1. Relating Confidence Intervals and Hypothesis Tests for <u>Means</u> and <u>Parameter Estimates</u>

Assumptions satisfied	**Property of estimators $\hat{\boldsymbol{\beta}}_j$**	**Why this is important**
MLR.1-MLR.4	$E(\hat{\beta}_j) = \beta_j$	OLS estimates are unbiased.
MLR.1-MLR.5 [MLR.5: $Var(u \mid x_1, \ldots, x_k) = \sigma^2$] ("Homoskedasticity")	$Var(\hat{\beta}_j) = \dfrac{\sigma^2}{SST_j(1 - R_j^2)}$	The standard errors we get from Stata regressions are right.
MLR.1-MLR.6 [MLR.6: $u \sim \mathcal{N}(0, \sigma^2)$]	$\dfrac{\hat{\beta}_j - \beta_j}{SE(\hat{\beta}_j)} \sim t_{n-k-1}$	We can do hypothesis tests (t, F) for the parameter values and construct confidence intervals.

We learned the first result a long time ago. The second doesn't give us a lot of insight, just that for this particular formula for the variance of the estimators to be right, we need the unobservables to have the same variance regardless of what the *x*'es are.

The third result is what interests us for hypothesis testing. It says that if the unobservables follow a normal distribution (again, an assumption, we can't explicitly verify this), then the $\hat{\beta}_j$'s follow a particular distribution too.

Let's compare confidence intervals and test statistics between population means and regression parameters:

	Confidence interval	**Test statistics**
Population mean	$CI = \left[\bar{x} - c(SE(\bar{x})), \bar{x} + c(SE(\bar{x}))\right]$	$\dfrac{\bar{x} - \mu}{SE(\bar{x})} \sim t_{n-1}$
Regression parameters	$CI = \left[\hat{\beta}_j - c\left(SE(\hat{\beta}_j)\right), \hat{\beta}_j + c\left(SE(\hat{\beta}_j)\right)\right]$	$\dfrac{\hat{\beta}_j - \beta_j}{SE(\hat{\beta}_j)} \sim t_{n-k-1}$

As you see, they are almost the same. In fact, once you have the appropriate standard error, you follow the exact same procedures for constructing a confidence interval or hypothesis testing.

2. Confidence Intervals and Hypothesis Tests for One Regression Parameter

The steps for these will look very familiar.

Example:
Here are the results of a regression of Michigan State University students' cumulative GPAs on their ACT score, high school GPAs, gender, and whether or not they were engineering or business majors.

```
      Source |       SS       df       MS              Number of obs =     141
-------------+------------------------------           F(  5,   135) =    6.63
       Model |  3.82623697     5  .765247395           Prob > F      =  0.0000
    Residual |  15.5798625   135  .115406389           R-squared     =  0.1972
-------------+------------------------------           Adj R-squared =  0.1674
       Total |  19.4060994   140  .138614996           Root MSE      =  .33972

------------------------------------------------------------------------------
      colGPA |      Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
         ACT |   .0102472   .0111827
       hsGPA |   .4666345   .1003923
    engineer |  -.2147078   .1686759
    business |   .0603423   .0764355
        male |   .0043776    .061257
       _cons |    1.17904   .3535725
------------------------------------------------------------------------------
```

Confidence interval for $\boldsymbol{\beta_{ACT}}$:
We know $CI_{95} = \left[\hat{\beta}_{ACT} - c_{95}\left(SE(\hat{\beta}_{ACT})\right), \hat{\beta}_{ACT} + c_{95}\left(SE(\hat{\beta}_{ACT})\right)\right]$, and Stata gave us $\hat{\beta}_{ACT}$ and $SE(\hat{\beta}_{ACT})$. We just need c_{95}:

n = ______________ k = ______________ n-k-1 = ______________ c_{95} = _____________

So $CI_{95} = \left[\qquad -1.98\left(\qquad \right), \qquad +1.98\left(\qquad \right)\right] = \left[\qquad , \qquad \right]$

Hypothesis test for $\boldsymbol{\beta_{ACT} = 0}$:

STEP 1. Define hypotheses:

$$H_0: \beta_{ACT} = 0$$
$$H_1: \beta_{ACT} \neq 0$$

Under the null hypothesis, β_{ACT} is 0, so we assume this is true for now. If true, then $\frac{\hat{\beta}_{ACT} - 0}{SE(\hat{\beta}_{ACT})} \sim t_{n-k-1}$.

STEP 2. Compute the test statistic:
Use the information from the Stata output:

$$\frac{\hat{\beta}_{ACT} - \beta_{ACT}}{SE(\hat{\beta}_{ACT})} = \frac{\qquad -0}{\qquad\qquad} =$$

STEP 3. Get the significance level of the test:
Here we'll choose the 5% significance level, meaning we'll wrongly reject a **true** H_0 5% of the time. We go to the t-table and find out what the corresponding critical value (c) is. It's about 1.98 for *n-k-1* = 141-5-1 = 135.

STEP 4. Reject the null hypothesis or fail to reject it:

Our critical value, c, is 1.98. Our t-statistic was ___________.

If our $|t_{n-k-1}| > 1.98$ then we reject the null hypothesis because our $\hat{\beta}_{ACT}$ was so far away from the null hypothesis of 0. If $|t_{n-k-1}| < 1.98$ then we can't reject the null hypothesis because $\hat{\beta}_{ACT}$ is close enough to the null hypothesis of 0 that we can't say it's wrong with enough confidence.

Did we reject H_0? YES NO

STEP 5. Interpret:
If reject:
There is statistical evidence at the 5% significance level that, **conditional upon high school GPA, gender, and type of major**, ACT score affects college GPA.

If fail to reject:
There is no statistical evidence at the 5% significance level that, **conditional upon high school GPA, gender, and type of major**, ACT score affects college GPA.

Note: here we tested for the β_j being equal to zero. This is the hypothesis test for which Stata reports the t-statistic and p-value, because it's the most common test for us to care about. But nothing stops us from doing tests for β_j being equal to some other number besides zero.

3. Hypothesis Testing for Multiple Parameters

Sometimes we want to do hypothesis tests involving more than one of the β_j's from a regression. We can do this by performing an F-test. F-tests can be used for many kinds of hypotheses, some of them being pretty complex and interesting. So far, we've just talked about testing that multiple β_j's are equal to zero.

If you know how to do a t-test for one β_j, then you already know how to do most of an F-test. You'll still follow the five steps of hypothesis tests. Here are some key differences between t-tests and F-tests for regressions:

t-test	**F-test**
Hypotheses are about **one** β	Hypotheses are about one **or more** β's
Calculate the **t-statistic**; compare to critical **t-value**	Calculate the **F-statistic**; compare to critical **F-value**
Choose one-tailed or two-tailed test	Do not choose tails – just reject if F > c
Has direct parallels to a confidence interval	No direct parallel to a confidence interval

Even though the F-test can involve hypotheses about many parameters, an F-test isn't more complicated to perform than a t-test.

Example:
A reasonable person might think that, even after controlling for initial student ability, a student's major in college will affect his/her GPA. For example, engineering courses might be harder and result in lower grades than humanities or business courses. Rather than speculate, let's use data to test this (at MSU at least):

Regression including dummy variables for majors (UNRESTRICTED):

```
      Source |       SS       df       MS              Number of obs =     141
-------------+------------------------------           F(  5,   135) =    6.63
       Model |  3.82623697     5  .765247395           Prob > F      =  0.0000
    Residual |  15.5798625   135  .115406389           R-squared     =  0.1972
-------------+------------------------------           Adj R-squared =  0.1674
       Total |  19.4060994   140  .138614996           Root MSE      =  .33972

------------------------------------------------------------------------------
      colGPA |      Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
         ACT |   .0102472   .0111827     0.92   0.361    -.0118687    .0323631
       hsGPA |   .4666345   .1003923     4.65   0.000     .2680895    .6651796
    engineer |  -.2147078   .1686759    -1.27   0.205    -.5482967    .1188812
    business |   .0603423   .0764355     0.79   0.431    -.0908236    .2115082
        male |   .0043776    .061257     0.07   0.943    -.1167699    .1255251
       _cons |    1.17904   .3535725     3.33   0.001     .4797823    1.878298
------------------------------------------------------------------------------
```

Regression excluding dummy variables for majors, i.e. restricting $\beta_{engineer} = 0, \beta_{business} = 0$ (RESTRICTED):

```
      Source |       SS       df       MS              Number of obs =     141
-------------+------------------------------           F(  3,   137) =    9.79
       Model |  3.42521974     3  1.14173991           Prob > F      =  0.0000
    Residual |  15.9808797   137  .116648757           R-squared     =  0.1765
-------------+------------------------------           Adj R-squared =  0.1585
       Total |  19.4060994   140  .138614996           Root MSE      =  .34154

------------------------------------------------------------------------------
      colGPA |      Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
         ACT |   .0097797   .0112388     0.87   0.386    -.0124442    .0320036
       hsGPA |   .4500655   .1005145     4.48   0.000      .251305     .648826
        male |  -.0070857   .0611797    -0.12   0.908    -.1280643    .1138929
       _cons |   1.293037   .3469184     3.73   0.000     .6070302    1.979045
------------------------------------------------------------------------------
```

STEP 1. Define hypotheses:

$$H_0: \beta_{engineer} = 0 \; AND \; \beta_{business} = 0$$
$$H_1: \beta_{engineer} \neq 0 \; OR \; \beta_{business} \neq 0$$

where "OR" can mean that *one* or *both* of the parameters is non-zero.

STEP 2. Compute the test statistic:

We can compute the F-statistic either using the residual sum of squares (RSS) or the R^2. Both give the same F-statistic but most people prefer to work with R^2 as it is a small number to put into the calculator.

METHOD 1: RSS.

$$F_{q,n-k-1} = \frac{(SSR_{restricted} - SSR_{unrestricted})/q}{SSR_{unrestricted}/(n-k-1)}$$

where: q is the number of restrictions being imposed (here, there are two: $\beta_{engineer} = 0$ and $\beta_{business} = 0$);
n is the number of observations (here, 141);
k is the number of explanatory variables in the *unrestricted* regression (here, 5).

$$F_{2,141-5-1} = \frac{(\qquad - \qquad)/2}{\qquad /(141-5-1)} = \_\_\_\_\_\_\_\_\_\_$$

METHOD 2: R^2.

$$F_{q,n-k-1} = \frac{(R^2_{unrestricted} - R^2_{restricted})/q}{(1 - R^2_{unrestricted})/(n-k-1)}$$

$$F_{2,141-5-1} = \frac{(\qquad - \qquad)/2}{(\qquad)/(141-5-1)} = \_\_\_\_\_\_\_\_\_\_$$

Notice that the order of the unrestricted and restricted terms in the numerator flipped between the methods. When in doubt, check the formula. You'll know you were wrong if you get a negative F-statistic.

STEP 3. Get the significance level of the test:

Here we'll choose the 10% significance level, meaning we'll wrongly reject a **true** H_0 10% of the time. We go to the F-tables and find out what the corresponding critical value (c) is when we have *2 degrees of freedom in the numerator (q)* and *135 degrees of freedom in the denominator (n – k – 1)*. It's about **2.35**.

Note on the F-tables: there is one table for each significance level. The rows and columns correspond to the degrees of freedom in the numerator and denominator (*q* and *n – k – 1*, respectively).

STEP 4. Reject the null hypothesis or fail to reject it:

Our critical value, c, is 2.35. Our F-statistic was ___________.

If our $F_{q,n-k-1} > 2.35$ then we reject the null hypothesis because our $\hat{\beta}$'s were collectively too far away from the null hypothesis of 0. If $F_{q,n-k-1} < 2.35$ then we can't reject the null hypothesis because our $\hat{\beta}$'s are close enough to the null hypothesis of 0 that we can't say the null is wrong with enough confidence.

Did we reject H_0? YES NO

STEP 5. Interpret:
If reject:
There is statistical evidence at the 10% significance level that, conditional upon high school GPA, gender, and ACT score, college major affects college GPA.

If fail to reject:
There is no statistical evidence at the 10% significance level that, conditional upon high school GPA, gender, and ACT score, college major affects college GPA.

FYI: Here is the distribution of the F-statistic with 2 and 135 degrees of freedom: