9b: Randomization Equivalent to AxB ANOVA

Let's take another look at the cholesterol data from Activity #8:

	Placebo	Drug	Total
Not obese	85.08 92.68 90.60 89.86 76.85 77.63 90.24 107.21 91.39 90.24 n = 10 mean = 89.2 std. dev = 8.5	92.11 85.14 107.13 87.59 91.20 79.83 82.05 103.61 101.25 99.96 n = 10 mean = 93.0 std. dev = 9.5	n = 20 mean = 91.1 std. dev = 9.0
Obese	103.13 95.17 116.50 100.63 102.87 115.14 101.10 100.16 109.20 94.27 n = 10 mean = 103.8 std. dev = 7.6	96.97 105.78 85.51 94.67 90.53 94.39 87.55 106.52 108.24 86.86 n = 10 mean = 95.7 std. dev = 8.5	n = 20 mean = 99.8 std. dev = 8.9
Total	n = 20 mean = 96.5 std. dev = 10.8	n = 20 mean = 94.3 std. dev = 8.9	n = 40 mean = 95.4 std. dev = 9.862

1) The following Stata output shows the results of conducting an AxB ANOVA on this dataset. What assumptions were made in conducting this test? Are these assumptions reasonable in this scenario? Verify the values for SStotal and the various degrees of freedom. What conclusions could we make from the output? Calculate an effect size and interpret.

```
. 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
```

2) I had Stata conduct a *robust equal variances test* to compare the variances among the 4 cells. This test resulted in a p-value of 0.70. What can we conclude from this? What can we conclude from the p-values reported below (from a *Shapiro-Wilk W Test for Normality* on the data within each cell)?

Placebo-Not Obese: p = 0.11 Drug-Not Obese: p = 0.59
Placebo-Obese: p = 0.26 Drug-Obese: p = 0.19

3) What options do we have if we have serious concerns about the assumptions necessary to conduct an AxB ANOVA? Well, if the sample sizes within each cell are large and equal, we could still choose to conduct an ANOVA (it's robust against violations of normality).

Throughout MATH 300 (and in the first unit of this course), we've briefly investigated alternative analyses based on *randomization methods*. Recall that randomization methods require us to:

(1) Pool all the data (ignoring group membership)
(2) Randomly assign observations to groups (assuming the groups have no impact on the observations)
(3) Calculate a test statistic
(4) Repeat steps 1-3 many, many times and record the test statistic each time
(5) Determine the likelihood of the observed data based on all these test statistics

In the last unit, we calculated an F-statistic after randomly assigning observations to groups. We then repeated that process 10,000 times to create a sampling distribution. We then used that sampling distribution to determine how likely we were to obtain our actual (observed) F-statistic.

We'll do something similar here. We already know our observed F-statistics (from the AxB ANOVA we conducted earlier) are:

F (interaction) = 4.85 (p = 0.0342)
F (drug) = 0.63 (p = 0.4317)
F (obesity) = 10.27 (p = 0.0028)

To use randomization methods, we'll need to deal with each of these F-statistics separately. Let's first focus on the interaction effect.

4) To test the interaction of drug and obesity on cholesterol levels using randomization methods, we will:

(1) Pool all 40 observations into a single group (ignoring the drug/obesity status of each subject)
(2) Randomly assign 20 observations to the "drug" group and 20 observations to the "placebo" group
(3) Randomly assign 20 observations to the "obese" group and 20 observations to the "not obese" group
(Note: We're assuming a null hypothesis that there's no drug- or obesity-effect on cholesterol. With this assumption, we can assume that each subject would have had the same cholesterol regardless of obesity or whether that subject was given the drug/placebo.)
(4) Conduct an AxB ANOVA on this randomized data and record the value of F-interaction
(5) Repeat steps 1-4 many, many times.
(6) Determine the likelihood of the observed F-interaction (F = 4.85) based on all these randomized values.

5) In this study, there are 4 705 360 871 073 570 227 520 possible randomizations. How did I calculate this? If it took you only 10 seconds to find each randomization, it would take you more than 1 491 040 000 000 000 years to find them all. Rather than trying to find them all, let's have the computer find 100,000 randomizations and see what we get.

In Stata, the code to get 100,000 randomizations is:

```
permute cholest permt=e(F_3), reps(100000) nodots saving(/Users/Brad/Desktop/Fs.dta): anova cholest drug##obese
```

- permute cholest: Instead of randomly assigning drug/obesity to subjects, I randomly assigned cholesterol levels to the existing subjects
- permt=e(F_3): Calculate the 3rd F statistic from AxB ANOVA (interaction)
- reps(100000): Repeat this 100,000 times
- saving(/Users/Brad/Desktop/Fs.dta): save all the F values from the 100,000 randomizations (so we can graph them)
- anova cholest drug##obese: For each randomization, conduct an AxB ANOVA with: Dependent var = cholesterol, Ind. var = drug, obesity

Granted, I don't have the fastest computer available, but it took around 20 minutes for Stata to run 100,000 randomizations. The output I obtained is pasted below.

```
Monte Carlo permutation results                   Number of obs   =         40

      command:  anova cholest drug##obese
        permt:  e(F_3)
  permute var:  cholest

------------------------------------------------------------------------------
T            |     T(obs)       c       n   p=c/n   SE(p) [95% Conf. Interval]
-------------+----------------------------------------------------------------
       permt |   4.846502    3460  100,000 0.0346  0.0006  .0334757   .0357513
------------------------------------------------------------------------------
Note:  confidence interval is with respect to p=c/n.
Note:  c = #{|T| >= |T(obs)|}
```

Notes:

- On the top-left, we can see Stata provided Monte Carlo permutation results. *Monte Carlo* methods use simulations based on random sampling (what we call randomization methods). Recall that *permutations* refer to rearrangements of objects in different orders.
- On the top-right, we see Stata used all 40 of our observations
- In the table, "T" represents "test statistic." So our observed test statistic -- T(obs) -- is 4.846502. This is our observed F-statistic of 4.85 that we calculated earlier in our AxB ANOVA.
- The n of 100,000 represents the number of F-statistics that were calculated from our randomizations
- "C" represents the number of randomizations that were greater or equal to our observed test statistic. In other words, 3460 of our 100,000 randomizations resulted in F-statistics greater or equal to 4.85.
- Our p-value is p = c/n = 3460 / 100000 = 0.0346. This is similar to the p-values we calculated in our original ANOVA and in the Kruskal-Wallis analysis.
- SE(p) is a standard error of our p-value. The last 2 columns represent a 95% confidence interval for our p-value.

6) Based on the Stata output, briefly write any conclusions you can make from this study. Do we have a significant interaction? Does this randomization approach agree with the results we obtained from the AxB ANOVA?

7) Here's a histogram of the 100,000 randomized F-statistics. Locate our observed F-statistic and shade-in the p-value.

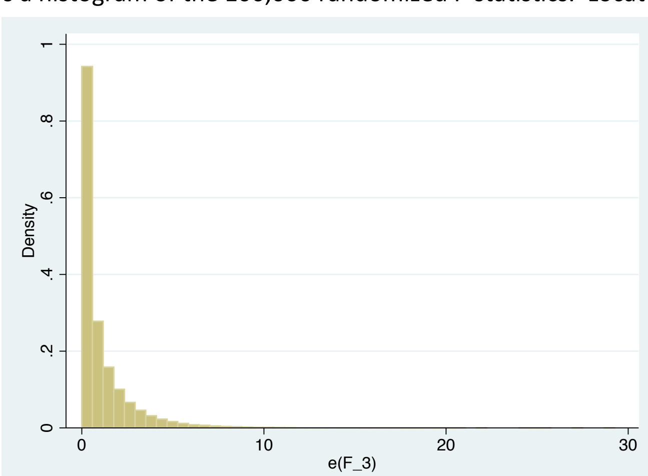

8) Just for this example, let's ignore the fact that we found a significant interaction in our analysis. Let's go ahead and use randomization methods to determine if the drug has a significant effect on cholesterol. To do this, we will need to:

(1) Pool the 40 observations into 2 existing groups: obese and not obese

(2) Randomly assign 10 observations within each group to "drug" and the remaining 10 observations to the "placebo"

(Note: We're assuming a null hypothesis that the drug has no effect on cholesterol.)

(3) Conduct an AxB ANOVA on this randomized data and record the value of F-drug

(4) Repeat steps 1-3 many, many times.

(5) Determine the likelihood of the observed F-drug (F = 0.63, p = 0.43) based on all these randomized values.

9) Look at step (1) above. Why do we keep observations in the existing obese / not obese groups?

10) Because I don't want to wait 20 minutes again, I'm going to have Stata calculate only 10,000 randomizations. Interpret the following Stata syntax:

```
permute cholest permt=e(F_1), strata(obese) reps(10000) nodots saving(/Users/Brad/Desktop/F1s.dta): anova cholest drug##obese
```

11) Interpret the following output:

```
Monte Carlo permutation results

Number of strata =           2                        Number of obs    =          40

      command:  anova cholest drug##obese
        permt:  e(F_1)
  permute var:  cholest

------------------------------------------------------------------------------
T            |     T(obs)       c       n   p=c/n   SE(p) [95% Conf. Interval]
-------------+----------------------------------------------------------------
       permt |   .6322496    4271   10000  0.4271  0.0049  .4173773   .4368652
------------------------------------------------------------------------------
Note:  confidence interval is with respect to p=c/n.
Note:  c = #{|T| >= |T(obs)|}
```

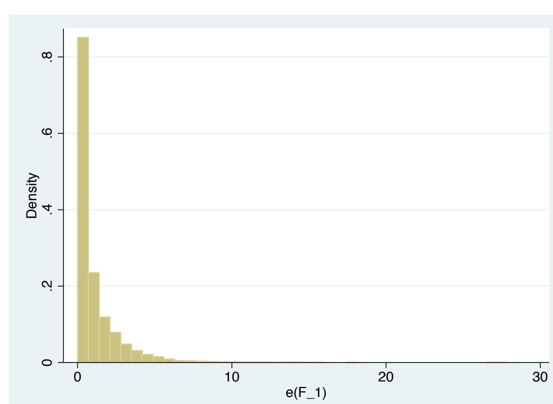

12) Locate the observed F-statistic on the randomization distribution.

13) Let's go ahead and use randomization methods to get the F-statistic for the effect of obesity on cholesterol. Fill-in the blanks to show what steps we would need to follow:

(1) Pool the 40 observations into 2 existing groups: ______________________ and ______________________

(2) Randomly assign 10 observations within each group to ____________________ and ____________________.

(Note: We're assuming ____)

(3) Conduct an AxB ANOVA on this randomized data and ____

(4) ____

(5) Use these randomized values to determine ____.

14) Fill-in the blanks to show the Stata syntax we would need to enter to get 10,000 randomizations:

permute ___________________ permt=e(F_2), strata(_________________________) reps(___________)

nodots saving(/Users/Brad/Desktop/F2.dta): ____

15) Fill-in the blanks and Interpret the output:

```
Number of strata =          2                      Number of obs    =         40

      command:  anova cholest drug##obese
        permt:  e(F_2)
  permute var:  cholest

------------------------------------------------------------------------------
T            |     T(obs)       c       n   p=c/n   SE(p) [95% Conf. Interval]
-------------+----------------------------------------------------------------
             |
       permt |   ________       32    10000  ________ 0.0006  .0021898   .0045145
------------------------------------------------------------------------------
Note:  confidence interval is with respect to p=c/n.
Note:  c = #{|T| >= |T(obs)|}
```

16) What are the advantages and disadvantages of this randomization approach versus a traditional AxB ANOVA?