Intermediate Econometrics 04

Chen Zhu

College of Economics and Management
China Agricultural University

Testing other hypotheses about β_j

1. We sometimes want to test whether β_j is equal to some other given constant.
2. Generally, if the null is stated as

$$H_0: \ \beta_j = a_j. \tag{1}$$

where a_j is our hypothesized value of β_j, then the appropriate t statistic is

$$t = \frac{\hat{\beta}_j - a_j}{se(\hat{\beta}_j)}. \tag{2}$$

Testing other hypotheses about β_j

1. The general t statistic is usefully written as

$$t = \frac{estimate - hypothesized\ value}{standard\ error}. \tag{3}$$

2. t measures how many estimated standard deviations $\hat{\beta}_j$ is away from the hypothesized value of β_j.

Computing p-values for t tests

The classical testing approach:

1. After stating the alternative hypothesis, we choose a significance level, which then determines a critical value.
2. Once the critical value has been identified, the value of the t statistic is compared with the critical value, and the null is either rejected or failed to reject at the given significance level.

Computing p-values for t tests

1. Rather than testing at different significance levels, it is more informative to answer the following question: Given the observed value of the t statistic, what is the smallest significance level at which the null hypothesis would be rejected?
2. This level is known as the p-value for the test, which is a probability.
   - If the null is not rejected at the 5% level, we know the p-value is greater than .05.
   - If the null is rejected at the 10% level, we know that the p-value is less than .10.

Computing p-values for t tests

1. Most modern regression packages have the capability to compute areas under the probability density function of the t distribution.
2. If a regression package reports a p-value along with the standard OLS output, the p-value in this case is

$$\text{p-value} = P(|T| > |t|), \tag{4}$$

where T denote a t distributed random variable with n-k-1 degrees of freedom and t denote the numerical value of the test statistic.

- The p-value is the probability of observing a t statistic as extreme as we did if the null hypothesis is true.
- This means that small p-values are evidence against the null; large p-values provide little evidence against H_0.

Computing p-values for t tests

1. e.g. In the example with df=40 and t=1.85, the p-value is computed as .0718.
    - This means that, if the null hypothesis is true, we would observe an absolute value of the t statistic as large as 1.85 about 7.2% of the time.
    - This provides some evidence against the null hypothesis, but we would not reject the null at the 5% significance level.
    - Once the p-value has been computed, a classical test can be carried out at any desired level.

Computing p-values for t tests

Obtaining the *p*-value against a two-sided alternative, when $t = 1.85$ and $df = 40$.

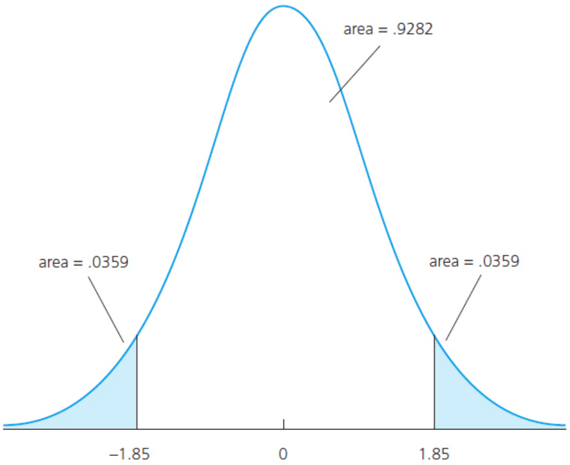

Confidence intervals

1. Under the classical linear model assumptions, we can easily construct a confidence interval (CI) for the population parameter β_j .
2. Using the fact that $(\hat{\beta}_j - \beta_j)/se(\hat{\beta}_j)$ has a t distribution with n-k-1 degrees of freedom, simple manipulation leads to a CI for the unknown β_j: a 95% confidence interval, given by

$$CI = \hat{\beta}_j \pm c \cdot se(\hat{\beta}_j), \tag{5}$$

   where the constant c is the 97.5% percentile in a t_{n-k-1} distribution.

   - As an example, for df=25, a 95% confidence interval for any β_j is given by $[\hat{\beta}_j - 2.06 \cdot se(\hat{\beta}_j),\ \hat{\beta}_j + 2.06 \cdot se(\hat{\beta}_j)]$.

Testing multiple linear restrictions: the F test

1. The t statistic associated with any OLS coefficient can be used to test whether the corresponding unknown parameter in the population is equal to any given constant.
2. Frequently, we wish to test multiple hypotheses about the underlying parameters $\beta_0, \beta_1, \cdots, \beta_k$.
3. We begin with the leading case of testing whether a set of independent variables has no partial effect on a dependent variable.

Testing exclusion restrictions

1. We want to test whether a group of variables has no effect on the dependent variable.
2. More precisely, the null hypothesis is that a set of variables has no effect on y, once another set of variables has been controlled.
3. The unrestricted model with k independent variables as

$$y = \beta_0 + \beta_1 x_1 + \cdots + \beta_k x_k + u; \tag{6}$$

the number of parameters in the unrestricted model is k+1.

Testing exclusion restrictions

1. Suppose that we have q exclusion restrictions to test: that is, the null hypothesis states that q of the variables have zero coefficients.
2. For notational simplicity, assume that it is the last q variables in the list of independent variables: $x_{k-q+1}, \cdots, x_k$.
3. The null hypothesis is stated as

$$H_0 := \beta_{k-q+1} = 0, \cdots, \beta_k = 0, \tag{7}$$

   which puts q exclusion restrictions on the model above.
4. The alternative H_1 is that at least one of the parameters listed is different from zero.

Testing exclusion restrictions

1. When we impose the restrictions under H_0, we are left with the restricted model:

$$y = \beta_0 + \beta_1 x_1 + \cdots + \beta_{k-q} x_{k-q} + u. \tag{8}$$

2. The F statistic (or F ratio) is defined by

$$F = \frac{(\text{SSR}_r - \text{SSR}_{ur})/q}{\text{SSR}_{ur}/(n-k-1)}, \tag{9}$$

where SSR_r is the sum of squared residuals from the restricted model and SSR_{ur} is the sum of squared residuals from the unrestricted model.

Testing exclusion restrictions

1. Since SSR_r can be no smaller than SSR_{ur}, the F statistic is always nonnegative (and almost always strictly positive).
2. The easiest way to remember where the SSRs appear is to think of F as measuring the relative increase in SSR when moving from the unrestricted to the restricted model.
3. q is the number of restrictions imposed in moving from the unrestricted to the restricted model (q independent variables are dropped):

$$q = \text{numerator degrees of freedom} = df_r - df_{ur}, \tag{10}$$

which also shows that q is the difference in degrees of freedom between the restricted and unrestricted models.

Testing exclusion restrictions

1. Notice that the restricted model has fewer parameters.
2. The SSR in the denominator of F is divided by the degrees of freedom in the unrestricted model:

$$n - k - 1 = \text{denominator degrees of freedom} = df_{ur}. \quad (11)$$

3. In fact, the denominator of F is just unbiased estimator of σ^2=Var(u) in the unrestricted model.

Testing exclusion restrictions

1. It is often more convenient to have a form of the F statistic that can be computed using the R-squareds from the restricted and unrestricted models.
2. The R-squared form of the F statistic is

$$F = \frac{(\mathrm{R}^2_{ur} - \mathrm{R}^2_{r})/q}{(1 - \mathrm{R}^2_{ur})/(n-k-1)} = \frac{(\mathrm{R}^2_{ur} - \mathrm{R}^2_{r})/q}{(1 - \mathrm{R}^2_{ur})/df_{ur}}. \quad (12)$$

Testing exclusion restrictions

1. Remark: A special set of exclusion restrictions is routinely tested by most regression packages: none of the explanatory variables has an effect on y.
   - The alternative is that at least one of the β_j is different from zero.
   - The F statistic in this case is

$$F = \frac{R^2/k}{(1-R^2)/(n-k-1)}. \tag{13}$$

   - Most regression packages report this F statistic automatically for testing joint exclusion of all independent variables.
   - This is sometimes called determining the overall significance of the regression.
   - If we fail to reject it, then there is no evidence that any of the independent variables help to explain y. This usually means that we must look for other variables to explain y.

F distribution

1. To use the F statistic, we must know its sampling distribution under the null in order to choose critical values and rejection rules.
2. It can be shown that, under H_0 (and assuming the CLM assumptions hold), F is distributed as an F random variable with (q, n-k-1) degrees of freedom.
3. We write this as

$$F \sim F_{q,\, n-k-1} \tag{14}$$

F distribution

1. The distribution of $F_{q,\, n-k-1}$ is available in statistical tables.
2. Similarly, we will reject H_0 in favor of H_1 when F is sufficiently "large".
3. If H_0 is rejected, then we say that $x_{k-q+1}, \cdots, x_k$ are jointly statistically significant at the appropriate significance level.
4. If the null is not rejected, then the variables are jointly insignificant.
5. The F statistic is often useful for testing exclusion of a group of variables when the variables in the group are highly correlated.

F distribution

1. Suppose we want to test whether firm performance affects the salaries of CEOs.
2. There are many ways to measure firm performance, e.g. revenue, profits.
3. Since measures of firm performance are likely to be highly correlated, individual coef. could be insignificant due to multicollinearity (t test).
4. An F test can be used to determine whether, as a group, the firm performance variables affect salary.

Homework 2

Wooldridge, 5th Edition, 2013:

Chapter 3 (pp. 106) Problem 5, 12, 13;
Chapter 4 (pp. 160) Problem 6. (i)-(iii).