

General linear Model Continued..

$$y = X\beta + u$$

X non random

$$u \sim N(0, \sigma^2 I_n)$$

$$\hat{\beta} = (X'X)^{-1}X'y$$

We know

$$E(\hat{\beta}) = \beta$$

$$\text{Var}(\hat{\beta}) = \sigma^2(X'X)^{-1}$$

Generalised Least Squares

Have assumed error independent with common variance, but this may not always be realistic. Relaxing this assumption

$$y = X\beta + u$$

$$E(uu') = \Omega$$

where Ω is an arbitrary positive definite matrix

Premultiply by $\Omega^{-1/2}$

$$y^* = X^*\beta + u^*$$

$$y^* = \Omega^{-1/2}y$$

$$x^* = \Omega^{-1/2}X$$

$$u^* = \Omega^{-1/2}u$$

$$E(u^*u^{*'}) = \Omega^{-1/2}E(uu')\Omega^{-1/2} = I$$

So BLUE of β in this case is

$$\begin{aligned}\beta_{GLS} &= (X^{*'}X^*)^{-1}(X^{*'}y^*) \\ &= (X'\Omega^{-1}X)^{-1}X'\Omega^{-1}y\end{aligned}$$

which is the generalised least squares (GLS) estimator. Generally, Ω is a diagonal matrix with the i^{th} diagonal σ_i^2

A feasible GLS estimator then results from replacing Ω with a consistent estimator.

Now

$$\text{var}(\hat{\beta}_{GLS}) = \sigma^2(X^{*'}X^*)^{-1} = \sigma^2(X'\Omega^{-1}X)^{-1}$$

while variance for OLS is

$$\text{var}(\hat{\beta}_{OLS}) = \sigma^2(X'X)^{-1}$$

If we were to estimate OLS when GLS is appropriate $E(u) = 0$ will still be true so the estimator is unbiased, but its variance is now

$$\begin{aligned}\text{var}(\hat{\beta}_{OLS}) &= (X'X)^{-1}X'E(uu')X'(X'X)^{-1} \\ &= (X'X)^{-1}X'\Omega X(X'X)^{-1}\end{aligned}$$

Generally this means it is inefficient, though it is not necessarily the case -if $\Omega \neq \sigma^2I$

Instrumental Variables

We know that if the regressors are not independent of the disturbances then OLS estimates are both biased and inconsistent

$$\begin{aligned}y &= X\beta + u \\ \Rightarrow \hat{\beta} &= \beta + (X'X)^{-1}X'u \\ \Rightarrow p\lim(\hat{\beta}) &= \beta + p\lim\left(\frac{1}{n}X'X\right)^{-1} p\lim\left(\frac{1}{n}X'u\right)\end{aligned}$$

Assume that

$$p\lim\left(\frac{X'X}{n}\right) = \Sigma_{XX}$$

a positive definite matrix of full rank

and

$$p\lim\left(\frac{X'u}{n}\right) = \Sigma_{Xu} \neq 0$$

then

$$p\lim(\hat{\beta}) = \beta + \Sigma_{XX}^{-1}\Sigma_{Xu}$$

So the correlation of the disturbance term with one or more of the regressors will make OLS inconsistent.

Such correlations can be caused by measurement error in one or more regressors, but there are other possibilities:

- lagged dependent variables
- autoregressive disturbance
- simultaneity

Can get consistent estimator by instrumental variables. Consider:

$$y = X\beta + u$$

with

$$\text{var}(u) = \sigma^2 I$$

but

$$p \lim \left(\frac{X'u}{n} \right) \neq 0$$

Suppose we can find a data matrix Z of order $n \times l$ where $l > k$

1. variables in Z correlated with variables in X

$$p \lim \left(\frac{Z'X}{n} \right) = \Sigma_{ZX}$$

a finite matrix of full rank

Variables in Z are in the limit uncorrelated with u

$$p \lim \left(\frac{Z'u}{n} \right) = 0$$

then take

$$y = X\beta + u$$

premultiply both sides

$$Z'y = Z'X\beta + Z'u$$

$$\text{var}(Z'u) = \sigma^2 Z'Z$$

If we use GLS

$$\begin{aligned} \hat{\beta}_{GLS} &= \hat{\beta}_{IV} = (X'Z(Z'Z)^{-1}Z'X)^{-1}X'Z(Z'Z)^{-1}Z'y \\ &= (X'PX)^{-1}X'Py \end{aligned}$$

$$\text{with } P = Z(Z'Z)^{-1}Z'$$

then

$$\text{Var}(\hat{\beta}_{IV}) = \sigma^2 (X'PX)^{-1}$$

$$\hat{\sigma}^2 = (y - \hat{\beta}_{IV})'(y - \hat{\beta}_{IV})/n$$

using $n - k$ or n does not matter asymptotically).

$$\hat{\beta}_{IV} = \beta + \left(\frac{1}{n} X'PX \right)^{-1} \left(\frac{1}{n} X'Pu \right)$$

Now

$$\frac{1}{n} X'PX = \left(\frac{1}{n} X'Z \right) \left(\frac{1}{n} Z'Z \right)^{-1} \left(\frac{1}{n} Z'X \right)$$

Assume middle term has $\text{plim } \Sigma_{ZZ}^{-1}$ so

$$p \lim \left(\frac{1}{n} X'PX \right) = \Sigma_{XZ} \Sigma_{ZZ}^{-1} \Sigma_{ZX}$$

which will be finite non singular matrix.

Similarly

$$p \lim \left(\frac{1}{n} Z'Pu \right) = \Sigma_{XZ} \Sigma_{ZZ}^{-1} \Sigma_{Zu} = 0$$

as instruments z are assumed uncorrelated with u in the limit. So the IV estimator is consistent

If $l = k$ so that Z has the same number of columns as X then $X'Z$ is $k \times k$ and nonsingular. Then all simplifies to

$$\hat{\beta}_{IV} = (Z'X)^{-1}Z'y$$

$$\text{var}(\hat{\beta}_{IV}) = \sigma^2(Z'X)^{-1}(Z'Z)(X'Z)^{-1}$$

if $l < k$ then singular

Two Stage Least Squares

A form of IV can be seen as result of double application of OLS

1. Regress each of the variables in the X matrix on Z to give \hat{X}

$$\hat{X} = Z(Z'Z)^{-1}Z'X$$

$$= PX$$

2. Regress y on \hat{X} to obtain β

$$\hat{\beta}_{TSLS} = (\hat{X}'\hat{X})^{-1}\hat{X}'y$$

$$= (X'PX)^{-1}(X'Py)$$

$$= \hat{\beta}_{IV}$$

As before

$$\text{var}(\hat{\beta}_{IV}) = \sigma^2(X'PX)^{-1}$$

$$\hat{\sigma}^2 = (y - X\hat{\beta}_{IV})'(y - X\hat{\beta}_{IV})/n$$

Choice of instruments

- can use variables from X
- any variable thought exogenous and independent of disturbances
- lagged values

When some of the X are used need to partition:

$$X = [X_1 \ X_2] \text{ and } Z = [X_1 \ Z_1]$$

X_1 is $n \times r$

X_2 is $n \times (k - r)$ with $r < k$

Z_1 is $n \times (p - r)$

can show

$$\hat{X} = [X_1 \ \hat{X}_2]$$

where

$$\hat{X}_2 = Z(Z'Z)^{-1}Z'X_2$$

So variables in X_1 serve as instruments for themselves and the remaining second stage regressors are the fitted values of X_2 obtained from regressing X_2 on the full set of instruments Z

Minimum number of instruments is k including any variables that serve as their own instruments.

Asymptotic efficiency increases with the number of instruments, but so does finite sample bias

if select n instruments then $P = I$ and get the biased and inconsistent OLS estimates

Multicollinearity

Have assumed in the past that the explanatory variables were linearly independent. This means that $(X'X)^{-1}$ exists.

Perfect multicollinearity implies that $(X'X)$ will be a singular matrix with rank less than k . Which implies do not have unique solutions to the normal equations

As

$$\hat{\beta} = (X'X)^{-1}X'y$$

then if $(X'X)$ is singular $\hat{\beta}$ cannot be estimated. In fact it means that not all regression parameters (the whole vector of parameters) are estimable, only certain linear functions of $\hat{\beta}_i$.

Testing

Again we have

$$y = X\beta + u$$

X non random

$$u \sim N(0, \sigma^2 I_n)$$

$$\hat{\beta} = (X'X)^{-1}X'y$$

We know

$$E(\hat{\beta}) = \beta$$

$$Var(\hat{\beta}) = \sigma^2(X'X)^{-1}$$

We saw that

$$\hat{\beta} - \beta = (X'X)^{-1}X'u$$

so $\hat{\beta} - \beta$ is a linear function of a normally distributed random vector u and so is itself normal.

$$\hat{\beta} \sim N(\beta, \sigma^2(X'X)^{-1})$$

From before

$$\hat{\sigma}^2 = \frac{\hat{u}'\hat{u}}{n-k} = \frac{u'Nu}{n-k}$$

$$\text{where } N = [I_n - X(X'X)^{-1}X']$$

Now we can show that for any symmetric idempotent matrix with $\varepsilon \sim N(0, \sigma^2 I_n)$

$$\frac{\varepsilon'Q\varepsilon}{\sigma^2} \sim \chi_{\text{trace}(Q)}^2$$

Since N is symmetric and idempotent

$$\frac{u'Nu}{\sigma^2} = \frac{(n-k)\hat{\sigma}^2}{\sigma^2} \sim \chi_{n-k}^2$$

Now

$$\begin{aligned}\hat{\beta} - \beta &= (X'X)^{-1}X'u \\ \Rightarrow X(\hat{\beta} - \beta) &= X(X'X)^{-1}X'u = Mu \\ \Rightarrow (\hat{\beta} - \beta)'X'X(\hat{\beta} - \beta) &= u'M^2u = u'Mu\end{aligned}$$

as M is idempotent. So

$$\frac{(\hat{\beta} - \beta)'X'X(\hat{\beta} - \beta)}{\sigma^2} = \frac{u'Mu}{\sigma^2} \sim \chi_k^2$$

As $MN = 0$ these two results are independent

$$\frac{u'Mu/k}{u'Nu/n-k} \sim \frac{\chi_k^2}{\chi_{n-k}^2} \sim F_{(k,n-k)}$$

That is

$$\frac{(\hat{\beta} - \beta)'X'X(\hat{\beta} - \beta)/n}{(y - X\hat{\beta})'(y - X\hat{\beta})/n-k} \sim F_{(k,n-k)}$$

which can be used to construct significance tests on β . replace the value of β given by the null and compute the ratio.

More usually we don't test the hypothesis on the whole parameter space. Consider the linear restriction

$$R\beta = d$$

Want to test

$$H_0 : R\beta = d$$

$$H_1 : R\beta \neq d$$

R is an $r \times k$ matrix of constants with $\text{Rank}(R) = r$ and d is an $r \times 1$ vector. Many economic hypotheses can be put on this form.

$$y_i = \beta_0 + \beta_1 x_{1i} + \beta_2 x_{2i} + \beta_3 x_{3i} + u_i$$

To test

$$\beta_1 = 0 : R = (0100) : d = 0$$

$$\beta_2 = -\beta_3 : R = (0011) : d = 0$$

$$\beta_1 + \beta_2 + \beta_3 = 1 : R = (0111) : d = 1$$

In each case there is one restriction and $\text{Rank}(r) = 1$. For

$$\beta_0 = 0; \beta_1 = 1; \beta_2 = -\beta_3$$

$$R = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{bmatrix} : d = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}$$

and the $\text{Rank}(R) = 3$. In all of these cases the restricted model can be estimated by imposing the restrictions on the data used to estimate the model

$$y_i = \beta_0 + \beta_1 x_{1i} + \beta_2 x_{2i} + \beta_3 x_{3i} + u_i$$

$$y_i = \beta_0 + x_{1i} + \beta_2 x_{2i} + \beta_3 x_{3i} + u_i \text{ when } \beta_1 = 1$$

$$y_i - x_{1i} = \beta_0 + \beta_2(x_{2i} - x_{3i}) + u_i \text{ when } \beta_2 = -\beta_3$$

$$y_i - x_{1i} = \beta_2(x_{2i} - x_{3i}) + u_i \text{ when } \beta_0 = 0$$

Note that in each case after you impose r restrictions only estimate $k-r$ parameters. So each restriction reduces the number of parameters by one. What we are doing is minimising $\hat{u}'\hat{u}$

subject to the restrictions $R\beta = d$

As a Lagrangian

$$\begin{aligned} L &= \hat{u}'\hat{u} + \lambda(R\beta - d) \\ &= (y - X\beta)'(y - X\beta) - \lambda(R\beta - d) \end{aligned}$$

The restricted estimator can be written

$$\beta^* = \hat{\beta} - (X'X)^{-1}R\lambda$$

So define unrestricted and restricted residual sum of squares

$$\begin{aligned} URSS &= (y - X\hat{\beta})'(y - X\hat{\beta}) \\ RRSS &= (y - X\beta^*)'(y - X\beta^*) \end{aligned}$$

Then if the null is true

$$\frac{(RRSS - URSS)/r}{URSS/n - k} \sim F_{(r, n-k)}$$

which is the familiar F test. This measures whether an increase in RSS per degree of freedom from imposing the restrictions is large relative to the variance. If it is large reject the restriction, if it is small accept the restriction. What is large depends on the significance level chosen. Note that it is distributed F only if the null is true. If $R\hat{\beta} \neq d$ then β^* will be a long way from $\hat{\beta}$.

So an alternative way of looking at this is that if the restrictions are not true $R\hat{\beta} - d$ will be far from zero. To measure how far we could standardise by the estimated variance of $R\hat{\beta} - d$ namely $\hat{V}(R\hat{\beta} - d)$. Then

$$\frac{1}{r} (R\hat{\beta} - d)' \hat{V}(R\hat{\beta} - d)^{-1} (R\hat{\beta} - d) \sim F_{(r, n-k)}$$

with

$$\hat{V}(R\hat{\beta} - d)^{-1} = \frac{1}{\sigma^2} [R(X'X)^{-1}R']^{-1}$$

An important special case is where the hypothesis to be tested is that a particular β equals a particular value, eg $\beta_1 = \bar{\beta}_1$

Then $R = [0100\dots] : d = \bar{\beta}_1$

$$R\hat{\beta} - d = \hat{\beta}_1 - \bar{\beta}_1$$

with $\hat{V}(R\hat{\beta} - d)$ the estimated variance of $\hat{\beta}_1$

So

$$\frac{(\hat{\beta}_1 - \bar{\beta}_1)^2}{\text{Var}(\hat{\beta}_1)} \sim F_{(1, n-k)}$$

Now if $w^2 \sim F_{(1, n-k)} \Rightarrow w \sim t_{n-k}$ so

$$\frac{\hat{\beta}_1 - \bar{\beta}_1}{\text{se}(\hat{\beta}_1)} \sim t_{n-k}$$

which is the commonly used t test

These are small sample tests and require the model be linear and the errors normal. If these don't hold then have to use asymptotic tests.

The 3 widely asymptotic tests are:

- - The likelihood ratio test: uses restricted and unrestricted

- The Wald test: uses only the unrestricted
- The Lagrange Multiplier (LM) test: uses only restricted

These are asymptotically equivalent and are each distributed χ_r^2 (where r is the number of restrictions) if the null hypothesis is true.

All can be written as measures of distance standardised by a variance covariance matrix, each differing by which distance is measured.

The ML estimates are those which maximise $LL(\theta)$, i.e. the $\hat{\theta}$, which make

$$\frac{\partial LL(\theta)}{\partial \theta} = S(\hat{\theta}) = 0$$

where $S(\hat{\theta})$ is the score vector, the derivatives of the LL with respect to each of the k elements of the vector θ evaluated at the values, $\hat{\theta}$, which make $S(\theta) = 0$. We will call these the unrestricted estimates and the value of the Log-likelihood at $\hat{\theta}$, $LL(\hat{\theta})$.

Suppose theory suggests $m \leq k$ prior restrictions of the form $R(\theta) = 0$. If $m = k$, theory specifies all the parameters and there are none to estimate. The restricted estimates maximise

$$L = LL(\theta) - \lambda' R(\theta)$$

where λ is a $m \times 1$ vector of Lagrange Multipliers. The solution to

$$\frac{\partial \mathcal{L}}{\partial \theta} = \frac{\partial LL(\theta)}{\partial \theta} - \frac{\partial R(\theta)}{\partial \theta} \lambda = 0$$

gives the restricted estimator θ^* , we can write this

$$S(\theta^*) - F(\theta^*) \lambda^* = 0$$

where $S(\theta^*)$ is the $k \times 1$ Score vector evaluated at the restricted estimates and $F(\theta^*)$ is the $k \times m$ matrix of the derivatives of the restrictions with respect to the parameters evaluated at the restricted estimates. Notice that at θ^* the derivative of the Log-likelihood function with respect to the parameters is not equal to zero but to $F(\theta^*) \lambda^*$. The value of the Log-likelihood at θ^* is $LL(\theta^*)$ which is less than or equal to $LL(\hat{\theta})$.

If the hypotheses (restrictions) are true:

1. the two log-likelihoods should be similar, i.e. $LL(\hat{\theta}) - LL(\theta^*)$ should be close to zero;
2. the unrestricted estimates that satisfy the restrictions $R(\hat{\theta})$ should be close to zero (note $R(\theta^*)$ is exactly zero by construction);
3. the restricted score, $S(\theta^*)$, should be close to zero (note $S(\hat{\theta})$ is exactly zero by construction) or equivalently the Lagrange Multipliers λ^* should be close to zero, the restrictions should not be binding.

ILLUSTRATE WITH DIAGRAM FROM KENNEDY

These implications are used as the basis for three types of test procedures. The issue is how to judge 'close to zero'? To judge this we use the asymptotic equivalents of the linear distributional results used above in the discussion of the properties of the Linear Regression Model.

Asymptotically the ML estimator is normal

$$\hat{\theta} \sim N(\theta, I(\theta)^{-1})$$

asymptotically the scalar quadratic form is chi-squared

$$(\hat{\theta} - \theta)' I(\theta) (\hat{\theta} - \theta) \sim \chi^2(k).$$

and asymptotically $R(\hat{\theta})$ is also normal

$$R(\hat{\theta}) \sim N(R(\theta), F(\theta)' I(\theta)^{-1} F(\theta))$$

This gives us three procedures for generating asymptotic test statistics for the m restrictions

$H_0 : R(\theta) = 0$; each of which are distributed $\chi^2(m)$, when the null hypothesis is true:

1. Likelihood Ratio Tests

$$LR = 2(LL(\hat{\theta}) - LL(\theta^*)) \sim \chi^2(m)$$

2. Wald Tests

$$W = R(\hat{\theta})' [F(\theta)' I(\theta)^{-1} F(\theta)]^{-1} R(\hat{\theta}) \sim \chi^2(m)$$

3. Lagrange Multiplier (or Efficient Score) Tests

$$LM = S(\theta^*)' I(\theta^*)^{-1} S(\theta^*) \sim \chi^2(m).$$

- The Likelihood ratio test is straightforward to calculate when both the restricted and unrestricted models have been estimated.
- The Wald test only requires the unrestricted estimates.
- The Lagrange Multiplier test only requires the restricted estimates.
- For the linear regression model, the inequality $W > LR > LM$ holds, so you are more likely to reject using W .
- In the LRM, the LM test is usually calculated using regression residuals.
- The Wald test is not invariant to how you write non-linear restrictions. Suppose $m = 1$, and $R(\theta)$ is $\theta_1\theta_2 - \theta_3 = 0$. This could also be written $\theta_1 - \theta_3/\theta_2 = 0$ and these would give different values of the test statistic. The former form, using multiplication rather than division, is usually better.

These formulas are not used to compute the test. The LR test is usually

$$-2 \log \lambda = 2 [LL(\hat{\theta}, y) - LL(\theta^*, y)] \sim \chi_r^2$$

while the LM and W are normally calculated from tests on the regression residuals.