Economics 311 | Fall 2017

Project #3 Assignment | Wednesday, September 27

This project uses a subset of a data set that accompanies Stock and Watson's popular econometrics textbook. A link to the data set is on the assignment page or you can download it from academic.reed.edu/economics/parker/311/asgns/datasets/Evals.dta.

Please copy and paste the relevant parts of your Stata output into your project report. You would not do this in a paper or thesis, but it is very helpful to me to see the command(s) you used to generate the results.

The data are taken from the following paper, to which the authors apply an amiable (Amy-able?) and appealingly alliterative appellation:

Daniel Hamermesh and Amy Parker, "Beauty in the Classroom: Instructors' Pulchritude and Putative Pedagogical Productivity," *Economics of Education Review*, 24(4), August 2005, 369–76.

(You are not expected to read this paper for this assignment; in fact, it's probably better if you don't because reading it may bias how you would approach or interpret the questions in the assignment. The reference is included to give appropriate credit and in case you are interested in the topic.)

The table below describes the variables in the data set:

Variable	**Definition**
eval	"Course overall" teaching evaluation score, on a scale of 1 (very unsatisfactory) to 5 (excellent)
beauty	Rating of instructor physical appearance by a panel of six students, averaged across the six panelists, shifted to have mean zero.
age	Instructor's age
female	= 1 if instructor is female, 0 otherwise

Exercises

1. Data and summary statistics

a. Examine the data using the data browser. Do there seem to be multiple observations from the same instructor? How might this violate the OLS "classical" assumptions?

Based on casual examination of the data, speculate on how large a problem this may be for this dataset.

b. Compute and present the summary statistics for all four variables. (The Stata command summarize is useful here.) The *female* variable is what we call a "dummy" (or indicator or binary) variable. How can we interpret the sample mean of this variable? The table above says that the *beauty* variable is "shifted to have mean zero." What does this mean (no pun intended … for once)? Is this verified by your table? What is the standard deviation of *beauty* and what does that mean? Is there anything else notable about the summary statistics?

2. Simple regression. Perform a simple (one-regressor) regression of course evaluations on beauty and use the results to answer the following questions:

a. What is the interpretation of the intercept term in the regression? What is the interpretation of the slope coefficient? If Professor Average has average beauty, what is our best prediction of her evaluation score? If Professor Gorgeous has a beauty score two standard deviations above the mean, what is our prediction of his score? If Professor Ugly has a beauty score two standard deviations below the mean and gets evaluations of 3.9, would you consider him to be a better or worse teacher than would be expected based solely on his appearance? Why?

b. If we increase the value of *beauty* by one standard deviation (of *beauty*), by how many standard deviations (of *eval*) does the conditional expectation (prediction) of *eval* go up along our regression line? (This measure is common in non-economics social sciences and is often misleadingly called a "beta" coefficient.) What is the advantage of this measure compared to the simple regression slope coefficient? What is the disadvantage?

c. Assess the strength of your regression relationship using the R-square statistic. Retrieve the fitted (predicted) values and residuals from this regression. (Use the predict command for this.) Compute the summary statistics of the predicted values and residuals.
   i. What is the mean of the residuals? Compare the mean of the fitted values to the mean of *eval*. What do you find? Interpret these results (which *always* happen in OLS).
   ii. What is the range of your fitted values compared to the actual values of *eval*? Is this a problem?
   iii. Examine the residual series to see if there are any obvious outliers. Create a graph with both the actual values of the dependent variable and the fitted values (measured on the vertical axis) plotted against the regressor (*beauty*) on the horizontal axis. Based on your graph, does the linear functional form seem appropriate?

d. Using this regression, what is the 95% confidence interval for the effect of beauty on evaluations? Can you say with 95% confidence that the effect of beauty is not zero?

3. Sub-sample regressions. Some have argued that our society seems to place more emphasis on women's appearance than on men's. Perform separate regressions for the sub-samples of female and male professors. Interpret the results. Do they support the supposition that appearance is more important for women's evaluations?

4. Averaging multiple courses by the same instructor. It seems fair to assume that any observations with the same values for *age*, *female*, and *beauty* are different courses taught by the same instructor. Sort the dataset by these three variables, then use the `collapse` command to create a new dataset with one observation per instructor and the evaluation variable set to the average evaluation for each instructor across all courses. (Be sure that you don't save this over your original data set!)

a. How many observations are in this dataset? What does this mean?

b. Re-run the *eval* on *beauty* regression and discuss how or whether the results change, and why. Is this what you expected?

c. Which regression do you think is more reliable, and why?