ECON 343
Lecture 10: Autocorrelation and Linear Time Series Models

Outline

1. Definition of autocorrelation
2. Some reminders of terminology
3. MA and AR models
4. Forms of autocorrelation
5. Causes and consequences of autocorrelation
6. Detecting autocorrelation
7. Estimation in the presence of autocorrelation

1. Definition of autocorrelation

- Relaxing assumption in OLS on the structure of the error term; allow the errors to be correlated, i.e.

$E(u_i u_j) \neq 0$ some $i \neq j$.

- The variance-covariance matrix of the errors have non-zero off-diagonal elements.

- Why might this happen? a data set where observations might be correlated with one another for reasons other than via X variables.

- The most striking examples of cross-observation correlations occur in time series data. Since many economic times series move slowly over time, these variables will always display correlations over time.

2. Some reminders of terminology

Terminology 1:Lags

Let Y_t be a time series variable. Then define Y_{t-1} , the value of Y one period before, as the lag of Y, or 'lagged Y', of the first lag of Y. Y_{t-2} is the 2nd lag of Y and Y_{t-k} is the kth lag of Y.

Terminology 2: Differences

The difference, $\Delta Y_t = Y_t - Y_{t-1}$ is called the first difference of Y, some times written $\Delta_1 Y_t$. Correspondingly $\Delta_j Y_t = Y_t - Y_{t-j}$ is the j- period difference. The expression second difference is usually reserved for the difference of a difference, i.e.

$\Delta^2 Y_t = \Delta Y_t - \Delta Y_{t-1} = (Y_t - Y_{t-1}) - (Y_{t-1} - Y_{t-2})$ and is occasionally written as $\Delta\Delta Y_t$.

Note also the relationship between the difference of the logarithm of a variable and the rate of change:

$$\Delta lnY_t = lnY_t - lnY_{t-1} = ln(Y_t/Y_{t-1}) = ln\left(1+\Delta Y_t \big/ Y_{t-1}\right) \cong \Delta Y_t \big/ Y_{t-1}$$

Terminology 3: Autocorrelation

Define the j^{th} autocorrelation =

$$\rho_j = \text{corr}\,(Y_t, Y_{t-j}) = \frac{\text{cov}(Y_t, Y_{t-j})}{\sqrt{Var(Y_t)Var(Y_{t-j})}}\,.$$

AUTOREGRESSIONS are models that decribe the evolution of a time series.

The first-order autoregression model is

$$Y_t = \alpha + \beta Y_{t-1} + u_t$$

The k^{th} order autoregression model is:

$$Y_t = \alpha + \beta_1 Y_{t-1} + \beta_2 Y_{t-2} + \ldots + \beta_\kappa Y_{t-k} + u_t$$

Note that these *AR* models contain serial dependence in the *Y* variable, but we haven't said much about the error term *u* yet.

3. MA and AR models

- It is perfectly possible that the shocks hitting *Y* have some persistence of their own. In the *AR1* model, a shock hits *Y* and then disappears, but persists inasmuch as *Y* is slow to react.

- What if the shocks themselves have persistent effects on *Y*, but *Y* has no persistence of its own? Than we get the moving-average (*MA)* time series model

- Here is a first-order version (*MA1)* where a shock has an immediate effect and then a one-period delayed effect.

$$Y_t = \alpha + u_t + \beta_1 u_{t-1}$$

- You might reasonably ask if the *MA* and *AR* models are really all that different. Well, there is a relationship. Take the *AR1* model.

$$Y_t = \alpha + \beta Y_{t-1} + u_t$$

Then substitute for Y_{t-1} and so on

$$Y_t = \alpha + \beta(\alpha + u_{t-1} + \beta\{\alpha + u_{t-2} + \beta[\ldots]\})$$

$$Y_t = \alpha(1 + \beta + \beta^2 + \ldots) + u_t + \beta u_{t-1} + \beta^2 u_{t-2} + \ldots$$

So, an *AR1* is an infinite *MA* process, though with β <1, the weights decline to zero. Same goes the other way, that is, a *MA1* is an infinite *AR1*.

- We can also have mixed processes. An *ARMA*(*p*, *q*) model is a p^{th} order *AR* combined with a q^{th} order *MA*.

- An *ARIMA*(*p*, *d*, *q*) is an *ARMA* model where the dependent variable has been differenced *d* times, so the model $\Delta Y_t = 0.017 + 0.303 \Delta Y_{t-1} + res_t$ is an *ARIMA*(*1*, *1*, *0*).

- There are ADL, (*autoregressive distributed lag*) models. In these models we allow and *AR*(*p*) in the dependent variable, but also *q* lags of one or more variables used to help predict.

- An example would be an ADL(1, 1) of changes in the rate of inflation (*infl*) with changes in the rate of unemployment (*u*) as the lagged X variable:

$$\Delta infl_t = \beta_o + \beta_1 \Delta infl_{t-1} + \beta_2 \Delta u_{t-1} + \varepsilon_t.$$

4. Forms of autocorrelation

- Let us imagine a time series regression model, and assume joint stationarity

$$Y_t = \beta_0 + \beta_1 X_{t-1} + u_t$$

- Several forms of autocorrelation are possible, the simplest being

$$u_t = \rho u_{t-1} + v_t$$

This is first order autocorrelation. The parameter ρ measures the correlation between u_t and u_{t-1} and v_t is another random error term that is IID (i.e. not autocorrelated).

- If $\rho > 0$ then we have positive autocorrelation (positive (negative) errors tend to be followed by positive (negative) errors); if $\rho < 0$ then we have negative autocorrelation (positive (negative) errors tend to be followed by negative (positive) errors).

- Higher order autocorrelation can also occur. Using quarterly data, it is likely that the error in quarter 4 is correlated with the error in quarter 4 the following year (Christmas spending). Hence we have

$$u_t = \rho u_{t-4} + v_t$$

or, more generally

$$u_t = \rho_1 u_{t-1} + \rho_2 u_{t-2} + \rho_3 u_{t-3} + \rho_4 u_{t-4} + v_t$$

This is fourth order autocorrelation. Some of the ρ values may not be significant.

5. Causes and consequences of autocorrelation

Causes of autocorrelation

- **an error in functional form of indeed systematic measurement error**
- **an omitted explanatory variable, which is itself autocorrelated, or more generally a dynamic mis-specification.**

Consequences of autocorrelation

These are similar to heteroscedasticity since the same assumption is relaxed:

- **OLS coefficient estimates are unbiased and consistent, but not BLUE nor asymptotically efficient. If $\rho = 0.8$ (the correlation between u_t and u_{t-1}) for example, then OLS is only about 20% efficient, i.e. the sampling variance is about five times as great as that of GLS - a big difference.**

- **the estimate of σ^2 is biased, affecting inference. If you have positive autocorrelation and the X variables are trended, then OLS formulae *underestimate* the true standard errors. If X ~ AR(1) with parameter 0.9 and ρ = 0.5 then the OLS estimate of the variance of b is about one-third of its true size - a big error. Hence one tends to reject H_0 when one should not.**

6. Detecting autocorrelation

Durbin-Watson test

- This tests for <u>first order autocorrelation</u> (u ~ AR(1)), $u_t = \rho u_{t-1} + v_t$ where v is iid. This implies the correlation between u_t and u_{t-1} is ρ. (Note here that $u_t = \rho(\rho u_{t-2} + v_{t-1}) + v_t = \rho^2 u_{t-2} + \rho v_{t-1} + v_t$ so ρ^2 is the correlation between u_t and u_{t-2}.). The test also assumes the X variables are independent.

- Consider the statistic $\frac{\sum(u_t - u_{t-1})^2}{\sum u_t^2}$. If $\rho > 0$ then $u_t \approx u_{t-1}$ and so this expression will be 'small' (near 0). If $\rho < 0$ then it will be 'large' (near 4). If $\rho = 0$ the expression will be near 2.

- We have to substitute e for the unobserved *u* to get Durbin and Watson's *d* statistic. Under H_0: $\rho = 0$, *d* varies around 2. Unfortunately, the distribution of *d* depends upon the X values, hence we can only get upper and lower bounds for the critical value. To interpret *d*

0 d_L d_U 2 $4\text{-}d_U$ $4\text{-}d_L$ 4

positive auto	uncertain region	OK	uncertain region	negative auto

From the *d* statistic we can get a good approx of the value of ρ since it can be shown that *d* ≈ 2 (1-r) where r = $\Sigma e_t e_{t-1} / \Sigma e_t^2$ is the first order sample autocorrelation coefficient and is an estimate of ρ. Hence we get

$\hat{\rho}$ **= r = 1-*d*/2**

Durbin's h-test

If there is a LDV then it can be shown that *d* is biased towards 2 (no autocorrelation). In this case we can use Durbin's *h* test. Again this only tests for AR(1).

$$h = (1-d/2)\sqrt{\frac{n}{1-ns_b^2}} \overset{a}{\sim} N(0,1) \quad \text{(asymptotic Normal)}$$

where s_b^2 is the estimated variance of the coefficient of the LDV. Here *h* = -0.13. So, we cannot reject H_0.

<u>What do we do if we find evidence of first order autocorrelation?</u> It means the following equations, for instance hold:

$$Y_t = \beta_0 + \beta_1 X_{t-1} + u_t$$

$$u_t = \rho u_{t-1} + v_t$$

Lag the first equation and multiply by ρ:

$$\rho Y_{t-1} = \rho\beta_0 + \rho\beta_1 X_{t-2} + \rho u_{t-1}$$

Now subtract from the original equation

$$Y_t - \rho Y_{t-1} = \beta_0 - \rho\beta_0 + \beta_1 X_{t-1} - \rho\beta_1 X_{t-2} + v_t$$

This model is an ADL(*1*, *2*), but with restricted coefficients. Count them, there are three: ρ, β_0, β_1. A freely-estimated version of this mode would have four parameters:

$$Y_t = \alpha_0 + \alpha_1 Y_{t-1} + \alpha_2 X_{t-1} + \alpha_3 X_{t-2} + v_t$$

So, a static model with autocorrelated errors is a restricted version of a more general autoregressive distributed lag model. Those restrictions are testable, but the big point is that evidence of autocorrelation in the residuals is a strong indicator of dynamic mis-specification.

More general tests

The above only test for AR(1). A more general test is the LM test due to Godfrey. This tests for higher order auto (1 with annual data, 4 with quarterly, etc.) *and* it works with stochastic regressors.

Suppose we suspect an AR(2) structure:

$$u_t = \rho_1 u_{t-1} + \rho_2 u_{t-2} + v_t$$

then (in rather sloppy notation)

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

(A) $y = X\beta + \rho_1 u_{t-1} + \rho_2 u_{t-2} + v_t$

Under H_0: $\rho_1 = \rho_2 = 0$ (no autocorrelation) we have

(B) $y = X\beta + v$.

B is the restricted version of A. Hence using the LM principle:

- **estimate B and obtain residuals e**
- **regress residuals on $X\beta$, u_{t-1}, u_{t-2} (but replace unobservable u by e)**
- **$nR^2 \sim \chi_q^2$ from latter regression, where q = 2 is the number of restrictions.**

8. Estimation in the presence of autocorrelation

What to do depends upon the cause. If it's due to omitted variables or wrong functional form then these should be cured. Often the cause is insufficient lagged values of the X's being included (theory doesn't tell us much about this). The RESET test for functional form should help here.

But if you have 'genuine' autocorrelation then you should use GLS, as for heteroscedasticity. The Ω matrix looks like (for AR(1))

$$\Omega = \begin{pmatrix} 1 & \rho & \rho^2 & \cdots & \rho^{T-1} \\ \rho & 1 & \rho & & \rho^{T-2} \\ \rho^2 & \rho & \ddots & & \\ \vdots & & & & \rho \\ \rho^{T-1} & \rho^{T-2} & & \rho & 1 \end{pmatrix}$$

so this could be put into the GLS formula $b_{GLS} = (X'\Omega^{-1} X)^{-1} X' \Omega^{-1} y$ *if* we knew *ρ*. This would be BLUE. This is equivalent to using OLS to estimate

(C) $y_t - \rho y_{t-1} = (1-\rho)\beta_0 + \beta_1 (x_{1t} - \rho x_{1t-1}) + ... + \beta_k (x_{kt} - \rho x_{kt-1}) + u_t - \rho u_{t-1}$

Note that the error term is now OK, since $u_t - \rho u_{t-1} = v_t$. This procedure loses the first sample observation, but we can construct it using $y'_1 = y_1 \sqrt{(1-\rho^2)}$. etc.