

5. Differences-in-Differences and A Brief Introduction to Panel Data


Differences-in-Differences

We have already come across the idea of ‘differencing’ as a way to deal with the problem of omitted variables. In the context of the analysis of experimental data the simple comparison of the mean of the outcome in treatment and control groups (the ‘differences’ estimator) is justified on the grounds that the randomization guarantees they should not have any systematic differences in any other pre-treatment variable.

This idea of trying to mimic an experiment suggests trying to find equivalents of ‘treatment’ and ‘control groups’ in which everything apart from the variable of interest (or other things that can be controlled for) are assumed to be the same. But this is often a very difficult claim to make as it is rarely possible to do this perfectly in which case observed differences between treatment and control groups may be the result of some other omitted factors.

But, even if one might not be prepared to make the assumption that the treatment and control groups are the same in every respect apart from the treatment one might be prepared to make the assumption that, in the absence of treatment, the unobserved differences between treatment and control groups are the same over time.

In this case one could use data on treatment and control group before the treatment to estimate the ‘normal’ difference between treatment and control group and then compare this with the difference after the receipt of treatment. Perhaps a graph will make the idea clearer



If one just used data from the post-treatment period then one would estimate the treatment effect as the distance AB – this estimate being based on the assumption that the only reason for observing a difference in outcome between treatment and control group is the receipt of treatment.

In contrast the ‘difference-in-difference’ estimator will take the ‘normal’ difference between the treatment and control group as the distance CB and estimate the treatment effect as the distance AC. Note that the validity of this is based on the assumption that the ‘trend’ in y is the same in both treatment and control group – if, for example, the trend was greater in the treatment group then AC would be an over-estimate of the treatment group. One can never test this identifying assumption of the same trend in the absence of treatment. But, if there are more than two observations on treatment and control group one can see whether in other periods the assumption of a common trend seems to be satisfied - I will give examples of this a bit later.

Let’s introduce some notation. Define μ_{it} to be the mean of the outcome in group i at time t . Define $i=0$ for the control group and $i=1$ for the treatment group. Define $t=0$ to be a pre-treatment period and $t=1$ to be the post-treatment period (though only the treatment group gets the treatment).

The difference estimator we have discussed so far simply uses the difference in means between treatment and control group post-treatment as the estimate of the treatment effect i.e. it uses an estimate of $(\mu_{11} - \mu_{01})$. However, this assumes that the treatment and control groups have no other differences apart from the treatment, a very strong assumption with non-experimental data. A weaker assumption is that any differences in the change in means between treatment and control groups is the result of the treatment i.e. to use an estimate of $(\mu_{11} - \mu_{01}) - (\mu_{10} - \mu_{00})$ as an estimate of the treatment effect – this is the *differences-in-differences* estimator.

How can one estimate this in practice? One way is to write the D-in-D estimator as $(\mu_{11} - \mu_{10}) - (\mu_{01} - \mu_{00})$ - note that the first term is the change in outcome for the treatment group and the second term the change in outcome for the control group and then simply estimate the model:

$$\Delta y_i = \beta_0 + \beta_1 X_i + \varepsilon_i \quad (5.1)$$

where:

$$\Delta y_i = y_{i1} - y_{i0} \quad (5.2)$$

Note that this is simply the differences estimator applied to differenced data.

To implement the difference-in-difference estimator in the form in (5.1) requires data on the same individuals in both the pre- and post- periods. But it might be the case that the individuals observed in the two periods are different so that those in the pre-period who are in the treatment group are observed prior to treatment but we do not observe their outcome after the treatment. If we use $t=0$ to denote the pre-period and $t=1$ to denote the post-period y_{it} to denote the outcome for individual i in period t then an alternative regression-based estimator that just uses the level of the outcome variable is to estimate the model:

$$y_{it} = \beta_0 + \beta_1 X_i + \beta_2 T_t + \beta_3 X_i * T_t + \varepsilon_{it} \quad (5.3)$$

where X_i is a dummy variable taking the value 1 if the individual is in the treatment group and 0 if they are in the control group, and T_t is a dummy variable taking the value 1 in the post-treatment period and 0 in the pre-treatment period.

The D-in-D estimator is going to be the OLS estimate of β_3 , the coefficient on the interaction between X_i and T_i . Note that this is a dummy variable that takes the value one only for the treatment group in the post-treatment period.

From what you have done already you should be able to prove the following Proposition.

Proposition 5.1: In the estimation of (5.3) we will have that:

$$\begin{aligned} p \lim \hat{\beta}_0 &= \mu_{00} \\ p \lim \hat{\beta}_1 &= \mu_{10} - \mu_{00} \\ p \lim \hat{\beta}_2 &= \mu_{01} - \mu_{00} \\ p \lim \hat{\beta}_3 &= (\mu_{11} - \mu_{01}) - (\mu_{10} - \mu_{00}) \end{aligned} \tag{5.4}$$

Proof: An Exercise ☺

Proposition 5.1 implies that $\hat{\beta}_3$ is a consistent estimate of the treatment effect.

Where one has repeated observations on the same individuals one can use both estimation methods - (5.1) and (5.3) – on the same data and they will give exactly the same estimate of the treatment effect. However the standard error of that estimate will be different in the two cases – the class exercise asks you about the reasons for that.


Other Regressors

You can include other regressors in either (5.1) or (5.3). Note that if you think it is the level of some variable that affects the level of y then you should probably include the change in that variables as one of the other regressors if one is estimating the model (5.1) i.e. in differenced form.

Differential Trends in Treatment and Control Groups

The validity of the differences-in-differences estimator is based on the assumption that the underlying ‘trends’ in the outcome variable is the same for both treatment and control group. This assumption is never testable and with only two observations one can never get any idea of whether it is plausible. But, with more than two observations we can get some idea of its plausibility.

To give an example, consider the paper “Vertical Relationships and Competition in Retail Gasoline Markets”, by Justine Hastings, published in the American Economic Review, 2004. She was interested in the effect on retail petrol prices as a result of an increase in vertical integration when a chain of independent Californian ‘Thrifty’ petrol stations were acquired by ARCO, who also have interests in petrol refining. She defined a petrol station as being in the ‘treatment’ group if it was within one mile of a Thrifty station (i.e. one can think of it as having a competitor that was a ‘Thrifty’) and in the ‘control’ group if it did not. Because there are likely to be all sorts of factors that causes petrol prices to differ across locations, this lends itself to a difference-in-difference approach. The basic conclusions can be summarized in the following graph.




(a) LOS ANGELES

Before the acquisition, prices in the ‘treatment’ group were, on average 2-3 cents lower than in the control group, but after the acquisition they were 2-3 cents higher. Hence, the difference-in-difference estimate of the effect of the acquisition is 5 cents.

This picture also presents information on prices not just in the periods immediately prior to the acquisition and immediately afterwards but also in other periods. One can see that the trends in prices in treatment and control groups are similar in these other periods suggesting that the assumption of common trends is a reasonable one.

Lets also give a famous example where the D-in-D assumption does not seem so reasonable. In “Estimating the Effect of Training Programs on Earnings”, Review of Economics and Statistics, 1978, Orley Ashenfelter was interested in estimating the effect of government-sponsored training on earnings. He took a sample of trainees under the Manpower Development and Training Act (MDTA) who started their training in the first 3 months of 1964. Their earnings were tracked both prior, during and after training from social security records. A random sample of the working population were used as a comparison group.

The average earnings for white males for the two groups in the years 1959-69 inclusive are shown in the following Figure.



There are several points worth noting

First the earnings of the trainees in 1964 are very low because they were training and not working for much of this year – we should not pay much attention to this.

Secondly the trainee and comparison groups are clearly different in some way unconnected to training as their earnings both pre- and post-training are different. This means that the differences estimator based, say, on 1965 data would be a very misleading measure of the impact of training on earnings. This suggests a differences-in-differences approach.

A simple-minded approach would be to use the data from 1963 and 1965 to give a difference-in-difference estimate the effect of training on earnings. However inspection of the figures suggests that the earnings of the trainees were rather low not just in the training year, 1964, but also in the previous year, 1963¹. This is what is known as ‘Ashenfelter’s Dip’ – the most likely explanation is that the trainees had a bad year that year (e.g. they lost a job) and it was this that caused them to enter training in 1964. Because the earnings of the trainees are rather low in 1963 a simple differences-in-differences comparison of earnings in 1965 and 1963 is likely to over-estimate the returns to training. Ashenfelter then describes a number of ways to deal with this problem that I am not going to discuss here – the point is that observations on multiple years can be used to shed light on whether the assumption underlying the use of differences-in-differences is a good one.

¹ Note that they are only slightly lower than 1962 earnings but the usual pattern is for earnings growth so it is quite a lot lower than what one might have expected to see.

Panel Data

I introduced the D-in-D estimator in a situation where we have two observations on treatment and control group. This is the most rudimentary form of panel data – a situation used to describe data where we have more than one observation on the same individual so the data set has both a cross-section and a time-series dimension. Whole books have been written about the analysis of panel data which sometimes give the impression that the analysis of such data requires very different ideas and estimation technologies. But, while there are differences, the basics are not very different from standard regression and I will emphasize those similarities here.

Let's start with some terminology. I will denote the number of individuals in the data set by N and the number of time periods over which we have information on the individual by T . I will restrict attention to balanced panels in which we have the same number of observations on every individual – the analysis of unbalanced panels in which we have more observations on some individuals than others is not much more difficult but the notation is messier. Denote by y_{it} the outcome variable for individual i in period t – similarly define x_{it} .

In total we will have NT observations. One difference from normal cross-section data is that when we do asymptotics and take the number of observations to infinity, we can do this in a number of ways – N can go to infinity with T fixed, T can go to infinity with N fixed or both could go to infinity. It is most common to see the 'large N , fixed T case' because that is felt to be the best approximation to the situation in which researchers in microeconomics find themselves. But in other parts of the subject e.g. macroeconomics, it is common to be in a 'fixed N , large T ' case for which the asymptotics can be very different.

A first approach to estimating models using panel data would be to ignore the panel nature of the data and simply estimate:

$$y_{it} = \beta' x_{it} + \varepsilon_{it} \quad (5.5)$$

OLS estimation of this will lead to a consistent estimate of β if $E(\varepsilon_{it} | x_{it}) = 0$.

However, unless STATA is told otherwise, the standard errors will be computed under the assumption that the ε_{it} are all independent of each other, something that is very unlikely. For example, it is very likely that the outcome variable for the same individual is strongly correlated over time. There are a number of ways in which one might capture this idea – I will discuss one of them.

This is to introduce an individual specific component into (5.5) and write:

$$y_{it} = \beta' x_{it} + \theta' D_i + \varepsilon_{it} \quad (5.6)$$

Where θ is an $N \times 1$ vector and D_i is a vector consisting of zeros everywhere except for a 1 in the i th position. (5.6) will often be written more commonly as:

$$y_{it} = \beta' x_{it} + \theta_i + \varepsilon_{it} \quad (5.7)$$

You should recognise these models as very similar to the ones we discussed when talking about clustered standard errors – the 'group' is now a particular individual.

I will discuss a number of different versions of (5.7) that are commonly estimated:

- the fixed effects model
- the random effects model
- the between-groups model

The Fixed Effects Model

The first, the *fixed effects model*, treats θ_i as a parameter to be estimated (similar to β). The θ_i is often called the individual fixed effect. Note that there will often be a very large number of these individual fixed effects to be estimated as N can often be very large. Although this may cause computational problems (much more in the past than now) there is nothing conceptually difficult about this estimator. In STATA one would write a command like

```
. xtreg y x, fe i(id)
```

Where id is a variable indicating the individual observation.

The assumption required for the fixed effect estimator to give a consistent estimate of β is contained in the following Proposition:

Proposition 5.2: The fixed effect estimator of β in (5.7) will be consistent under the following assumptions:

- $E(\varepsilon_{it} | x_{it}) = 0$ (5.8)
- $\text{rank}(X, D) = N + K$

Proof: This is simple application of what you should know about the linear regression model. ☺

The intuition for the first condition should be obvious but the second may need more explanation. It amounts to the usual condition that the matrix of regressors must be of full rank i.e. we do not have perfect multicollinearity among them. In the context of panel data models the main way in which this will fail when using the fixed effects estimator is if some regressor varies only across individuals and not over time for a given individual. The fixed effect estimator cannot provide a consistent estimate of the coefficient on such a variable as it will be perfectly multicollinear with the individual fixed effect.

One could implement the fixed effects estimator by brute force, generating a dummy variable for each observation and then including these dummy variables in the regression. But, in practice this estimator is often estimated in mean deviation form – define the mean of y for individual i as:

$$\bar{y}_i = \frac{1}{T} \sum y_{it} \quad (5.9)$$

And then the deviation of each observation from this mean:

$$\tilde{y}_{it} = y_{it} - \bar{y}_i \quad (5.10)$$

And similarly for x. Taking means of (5.7) and taking these means away from (5.7), we can write the model as:

$$\tilde{y}_{it} = \beta' \tilde{x}_{it} + \tilde{\varepsilon}_{it} \quad (5.11)$$

And the θ_i disappear. The model in (5.11) can simply be estimated by OLS albeit without a constant (as all variables will, by construction, have mean zero). That this will give exactly the same estimate of β (and the same standard errors) as direct estimation of (5.7) follows from the Frisch-Waugh Theorem.

This way of writing the fixed effects estimator is also helpful in showing that the only variation that is exploited is variation for the same individual – this is the reason it is sometimes called the ‘within-groups’ estimator. Because this variation is often a quite small part of the total variation, fixed effect estimators can often have quite low precision.

This is also an alternative way to show that the fixed effects estimator cannot be used if one wants to estimate the effect of a regressor that is always constant for the same individual e.g. gender or (in many applications) education. The easiest way to see this is to think that the deviation of such a variable from the individual-specific mean will always be zero.

For these reasons among others, one often sees another estimator – the *random effects estimator*.

The Random Effects Estimator

In this case θ_i is viewed not as a parameter to be estimated but as a component in the error.

The presence of θ_i in the error now allows the residuals to be correlated for individual i though still uncorrelated for different individuals. You should recognise this model as essentially the same as the model we introduced for clustered standard errors.

Given the assumption of the independence of the θ_i one could simply estimate the model (5.7) by OLS and correct the standard errors (essentially using the same formula as used for clustered standard errors).

However, this is not the random effects estimator. This is the feasible GLS estimator which can be written as:

$$\hat{\beta}^{RE} = \left(X' \hat{\Omega}^{-1} X \right)^{-1} X' \hat{\Omega}^{-1} y \quad (5.12)$$

Where $\hat{\Omega}$ is an estimate of the covariance matrix of the residuals. I am not going to go into the detail of how $\hat{\Omega}$ is computed – a book like Wooldridge “Econometric Analysis of Cross-Section and Panel Data”, ch10, has a good discussion.

Implementation using STATA is straightforward. Simply type

```
. xtreg y x, re i(id)
```

When will the random effects estimator give a consistent estimate of β ? – the assumptions are contained in the following Proposition.

Proposition 5.3: The random effects estimator of β in (5.7) will be consistent under the following assumptions:

$$a. \quad E(\varepsilon_{it} | x_{i1}, x_{i2}, \dots, x_{iT}) = 0 \quad (5.13)$$

$$b. \quad E(\theta_i | x_{i1}, x_{i2}, \dots, x_{iT}) = 0 \quad (5.14)$$

$$c. \quad \text{rank}(X' \Omega^{-1} X) = k$$

Proof: The random effects estimator is a special case of the feasible GLS estimator and the conditions for the consistency of that estimator apply also here. The two conditions (5.13) and (5.14) are needed because the error term has two components.

In comparing Proposition 5.3 with Proposition 5.2 there are a number of features to note. First the assumptions about the exogeneity of the errors required for the consistency of the random effects estimator are stronger than those for the consistency of the fixed effects estimator. There are two parts to this. First, we need to assume that the individual component θ_i is uncorrelated with the regressors for individual i , something we did not need in the case of the fixed effects estimator. One way of understanding this result is to think of the individual fixed effects being omitted variables in the random effects model and this leads to inconsistency unless the omitted variables are uncorrelated with the regressors.

The second aspect to the strengthening of the exogeneity conditions is (5.13) – it is no longer enough to assume that ε_{it} is uncorrelated with the contemporaneous values of the regressors – it must be uncorrelated with the whole history for the individual – what is sometimes called strict exogeneity. The reason this is required is that the feasible GLS estimator in (5.12) is exploiting the fact that the residuals for an individual are serially correlated and when residuals are serially correlated one needs to make stronger exogeneity assumptions.. One could weaken (5.13) to (5.8) (though one has to maintain (5.14)) if one estimated the model using OLS though one would have to use the cluster option to get correct standard errors.

One might wonder from this discussion why anyone ever uses the random effects rather than the fixed effects model – there are two possible reasons.

First, the rank condition for consistent estimation of β is weaker in the random effects case than the fixed effects case – essentially this is because one can estimate the effects of variables that do not vary for a given individual.

Secondly, if the assumptions needed for the consistency of the random effects estimator are satisfied, it will be more efficient than the fixed effects estimator i.e. the standard errors will typically be smaller. Essentially this is an example of the general principle that imposing a true restriction on the data leads to greater efficiency. Intuitively this is because the random effects estimator uses all the variation in the regressors not just the within-group variation as is the case with the fixed effects estimator.

Before we leave the discussion of the random effects estimator, we will discuss another useful result. It can be shown (I am not going to do it but a good place to find

more information is Wooldridge “Econometric Analysis of Cross-Section and Panel Data”, ch10, pp286-7) that the RE estimator can be thought of as is an OLS regression of $\tilde{y}_{it} = (y_{it} - \lambda \bar{y}_i)$ on $\tilde{x}_{it} = (x_{it} - \lambda \bar{x}_i)$ where

$$\lambda = 1 - \sqrt{\frac{\sigma_{\varepsilon}^2}{\sigma_{\varepsilon}^2 + T\sigma_{\theta}^2}} \quad (5.15)$$

This is sometimes referred to as quasi-time demeaning. If λ is close to one the random and fixed effects estimates will be very similar – inspection of (5.15) shows that this is the case when $\lambda = T\sigma_{\theta}^2 / \sigma_{\varepsilon}^2$ is large which is when the number of observations per individual is large and when the variance in the individual fixed effect is large relative to the variance in ε .

The Between-Groups Estimator

A third estimate is based on using the individual means as defined in (5.9). If we take individual means of (5.7) we end up with:

$$\bar{y}_i = \beta' \bar{x}_i + \theta_i + \bar{\varepsilon}_i \quad (5.16)$$

The between-groups estimator is simply the OLS estimate of this model. It can be implemented in stata using the command:

```
. xtreg y x, be i(id)
```

The exogeneity assumptions required for the consistency of the between-groups estimator are the same as described for the random effects estimator in Proposition 5.3. This is simple to understand as consistency requires no correlation between the error in (5.16) and the regressors.

But the between-groups estimator does not exploit the variation in the regressors over time for a given individual so is less efficient than the random effects estimator. And for variables that do not vary across individuals e.g. time trend, the between-groups estimator cannot identify the effects. Hence, one might wonder why the between-groups estimator is of any interest.

One possible reason comes from thinking about the effects of measurement error in panel data models.

Measurement Error in Panel Data Models

Suppose the true model is:

$$y_{it} = \beta_0 + \beta_1 x_{it}^* + \theta_i + \varepsilon_{it} \quad (5.17)$$

Where x^* is the true value of x . Let's assume that $E(\theta_i | x_{i1}^*, \dots, x_{iT}^*) = 0$ and that

$E(\varepsilon_{it} | x_{i1}^*, \dots, x_{iT}^*) = 0$ so that both random and between-groups estimators are consistent. To keep things simple assume x is one-dimensional.

However, x^* is only observed with measurement error. Assume the observed value of x is given by:

$$x_{it} = x_{it}^* + u_{it} = x_i^* + \eta_{it} + u_{it} \quad (5.18)$$

Where u is classical measurement error x_i^* is average value of x^* for individual i and η_{it} is variation around the true value which is assumed to be uncorrelated with x_i^* and u and iid.

We know that this type of measurement error is likely to cause attenuation bias. The important point here is that the extent of attenuation bias is likely to vary between fixed, random and between-groups estimators. The results are contained in the following Proposition.

Proposition 5.4:

a. For the fixed effects model the attenuation bias is given by:

$$p \lim \hat{\beta}_1^{FE} - \beta_1 = -\beta_1 \frac{Var(u)}{Var(\eta) + Var(u)} \quad (5.19)$$

b. For the between-groups model the attenuation bias is given by:

$$p \lim \hat{\beta}_1^{BE} - \beta_1 = -\beta_1 \frac{Var(u)}{TVar(x^*) + Var(\eta) + Var(u)} \quad (5.20)$$

c. For the random-effects model the attenuation bias is given by:

$$p \lim \hat{\beta}_1^{RE} - \beta_1 = -\beta_1 \frac{Var(u)}{\kappa Var(x^*) + Var(\eta) + Var(u)} \quad (5.21)$$

Where:

$$\kappa = \frac{(1-\lambda)^2}{\left(1 + \frac{\lambda(\lambda-2)}{T}\right)} = \frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + \sigma_\theta^2} < 1$$

Proof:

The way to prove this is to invoke the result from our earlier discussion of the attenuation bias caused by classical measurement error (summarized in Proposition 4.2) that the bias is measured by the difference in variance between the variance in the true variable and the observed (which will be the variance in the measurement error) divided by the observed.

a. We know from the earlier discussion that the fixed effects estimator can be written as the OLS estimate on de-meaned data. Using (5.18) we have that:

$$\bar{x}_i = x_i^* + \frac{1}{T} \sum \eta_{it} + \frac{1}{T} \sum u_{it} \quad (5.22)$$

And that:

$$\bar{x}_i^* = x_i^* + \frac{1}{T} \sum \eta_{it} \quad (5.23)$$

This means that we have:

$$x_{it} - \bar{x}_i = \eta_{it} + u_{it} - \frac{1}{T} \sum \eta_{it} - \frac{1}{T} \sum u_{it} \quad (5.24)$$

And that:

$$x_{it}^* - \bar{x}_i^* = \eta_{it} - \frac{1}{T} \sum \eta_{it} \quad (5.25)$$

Taking variances we have that:

$$\text{Var}(x_{it} - \bar{x}_i) = \frac{T-1}{T} \text{Var}(\eta) + \frac{T-1}{T} \text{Var}(u) \quad (5.26)$$

and that:

$$\text{Var}(x_{it}^* - \bar{x}_i^*) = \frac{T-1}{T} \text{Var}(\eta) \quad (5.27)$$

Using (5.26) and (5.27) gives (5.19) as the standard formula for the attenuation bias.

b. Now consider the attenuation bias in the between-groups estimator. Using (5.22) and (5.23) we have that:

$$\text{Var}(\bar{x}_i) = \text{Var}(x^*) + \frac{1}{T} \text{Var}(\eta) + \frac{1}{T} \text{Var}(u) \quad (5.28)$$

And that:

$$\text{Var}(\bar{x}_i^*) = \text{Var}(x^*) + \frac{1}{T} \text{Var}(\eta) \quad (5.29)$$

Using (5.28) and (5.29) gives (5.20) as the standard formula for the attenuation bias.

c. To derive the attenuation bias for the random effects model is bit messier but we will use the earlier result of thinking of this estimator as a quasi-time de-meaned i.e. an OLS regression of $\tilde{y}_{it} = (y_{it} - \lambda \bar{y}_i)$ on $\tilde{x}_{it} = (x_{it} - \lambda \bar{x}_i)$ where λ is given in (5.15)

The attenuation bias can be written as:

$$\frac{\text{Var}(\tilde{u}_{it})}{\text{Var}(\tilde{x}_{it})} = \frac{\text{Var}\left(u_{it} - \frac{\lambda}{T} \sum u_{is}\right)}{\text{Var}\left(x_{it} - \frac{\lambda}{T} \sum x_{is}\right)} = \frac{\text{Var}\left(u_{it} - \frac{\lambda}{T} \sum u_{is}\right)}{\text{Var}\left(x_{it}^* - \frac{\lambda}{T} \sum x_{is}^*\right) + \text{Var}\left(u_{it} - \frac{\lambda}{T} \sum u_{is}\right)} \quad (5.30)$$

Now:

$$\text{Var}\left(u_{it} - \frac{\lambda}{T} \sum u_{is}\right) = \sigma_u^2 \left(1 + \frac{\lambda^2}{T} - \frac{2\lambda}{T}\right) = \sigma_u^2 \left(1 + \frac{\lambda(\lambda-2)}{T}\right) \quad (5.31)$$

And we have:

$$\begin{aligned} \text{Var}\left(x_{it}^* - \frac{\lambda}{T} \sum x_{is}^*\right) &= \text{Var}\left((1-\lambda)x_{it}^* + \eta_{it} - \frac{\lambda}{T} \sum \eta_{is}\right) \\ &= (1-\lambda)^2 \sigma_{x^*}^2 + \sigma_\eta^2 \left(1 + \frac{\lambda(\lambda-2)}{T}\right) \end{aligned} \quad (5.32)$$

Hence the attenuation bias is related to:

$$\frac{\text{Var}(u)}{\kappa \text{Var}(x^*) + \text{Var}(\eta) + \text{Var}(u)} \quad (5.33)$$

Where:

$$\kappa = \frac{(1-\lambda)^2}{\left(1 + \frac{\lambda(\lambda-2)}{T}\right)} = \frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + \sigma_\theta^2} < 1 \quad (5.34)$$

☺

This is all rather complicated and you should not worry about the details. But you should note and understand the following. First, the fixed effects estimator is going to

be more vulnerable to attenuation bias than the random or between-groups estimator. One can see this because the $\text{Var}(x^*)$ does not appear in the denominator of the expression for attenuation bias. The intuition for this is simple: the fixed effects estimator does not use the variation in x^* across individuals – hence a greater part of the variation that remains is measurement error.

Secondly, the random effects estimator has more attenuation bias than the between-groups estimator as we will have $T > \kappa$. The intuition is that the averaging in the between-groups estimator reduces the importance of measurement error.

One should realize that these results are dependent on the particular form of measurement error assumed – things would be very different if measurement error was an individual effect that did not vary over time. But the example does make the point that measurement error considerations should be taken into account when choosing a panel data estimator.

Time Effects

The discussion so far has treated the individual and time dimensions of panel data very differently. But one might also expect the outcomes of individuals at a particular point in time to be correlated because, for example, of the presence of aggregate shocks. The most common way to do this is to include a set of time effects, one dummy variable for each period, in the model. If one does this with individual fixed effects as well the only variation in a regressor one will be exploiting is that after taking individual and time effects out which effectively means individual-specific time trends.

The Fixed Effects Estimator in Difference Form

When estimating the fixed effect estimator, we got ‘rid’ of the individual fixed effects by de-meaning every variable. But an alternative is to take differences. Write (5.7) for the previous period as:

$$y_{it-1} = \beta' x_{it-1} + \theta_i + \varepsilon_{it-1} \quad (5.35)$$

And then take this away from (5.7) to give:

$$\Delta y_{it} = \beta' \Delta x_{it} + \Delta \varepsilon_{it} \quad (5.36)$$

Note that θ_i have disappeared and one could simply estimate this by OLS. For this to deliver consistent estimates you should find it obvious that one requires Δx and $\Delta \varepsilon$ to be uncorrelated. This is stronger than the exogeneity assumption required for consistency of the fixed effects estimator (see (5.8)) but weaker than that required for consistency of the random effects estimator (see (5.13)).

If there are only two periods this will give exactly the same estimates of the coefficients as the ‘de-meaned’ method – the exercise asks you to show this. If there are more than two periods then the estimated coefficients will be different.

But, in all cases the reported standard errors will be different. The reason is one issue we have not touched upon so far in panel data that is very important. We have worried about the fact that observations in the same cross-sectional unit might have errors that are correlated but we have not worried about the fact that the errors for a

given observation might be correlated over time – this is autocorrelation. Yet this is very likely to be the case.

Indeed the difference in the standard errors for the two methods of computing the standard errors comes down to this difference.

In the ‘levels’ version the standard errors will be computed using the assumption that:

$$Cov(\varepsilon_{it}, \varepsilon_{is}) = 0, \quad t \neq s \quad (5.37)$$

whereas in the ‘difference’ version they will assume that:

$$Cov(\Delta \varepsilon_{it}, \Delta \varepsilon_{is}) = 0, \quad t \neq s \quad (5.38)$$

These are inconsistent as can be seen if we take $s=t-1$. Then the covariance on the left-hand side of (5.38) can be written as:

$$Cov(\Delta \varepsilon_{it}, \Delta \varepsilon_{it-1}) = Cov(\varepsilon_{it}, \varepsilon_{it-1}) - Cov(\varepsilon_{it}, \varepsilon_{it-2}) + Cov(\varepsilon_{it-1}, \varepsilon_{it-2}) - Var(\varepsilon_{it-1}) \quad (5.39)$$

If (5.37) is satisfied then we will have that:

$$Cov(\Delta \varepsilon_{it}, \Delta \varepsilon_{it-1}) = -Var(\varepsilon_{it-1}) \neq 0 \quad (5.40)$$

i.e. (5.38) cannot be true. The converse can also be shown – if (5.38) is satisfied then (5.37) cannot be. Which is the ‘correct’ estimate of the standard error depends on which assumption about the errors is correct though it is possible that neither is and one would want to compute the standard errors in some other way.

To think about these issues we are going to drop the ‘cross-section’ aspect of our data and just focus on the time series. Although we have arrived at a discussion of time series from a discussion of the analysis of cross-sectional and panel data, one should not think of this as anything other than an arbitrary route through econometrics. One could just as easily have started from time series and then introduced the other ideas.