Gender Bias in Pay: *Dummy variables in action*

In theory, you might think that dummy variables would facilitate a simple and compelling test for bias or discrimination.

For example, suppose you wanted to test for gender bias in pay. **It's really very simple! Yes?**

Grab a fabulously complete dataset, which includes pay and gender variables as well as a gazillion other explanatory/ control variables (all the non-gender factors that might otherwise explain the variation in pay). Build your most amazing MLR model explaining the variation in pay, controlling for all of the gazillion possible explanatory factors other than gender.[1]

And then you have a couple choices:

1) *Estimate one model*: After you've built the world's most amazing MLR model, you have just one additional task: Add a binary *gender* dummy variable to the model.

   The estimated *gender* coefficient will capture average differences in pay (across gender) not explained by the rest of the model (so it captures average residuals).

   But since your awesome model has completely controlled for all possible explanatory factors, there's only one conclusion: Those residuals are driven by *Gender Bias!*

   And the p value and t stat on the *gender* variable tell you if you have statistical significance.[2] So doing the test for statistical significance is a breeze. Done!

2) *Estimate two models*: This is a slightly more complicated approach.

   Estimate two pay models, one for males and the other for females, with all of your fabulous explanatory variables on the RHS. This will explicitly allow for different SRFs for the two populations. The female SRF generates predicted pay for females as a function of a bunch of explanatory factors, presumably including education, experience, job tenure and so forth. And the male SRF similarly generates predicted pay for males (as a function of a bunch of a fairly similar set explanatory factors… after all, those are the factors that drive compensation!).

   Then apply the female SRF to the males' RHS data to predict female-model driven pay for males, and do the reverse, applying the males SRF to the females' RHS data to predict male model-driven pay for females. You could do this on a case by case basis, or

[1] Don't worry one bit about multicollinearity, because at the end of the day, you're going to be working with a favorite coefficient model.

[2] In Hazelwood Sch. Dist. v. United States, 433 U.S. 299 (1977), Footnote 17, the Supreme Court appeared to endorse the proposition that a t stat of at least two provides (statistical) evidence of discrimination.

look at population means… but either way, the differences will tell you something about gender bias.[3]

Sounds simple, Yes? But sorry, it is anything but simple!

There are extraordinary opportunities for biases with these models, including sample selection bias and the dreaded omitted variable bias. Your model is only as good as the data you chose to work with… and as bad as the data you left out. The coefficient on the dummy variable will capture average gender differences for effects not otherwise captured/explained by the model. So if your model is incomplete, you will attribute those excluded effects to gender, when in fact they might very well have everything to do with, say, omitted variables, and in fact, have nothing to do with gender discrimination.

And so the pressure is on: What important explanatory factors did you leave out of your model? How representative/unbiased is your sample? And how are those factors biasing your conclusions about gender discrimination/bias?

Application: *Working with the wage1 dataset*

Let's explore. You'll be working with the *wage1* dataset, which was assembled by Geoff Wooldridge and Hank Farber in 1988 (the data are from the 1976 Current Population Survey). I know it's ancient history, but *wage1* is easily accessed through *bcuse*… and it will nicely illustrate the various dimensions of the challenge.

Here are brief descriptions of the variables in the dataset (many of which are (0,1) dummies)… and a comparison of the means by gender:

```
bcuse wage1

 1. wage     average hourly earnings
 2. educ     years of education
 3. exper    years potential experience
 4. tenure   years with current employer
 5. nonwhite =1 nonwhite
 6. female   =1 female
 7. married  =1 married
 8. numdep   number of dependents
 9. smsa     =1 live in SMSA
10. northcen =1 live in north central US
11. south    =1 live in southern region
12. west     =1 live in western region
13. construc =1 work in construc. indus.
14. ndurman  =1 in nondur. manuf. indus.
15. trcommpu =1 in trans, commun, pub ut
16. trade    =1 in wholesale or retail
17. services =1 in services indus.
18. profserv =1 in prof. serv. indus.
19. profocc  =1 in profess. occupation
20. clerocc  =1 in clerical occupation
21. servocc  =1 in service occupation
```

[3] And maybe you want to know about statistical significance? Me too:) Not sure how to handle that (other than maybe empirical distributions?).

Gender Bias in Pay: Dummy variables in action

		Males	Females	All		*Female Delta*
nobs		274	252	526		***(22)***
wage		$7.10	$ 4.59	$5.90		***$ (2.51)***
educ		12.79	12.32	12.56		***(0.47)***
exper		17.56	16.43	17.02		***(1.13)***
tenure		6.47	3.62	5.10		***(2.86)***
nonwhite		11%	10%	10%		***-1%***
married		69%	52%	61%		***-16%***
numdep		1.00	1.09	1.04		***0.08***
smsa		72%	73%	72%		***1%***
northcen		24%	26%	25%		***1%***
south		38%	33%	36%		***-4%***
west		15%	19%	17%		***4%***
construc		6.2%	2.8%	4.6%		***-3.4%***
ndurman		14.2%	8.3%	11.4%		***-5.9%***
trcommpu		4.7%	4.0%	4.4%		***-0.8%***
trade		31%	26%	29%		***-5%***
services		7%	13%	10%		***7%***
profserv		17%	36%	26%		***19%***
profocc		45%	28%	37%		***-17%***
clerocc		4%	31%	17%		***27%***
servocc		9%	20%	14%		***11%***

The difference in average wages is about $2.50 … but lots of other things differ as well: females have 5-10% less education and experience than males, and about half as much job tenure. You would normally expect that those three differences would alone and collectively imply lower wages for females… but $2.51 lower?

Let's investigate.

We'll eventually get to the more interesting issues of differences in education, experience and tenure … but let's start with some simple applications of dummy variables

Gender Bias in Pay: Dummy variables in action

1. First regress *wage* on a constant term to find the overall wage average.

Average wages: You can use regression models to calculate sample means by category.

```
. summm wage

    Variable |        Obs        Mean    Std. Dev.       Min        Max
-------------+---------------------------------------------------------
        wage |        526    5.896103    3.693086        .53      24.98

. reg wage

      Source |       SS           df       MS      Number of obs   =       526
-------------+----------------------------------   F(0, 525)       =      0.00
       Model |           0         0           .   Prob > F        =         .
    Residual |  7160.41429       525  13.6388844   R-squared       =    0.0000
-------------+----------------------------------   Adj R-squared   =    0.0000
       Total |  7160.41429       525  13.6388844   Root MSE        =    3.6931

------------------------------------------------------------------------------
        wage |      Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
       _cons |   5.896103   .1610262    36.62   0.000     5.579768    6.212437
------------------------------------------------------------------------------
```

Note that the reported Std. Err. is in fact the standard error associated with the sample mean estimator, $S_y / \sqrt{n}$.

```
. di 3.693086/526^.5
.1610262
```

So dummies in regressions provide an easy way to generate sample means and test the Null Hypothesis that the true mean is zero. In the results above, the t stat is 36.62 and the p value is 0…. and so it's easy to reject $H_0 : \mu = 0$ at any standard level of statistical significance.

2. Add in the female dummy variable to find the average wages for females and males

Average wages by gender:

```
. tabstat wage, by(female)

  female |      mean
---------+----------
       0 |  7.099489
       1 |  4.587659
---------+----------
   Total |  5.896103

. reg wage female

      Source |       SS       df       MS              Number of obs =     526
-------------+------------------------------           F(  1,   524) =   68.54
       Model |  828.220467     1  828.220467           Prob > F      =  0.0000
    Residual |  6332.19382   524  12.0843394           R-squared     =  0.1157
-------------+------------------------------           Adj R-squared =  0.1140
       Total |  7160.41429   525  13.6388844           Root MSE      =  3.4763

------------------------------------------------------------------------------
        wage |      Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
      female |   -2.51183   .3034092    -8.28   0.000    -3.107878   -1.915782
       _cons |   7.099489   .2100082    33.81   0.000     6.686928     7.51205
------------------------------------------------------------------------------
```

Predicted wages are: $\hat{w} = 7.10 - 2.51\, female$.

- For females, $female = 1$, and so $\hat{w} = 7.10 - 2.51\,(1) = 4.59$ … the average female wage!
- For males, $female = 0$, and so $\hat{w} = 7.10 - 2.51\,(0) = 7.10$ … the average male wage!

These predicted wages are just the average wages for males and females.

The _cons coefficient (7.10) is the average wage for males (the predicted wage when female=0), and the female coefficient is the **difference** (-2.51) in average wages between males and females.

So you can read the **difference** in mean wages right off the regression results with no further calculation… it's just the coefficient on the dummy variable. Since the t-stat is -8.28 and the p-value is 0.000, we reject the hypothesis that there is no difference between wages for men and women (at any usual level of statistical significance). But of course, we haven't yet controlled for any of the other factors that might explain differences in wages.

We could instead use a male dummy variable… and we'd get the same results.

```
. gen male=(female==0)
. reg wage male
      Source |       SS       df       MS              Number of obs =     526
-------------+------------------------------           F(  1,   524) =   68.54
       Model |  828.220467     1  828.220467           Prob > F      =  0.0000
    Residual |  6332.19382   524  12.0843394           R-squared     =  0.1157
-------------+------------------------------           Adj R-squared =  0.1140
       Total |  7160.41429   525  13.6388844           Root MSE      =  3.4763

------------------------------------------------------------------------------
        wage |      Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
        male |    2.51183   .3034092     8.28   0.000     1.915782    3.107878
       _cons |   4.587659   .2189834    20.95   0.000     4.157466    5.017852
------------------------------------------------------------------------------
```

Gender Bias in Pay: Dummy variables in action

For this model, predicted wages are: $\hat{w} = 4.59 + 2.51\, male$.

- For females, $male = 0$, and so $\hat{w} = 4.59 + 2.51\,(0) = 4.59$ … the average male wage
- For males, $male = 1$, and so $\hat{w} = 4.59 + 2.51\,(1) = 7.10$ … the average female wage

Note: The two models (w/ female or male dummies) are virtually identical… except that they have different *benchmarks* (sometimes called the *excluded dummy* or *excluded other*). The benchmark is the case in which the dummy variable is 0. Sometimes the benchmark is obvious; sometimes it's not so obvious… and you need to understand your data better to identify the actual benchmark.

If you put both dummies (*male* and *female*) in the model, Stata will reject one dummy due to perfect multicollinearity. The error message will say "(omitted)" … but variables are dropped for only one reason, perfect collinearity… and of course *male* and *female* are perfectly collinear: $male = 1 - female$.

```
. reg wage male female
note: female omitted because of collinearity

      Source |       SS           df       MS      Number of obs   =       526
-------------+----------------------------------   F(1, 524)       =     68.54
       Model |  828.220467         1  828.220467   Prob > F        =    0.0000
    Residual |  6332.19382       524  12.0843394   R-squared       =    0.1157
-------------+----------------------------------   Adj R-squared   =    0.1140
       Total |  7160.41429       525  13.6388844   Root MSE        =    3.4763

------------------------------------------------------------------------------
        wage |      Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
        male |    2.51183   .3034092     8.28   0.000     1.915782    3.107878
      female |          0  (omitted)
       _cons |   4.587659   .2189834    20.95   0.000     4.157466    5.017852
------------------------------------------------------------------------------
```

It would be an egregious error to claim that you have any evidence whatsoever of gender bias … since your analysis is based only on the difference in mean wages for males and females. Especially since you believe that *educ*, *exper tenure* and *married* are all correlated with gender… so you'd want to control for those effects. Otherwise your estimated *gender* coefficient may be biased by the omission of those variables from the analysis

So let's worry about all that.

3. Look at correlations in the hopes of identifying possible omitted variable bias.

```
. corr wage female educ exper tenure nonwhite married numdep

           |     wage   female     educ    exper   tenure nonwhite  married   numdep
-----------+------------------------------------------------------------------------
      wage |   1.0000
    female |  -0.3401   1.0000
      educ |   0.4059  -0.0850   1.0000
     exper |   0.1129  -0.0416  -0.2995   1.0000
    tenure |   0.3469  -0.1979  -0.0562   0.4993   1.0000
  nonwhite |  -0.0385  -0.0109  -0.0847   0.0144   0.0116   1.0000
   married |   0.2288  -0.1661   0.0689   0.3170   0.2399  -0.0623   1.0000
    numdep |  -0.0538   0.0331  -0.2153  -0.0563  -0.0270   0.0777   0.1545   1.0000
```

Omitted variable bias appears to be lurking, as wages are correlated with *educ*, *exper*, *tenure* and *married*, as is *female*. Leave any one of these explanatory variables out of your model at the peril of omitted variable bias!

4. Controlling for *tenure* effects.

Let's start looking at the other explanatory factors (other than gender), and start with the variable most highly correlated with female, *tenure*. Here's a look at the overall relationship between *tenure* and *wage*:

```
. reg wage tenure

      Source |       SS           df       MS      Number of obs   =       526
-------------+----------------------------------   F(1, 524)       =     71.68
       Model |   861.62965         1   861.62965   Prob > F        =    0.0000
    Residual |  6298.78464       524  12.0205814   R-squared       =    0.1203
-------------+----------------------------------   Adj R-squared   =    0.1187
       Total |  7160.41429       525  13.6388844   Root MSE        =    3.4671

------------------------------------------------------------------------------
        wage |      Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
      tenure |   .1773271   .0209449     8.47   0.000     .1361809    .2184733
       _cons |   4.990925    .185158    26.95   0.000     4.627182    5.354669
------------------------------------------------------------------------------
```

So, yes, tenure matters (t = 8.47; p = 0.000)… maybe we should control for tenure effects when looking at the wage discrimination question.

```
. reg wage female tenure

      Source |       SS           df       MS      Number of obs   =       526
-------------+----------------------------------   F(2, 523)       =     64.16
       Model |  1410.73305         2  705.366525   Prob > F        =    0.0000
    Residual |  5749.68124       523  10.9936544   R-squared       =    0.1970
-------------+----------------------------------   Adj R-squared   =    0.1939
       Total |  7160.41429       525  13.6388844   Root MSE        =    3.3157

------------------------------------------------------------------------------
        wage |      Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
      female |  -2.086511   .2952329    -7.07   0.000    -2.666499   -1.506523
      tenure |   .1487455   .0204344     7.28   0.000     .1086019    .1888891
       _cons |   6.136443   .2400553    25.56   0.000     5.664852    6.608034
------------------------------------------------------------------------------
```

Gender Bias in Pay: Dummy variables in action

Controlling for tenure effects, women are now on average paid $2.09 less than men … and the effect is highly statistically significant. So adding *tenure* to the model caused the gender bias estimate to drop by $0.43. So almost 20% of the originally estimate bias was in fact driven by omitted variable bias associated with tenure and differences in tenure between males and females (a factor we might initially not be inclined to relate to gender bias).

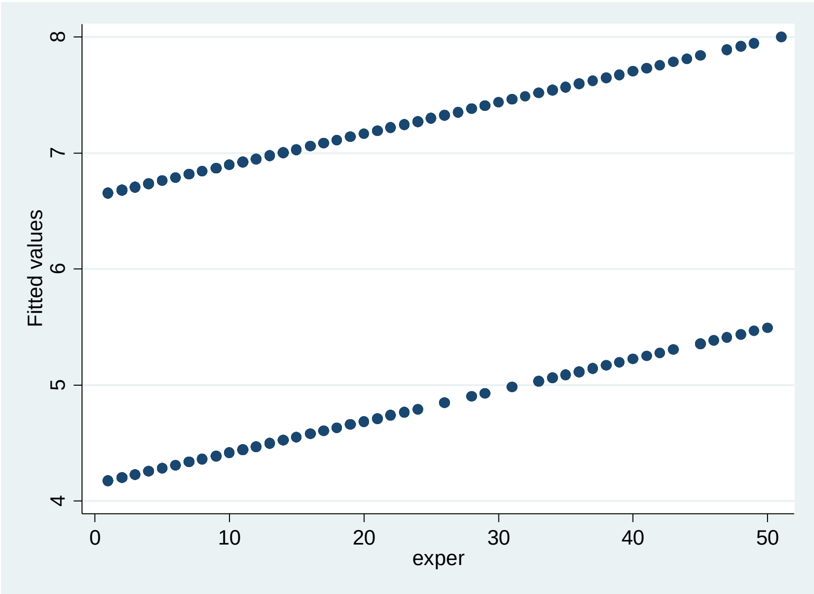

The predicted values from the model are to the right, where the higher line is for males; for any level of tenure, predicted wages for males are $2.08 higher. The slopes of the two lines are identical, because we restricted the model to have a common incremental *tenure* effect. We'll be dropping that assumption below.

So the female dummy allows for different intercepts for males and females… but the slopes of male and female SRFs are identical for this specification of the model.

Intercept dummies… and ***slope dummies***

For this reason, we sometimes refer to variable like *female* as *intercept dummies*… as they allow for different SRF intercepts for different categories. Later, we'll look at *slope dummies*, which will allow for different SRF slopes for different categories. And not surprisingly, if ou have interecepta dn slope dummies in your model, you allow for different SRFs slopes and intercepts fordifferent categories

5. Controlling as well for *educ, exper* and *married* effects.

Adding *educ*, *exper* and *married* to the model, we get:

```
--------------------------------------------------------------------------------------
                 (1)          (2)          (3)          (4)          (5)          (6)
                wage         wage         wage         wage         wage         wage
--------------------------------------------------------------------------------------
female        -2.512***    -2.481***    -2.294***    -2.273***    -2.087***    -1.741***
             (-8.28)      (-8.21)      (-7.58)      (-8.15)      (-7.07)      (-6.53)

exper                      0.0269*                                             0.0187
                          (2.42)                                              (1.56)

married                                 1.339***                               0.559
                                       (4.32)                                 (1.96)

educ                                                 0.506***                  0.556***
                                                    (10.05)                  (11.14)

tenure                                                            0.149***     0.139***
                                                                 (7.28)       (6.57)

_cons          7.099***     6.627***     6.180***     0.623        6.136***    -1.618*
             (33.81)      (23.15)      (20.86)       (0.93)      (25.56)      (-2.24)
--------------------------------------------------------------------------------------
N                526          526          526          526          526          526
R-sq           0.116        0.125        0.146        0.259        0.197        0.368
adj. R-sq      0.114        0.122        0.143        0.256        0.194        0.362
--------------------------------------------------------------------------------------
t statistics in parentheses
* p<0.05, ** p<0.01, *** p<0.001
```

Model (1) is the original model. In Models (2)-(5), *exper*, *married*, *educ* and *tenure* have been individually added to Model (1). Allowing for *experience* effects leads to a minimal decline ($.03) in the gender bias estimate; *married* alone drops the estimate by $.22; *education* alone drop the bias estimate by a few more pennies ($.24), and *tenure* effects easily have the greatest impact on the estimated bias, ($.42).

All of the estimated coefficients in models (2)-(5) are highly statistically significant. And when all four explanatory variables are included in the analysis, the estimated gender bias is ($1.74), a drop of about a third from the estimated bias in the first model.

To assess the joint statistical significance of the four additional explanatory variables in Model (5), we just do an F test after running the model:

```
. reg wage female exper married educ tenure

. test exper married educ tenure

 ( 1)  exper = 0
 ( 2)  married = 0
 ( 3)  educ = 0
 ( 4)  tenure = 0

       F(  4,   520) =   51.96
            Prob > F =    0.0000
```

Since F=51.96 and the p value is 0, *Reject Reject Reject* the joint Null hypothesis the set of associated parameters are all zero. Put differently, the collection of the four RHS variables is statistically significant at all standard levels of significance… even though *exper* and *married* are not individually statistically significant at the 5% level.

6. Tenure effects with slope dummies

We have used intercept dummies to capture average differences in wages controlling for whatever else was in the model. We now turn to slope dummies, and allow for different marginal relations (slopes) between wages and, say, tenure… for males and females. To generate slope dummies for , say, tenure, we interact the tenure variable with the female dummy variable:

```
. gen ftenure = female*tenure
. reg wage tenure ftenure
. gen ftenure = female*tenure

. reg wage tenure ftenure

      Source |       SS       df       MS              Number of obs =     526
-------------+------------------------------           F(  2,   523) =   56.11
       Model |  1264.97005     2  632.485027           Prob > F      =  0.0000
    Residual |  5895.44424   523  11.2723599           R-squared     =  0.1767
-------------+------------------------------           Adj R-squared =  0.1735
       Total |  7160.41429   525  13.6388844           Root MSE      =  3.3574

------------------------------------------------------------------------------
        wage |      Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
      tenure |   .2262657   .0218704    10.35   0.000      .183301    .2692303
     ftenure |  -.2292762   .0383293    -5.98   0.000    -.3045744   -.1539779
       _cons |   5.138208   .1809855    28.39   0.000      4.78266    5.493756
------------------------------------------------------------------------------
```

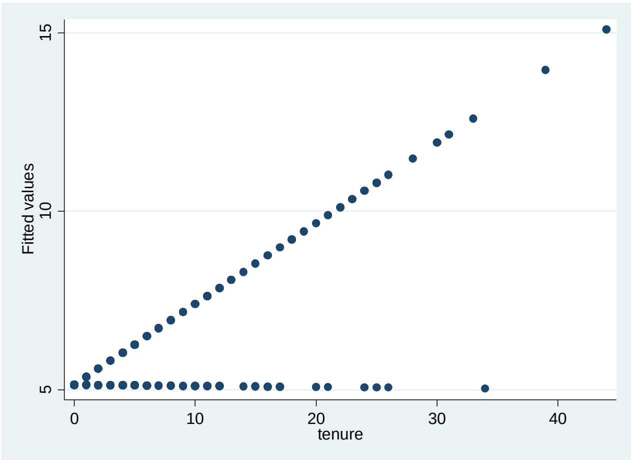

Predicted wages are: $\hat{w} = 5.14 + .227\ tenure - .229\ \ ftenure$.

- For females, $female = 1$, and so $\hat{w} = 5.14 + .227\ tenure - .229\ tenure = 5.14 - .002\ tenure$
- For males, $female = 0$, and so $\hat{w} = 5.14 + .227\ tenure - .229\ (0) = 5.14 + .227\ tenure$

The *ftenure* coefficient (-.229) is the **<u>difference</u>** in slopes (average incremental relationships between tenure and wages) between males and females. Since the t stat for *ftenure* is -5.98 and the p value is 0, it is easy to reject the null hypothesis that wages respond differently to changes in tenure for males and females.

If you want to test the null hypothesis that female wage do not respond to changes in tenure, just run the F test (of the Null hypothesis that the female tenure slope is zero)::

```
. test tenure+ftenure = 0

 ( 1)  tenure + ftenure = 0

       F(  1,   523) =    0.01
            Prob > F =    0.9340
```

As you can see, we cannot reject the null hypothesis as anywhere close to an attractive significance level. But let's not get too carried away with this… as we have yet to control for other explanatory factors.

7. Estimate two wage models and compare predictions

Working with only the female data, estimate a model that seeks to explain the variation in wages for females, and call the SRF from the model the femaleSRF. Do the same for males and then compare predictions. One measure of gender bias might be to compare predicted wages for females according to the two models. So for example, comparing average female wages if they are paid according to the female and male SRFs. And you might do something similar looking at males. The differences in average wages tell you something about gender bias.

Let's do that! … and start with the simple model with *tenure* as the single explanatory variable in the SLR models:

```
. *Female SRF:
. reg wage tenure if female==1
. predict fwhat

. *Male SRF:
. reg wage tenure if female==0
. predict mwhat
```

	(females) wage	(males) wage
tenure	0.0652*	0.180***
	(2.21)	(6.41)
_cons	4.352***	5.933***
	(22.80)	(19.98)
N	252	274
R-sq	0.019	0.131
adj. R-sq	0.015	0.128

t statistics in parentheses
* p<0.05, ** p<0.01, *** p<0.001

Gender Bias in Pay: Dummy variables in action

Here are the SRFs from the two models:

```
twoway (scatter mwhat tenure if female==0) (scatter fwhat tenure if female==1)
```

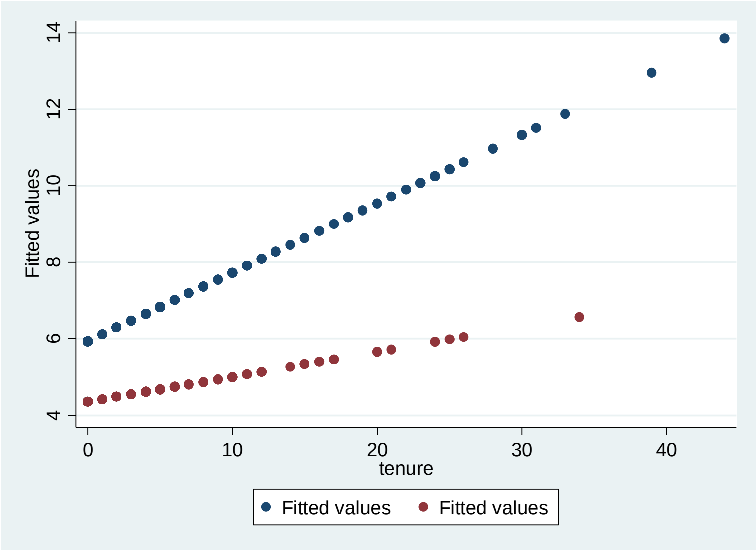

Predicted male wages are always above predicted female wages… and predicted male wages seem to increase at a more rapid pace with increases in tenure.

To estimate gender bias then we might just compare what men and women would be paid on average, under the two estimated SRFs:

```
. tabstat wage mwhat fwhat, by(female)

Summary statistics: mean
  by categories of: female

  female |      wage     mwhat     fwhat
---------+------------------------------
       0 |  7.099489  7.099489  4.774143
       1 |  4.587659  6.584172  4.587659
---------+------------------------------
   Total |  5.896103  6.852607  4.684801
----------------------------------------.
. tabstat wage mwhat fwhat, by(female)
```

Looking at females (*female*=1; the second row in the table): The average female wage is $4.59, which is also what is predicted under the femaleSRF (no surprise there). However if they were paid according to the maleSRF, their average predicted wage would be $6.58, implying an averaege gender gap of $1.99.

Looking at males (*female*=0; the first row in the table): The average male wage is $7.10, which is also what is predicted under the maleSRF (again, no surprise there). However if they were paid according to the femaleSRF, their average predicted wage would be $4.77, implying an average gender gap of $2.33.

8. Combining slope and intercept dummies

In the previous analysis, we effectively allowed for different slopes and intercepts for males and females. We can do that in one model if we interact *female* with the RHS variable *tenure*:

```
. gen ftenure = female*tenure
. reg wage female tenure ftenure

      Source |       SS           df       MS      Number of obs   =       526
-------------+----------------------------------   F(3, 522)       =     45.33
       Model |  1479.94868         3  493.316228   Prob > F        =    0.0000
    Residual |  5680.46561       522   10.882118   R-squared       =    0.2067
-------------+----------------------------------   Adj R-squared   =    0.2021
       Total |  7160.41429       525  13.6388844   Root MSE        =    3.2988

------------------------------------------------------------------------------
        wage |      Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
      female |  -1.580774    .355655    -4.44   0.000    -2.279465   -.8820832
      tenure |   .1802201   .0238554     7.55   0.000     .1333557    .2270844
     ftenure |  -.1150015   .0455993    -2.52   0.012    -.2045822   -.0254208
       _cons |   5.932663   .2521324    23.53   0.000     5.437344    6.427981
------------------------------------------------------------------------------
```

Predicted wages are: $\hat{w} = (5.93 - 1.58\ female) + (.180 - .115\ female)\ tenure$.

The implied intercepts and slopes from the model are:

Intercepts: female: 5.93 - 1.58 (1) = 4.35 male: 5.93 - 1.58 (0) = 5.93

Slopes: female: .180 - .115 (1) = .065 and male .180 -.115 (0) = .180

And so the estimated *female* coefficient (-1.58) is the difference in the intercepts, and the estimated *ftenure* coefficient (-.115) is the difference in slopes. And so as before, intercept and slope dummies capture differences between intercepts and slopes between males and females.

An F test allows us to test the joint null hypothesis that the male and female intercepts are the same, as are the two slopes… or put differently, that the differences in slopes and intercepts are zero:

```
. test female ftenure

 ( 1)  female = 0
 ( 2)  ftenure = 0

       F(  2,   522) =   28.41
            Prob > F =    0.0000
```

Reject, Reject, Reject!

This last test is called the **Chow Test**…. to which we will return later in the course.

Gender Bias in Pay: Dummy variables in action

Here's the SRF from the previous model:

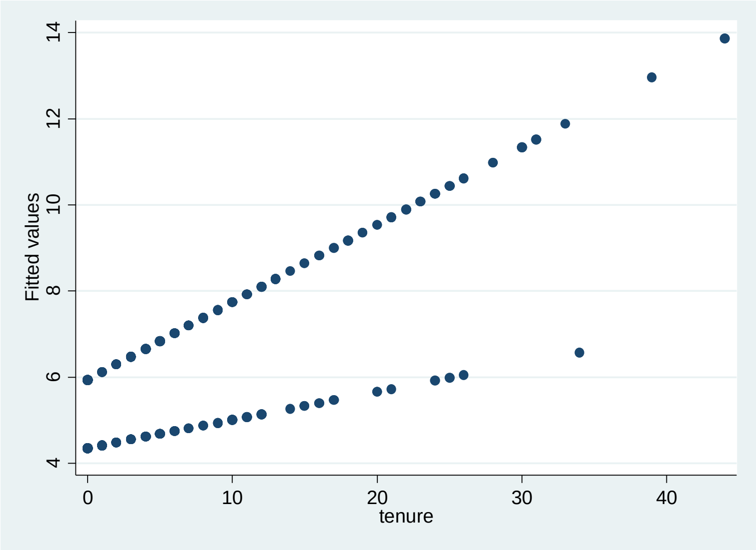

And Yes, you've seen this figure before. Perhaps not surprisingly, since this model effectively allows for different intercepts and slopes for males and females, the SRFs in this model are the same as the two SRFs in the previous approach in which we estimated two separate models (thereby allowing for separate intercepts and slopes for males and females).

Diff-in-Diff: As you can see in the SRFs, the estimated gender gap is expanding with increases in tenure. When *tenure* = 1, the predicted gender gap is about $1.70, and when *tenure* is 15, it is $3.30, almost twice the gap observed at tenure=1. The focus on how the difference in predicted wages, the estimated gender bias in this way-too-simplistic model, responds to, or is exacerbated by, changes in tenure levels is sometimes referred to as ***differences-in-differences***, or ***diff-in-diff*** for short. It can be a useful and powerful tool for understanding the impact of various factors, in this case tenure, on estimated bias.

9. The Kitchen Sink

Let's return to the first approach and since we want to control for everything else that might explain the wage differential, bring on the *Kitchen Sink* … adding in other variables including regional dummies, and second order terms for *tenure* and *exper*:

```
. reg wage female tenure tenure2 educ exper exper2 married nonwhite smsa south
northcen west

      Source |       SS           df       MS      Number of obs   =       526
-------------+----------------------------------   F(12, 513)      =     31.71
       Model |  3049.17602        12  254.098001   Prob > F        =    0.0000
    Residual |  4111.23827       513   8.0141097   R-squared       =    0.4258
-------------+----------------------------------   Adj R-squared   =    0.4124
       Total |  7160.41429       525  13.6388844   Root MSE        =    2.8309

------------------------------------------------------------------------------
        wage |      Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
      female |  -1.836211   .2565784    -7.16   0.000    -2.340285   -1.332138
      tenure |   .2037243   .0487076     4.18   0.000     .1080335    .2994152
     tenure2 |  -.0028757     .00167    -1.72   0.086    -.0061565    .0004052
        educ |    .486951   .0496408     9.81   0.000     .3894268    .5844752
       exper |   .1896114    .038007     4.99   0.000     .1149428      .26428
      exper2 |  -.0037989   .0008011    -4.74   0.000    -.0053727   -.0022251
     married |   .1816321   .2911276     0.62   0.533     -.390317    .7535812
    nonwhite |  -.1751238   .4122931    -0.42   0.671    -.9851145    .6348669
        smsa |   .7828516   .2917532     2.68   0.008     .2096736     1.35603
       south |   -.624951   .3438007    -1.82   0.070    -1.300382    .0504795
    northcen |  -.5810151   .3604689    -1.61   0.108    -1.289192    .1271617
        west |   .4038231   .3998333     1.01   0.313     -.381689    1.189335
       _cons |  -1.943356   .7590851    -2.56   0.011    -3.434654   -.4520579
------------------------------------------------------------------------------
```

And the conclusion remains virtually the same as before…the gender wage gap is now $1.84, and highly statistically significant, with a t stat well above the Supreme Court's threshold of two (2).

Gender bias in pay! No doubt!

… or maybe the model is seriously flawed and there's some logical explanation other than bias/discrimination. You never know until you've looked at everything!

Remember that the *female* coefficient picks up differences in average residuals between males and females, where the residuals are driven entirely by the rest of the model. If your model is A+, then maybe that estimated difference is worth paying attention to…. but if you have a crummy model, then no one should pay any attention to your estimate of the gender wage gap.

Or put differently: The quality of your estimate of the gender wage gap is entirely dependent on the quality of your model… and especially dependent on the extent to which you may not have accounted for important explanatory factors that drive compensation levels.

Think Endogeneity!… and worry as well about sample selection bias!

Speaking of which…

10. SINKS: Single Income No Kids/dependents

Suppose we focus on ***Single Income No Kids/dependents***… so let's select individuals with *numdep*=0 and *married* = 0, and rerun the *Kitchen Sink* model:

```
. reg wage female tenure tenure2 educ exper exper2 married nonwhite smsa south
northcen west if (numdep == 0 & married ==0)
note: married omitted because of collinearity

      Source |       SS           df       MS      Number of obs   =       123
-------------+----------------------------------   F(11, 111)      =      4.83
       Model |   401.13671        11  36.4669736   Prob > F        =    0.0000
    Residual |  838.079402       111  7.55026488   R-squared       =    0.3237
-------------+----------------------------------   Adj R-squared   =    0.2567
       Total |  1239.21611       122  10.1575091   Root MSE        =    2.7478

------------------------------------------------------------------------------
        wage |      Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
      female |  -.1616863   .5262951    -0.31   0.759    -1.204575    .8812025
      tenure |   .2745529   .1343989     2.04   0.043     .0082326    .5408732
     tenure2 |   -.007322   .0053599    -1.37   0.175    -.0179431    .0032991
        educ |   .4232028   .0998875     4.24   0.000     .2252691    .6211364
       exper |   .2026042   .0786833     2.57   0.011      .046688    .3585205
      exper2 |  -.0040294   .0016582    -2.43   0.017    -.0073153   -.0007435
     married |          0  (omitted)
    nonwhite |  -.5179759   .7993828    -0.65   0.518    -2.102006    1.066054
        smsa |    1.33732    .745724     1.79   0.076    -.1403819    2.815022
       south |  -.1659788   .7234019    -0.23   0.819    -1.599448     1.26749
    northcen |  -.2303097   .7090453    -0.32   0.746     -1.63533    1.174711
        west |  -.5208943   .7691747    -0.68   0.500    -2.045065    1.003277
       _cons |  -2.432175   1.531665    -1.59   0.115    -5.467271    .6029213
------------------------------------------------------------------------------
```

And the *female* dummy in the wage equation goes from -$1.84 and highly statistically significant to -$0.16 and having a p-value of .76. … *Interesting!*

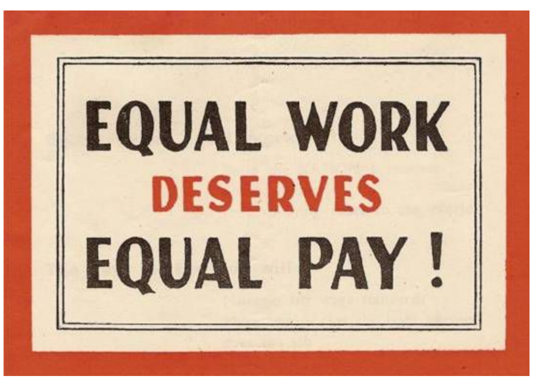

More work clearly needs to be done before any conclusions are reached. It's one thing to observe a wage gap…. and quite another to attribute that gap to gender bias/discrimination.

I'm not saying that there is no gender bias/discrimination in wages/compensation.

But I am saying: ***This is indeed very complicated!***

It's really very simple! ~~Yes?~~ No!

If you want to learn more about the topic, I recommend the 2017 *Journal of Economic Literature* survey piece by Blau and Kahn (a copy of this paper has been posted to Canvas):

Blau, Francine D., and Lawrence M. Kahn. 2017. "The Gender Wage Gap: Extent, Trends, and Explanations." Journal of Economic Literature, 55 (3): 789-865. https://www.aeaweb.org/articles?id=10.1257/jel.20160995