Computer Exercise 3 Answers – Hypothesis Testing

```
. reg lnhpay xper yearsed tenure

  Source |       SS       df       MS              Number of obs =    6225
---------+------------------------------           F(  3,  6221) =  512.58
   Model |  457.732594     3  152.577531           Prob > F      =  0.0000
Residual |  1851.79026  6221  .297667619           R-squared     =  0.1982
---------+------------------------------           Adj R-squared =  0.1978
   Total |  2309.52285  6224  .371067296           Root MSE      =  .54559

------------------------------------------------------------------------------
  lnhpay |      Coef.   Std. Err.       t     P>|t|       [95% Conf. Interval]
---------+--------------------------------------------------------------------
    xper |   .0054206    .000663      8.176   0.000       .0041209    .0067202
 yearsed |   .0740664   .0023269     31.830   0.000       .0695048    .0786279
  tenure |   .0171891   .0009445     18.199   0.000       .0153375    .0190406
   _cons |   .6160226   .0372179     16.552   0.000       .5430626    .6889826
------------------------------------------------------------------------------
```

You should by now be able to look at the regression output and work out that a test of a hypothesis for a single coefficient can be done by calculating the appropriate t statistic. The regression output always gives the t value for the hypothesis that any coefficient is zero.

In this case the t value for the hypothesis that the true coefficient on the experience variable is 8.176 (= .0054/.00066) which given degrees of freedom 6225-5 =6221 means can reject null hypothesis at the 5% level of significance. Note the p value in the output is the exact level of significance at which the hypothesis can be rejected. In the above case this p value is so small that it registers as zero, (see the column after the t values above).

To test b=.001, calculate t=(.0054-.001)/.00066 = 6.67
Again can reject null that true value is zero

(Often when hypothesised values are close together will often reject or not reject both hypotheses).

Stata does the above tests as follows:

```
. test xper=0

 ( 1)  xper = 0.0

       F(  1,  6221) =   66.85
            Prob > F =    0.0000

. test xper=0.001

 ( 1)  xper = .001

       F(  1,  6221) =   44.46
            Prob > F =    0.0000
```

The difference is that Stata uses the F version of the test (see lecture notes). You should verify that these F values are the square of the t values you calculated earlier. The numbers on the 2nd line are the p values of the test (values below .05 (5%) mean the null can be rejected at the 5% level of significance). In these cases the predicted F values are so large that they easily reject the null at the 5% level of significance (the p value is so low that it does not register)

To test equality of coefficients (again a single linear hypothesis), use

```
. test xper=tenure

 ( 1)  xper - tenure = 0.0

       F(  1,  6221) =   73.57
            Prob > F =    0.0000
```

Predicted F value exceeds critical level so reject null.

Again you should check with your lecture notes to confirm that this test uses the hypothesised coefficients of the two variables weighted by their respective variance and covariances,
(in this case Var(xper) + Var(tenure) - Cov(xper,tenure)

```
. test xper+tenure=.02

 ( 1)  xper + tenure = .02

       F(  1,  6221) =    8.72
            Prob > F =    0.0032
```

Predicted F value exceeds critical level so reject null. (p value = 0.3%)
(in this case the weighting used is Var(xper) + Var(tenure) + Cov(xper,tenure)

The next test is a test that all the coefficients in the model (except the constant) are zero. This is a test of the goodness of fit of the model as a whole and so can be calculated directly from the R^2 value in the regression output above.

Using $F = (R^2 /k-1)/((1- R^2)/N-k) \sim F(k-1, N-k)$

$= (0.198/4-1)/((1-.198)/6225-4) \sim F(4-1, 6225-4)$

$= 512$

Stata does this using the command

```
. test xper yearsed tenure

 ( 1)  xper = 0.0
 ( 2)  yearsed = 0.0
 ( 3)  tenure = 0.0

       F(  3,  6221) =  512.58
            Prob > F =    0.0000
```

Again easily reject null that all coefficients are zero, so model has statistically significant explanatory power.

The next test is a test of subsets of the regression coefficients so we can use the form of the F test

$$F = \frac{RSSrestrict - RSSunrestrict /J}{RSSunrestrict /N- Kunrestrict} \sim F(J, N\text{-}Kunrestrict)$$

Where j is the number of restricted coefficients

The unrestricted RSS can be found in the regression output above. The restricted RSS can be found by running the following restricted regression.

```
. reg lnhpay yearsed

  Source |       SS       df       MS              Number of obs =    6225
---------+------------------------------           F(  1,  6223) =  797.86
   Model |   262.45794     1   262.45794           Prob > F      =  0.0000
Residual |  2047.06491  6223  .328951456           R-squared     =  0.1136
---------+------------------------------           Adj R-squared =  0.1135
   Total |  2309.52285  6224  .371067296           Root MSE      =  .57354

------------------------------------------------------------------------------
  lnhpay |      Coef.   Std. Err.       t     P>|t|       [95% Conf. Interval]
---------+--------------------------------------------------------------------
 yearsed |   .0637578   .0022572     28.246   0.000       .0593329    .0681827
   _cons |   1.015405    .029999     33.848   0.000       .9565962    1.074213
------------------------------------------------------------------------------
```

So F = $\frac{2047 - 1852 /2}{1852 /6225 - 4}$ ~ F(2, 6225-4)

= 328 ~ F(2, 6221)

The automated version of the test is given by

```
. test xper tenure

 ( 1)  xper = 0.0
 ( 2)  tenure = 0.0

       F(  2,  6221) =  328.01
            Prob > F =    0.0000
```

Again the F value easily exceeds the critical 5% value so we can reject the null (that the 2 variables have no explanatory power) at the 5% level of significance.

The final test is a form of the structural break test (testing that the coefficients are equal for the male and female sub-samples)

Regression for men:

```
. reg lnhpay xper yearsed tenure if female==0

  Source |       SS       df       MS              Number of obs =    3052
---------+------------------------------           F(  3,  3048) =  254.41
   Model |  214.059947     3  71.3533156           Prob > F      =  0.0000
Residual |  854.863086  3048  .280466892           R-squared     =  0.2003
---------+------------------------------           Adj R-squared =  0.1995
   Total |  1068.92303  3051  .350351699           Root MSE      =  .52959

------------------------------------------------------------------------------
  lnhpay |      Coef.   Std. Err.       t     P>|t|       [95% Conf. Interval]
---------+--------------------------------------------------------------------
    xper |   .0100297   .0009035     11.101   0.000       .0082582    .0118013
 yearsed |   .0706268   .0030956     22.815   0.000       .0645572    .0766964
  tenure |   .0116025   .0011813      9.822   0.000       .0092863    .0139188
   _cons |   .7426208   .0497533     14.926   0.000       .6450674    .8401741
------------------------------------------------------------------------------
```

regression for women:

```
. reg lnhpay xper yearsed tenure if female==1

  Source |       SS       df       MS              Number of obs =    3173
---------+------------------------------           F(  3,  3169) =  263.51
   Model |  219.977658     3  73.3258859           Prob > F      =  0.0000
Residual |  881.832729  3169  .278268453           R-squared     =  0.1997
---------+------------------------------           Adj R-squared =  0.1989
   Total |  1101.81039  3172  .347355103           Root MSE      =  .52751

------------------------------------------------------------------------------
  lnhpay |      Coef.   Std. Err.       t     P>|t|       [95% Conf. Interval]
---------+--------------------------------------------------------------------
    xper |   .0016803   .0009192      1.828   0.068      -.0001219    .0034826
 yearsed |   .0745759   .0032974     22.617   0.000       .0681107    .0810412
  tenure |   .0196738   .0014838     13.259   0.000       .0167645    .0225832
   _cons |   .5477843   .0524783     10.438   0.000       .4448894    .6506791
------------------------------------------------------------------------------
```

Now this is also an f test of the form

$$F = \frac{RSSrestrict - RSSunrestrict / J}{RSSunrestrict / N - Kunrestrict} \sim F(J, N\text{-}Kunrestrict)$$

Where now:
RSSunrestrict = RSSmen + RSSwomen = 855 + 882 (from above output)
RSSrestrict = RSS from the original regression which combines men & women = 1852
J = no. restricted coefficients = 4 (constant, xper, yearsed & tenure)
Kunrestrict = 8

So F = (1852 – (855+882)/4)/((855+882)/6225-8) ~ F(4,6217)

= 103

which again exceeds the 5% critical value = 2.37(you should be able to find these critical values from statistical tables). So can again reject null that coefficients are equal in the 2 sub-samples.
From the regression output you should be able to see that the coefficient on experience, for example, is much larger in the male regression.

You can do a Chow test in Stata by recognising that running a single unrestricted regression amounts to letting there be separate coefficients on each variable (including the constant) for men & women.

$$Ln(Hourpay)_i = b_0 + b_1 xper_i + b_2 yearsed_i + b_3 tenure_i$$
$$+ b_4 female + b_5 xper*female_i + b_6 yearsed*female_i + b_7 tenure*female_i + e_i$$

You can do this by creating interactions of the female dummy variable with the other variables and adding them to the original regression

```
. g femxp=female*xper
. g femed=female*yearsed
. g femten=female*tenure
```

```
. reg lnhpay xper yearsed tenure female femxp femed femten

  Source |       SS       df       MS                  Number of obs =    6225
---------+------------------------------               F(  7,  6217) =  292.94
   Model |  572.827038     7  81.8324341               Prob > F      =  0.0000
Residual |  1736.69581  6217  .279346279               R-squared     =  0.2480
---------+------------------------------               Adj R-squared =  0.2472
   Total |  2309.52285  6224  .371067296               Root MSE      =  .52853

------------------------------------------------------------------------------
  lnhpay |      Coef.   Std. Err.       t     P>|t|       [95% Conf. Interval]
---------+--------------------------------------------------------------------
    xper |   .0100297   .0009017     11.123   0.000       .0082621    .0117974
 yearsed |   .0706268   .0030894     22.861   0.000       .0645705     .076683
  tenure |   .0116025    .001179      9.841   0.000       .0092914    .0139137
  female |  -.1948365   .0723197     -2.694   0.007      -.3366081   -.0530649
   femxp |  -.0083494   .0012889     -6.478   0.000      -.0108761   -.0058228
   femed |   .0039491   .0045232      0.873   0.383      -.0049179    .0128161
  femten |   .0080713   .0018974      4.254   0.000       .0043517    .0117909
   _cons |   .7426208   .0496538     14.956   0.000       .6452822    .8399593
------------------------------------------------------------------------------
```

Then test the hypothesis that the coefficients on all the terms involving the female dummy are jointly zero.

```
. test female femxp femed femten

 ( 1)  female = 0.0
 ( 2)  femxp = 0.0
 ( 3)  femed = 0.0
 ( 4)  femten = 0.0

       F(  4,  6217) =  103.00
            Prob > F =    0.0000
```