Testing statistical hypotheses is one of the main tasks of an applied researcher.

We have derived the sampling distribution of the $\hat{\beta}$s because they are essential to hypothesis testing. Recall, for $\hat{\beta}_1$ (in a model with only one X variable) we had:

$$\hat{\beta}_1 \sim N\left(\beta_1, Var(\hat{\beta}_1)\right), \text{ where } Var(\hat{\beta}_1) = \frac{\sigma^2}{\sum_{i=1}^{n}(X_i - \bar{X})^2}$$

In the general model with k independent variables,

$$Y_i = \beta_0 + \beta_1 X_{1i} + \beta_2 X_{2i} + \dots + \beta_k X_{ki} + \varepsilon_i$$

we had:

$$\hat{\beta}_j \sim N\left(\beta_j, Var(\hat{\beta}_j)\right), \text{ for } j = 0,1,\dots,k$$

Remember we are assuming $\varepsilon_i \sim N(0, \sigma^2)$.

Our goal is to learn something about the population with our single sample. With hypothesis testing what we are trying to do is figure out how likely our sample result could have been obtained by chance.

We reviewed testing hypotheses in Chapter 17.

All types of hypothesis testing consist of three main steps:

1. Specify two opposing hypotheses.
   - Specify H_0 and H_A
   - Choose level of significance α

2. Derive a test statistic and its statistical distribution under the null hypothesis.

3. Derive a decision rule to use in deciding whether to reject the null hypothesis.
   - Compare the value of the test statistic to the critical value
   - If test statistic is larger (in absolute value) than the critical value, we reject H_0, otherwise we do not reject H_0
   - Type I and Type II errors

Alternative testing methods: p-values, confidence intervals

We worked all this out for testing the population mean μ in Chapter 17. Here our interest is on testing the individual βs.

Examples of null and alternative hypotheses for the βs:

	(a)	(b)	(c)
H_0	$\beta = \beta^*$	$\beta \leq \beta^*$	$\beta \geq \beta^*$
H_A	$\beta \neq \beta^*$	$\beta > \beta^*$	$\beta < \beta^*$

How do we construct a test statistic?

We know $\hat{\beta}_j \sim N\left(\beta_j, Var(\hat{\beta}_j)\right)$. We can standardize $\hat{\beta}_j$ so that

$$Z = \frac{\hat{\beta}_j - \beta_j}{\sqrt{Var(\hat{\beta}_j)}} \sim N(0,1)$$

Suppose we want to test the null and alternative hypotheses given in (b). If the null is true, then

$$Z = \frac{\hat{\beta}_j - \beta_j^*}{\sqrt{Var(\hat{\beta}_j)}} \sim N(0,1)$$

Using the critical value, we would reject H_0 is the value of Z was in the rejection region.

Unfortunately, we typically do not know $Var(\hat{\beta}_j)$ so we can't use Z. In the previous lecture, we used $\widehat{Var}(\hat{\beta}_j)$ as an unbiased estimator of $Var(\hat{\beta}_j)$ and se($\hat{\beta}_j$) as an unbiased estimator of $\sqrt{Var(\hat{\beta}_j)}$. If we plug se($\hat{\beta}_j$) into Z, we obtain:

$$t = \frac{\hat{\beta}_j - \beta_j}{se(\hat{\beta}_j)}$$

If $\varepsilon_i \sim N(0, \sigma^2)$ then $t \sim t_{n-k-1}$.

We could have conducted the hypothesis test using a p-value. Alternatively, for a two-sided hypothesis, we could have used a confidence interval:

$$\Pr[\hat{\beta}_j - se(\hat{\beta}_j)t_{n-k-1,\alpha/2} \leq \beta_j \leq \hat{\beta}_j + se(\hat{\beta}_j)t_{n-k-1,\alpha/2}]$$

Of particular interest is testing a regression coefficient is not equal to zero:

	(a)	(b)	(c)
H_0	$\beta = 0$	$\beta \leq 0$	$\beta \geq 0$
H_A	$\beta \neq 0$	$\beta > 0$	$\beta < 0$

Using a *t*-test:

$$t = \frac{\hat{\beta}_j}{se(\hat{\beta}_j)}$$

The value of this statistic is what is reported in EViews output. Two-sided p-values are also reported for this test.

Some general caution about the *t*-test:

- Statistical significance does not mean theoretical significance
- *t*-test does not test importance
- *t*-test not intended for tests of entire population

Testing for Overall Goodness of Fit

The *t*-test is used to test the significance of **individual** regression coefficients. Often, there is interest in testing the **joint significance** of **ALL** slope coefficients.

The ***F*-test** provides us with a formal statistical test for overall fit. Recall, we previously used R^2 and $\bar{R}^2$ to assess fit but these were descriptive measures and not formal tests.

To test the overall goodness of fit of a model, we test whether all the slope coefficients are simultaneously equal to zero:

$$H_0: \beta_1 = \beta_2 = \cdots = \beta_k = 0$$

$$H_A: \text{at least one } \beta_j \neq 0 \text{ for } j = 1,2,\ldots,k$$

Think about what a true null hypothesis implies. If all the slope coefficients are equal to zero then the new regression model would be

$$Y_i = \beta_0 + \varepsilon_i$$

for which the OLS estimator $\hat{\beta}_0 = \bar{Y}$.

So when H_0 is true we have

$$\hat{Y}_i = \bar{Y}.$$

The F-test will test then whether the fit of the regression equation is significantly better than that provided by the mean alone.

The F-statistic used for testing is

$$F = \frac{ESS/k}{RSS/(n-k-1)} \sim F_{k,n-k-1}$$

Critical values can be found in Tables B-2 and B-3.

This value of the F-statistic is reported in EViews regression output, along with its associated p-value.