

Regression Basics in Matrix Terms

1 The Normal Equations of least squares

Let y denote the dependent variable, a $n \times 1$ vector, and let X denote the $n \times k$ matrix of regressors (independent variables). Write b for the k -vector of regression coefficients, and write e for the n -vector of *residuals*, such that $e_i = y_i - X_i b$.

The vector b is the ordinary least squares (OLS) solution if and only if it is chosen such that the sum of squared residuals,

$$\text{SSR} = \sum_{i=1}^n e_i^2,$$

is at a minimum.

Attaining the minimum SSR can be approached as a calculus problem. In matrix notation, we can write the SSR as

$$\begin{aligned} e'e &= (y - Xb)'(y - Xb) \\ &= (y' - (Xb)')(y - Xb) \\ &= y'y - y'Xb - (Xb)'y + (Xb)'Xb \\ &= y'y - 2b'X'y + b'X'Xb \end{aligned}$$

The manipulations above exploit the properties of the matrix transpose, namely $(A + B)' = A' + B'$ and $(AB)' = B'A'$. In addition we rely on the point that $y'Xb$ and $(Xb)'y$ (which can also be written as $b'X'y$) are identical by virtue of the fact that they are 1×1 : they are transposes of each other, but transposing a 1×1 matrix gives back the same matrix.

Now take the derivative of the last expression above with respect to b and set it to zero. This gives

$$\begin{aligned} -2X'y + 2X'Xb &= 0 \\ \Rightarrow -X'y + X'Xb &= 0 \\ \Rightarrow X'Xb &= X'y \end{aligned}$$

(In general, the derivative of $b'Cb$ with respect to b is $(C + C')b$, but $X'X$ is a symmetric matrix and the derivative simplifies to that shown, $2X'Xb$.)

If the matrix $X'X$ is non-singular (i.e. it possesses an inverse) then we can multiply by $(X'X)^{-1}$ to get

$$b = (X'X)^{-1}X'y \tag{1}$$

This is a classic equation in statistics: it gives the OLS coefficients as a function of the data matrices, X and y .

Orthogonality of residuals and regressors. Here's another perspective. As we said above, the regression residuals are $e = y - Xb$. Suppose we require that the residuals be *orthogonal* to the regressors, X . We then have

$$\begin{aligned} X'e &= 0 \\ \Rightarrow X'(y - Xb) &= 0 \\ \Rightarrow X'Xb &= X'y \end{aligned}$$

which replicates the solution above. So another way of thinking about the least squares coefficient vector, b , is that it satisfies the condition that the residuals are orthogonal to the regressors. Why might that be "a good thing"? Well, intuitively, if this orthogonality condition were not satisfied, that would mean that the predictive power of X with regard to y is not exhausted. The residuals represent the "unexplained" variation in y ; if they were not orthogonal to X that would mean that more explanation could be squeezed out of X by a different choice of coefficients.

2 The expected value of the OLS estimator

Up to this point we have been concerned with the arithmetic of least squares. From an econometric viewpoint, however, we're interested in the properties of least squares as an *estimator* of some data generating process (DGP) that we presume to be operating in the economy. Let us write the DGP as

$$y = X\beta + u$$

where u is the error or disturbance term, satisfying $E(u) = 0$.

From this standpoint we consider the least squares coefficient vector b as an estimator, $\hat{\beta}$, of the unknown parameter vector β . By the same token our regression residuals, e , can be considered as estimates of the unknown errors u , or $e = \hat{u}$. In this section we consider the expected value of $\hat{\beta}$, with a question of particular interest being: under what conditions does $E(\hat{\beta})$ equal β ? In the following section we examine the variance of $\hat{\beta}$.

Since, as we saw above, the least squares $\hat{\beta}$ equals $(X'X)^{-1}X'y$, we can write

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

The next step is to substitute the presumed DGP for y , giving

$$\begin{aligned} E(\hat{\beta}) &= E[(X'X)^{-1}X'(X\beta + u)] \\ &= E[(X'X)^{-1}X'X\beta + (X'X)^{-1}X'u] \\ &= E[\beta + (X'X)^{-1}X'u] \\ &= \beta + E[(X'X)^{-1}X'u] \end{aligned}$$

(Since β itself is not a random variable, it can be moved through the expectation operator.) We then see that $E(\hat{\beta}) = \beta$ —in other words, the least squares estimator is unbiased—on condition that

$$E[(X'X)^{-1}X'u] = 0$$

What does it take to satisfy this condition? If we take expectations conditional on the observed values of the regressors, X , the requirement is just that

$$E(u|X) = 0 \tag{2}$$

i.e., the error term is independent of the regressors.

So: the least squares estimator is unbiased provided that (a) our favored specification of the DGP is correct and (b) the independent variables in X offer no predictive power over the errors u . We might say that the errors must be *uncorrelated* with the regressors, but the requirement in (2) is actually stronger than that.

3 The variance of the OLS estimator

Recall the basic definition of variance:

$$\text{Var}(X) = E[X - E(X)]^2 = E[(X - E(X))(X - E(X))]$$

The variance of a random variable X is the expectation of the squared deviation from its expected value. The above holds good for a scalar random variable. For a random *vector*, such as the least squares $\hat{\beta}$, the concept is similar except that squaring has to be replaced with the matrix counterpart, multiplication of a vector into its own transpose.

This requires a little thought. The vector $\hat{\beta}$ is a column vector of length k , and so is the difference $\hat{\beta} - E(\hat{\beta})$. Given such a k -vector, v , if we form $v'v$ (the inner product) we get a 1×1 or scalar result, the sum of squares of the elements of v . If we form vv' (the outer product) we get a $k \times k$ matrix. Which do we want here? The latter: by the variance of a random vector we mean the (co-)variance matrix, which holds the variances of the elements of the vector on its principal diagonal, and the covariances in the off-diagonal positions.

Therefore, from first principles,

$$\text{Var}(\hat{\beta}) = E \left[\left(\hat{\beta} - E(\hat{\beta}) \right) \left(\hat{\beta} - E(\hat{\beta}) \right)' \right]$$

If we suppose that the conditions for the estimator to be unbiased are met, then $E(\hat{\beta}) = \beta$ and, substituting in the least squares formula for $\hat{\beta}$, we get

$$\text{Var}(\hat{\beta}) = E \left[\left((X'X)^{-1} X'y - \beta \right) \left((X'X)^{-1} X'y - \beta \right)' \right]$$

Then, once again, substitute the DGP, $X\beta + u$, for y , to get

$$\begin{aligned} \text{Var}(\hat{\beta}) &= E \left[\left((X'X)^{-1} X'(X\beta + u) - \beta \right) \left((X'X)^{-1} X'(X\beta + u) - \beta \right)' \right] \\ &= E \left[\left(\beta + (X'X)^{-1} X'u - \beta \right) \left(\beta + (X'X)^{-1} X'u - \beta \right)' \right] \\ &= E \left[\left((X'X)^{-1} X'u \right) \left((X'X)^{-1} X'u \right)' \right] \\ &= E \left[(X'X)^{-1} X'uu'X(X'X)^{-1} \right] \end{aligned} \quad (3)$$

(Note that the symmetric matrix $(X'X)^{-1}$ is equal to its transpose.) As with figuring the expected value of $\hat{\beta}$, we will take expectations conditional on X . So we can write

$$\text{Var}(\hat{\beta}) = (X'X)^{-1} X' E(uu') X (X'X)^{-1} \quad (4)$$

The awkward bit is the matrix in the middle, $E(uu')$. This is the covariance matrix of the error term, u ; note that it is $n \times n$ and so potentially quite large. However, under certain conditions it simplifies greatly—specifically, if the error term is *independently and identically distributed* (i.i.d.). In that case $E(uu')$ boils down to an $n \times n$ version of the following,

$$\begin{bmatrix} \sigma^2 & 0 & 0 \\ 0 & \sigma^2 & 0 \\ 0 & 0 & \sigma^2 \end{bmatrix},$$

that is, a matrix with a constant value on the diagonal and zeros elsewhere. The constant σ^2 reflects the “identical” part of the i.i.d. condition (the error variance is the same at each observation) and the zeros off the diagonal reflect the “independent” part (there’s no correlation between the error at any given observation and the error at any other observation). Such an error term is sometimes referred to as “white noise”.

Under the assumption that u is i.i.d., then, we can write $E(uu') = \sigma^2 I$, where I is the $n \times n$ identity matrix. And the scalar σ^2 can be moved out, giving

$$\begin{aligned} \text{Var}(\hat{\beta}) &= \sigma^2 (X'X)^{-1} X' I X (X'X)^{-1} \\ &= \sigma^2 (X'X)^{-1} X' X (X'X)^{-1} \\ &= \sigma^2 (X'X)^{-1} \end{aligned} \quad (5)$$

The somewhat nasty equation (4) therefore gives the variance of $\hat{\beta}$ in general case, while equation (5) gives the nice and simple version that arises when the error term is assumed to be i.i.d.

How do we actually get a value for (that is, an estimate of) the variance of $\hat{\beta}$? The true error variance, σ^2 , is unknown, but we can substitute an estimate, s^2 , based on the regression residuals:

$$s^2 = \frac{\sum_{i=1}^n e_i^2}{n - k}$$

and our estimated variance is then $s^2 (X'X)^{-1}$ —provided, of course, that we are willing to assume that the error term is i.i.d.