ECON 523: Program Evaluation for International Development
In-Class Activity 11

In this exercise, we will simulate a data-generating process where outcomes are correlated within clusters. The code below generates a an outcome, `y` that is only trivially clustered (because each cluster contains $n_g = 1$ observations.) We will modify the code below to to better understand test size and statistical power in cluster-randomized trials.

Begin by creating a do file that contains the code below:

```
// preliminaries

clear all
set seed 24601

local numclusters = 1000
local obspercluster = 1
local effect = 0

// create an empty matrix to save results

local loopmax=100
matrix pval=J(‘loopmax’,1,.)

// create data sets w/ clusters

forvalues i  =1/‘loopmax’ {
display "Loop iteration ‘i’"
quietly set obs ‘numclusters’
quietly gen clustid = _n
quietly gen treatment=cond(_n>‘numclusters’/2,1,0)
quietly gen clusteffect = rnormal()
quietly expand ‘obspercluster’
quietly gen y = ‘effect’*treatment + clusteffect + rnormal()
quietly reg y treatment
mat V = r(table)
matrix pval[‘i’,1]=V[4,1]
drop clustid treatment clusteffect y
}

// save your results

svmat pval
summarize
```

```
gen significant = pval<0.05
tab significant
```

Extend your do file as you answer the questions below:

1. Look carefully at the program above. Make sure you understand what is happening in every line. What is the mean of the `treatment` variable (in each iteration of the loop)? You should be able to figure this out without running the code.

2. Looking at the current values of the local macros, would you say that the null hypothesis is true or false?

3. When you run the code and it tabulates the significant variable at the end, what is the expected number of times that you will reject the null hypothesis?

4. Now run the code. How many times do you actually reject the null hypothesis?

5. Change the number of iterations (the number of times the loop runs) to 1,000, and then run the code again. How many times do you reject the null hypothesis now? Was this what you expected?

6. What is the variance of `y`?

7. Use the MDE formula to calculate the expected standard error in the regression of `y` on treatment. What is it?

8. Now modify the program so that you also save the standard error from the regression of `y` on treatment. What is the average standard error across your 1,000 simulations?

9. You can multiply the expected standard error by 2.8 to calculate the minimum detectable effect given a test size of 0.05 and a power of 0.8. This tells us that the MDE is approximately 0.25. Change the local macro effect to 0.25 and run your code again. In how many of your 1,000 simulations do you reject the null hypothesis?

10. Now consider a case where treatment is assigned at the cluster level, and there are multiple observations per cluster. Change the local macro `numclusters` to 50 and the local macro `obspercluster` to 20. The expand command will make `obspercluster` identical copies of all of your observations (within your data set), so that you will have 1000 observations in total. Now, set effect to 0 again, so that the null hypothesis (that treatment has no effect) is correct. Run your code. How many times (out of 1,000) do you reject the null?

11. When treatment is assigned at the cluster level and outcomes are correlated within clusters, hypothesis tests are incorrectly sized unless we use the `cluster` option at the end of our regression. Run your code again, but add `, cluster(clustid)` to the end of your regression. How many times to you reject the null hypothesis now?