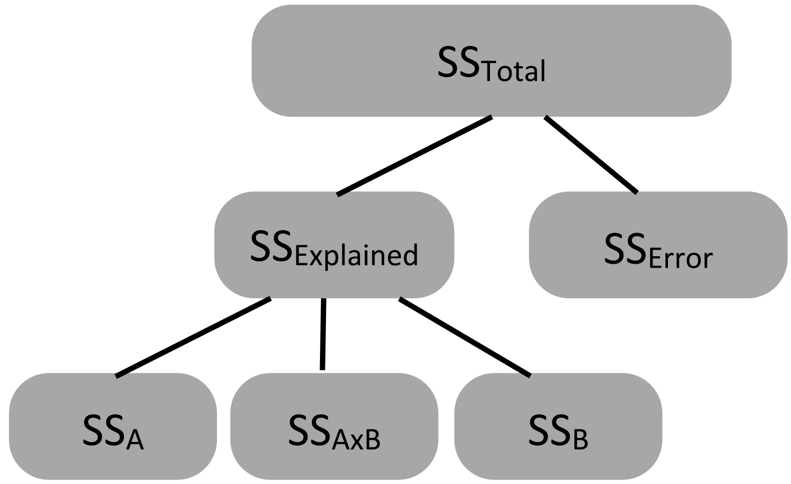

9) Before we get into the formulas to calculate SS, df, and MS values, let's take a quick moment to graph the means in our study. This graph will give us an idea of which effects we should expect to be significant. Based on the graph, do you expect to find a significant drug effect? obesity effect? interaction effect? Explain.

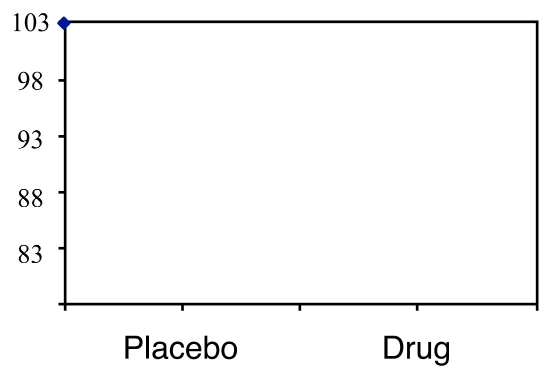

Source	SS	df	MS	MSR
A (Drug)	$\sum n_{Aj}\left(\bar{X}_{Aj} - M\right)^2$	a – 1	$SS\ /\ df$	$MS_A\ /\ MS_E$
B (Obesity)	$\sum n_{Bk}\left(\bar{X}_{Bk} - M\right)^2$	b – 1	$SS\ /\ df$	$MS_B\ /\ MS_E$
AxB (Interaction)	$\sum\sum n_{AjBk}\left(\bar{X}_{AjBk} - \bar{X}_{Aj} - \bar{X}_{Bk} + M\right)^2$	(a – 1)(b – 1)	$SS\ /\ df$	$MS_{AxB}\ /\ MS_E$
Error	$\sum\sum\sum\left(X_{ijk} - \bar{X}_{AjBk}\right)^2 = \sum\left(n_{AjBk} - 1\right)s^2_{AjBk}$	N – ab	$SS\ /\ df$	
Total	$SS_{\text{A}} + SS_{\text{B}} + SS_{\text{AxB}} + SS_{\text{Error}} = \sum\left(X_{ijk} - M\right)^2$	N – 1	(total variance of all N observations)	

10) I had Stata compute everything for me in the following table. Verify the calculations and interpret the output. How would we determine which MSR values are significant? What conclusions can we make from this study? Calculate an effect size.

Source	SS	df	MS	MSR	p-value
A (Drug)	1030.784	1	1030.784	14.063	0.001
B (Obesity)	1606.936	1	1606.936	21.923	<.001
AxB (Interaction)	38.426	1	38.426	0.524	0.474
Error	2638.774	36	73.299		
Total	5314.919	39			

At this point, you should attempt the cow and haircut problems on the assignment

Part B: Interaction effects

Recall our study of the effects of obesity and a new drug on the cholesterol levels of 40 subjects. Ten of the 20 non-obese individuals were given a placebo while the other ten were given an experimental new drug. The 20 obese individuals were also equally assigned to the placebo and treatment groups. After 6 months, the cholesterol levels of individuals in each group were measured. Run an AxB ANOVA on the following data (to simplify calculations, only the W_1D_2 group of cholesterol levels has changed):

	D_1 (placebo)	**D_2 (drug)**	
W_1 (non-obese)	**85.08** 92.68 90.60 89.86 76.85 77.63 90.24 107.21 91.38 90.24 Mean = 89.2 StDev = 8.5	92.11 85.14 107.13 87.59 **91.20** 79.83 82.05 103.61 101.25 99.96 Mean = 93.0 StDev = 9.5	Overall Non-Obese 20 subjects Avg. Cholesterol = 91.1 St. Deviation = 9.0
W_2 (obese)	103.13 95.17 116.50 100.63 102.87 115.14 101.10 100.16 109.19 94.27 Mean = 103.8 StDev = 7.6	96.97 105.78 85.51 94.67 90.53 94.39 87.55 106.52 108.24 86.86 Mean = 95.7 StDev = 8.5	Overall Obese 20 subjects Avg. Cholesterol = 99.8 St. Deviation = 8.9
	Overall Placebo 20 subjects Avg. Cholesterol = 96.5 St. Deviation = 10.8	Overall Drug 20 subjects Avg. Cholesterol = 94.3 St. Deviation = 8.9	Overall Study 40 subjects Avg. Cholesterol = 95.4 St. Deviation = 9.9

1) What assumptions must be met before running an AxB ANOVA? Are these assumptions met?

2) Graph the group means, using separate lines for the obese and non-obese groups. Does it appear as though obesity has a significant impact on cholesterol (a significant row effect)? Does it look as though there is a significant column effect (the drug impacts cholesterol)? Does it look as though we have a significant interaction between the drug and obesity treatments?

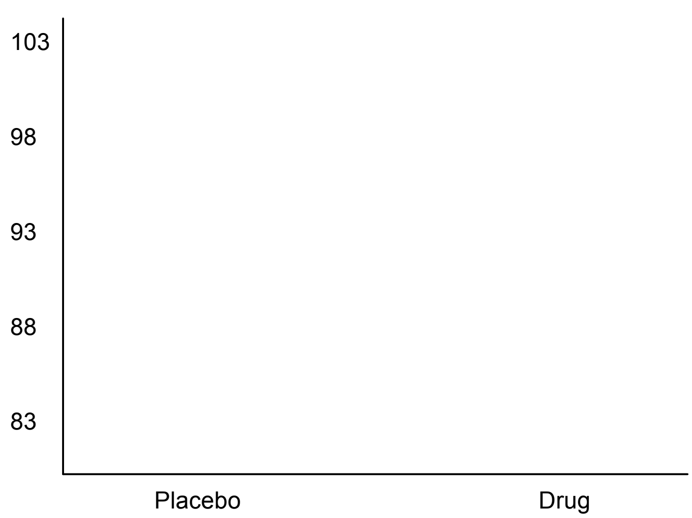

3) Write out the model for an AxB ANOVA. Select a couple cholesterol measurements and calculate the parameters in the model. What impact does being in the placebo, drug, obese, or non-obese groups have? What is the value of the interaction parameter for each group?

$$x_{ijk} = \mu + \alpha_j + \beta_k + \alpha\beta_{jk} + \varepsilon_{ijk}$$

Given	**Overall Mean**	**Alpha**	**Beta**	**Interaction**	**Error**
x_{ijk}	μ	$\alpha_j = (\mu_{Aj} - \mu)$	$\beta_k = (\mu_{Bk} - \mu)$	$\alpha\beta_{jk} = (\mu_{AjBk} - \mu_{Aj} - \mu_{Bk} + \mu)$	

4) Let's run the computations on this AxB ANOVA. The SS values have been filled in (*verify these when you have time*). Fill-in the missing degrees of freedom and calculate the MS values. Finally, calculate the three MSRs.

Source	Sums of Squares	Degrees of freedom	Mean Square	Mean Square Ratio
Drug (A)	46.343			
Obesity (B)	752.816			
Interaction (AB)	355.245			
Error	2638.774			
Total	3793.178			

5) Remember that the total variance in an AxB ANOVA is partitioned into four components. Make a diagram of this partitioning (write each SS value next to its place in the diagram). Explain what all of this means.

6) In this type of analysis, our first step is to always check for significant interaction. Find the critical F-value to compare to your MS-interaction. Is the interaction significant? What does this mean?

7) We now know that the efficacy of the drug depends on whether or not an individual is obese. Our original intent was to determine the overall efficacy of the drug. What additional analyses could we conduct to determine the drug's efficacy?

Once we find significant interaction, we should look at the ***simple effects*** of our variable of interest. To do this, we split our study into two groups (If variable A is important, we split the study into groups B1 and B2; if variable B is important, we split the study into groups A1 and A2).

8) Our variable of interest is "drug," so we need to split our study into two groups: obese and non-obese. The following two tables demonstrate this "split." What kind of analysis can we run on each of these tables?

	D_1 (placebo)	**D_2 (drug)**	
W_1 (non-obese)	n=10 Mean = 89.2 StDev = 8.5	n=10 Mean = 93.0 StDev = 9.5	<u>Overall Non-Obese</u> 20 subjects Avg. Cholesterol = 91.1 St. Deviation = 9.0

	D_1 (placebo)	**D_2 (drug)**	
W_2 (obese)	n=10 Mean = 103.8 StDev = 7.6	n=10 Mean = 95.7 StDev = 8.5	<u>Overall Obese</u> 20 subjects Avg. Cholesterol = 99.8 St. Deviation = 8.9

9) If we're now only comparing 2 group means, we can run a standard t-test. We would use the overall MSE as our best estimate of the pooled variance. If we're comparing 2+ group means, we can conduct a one-way ANOVA on our split data. In this ANOVA, we would use the MSE from the original analysis. Use the formulas for this type of analysis to verify the following calculations. Use the summary tables to draw conclusions from your study.

$$t_{df_E} = \frac{\bar{X}_1 - \bar{X}_2}{\sqrt{MS_E}\sqrt{\frac{1}{n_1} + \frac{1}{n_2}}}$$

Source	Sums of Squares	Degrees of freedom	Mean Square	Mean Square Ratio
SSA Treatment/Factor (Between)	$SS_A = \sum n_a \left(\bar{X}_a - M\right)^2$	$a-1$	$\frac{SS_A}{df_A}$	
SSE Error (Within)	SSE from ANOVA or $\sum_{i=1}^{n} (n_a - 1)s_a^2$	df_E from ANOVA or $N-a$	$\frac{SS_E}{df_E}$	

10) Let's run through the calculations for our example:

	D_1 (placebo)	D_2 (drug)	
W_1 (non-obese)	**85.08** 92.68 90.60 89.86 76.85 77.63 90.24 107.21 91.39 90.24 Mean = 89.2 StDev = 8.5	92.11 85.14 107.13 87.59 **91.20** 79.83 82.05 103.61 101.25 99.96 Mean = 93.0 StDev = 9.5	Overall Non-Obese 20 subjects Avg. Cholesterol = 91.1 St. Deviation = 9.0
W_2 (obese)	103.13 95.17 116.50 100.63 102.87 115.14 101.10 100.16 109.20 94.27 Mean = 103.8 StDev = 7.6	96.97 105.78 85.51 94.67 90.53 94.39 87.55 106.52 108.24 86.86 Mean = 95.7 StDev = 8.5	Overall Obese 20 subjects Avg. Cholesterol = 99.8 St. Deviation = 8.9
	Overall Placebo 20 subjects Avg. Cholesterol = 96.5 St. Deviation = 10.8	Overall Drug 20 subjects Avg. Cholesterol = 94.3 St. Deviation = 8.9	Overall Study 40 subjects Avg. Cholesterol = 95.4 St. Deviation = 9.9

Source	Sums of Squares	Degrees of freedom	Mean Square	Mean Square Ratio
Drug (A)	46.343	1	46.343	0.632
Obesity (B)	752.816	1	752.816	10.27
Interaction (AB)	355.245	1	355.245	4.85
Error	2638.774	36	73.299	
Total	3793.178	39		

Non-Obese Analysis

Obese Analysis

$$t_{36} = \frac{(103.8 - 95.7)}{\sqrt{73.3}\sqrt{\frac{1}{10} + \frac{1}{10}}} = \frac{8.1}{3.829} = 2.11 \sim t_{36}^{\alpha=0.05}(1.688)$$

If we can assume equal variances, we can use SSE and MSE from the AxB ANOVA:

$$SS_A = 10(103.8 - 99.8)^2 + 10(95.7 - 99.8)^2 = 328.1$$

Source	Sums of Squares	Degrees of freedom	Mean Square	Mean Square Ratio
Drug	328.1	1	328.1	4.476
Error	2638.774	36	73.3	
Total	N/A	N/A		

If we cannot assume equal variances, we can calculate SSE with our data:

$$SS_E = (10-1)(7.6^2) + (10-1)(8.5^2) = 1170.09$$

Source	Sums of Squares	Degrees of freedom	Mean Square	Mean Square Ratio
Drug	328.1	1	328.1	5.05
Error	1170.09	18	65.0	
Total	N/A	19		

11) Write out a paragraph explaining the results of this study.

Activity #5: Part C – Randomization-based approach

Here's the code for, and output from, Stata with regards to our AxB ANOVA on our cholesterol data:

```
. anova cholesterol i.obese##i.drug

                         Number of obs =         40     R-squared     =  0.3043
                         Root MSE      =     8.5615     Adj R-squared =  0.2464

              Source |  Partial SS    df       MS            F     Prob > F
          -----------+----------------------------------------------------
               Model |  1154.40398     3  384.801327       5.25     0.0041
                     |
                drug |  46.3434248     1  46.3434248       0.63     0.4317
               obese |  752.815519     1  752.815519      10.27     0.0028
          obese#drug |  355.245038     1  355.245038       4.85     0.0342
                     |
            Residual |  2638.77331    36  73.2992587
          -----------+----------------------------------------------------
               Total |   3793.1773    39  97.2609563
```

12. We found a significant interaction and split the data to investigate the simple effects of the study. Calculate and interpret an effect size for this study. What assumptions did we make when conducting this AxB ANOVA?

    Effect size = eta-squared = ______________________________

    Assumptions = ____

13. ANOVA is robust to violations of the assumptions, especially when we have an equal number of observations in each group. If we have obvious departures from normality and equal variances, we should consider conducting a different type of analysis. For example, the *Kruskal-Wallis* test is an ANOVA conducted on the data after we convert the data to ranks.

    We might also try a randomization-based approach. Recall that in randomization-based approaches, we:
    (1) Pool all the data together, ignoring group membership
    (2) Randomly assign observations to groups (assuming the treatment groups don't matter)
    (3) Calculate a test statistic from this randomized data
    (4) Repeat this process many times
    (5) Determine how unlikely it was for us to observe the test statistic we actually observed from our data.

    In activity #3, we saw how to conduct a randomization-based ANOVA using the MAD or F-statistic. In this example, we'll use the eta-squared statistic. Above, you've recorded the eta-squared statistic for the observed data. All we need to do is calculate that same statistic after randomizing the data.

    Our goal in this example will be to demonstrate how this randomization-based approach can replicate the results we got from our AxB ANOVA.

    Let's first focus on that interaction effect. In the ANOVA summary table, you can see the p-value for the interaction term was **0.0342**.

14. On the course website, I've uploaded code to conduct this randomization-based AxB ANOVA in either R or Stata. In this activity, I'll focus on the process and output; not the code.

    Our model is: **`Cholesterol = f(drug, obesity, drug*obesity)`**

    The data from this study yielded an eta-squared of 0.3043.

    To use our randomization approach, I simply had the computer shuffle the interaction labels. In other words, my model was: **`Cholesterol = f(drug, obesity, shuffle(drug)*shuffle(obesity))`**

    I had the computer calculate eta-squared from 10,000 replications of that shuffled model. Here's the distribution of those eta-squared values:

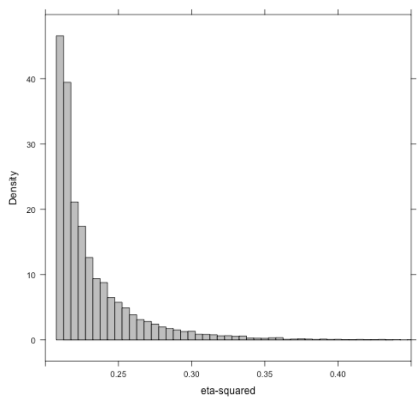

If the interaction term doesn't matter (in other words, if there is no interaction between the drug and obesity), the eta-squared value we calculated from our study came from this distribution.

Locate our observed eta-squared on this distribution. The computer tells me that of the 10,000 values on that distribution, our observed eta-squared is in the top 3.59%. Therefore, our p-value is 0.0359.

That is extremely similar to the p-value of 0.0342 we obtained from our AxB ANOVA. What can we conclude from this?

15. To replicate the p-value for our drug effect, ($p = 0.4317$) I simply had the computer run this model 10,000 times:

    `Cholesterol = f(shuffle(drug), obesity, drug*obesity)`

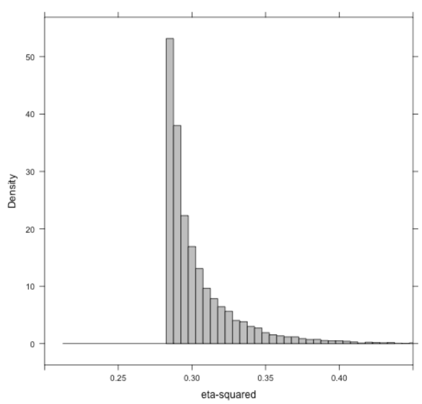

If the drug doesn't matter (in other words, if there is no drug effect), the eta-squared value we calculated from our study came from this distribution.

Locate our observed eta-squared on this distribution. The computer tells me that of the 10,000 values on that distribution, our observed eta-squared is in the top 32.25%. Therefore, our p-value is 0.3225.

That is similar to the p-value of 0.4317 we obtained from our AxB ANOVA. What can we conclude from this?

16. To replicate the p-value for our obesity effect, ($p = 0.0028$) I had the computer run this model 10,000 times:

    `Cholesterol = f(drug, shuffle(obesity), drug*obesity)`

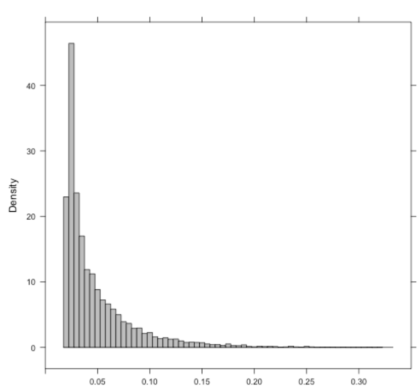

If obesity doesn't matter (in other words, if there is no obesity effect), the eta-squared value we calculated from our study came from this distribution.

Locate our observed eta-squared on this distribution. The computer tells me that of the 10,000 values on that distribution, our observed eta-squared is in the top 0.2%. Therefore, our p-value is 0.002.

That is similar to the p-value of 0.0028 we obtained from our AxB ANOVA. What can we conclude from this?

Go through the R Code for this activity to see more examples