Lecture 6-a
Testing in the CLM & Model Specification

Brooks (4th edition): Chapters 3 & 4

 1

Review: Testing H_0: $R\beta - q = 0$

- Q: Is **Rb** – **q** close to **0**? Two different approaches to this question.

Approach (1): We base the answer on the discrepancy vector:

$$\mathbf{m} = \mathbf{Rb} - \mathbf{q}.$$

Then, we construct a Wald statistic:

$$W = \mathbf{m}' \,(\text{Var}[\mathbf{m}\,|\,\mathbf{X}])^{-1}\, \mathbf{m}$$

Under H_0 and assuming (**A5**) & estimating σ^2 with $s^2 = \mathbf{e}'\mathbf{e}/(T - k)$:

$$W^* = (\mathbf{Rb} - \mathbf{q})' \,\{\mathbf{R}[s^2(\mathbf{X}'\mathbf{X})^{-1}]\mathbf{R}\}^{-1}\, (\mathbf{Rb} - \mathbf{q})$$

$$F = W^*/J \sim F_{J,T-k}.$$

If (**A5**) is not assumed, the results are only asymptotic: $J * F \xrightarrow{d} \chi^2_J$

Review: Testing H_0: $R\beta - q = 0$ with an F-Test

Approach (2): We base the answer on a model loss of fit when restrictions are imposed. Then, we construct an F test to check if the unrestricted RSS (RSS_U) is different from the restricted RSS (RSS_R). Does it go down a lot? –i.e., significantly?

Steps:

1. Estimate Restricted Model, get RSS_R
2. Estimate Unrestricted Model, get RSS_U

$$F = \frac{\frac{RSS_R - RSS_U}{(k_U - k_R)}}{\frac{RSS_U}{(T - k_U)}} \sim F_{J,T-k}. \quad (\text{where } J = k_U - k_R)$$

• The F-test constructed using a variable that can divide the data into 2 categories to compute RSS_R & RSS_U is usually referred as *Chow test.*

Review: Non-nested Models and Tests

• So far, all our tests (t-, F- & Wald tests) have been based on nested models, where the R model is a restricted version of the U model.

Example:

Model U	$\mathbf{Y} = \mathbf{X}\beta + \mathbf{W}\delta + \varepsilon$	(Unrestricted)
Model R	$\mathbf{Y} = \mathbf{X}\beta + \xi$	(Restricted)

Model U becomes Model R under H_0: $\boldsymbol{\delta} = \mathbf{0}$.

• Sometimes, we have two rival non-models -i.e., neither is a restricted version of the other. How do we choose a model?

Example:

Model 1	$\mathbf{Y} = \mathbf{X}\beta + \mathbf{W}\delta + \varepsilon$
Model 2	$\mathbf{Y} = \mathbf{X}\beta + \mathbf{Z}\gamma + \xi$

Review: Non-nested Models and Tests

- **Encompassing Test**

We have:

Model 1 $\mathbf{Y = X\beta + W\delta + \varepsilon}$

Model 2 $\mathbf{Y = X\beta + Z\gamma + \xi}$

Then, the **Encompassing Model** (**ME**) is:

ME: $\mathbf{Y} = \mathbf{X}\beta + \mathbf{W}\delta + \mathbf{Z}\gamma + \varepsilon$

Now test, separately, the hypotheses (1) $\boldsymbol{\delta} = \mathbf{0}$ and (2) $\boldsymbol{\gamma} = \mathbf{0}$. That is,

F-test for H_0: $\boldsymbol{\gamma} = \mathbf{0}$: **ME** (U Model) vs Model 1 (R Model).

F-test for H_0: $\boldsymbol{\delta} = \mathbf{0}$: **ME** (U Model) vs Model 2 (R Model).

Assuming the restrictions cannot be rejected, we prefer the model with the lower F statistic for the test of restrictions.

Review: Non-nested Models and Tests

- **J-Test**

Two non-nested models:

Model 1: $\mathbf{Y = X\beta + \varepsilon}$

Model 2: $\mathbf{Y = Z\gamma + \xi}$

- Steps:

(1) Estimate **Model 1** $\Rightarrow$ obtain fitted values: **Xb**.

(2) Add **Xb** to the list of regressors in Model 2: $\mathbf{Y} = \mathbf{Z}\gamma + \lambda\mathbf{Xb} + \xi$

(3) Do a *t-test* on λ. Rejecting H_0: λ=0 is evidence against **Model 2** & in favour of **Model 1**.

(4) Repeat the procedure (1)-(3) the other way round and do a *t-test* on λ. Rejecting H_0: λ=0 is evidence against **Model 1** & for **Model 2**.

(5) Rank the models on the basis of this test.

OLS Estimation - Assumptions

• CLM Assumptions

(**A1**) DGP: $\boldsymbol{y} = \mathbf{X}\,\beta + \boldsymbol{\varepsilon}$ is correctly specified.

(**A2**) $E[\boldsymbol{\varepsilon}|\mathbf{X}] = 0$

(**A3**) $Var[\boldsymbol{\varepsilon}|\mathbf{X}] = \sigma^2\,\mathbf{I}_T$

(**A4**) $\mathbf{X}$ has full column rank –rank($\mathbf{X}$)=k-, where $T \geq k$.

Q: What happens when (**A1**) is not correctly specified?

• We look at (**A1**). We have already studied what happens when we impose restrictions in the DGP: If we impose a true restrictions, estimation is unbiased & more efficient; a false restriction causes bias!

<u>This short lecture</u>: Are we omitting a relevant regressor? Are we including an irrelevant variable? Can we test for omitted variables?

Specification Errors: Omitted & Irrelevant X's

• Omitting relevant variables: Suppose the correct model (DGP) is

$\boldsymbol{y} = \mathbf{X}_1\boldsymbol{\beta}_1 + \mathbf{X}_2\boldsymbol{\beta}_2 + \boldsymbol{\varepsilon}$ –the "long regression," with $\mathbf{X}_1$ & $\mathbf{X}_2$.

But, we compute OLS omitting $\mathbf{X}_2$, a true driver of **y**. That is,

$\boldsymbol{y} = \mathbf{X}_1\boldsymbol{\beta}_1 + \boldsymbol{\varepsilon}$ –the "short regression."

<u>Implication</u>: Restricted estimator **b*** is **biased**, but **more efficient**.

• Irrelevant variables . Suppose the correct model is

$\boldsymbol{y} = \mathbf{X}_1\boldsymbol{\beta}_1 + \boldsymbol{\varepsilon}$ –the "short regression," with $\mathbf{X}_1$

But, we estimate, ignoring the true restriction $\boldsymbol{\beta}_2 = 0$:

$\boldsymbol{y} = \mathbf{X}_1\boldsymbol{\beta}_1 + \mathbf{X}_2\boldsymbol{\beta}_2 + \boldsymbol{\varepsilon}$ –the "long regression."

<u>Implication</u>: Estimator **b** is **unbiased**, but **inefficient**.

Trilogy of Asymptotic Tests: LR, Wald, and LM

• In practice, we rely on the asymptotic distribution of the Wald test. That is, $W \xrightarrow{d} \chi_J^2$.

• There are two other popular tests that are asymptotically equivalent –i.e., with same asymptotic distribution: the Likelihood Ratio (**LR**) and the Lagrange Multiplier (**LM**) tests.

• **LR test**: Based on the (log) *Likelihood*. It needs two ML estimations:

- The unrestricted estimation, producing $\hat{\theta}_{ML}$,
- The restricted estimation, producing $\hat{\theta}^R$.

Then, the LR test:

$$LR = 2[\log(L(\hat{\theta}_{ML})) - \log(L(\hat{\theta}^R))] \xrightarrow{d} \chi_J^2$$

Note: MLE requires assuming a distribution, usually, a normal.

The F Test: Are SMB and HML Priced Factors?

Example: We do a LR test to test if the SMB & HML FF factors are significant, using monthly data 1973 – 2020 (T=569). That is,

H_0: $\beta_{SMB} = \beta_{HML} = 0$

We use the function *lrtest* from the R package *lmtest*.

```
library(lmtest)
fit_ibm_ff3 <- lm (ibm_x ~ Mkt_RF + SMB + HML)
fit_ibm_capm <- lm (ibm_x ~ Mkt_RF)
lrtest(fit_ibm_ff3, fit_ibm_capm)

Likelihood ratio test

Model 1: ibm_x ~ Mkt_RF + SMB + HML
Model 2: ibm_x ~ Mkt_RF
  #Df LogLik Df  Chisq Pr(>Chisq)
1   5 810.03
2   3 805.30 -2 9.4616   0.008819 **
```

$\Rightarrow$ p-value is small: Reject H_0

Trilogy of Asymptotic Tests: LR, Wald, and LM

Technical note: The LR test is a *consistent test.* An asymptotic test which rejects H_0 with probability one when the H_1 is true is called a *consistent test.* That is, a consistent test has asymptotic power of 1. The LR test is a consistent test.

• **LM Test:** It needs only one estimation: the restricted estimation, producing $\hat{\theta}^R$. If the restriction is true, then the slope of the objective function (say, the Likelihood) at $\hat{\theta}^R$ should be zero. The slope is called the Score, S($\hat{\theta}^R$). The LM test is based on a Wald test on S($\hat{\theta}^R$) = 0.

$$LM = S(\hat{\theta}^R)'[Var(S(\hat{\theta}^R)]^{-1}S(\hat{\theta}^R) \xrightarrow{d} \chi^2_J$$

It turns out that there is a much simpler formulation for the LM test, based on the residuals of the restricted model. We will present this version of the test next.

Trilogy of Asymptotic Tests: LR, Wald, and LM

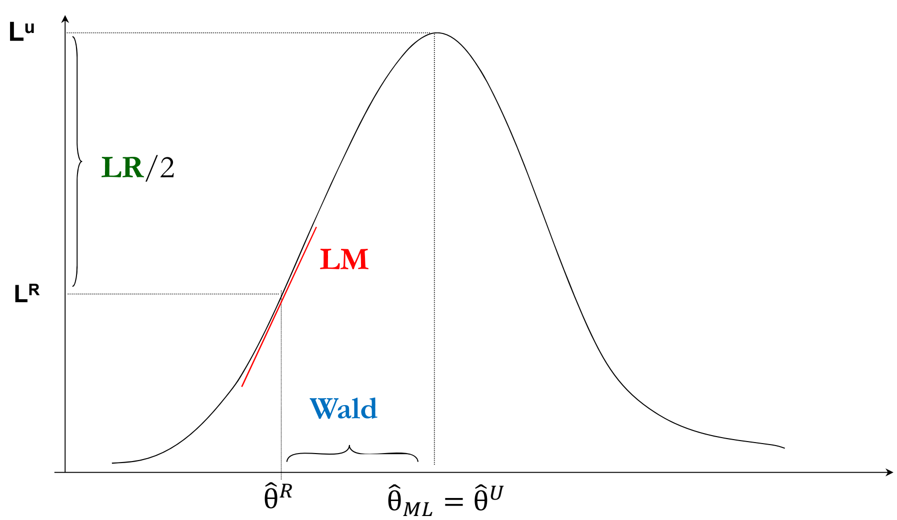

In general, W > LR > LM.

Testing Model Specification with an LM Test

• We can use the LM test to check for omitted variables, H_0: $\boldsymbol{\beta} = \mathbf{0}$. We have already presented LR & F tests of H_0: $\boldsymbol{\beta} = \mathbf{0}$. Why use an LM test? LM tests only use the restricted estimation, producing $\hat{\theta}^R$.

• The simpler formulation of the LM test is based on the residuals of the restricted model, $\mathbf{e}_R$.

Simple intuition. Everything that is omitted from (& belongs to!) a model should appear in the residuals ($\mathbf{e}_R$). Suppose the true model is:

$$\boldsymbol{y} = \mathbf{X}_1\boldsymbol{\beta}_1 + \mathbf{Z}\boldsymbol{\beta}_2 + \boldsymbol{\varepsilon}$$

But, we omit the J variables, $\mathbf{Z}$:

$$\boldsymbol{y} = \mathbf{X}_1\boldsymbol{\beta}_1 + \boldsymbol{\varepsilon}$$

An LM test checks if $\mathbf{e}_R$ can be explained by the J omitted variables $\mathbf{Z}$ 13

Testing Model Specification with an LM Test

• We use a simple regression of $\mathbf{e}_R$ against $\mathbf{Z}$ (dimension JxT) to check for the misspecification.

• LM test steps:

(1) Run restricted model ($\mathbf{y} = \mathbf{X}\,\boldsymbol{\beta}_1 + \boldsymbol{\varepsilon}$). Get restricted residuals, $\mathbf{e}_R$.

(2) (Auxiliary Regression). Run the regression of e_R on all the omitted J variables, $\mathbf{Z}$, and the k included variables, $\mathbf{X}$. In our case:

$$e_{R,i} = \alpha_0 + \alpha_1\, x_{i,1} + ... + \alpha_k\, x_{i,k} + \gamma_1\, z_{i,1} + ... + \gamma_J\, z_{i,J} + v_i$$

$\Rightarrow$ Keep the R^2 from this regression, R^2_{eR}.

(3) Compute LM-statistic:

$$\text{LM} = T * R^2_{eR} \overset{d}{\rightarrow} \chi^2_J.$$

14

Testing Model Specification with an LM Test

Technical Note: We include the original variables in (2), **X**, in the auxiliary regression to get the convenient form for the LM-test, as shown by Engle (1982).

• The LM Test is very general. It can be used in many settings, for example, to test for nonlinearities, interactions among variables, autocorrelation or heteroscedasticity (discussed later).

• Asymptotically speaking, the LM Test, the LR Test and the Wald Test are equivalent –i.e., they have the same limiting distribution, χ^2_J. In small T, they can have different conclusions. In general, however, we find: W > LR > LM. That is, the LM test is more conservative (cannot reject more often) and the Wald test is more aggressive.

15

Testing Model Specification with an LM Test

Example: We use an LM test to check if the standard CAPM for IBM returns omits **SMB** and **HML**. (J = 2)

```
fit_ibm_capm <- lm (ibm_x ~ Mkt_RF)                # Restricted Model
resid_r <- fit_ibm_capm$residuals                  # extract residuals from R model
fit_lm <- lm (resid_r ~ Mkt_RF + SMB + HML)        # auxiliary regression
> summary(fit_lm)

Coefficients:
             Estimate   Std. Error t value Pr(>|t|)
(Intercept)  0.0007021  0.0024875   0.282   0.7779
Mkt_RF       0.0125253  0.0567221   0.221   0.8253
SMB         -0.2124596  0.0841119  -2.526   0.0118 *
HML         -0.1715002  0.0846817  -2.025   0.0433 *
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
Residual standard error: 0.05848 on 565 degrees of freedom
Multiple R-squared:  0.01649,   Adjusted R-squared:  0.01127
F-statistic: 3.158 on 3 and 565 DF,  p-value: 0.02438
```

16

Testing Model Specification with an LM Test

Example (continuation):

```
R2_r <- summary(fit_lm)$r.squared          # extracting R^2 from fit_lm
> R2_r
[1] 0.01649104

LM_test <- R2_r * T
> LM_test
[1] 9.383402                               ⇒ LM_test > qchisq (.95,df=2) ⇒ Reject H0.

> qchisq(.95, df = 2)                      # chi-squared (df=2) value at 5% level
[1] 5.991465
p_val <- 1 - pchisq(LM_test, df = 2)       # p-value of LM_test
> p_val
[1] 0.009171071                            ⇒ p-value is small ⇒ Reject H0.
```

Note: In Lecture 5 we performed the same test with the Wald test (using the F distribution), the p-value was **0.0091175**. (This almost exact coincidence is not always the case.)

17