Advanced MLE

One-Way Unit Effects: Fixed, Random, and Between

July 25, 2007

Introduction

Last time, we noted that a key point in panel data models is the extent to which units may be seen as having the same coefficients. The most basic way of thinking of unit effects is in terms of an intercept, i.e., that some units/clusters have higher or lower levels of the outcome variable Y_{it} than others.

Consider a general model in which we have these individual (unit)-level effects:

$$Y_{it} = \alpha_i + \beta X_{it} + u_{it} \tag{1}$$

Here, we represent unit-level effects as separate intercept terms – i.e., some clusters tend to have higher values of Y than others. This is thus known as a **variable-intercepts model** (a la Hsiao 2002, Ch. 1) or a **one-way error components model** (Baltagi 2005, Ch. 2).

Variable-intercept models can be motivated in a bunch of different ways. One is as a model of individual-level heterogeneity (or, put differently, as a means of addressing the issue of omitted variable bias). Suppose for example that we have a model that considers the effects of unit varying, time-constant covariates (V), unit-constant, time-varying covariates (W) and covariates which vary across both units and time (X):

$$Y_{it} = \alpha + \beta X_{it} + \gamma V_i + \delta W_t + u_{it} \tag{2}$$

If we don't have measurements of the factors in V and W, we can consider instead their combined product; this leads naturally to a model where we have unit- and time-specific intercepts:

$$Y_{it} = \alpha + \beta X_{it} + \gamma_i + \delta_t + u_{it} \tag{3}$$

This model has both time- and unit-specific effects, and is known as a two-way error components model (e.g., Baltagi 2005, Ch. 3). Such models are a more general case of that in (1), and are used relatively little in the social sciences.

As a practical matter, there are a few different ways we can estimate the model in (1). The first we'll talk about is the *fixed-effects* approach.

The Fixed Effects Model

Treating the unit effects α_i in (1) as fixed values is, in many respects, the simplest thing we can do. That is, we simply estimate (1) by including $N - 1$ separate indicator variables, one for each unit, in the model along with the Xs. This is mathematically the same thing as *analysis of covariance* (and is the same as ANOVA if we drop the Xs); if we include both unit- and time-specific dummies, its equivalent to *two-way ANCOVA*. Stimson calls this approach the *least-squares dummy variables* (LSDV) method.

In panel data, when N is large relative to T, we have lots of effects to estimate; this means that our individual $\hat{\alpha}$s aren't going to be estimated very precisely. But, we shouldn't really care about that, since they're only "nuisance parameters" anyway. We might call this the "brute force dummy variables" (BFDV) method.

In practice, however, if we have lots of units, we can run into trouble calculating all those $\hat{\alpha}$s. Moreover, Neyman and Scott (1948) show that, for ML models, one's estimates can be inconsistent if the model contains parameters that grow with the number of observations. (The intuition is that, if the number of αs are growing in N, we are never able to "catch up" and get consistent estimates). In addition, estimating a fixed effects model requires inverting a large coefficient matrix (of size $N + k$, possibly up to several thousand square – remember, $\hat{\boldsymbol{\beta}} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y}$). Thankfully, there's another alternative.

Estimating Fixed-Effects Models

There's a useful way to think about the alternative to the BFDV approach. Recall the discussion of unit-varying versus time-varying variables from last time. We can rewrite the model in (1) as:

$$Y_{it} = \alpha_i + \beta_B \bar{X}_i + \beta_W (X_{it} - \bar{X}_i) + u_{it} \tag{4}$$

Now, to the extent that there is variation in the X's across units, it will be captured in the first term, while within–unit variation is in the second term. Moreover, the model in (4) reduces to the model in (1) when (that is, iff) $\beta_B = \beta_W$ (we'll talk more about this in a bit...).

The problem with (4) is that it is unidentified; the α_is and the $\bar{X}_i$s are the same thing (i.e., constants within units that vary only across units). In fact, as we'll see in a minute, they're *exactly* the same thing, which means that the β in the fixed-effects version of (1) is the same as β_W in (4).

Why is this important to estimation? Because it allows us to estimate these models in a much easier fashion. Recall that one characteristic of OLS is that, in the model $Y = X\beta + Z\gamma$ we can get estimates of $\hat{\beta}$ and $\hat{\gamma}$ by first regressing Y on X, and then regressing the residuals from that model in Z. Here, we could do the same thing:

1. Regress Y_{it} on α_i,
2. do the same thing for each of the Xs (that is, regress each X_{it} on α_i),
3. take the residuals from those equations and regress them on each other.

In fact, we don't even have to do all that. If you think about what you'll get when you undertake steps (1) and (2), its clear that the "regression line" for each unit will simply be the within-unit mean $\bar{X}_i$, and that the "residual" will be the difference between the observation and this mean, $X_{it} - \bar{X}_i$. So, in fact, we can estimate the βs by simply creating:

$$\begin{aligned} Y_{it}^* &= Y_{it} - \bar{Y}_i \\ X_{it}^* &= X_{it} - \bar{X}_i \end{aligned}$$

and then estimating

$$Y_{it}^* = \beta_{FE} X_{it}^* + u_{it} \tag{5}$$

This is what Stata's `areg` (stands for <u>a</u>bsorbing <u>reg</u>ression) procedure does. It provides estimates of the β_{FE}s, but not the αs; the latter are said to be "absorbed" (or "swept out", or "conditioned out") of the model. Since we're usually not interested in the αs anyway, this is no big deal.

A key thing to remember about (5) is that it makes use of only the *within-unit variation* in the Ys and the Xs; compared to (4), it's clear that (5) ignores cross-unit variation, since the variables Y_{it}^* and X_{it}^* are "de-meaned." This will be important in just a bit...

Issues with Fixed-Effects Models

Fixed effects models are intuitively appealing, and are fairly widely used, especially in economics. They do, however, have some particular issues that need to be considered in their use.

<u>*Detecting Fixed Effects*</u>

Unsurprisingly, since one way of thinking about fixed effects is as unit-specific independent variables, then an intuitive test for whether they are necessary or not is an F-test:

$$H_0 : \alpha_i = \alpha_j \forall i \neq j; H_1 : \alpha_i \neq \alpha_j \text{ for some } i \neq j \tag{6}$$

where we calculate the F by comparing the sums of squares from the models with and without fixed effects, and F has $(N-1, NT-(N-1))$ degrees of freedom (note that both `areg` and `xtreg, fe` do this for you automatically). A significant F is indicative of a need for fixed effects (or, at least, some kind of unit effects).

In practice, however, this test will almost always be significant, especially if NT is fairly large. If so, you then have to decide whether or not to include them. Which leads to...

Pros and Cons of Fixed Effects

There are good reasons both for and against using fixed effects models...

Reasons to Include Fixed Effects

- If you don't, you have specification bias, and there's nothing worse than that.
- Unit effects have an easy, intuitive explanation for non-methods types; this can be very useful.
- They're generally widely used, lots of people have heard of them; they're established and (almost always) non-controversial.

Reasons NOT to Include Fixed Effects

- They can kill your covariates. Fixed effects models can't include covariates that are constant within units (because they are perfectly collinear with the fixed effects). (Recall that, if you calculate the X^* variables above for unit-constant covariates, you get a uniform zero). Since we often care about the effects of variables that don't change over time (e.g., demographics like race and age in panel studies, or things like contiguity in IR), that can be a big problem (e.g., Green et al. 2001 *IO*).
- Relatedly, if you have variables that change slowly / only a little, their effects will be hard to estimate precisely. This is because they will be highly collinear with the fixed effects (that is, their X^* values will not vary much, and so will all be close to zero). But, see Plumper and Troeger (2007)...
- *Inefficiency.* Fixed effects models use up a lot of degrees of freedom (usually something like $\frac{1}{T}$, which, in short panels, is a lot). This means that one's standard error estimates can be pretty badly affected.

The moral of the story: I'm not going to tell you not to use fixed effects, but be careful...

An Example: Refugee Flows in Africa, 1992-2001

Consider a model of the refugee flows between two countries $A \rightarrow B$. A simple "gravity-type" model is:

$$\begin{aligned} ln(\text{Refugees})_{A\rightarrow Bt} &= \beta_0 + \beta_1 \text{Population Difference}_{ABt} + \beta_2 \text{Distance}_{AB} + \\ & \quad \beta_3 \textsf{POLITY} \text{ Difference}_{ABt} + \beta_4 \text{War Difference}_{ABt} + u_{ABt} \end{aligned} \tag{7}$$

Here, we have four variables: *Population*, ***POLITY** Difference*, and *War Difference* all vary both across dyads and across time, while *Distance* varies only across dyads, but not over

Variable	OLS	Fixed Effects
Constant	-0.32	-0.60
	(0.01)	(0.004)
Population Difference	-0.17	6.86
	(0.22)	(2.55)
Distance	-0.13	(dropped)
	(0.005)	
POLITY Difference	-0.0002	0.005
	(0.016)	(0.022)
War Difference	0.074	0.010
	(0.007)	(0.007)
$\hat{\rho}$	-	0.61

Note: $NT = 23618$ ($N = 2450$, $\bar{T} = 9.6$)

time (why?). Our data are for all "directed dyads" (that is, we include both Zaire → Angola and Angola → Zaire) in Africa (so $N = 2450$) for ten years (1992-2001). We compare OLS and FE models below.

Note several things about the two models:

- The *Distance* variable gets dropped in the fixed-effects model; this is because it doesn't vary "within" a given dyad.
- The effect of the *Population Difference* variable increases substantially in the fixed–effects model, and also changes sign (to be in the "right" direction). This suggests that within-dyad changes in population have large effects on refugee flows.
- The effect of the *POLITY Difference* variable grows a bit, but never attains statistical significance.
- The influence of the *War Difference* variable shrinks considerably in the fixed effects model; we'll talk about this more below.
- The value for the F-test is 12.1 (with 2449 and 21165 *df*) ($p < .001$). This indicates that the fixed-effects specification is "better" than the one without fixed effects. Whether one would want to use fixed effects in this case, however, would depend on how important the *Distance* variable was to you (more on this below...).

The One You Never Hear About: Between Effects

An important aspect of fixed effects models is that they only use information "'within" observations (that is, over time, or whatever else defines repeated observations on the same units). Go back to equation (4):

$$Y_{it} = \alpha_i + \beta_B \bar{X}_i + \beta_W (X_{it} - \bar{X}_i) + u_{it}$$

We said before that this model is unidentified, because the α_is and the $\bar{X}_i$s are perfectly collinear, because each only varies across (that is, between) units i. Because of this collinearity between the unit effects and any variables that vary only across units, fixed effects models are unable to estimate the influence of such covariates (such as the *Distance* variable in the example above).

Thus, as we noted above, an alternative way of estimating a fixed effects model is simply to "de-mean" each of the independent variables and use OLS:

$$Y_{it} - \bar{Y}_i = \alpha + \beta_W (X_{it} - \bar{X}_i) + u_{it}$$

Estimating this model will give you precisely the same estimates of $\hat{\beta}$ as the fixed–effects estimator. For this reason, the fixed effects model is often (and, in my opinion, more accurately) referred to as the "'within" or "within effects" estimator.

Logically, then, we could also consider an estimator that only considered information between units i:

$$\bar{Y}_i = \alpha + \beta_B \bar{X}_i + u_{it} \tag{8}$$

This is often known as the "between" estimator, because it uses only information *between* units i. One can estimate this model by calculating the within-unit mean of each X variable, and then regressing $\bar{Y}_i$ on those "means" using OLS; one can also use `Stata`'s `xtreg ..., be`. If we do this for the refugee data we've been taking as our example, we get the following:

Variable	OLS	Fixed/"Within" Effects	Between Effects
Constant	-0.32	-0.60	-0.30
	(0.01)	(0.004)	(0.03)
Population Difference	-0.17	6.86	-0.25
	(0.22)	(2.55)	(0.53)
Distance	-0.13	(dropped)	-0.13
	(0.005)		(0.01)
POLITY Difference	-0.0002	0.005	0.01
	(0.016)	(0.022)	(0.05)
War Difference	0.074	0.010	0.12
	(0.007)	(0.007)	(0.02)
$\hat{\rho}$	-	0.61	-

Note: $NT = 23618$ ($N = 2450$, $\bar{T} = 9.6$).

Note a few things:

- The between effects are, in general, closer to the OLS (pooled) estimates than are the fixed/within effects estimates; this is particularly true for the *Population Difference* variable (why?).
- The interpretation of these two models is very different.
  - The fixed/within effects model estimates the impact of a one–unit change in X on Y *within a particular unit* i; that is, the impact on Y of a particular unit's moving from X to $X + 1$.
  - The between effects model estimates the impact of a difference of one in X *across two different units* on the expectation of Y; that is, how much would we expect Y to be different between two different units, one of which had $X = k$ and the other had $X = k + 1$.
- This difference in interpretation has some substantive importance. Consider the *War Difference* variable...
  - Within a particular dyad, an increase in the *War Difference* variable (that is, country A experiencing greater conflict and/or country B experiencing less war) has a relatively small, marginally significant impact on refugee flows from A to B.
  - By looking at the between effects model, we can see that our expectation about the difference in refugee flows between two different (comparable) dyads, one with *War Difference* $= k$ and the other with *War Difference* $= k + 1$ is roughly ten times as great.

- Finally, given all this, the OLS results can be seen as something of an "average" of the two other models.

But the real value of the between effects model lies in it relationship to a third model that gets talked about a lot...

Random Effects Models

There's another way to think about individual-level heterogeneity in a panel context: as different components of the error term u_{it}. Suppose we start with a basic model:

$$Y_{it} = X_{it}\beta + u_{it} \tag{9}$$

and we think of the error term as having a number of different components, which vary either over time, across units, or both:

$$u_{it} = \alpha_i + \lambda_t + \eta_{it} \tag{10}$$

where the following conditions all hold:

$$\begin{aligned}
E(\alpha_i) = E(\lambda_t) = E(\eta_{it}) &= 0, \\
E(\alpha_i\lambda_t) = E(\alpha_i\eta_{it}) = E(\lambda_t\eta_{it}) &= 0, \\
E(\alpha_i\alpha_j) &= \sigma^2_\alpha \text{ if } i = j,\ 0 \text{ otherwise}, \\
E(\lambda_t\lambda_s) &= \sigma^2_\lambda \text{ if } t = s,\ 0 \text{ otherwise}, \\
E(\eta_{it}\eta_{js}) &= \sigma^2_\eta \text{ if } i = j,\ t = s,\ 0 \text{ otherwise}, \\
E(\alpha_i X_{it}) = E(\lambda_t X_{it}) = E(\eta_{it} X_{it}) &= 0.
\end{aligned}$$

If all these conditions hold, then the variance of Y_{it} conditional on X_{it} is simply $\sigma^2_\alpha + \sigma^2_\lambda + \sigma^2_\eta$. These are often referred to as *variance components*; each is a party of the total variance of u_{it}.

Note that, if we assume for the moment that $\lambda_t = 0$ (we'll come back to this assumption), we get a model that looks like:

$$Y_{it} = X_{it}\beta + \alpha_i + \eta_{it} \tag{11}$$

which looks an awful lot like the fixed effects model we talked about above. In fact, it is the same model, albeit with different assumptions about the αs, and about the relationship between the αs and the X_{it}s. (More on this in a minute...). In fact, there are a couple different ways we can motivate random-effects models.

Random-effects as Draws from a Distribution

The canonical way of thinking about random effects is to say that, rather than treating the α_is as fixed (and estimating them), we treat each as a random draw from some single distribution. We can then estimate the parameters of that distribution, which (relative to fixed effects) reduces the number of estimable parameters a whole lot. Typically, we assume that the distribution is Normal, i.e., that $\alpha_i \sim i.i.d.\ N(0, \sigma_\alpha^2)$. We then want to estimate $\hat{\sigma}_\alpha^2$, along with $\hat{\sigma}_\eta^2$. Intuitively, think about what this means:

- From what we said above, we know that the total error variance is just $\sigma_u^2 = \sigma_\alpha^2 + \sigma_\eta^2$.
- This means that we need to be able to separate out the unit-specific error component from the unit-and-observation-specific part, or
- at a minimum, figure out the relative contribution of the two components to the overall variability of the errors.

Put differently, we now have a model in which any given error u_{it} is, by necessity, correlated with every other error from the same unit u_{is}, by virtue of the fact that they share a common component α_i. How should/can we estimate this model?

One alternative is OLS. Since the "problems" with the model are all in the error terms, OLS will produce unbiased, consistent estimates for the $\hat{\beta}$s. The problem is that the standard errors will be too small (because we're acting as if we have information on NT separate observations, rather than on T observations on N units). This is a problem not unlike autocorrelation, in some respects.

Thus, if we are to estimate the model in (11), we need to account for the fact that the within-unit errors are correlated. The logical way to do this is through GLS, in the same way that we use GLS to deal with, e.g., heteroscedasticity. To do this, we need to consider the nature of the correlation within units. Let's think for a minute about what the variance of the errors are *within a particular unit* in this setup:

$$
\begin{aligned}
E(\mathbf{u}_i\mathbf{u}_i') \equiv \mathbf{\Sigma}_i &= \sigma_\eta^2 \mathbf{I}_T + \sigma_\alpha^2 \mathbf{i}\mathbf{i}' \\
&= \begin{pmatrix} \sigma_\eta^2 + \sigma_\alpha^2 & \sigma_\alpha^2 & \cdots & \sigma_\alpha^2 \\ \sigma_\alpha^2 & \sigma_\eta^2 + \sigma_\alpha^2 & \cdots & \sigma_\alpha^2 \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_\alpha^2 & \sigma_\alpha^2 & \cdots & \sigma_\eta^2 + \sigma_\alpha^2 \end{pmatrix}
\end{aligned} \tag{12}
$$

where $\mathbf{i}$ is a $T \times 1$ vector of ones. Assuming (as we do) that the errors across observations are uncorrelated, we can then "stack" these unit-level covariance matrices to get the overall variance-covariance matrix for $\mathbf{u}$:

$$\text{Var}(\mathbf{u}) \equiv \boldsymbol{\Omega} = \begin{pmatrix} \boldsymbol{\Sigma}_1 & 0 & \cdots & 0 \\ 0 & \boldsymbol{\Sigma}_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \boldsymbol{\Sigma}_N \end{pmatrix} \tag{13}$$

Now, we can use an estimate of the inverse of the "square root" of this $\hat{\boldsymbol{\Omega}}$ matrix as weights. Its possible (but a pain) to show that:

$$\boldsymbol{\Sigma}^{-1/2} = \frac{1}{\sigma_\eta}\left[\mathbf{I}_T - \left(\frac{\theta}{T}\mathbf{ii}'\right)\right] \tag{14}$$

where

$$\theta = 1 - \sqrt{\frac{\sigma_\eta^2}{T\sigma_\alpha^2 + \sigma_\eta^2}} \tag{15}$$

is an unknown quantity to be estimated.

This all seems really hard, but a close look at (15) will reveal a few useful things:

- If $\sigma_\alpha^2 = 0$, then
  - $\theta = 0$ as well,
  - there are no unit-level effects, and
  - the model reduces to a plain-vanilla pooled OLS model.
- If $\sigma_\eta^2 = 0$, then
  - $\theta = 1$,
  - *all* of the variation in the variables is "within" units, and
  - OLS will yield an $R^2 = 1$ regression.
- The term θ thus gives a measure of the relative sizes of the within- and between-unit variances.
- As T gets bigger, this term goes to zero; moreover, (14) converges to the standard OLS model. (As a practical matter, this means that there will be few differences between fixed- and random-effects model estimates when T is large; more on this in a bit...).

Once we have a consistent estimate of θ, we can use it (à la feasible GLS) to transform the observations in the data:

$$
\begin{aligned}
Y^*_{it} &= Y_{it} - \theta\bar{Y}_i \\
X^*_{it} &= X_{it} - \theta\bar{X}_i
\end{aligned}
$$

This gives the following feasible generalized least squares (FGLS) equation for the random-effects model (11):

$$
Y^*_{it} = (1-\theta)\alpha + X^*_{it}\beta_{RE} + [(1-\theta)\alpha_i + (\eta_{it} - \theta\bar{\eta}_i)] \tag{16}
$$

Note that the random-effects model in (16) can be seen to be a (matrix) weighted average of two other models: the "within" (i.e., fixed-effects) model (the one where we simply treat θ as equal to 1.0) and a "between" model, in which we consider only the variation between units (i.e., where we regress the means of Y on those of X). GLS is thus a "compromise" between fixed effects (which ignores the between-unit variation) and OLS.

This GLS approach to estimating (11) is the standard. One can also estimate the model via MLE, but (a) few do so, and (b) the documentation (at least, for doing so in Stata) isn't so good...

An Alternative View: Random Effects as a Restriction on Cov(X_{it}, α_i)

A different way of thinking about random versus fixed effects is solely in terms of the restrictions they put on $\text{Cov}(X_{it}, \alpha_i)$. Johnston and DiNardo (1997) point out, quite rightly, that the important distinction between "fixed" and "random" effects models isn't whether the effects are fixed or random; instead its whether the individual-level disturbance α_i is correlated with the covariates X_{it} or not. In the "random-effects" model, we assume that the unit effects are uncorrelated with the explanatory variables. Conversely, we need make no such assumption in the "fixed-effects" model.

This intuition makes it possible to do some specification tests, which we'll get to in a little bit.

Another Alternative View: Random Effects as Weighted Combinations of FE and BE Estimators

Equation (16) emphasizes a point that I mentioned earlier: *the random effects estimator is a matrix-weighted average of the between and within (fixed) effects estimators*, with the weight being determined by (among other things) the relative amount of within-unit and between-unit variability in the data. In the limit,

- If all the variation in the model is across-unit – e.g., if all the Xs in the data are non-time varying – then the BE and RE estimators will give the same results, asymptotically (and the FE model will be inestimable).

- Likewise, if all the variation is within-unit – e.g., if we are using "de-meaned" data – then FE and RE will give the same results asymptotically (and the BE model will blow up).

Graphically, we can think of the relationship among the models thusly:

Figure 1: Fixed, Between, and Random Effects Estimators

Thus, we can think of the RE estimator as something of a compromise... However, that compromise comes at something of a price, in that it requires that we assume that $Cov(X_{it}, \alpha_i) = 0$. More on that below.

Random Effects Models: Practical Stuff

Stata will, of course, estimate random effects models; it also provides easy ways to assess the relative "fit" of fixed versus random effects.

Estimation in Stata

The Stata command for fixed-effects models is simply `-xtreg-`; you can add the `re` option if you want, but the default is random effects. To obtain the MLE estimates (rather than the FGLS ones, which are the default), add the `mle` option to the end of the command (though I don't know why one would care to do this... details can be found at `http://www.stata.com/support/faqs/stat/mle_vs_gmm.html`).

Stata reports the standard coefficients and standard error estimates. In addition, it also reports estimates of $\hat{\sigma}_\alpha$ and $\hat{\sigma}_\eta$, as well as an estimate $\hat{\rho} = \frac{\hat{\sigma}_\alpha}{\hat{\sigma}_\eta + \hat{\sigma}_\alpha}$, which is simply the proportion of the total variance in u_{it} due to the α_is.

In addition, Stata provides you with estimates of three different "goodness of fit" measures: the "within", "between" and "overall" R^2s.

- The "overall" R^2 is simply the standard R^2 from the regression of Y on X.
- The "between" R^2 is the R^2 from the regression of the means of Y on the means of X (i.e., the R^2 of the "between" estimator).
- The "within" R^2 is similar, and amounts to the same thing as the R^2 from the "prediction equation" $(\hat{Y}_{it} - \hat{\bar{Y}}_i) = (X_{it} - \bar{X}_i)\hat{\beta}$.

The Refugee Example, Redux

Let's reconsider the refugee example from earlier. Specifically, let's compare the results of OLS, fixed, between, and random effects models...

Variable	OLS	Fixed Effects	Between Effects	Random Effects
Constant	-0.32	-0.60	-0.30	-0.31
	(0.01)	(0.004)	(0.03)	(0.03)
Population Difference	-0.17	6.86	-0.25	0.09
	(0.22)	(2.55)	(0.53)	(0.52)
Distance	-0.13	(dropped)	-0.13	-0.13
	(0.005)		(0.01)	(0.01)
POLITY Difference	-0.0002	0.005	0.01	0.0008
	(0.016)	(0.022)	(0.05)	(0.0199)
War Difference	0.074	0.010	0.12	0.022
	(0.007)	(0.007)	(0.02)	(0.007)
$\hat{\rho}$	-	0.61	-	0.56

Note: $NT = 23618$ ($N = 2450$, $\bar{T} = 9.6$).

Things to note:

- The effects of *Distance* are the same across all three models; this is because this variable only varies between / across units, so that the estimates are all bound to be the same.
- The *Population Difference* variable has very different effects across the three models; here it's clear that the RE results are averages of the others. The same is true of the *War Difference* variable.
- The random-effects model can estimate the influence of time-constant variables (which is a real advantage).

The biggest potential problem with random effects is the requirement that $\text{E}(\alpha_i, X_{it}) = 0$. Think about this: There are a bunch of unmeasured factors that all go into α_i and determine whether, on average, observation i will have a higher or lower value of Y_{it}. There are also some Xs that influence the same thing, which *are* in the model. What are the odds that those two sets of things will be uncorrelated? In some circumstances – e.g., field experiments, where we can randomize over important treatments (but still want to allow for unit-level heterogeneity), the answer might be "very likely." But, in most cases where we're using observational data, the odds are slim.

Choosing...

So, how does one choose between fixed- and random-effects models?

Statistical Criteria

There are a couple easy-to-use tests for assessing whether random effects, in particular, are a good idea.

The Breusch/Pagan LM Test

Breusch and Pagan (1980) developed a Lagrangian multiplier test for random effects. The basic idea is that, in the model

$$Y_{it} = X_{it}\beta + \alpha_i + u_{it} \tag{17}$$

if $\text{Var}(\alpha_i) = 0$, then random effects aren't needed. The B-P test is a Lagrangian multiplier test of this hypothesis. (An alternative way of thinking about it is that there is no within-unit correlation; this is also thus a test for that). The command for this test in Stata is `-xttest0-`; an example is presented below. In many respects, this is like an F-test; it's exceedingly unlikely that you'll not reject the null, especially if your NT is large.

The Hausman Test

Recall that one can think of fixed- and random-effects as simply the same model, with different assumptions about $\text{Cov}(\alpha_i, X_{it})$, and that the random effects estimator is nothing more than a weighted combination of $\hat{\beta}_B$ and $\hat{\beta}_W$. Moreover, the fixed and random effects estimators will differ only insofar as $\frac{\sigma^2_\eta}{T\sigma^2_\alpha + \sigma^2_\eta}$ is different from zero (thus, they will differ less and less as $T \to \infty$). Realistically, the former assumption is untestable; but Hausman (1978) gives us a test for whether the coefficients between the two models are the same.

Here's the logic:

- Under the null (i.e., that the data are generated by a random-effects DGP – or, equivalently, that the Xs and the αs are uncorrelated), both fixed- and random-effects are consistent, but the random-effects model is fully efficient. This is because it makes efficient use of both the within- and across-unit information in the data, whereas the fixed effects model only uses the within-unit information (and the between effects model only uses the across-unit information).

- Under the alternative assumption (i.e., that $\text{Cov}(\alpha_i, X_{it}) \neq 0$), the fixed-effects estimator is consistent, while the random-effects one isn't (i.e., they don't converge to the same values of $\hat{\beta}$).

This means that, asymptotically, we can get a handle on whether the random-effects assumption is valid or not by comparing $\hat{\beta}_{FE}$ with $\hat{\beta}_{RE}$. Formally, we do this by calculating:

$$\hat{W} = (\hat{\beta}_{\mathbf{FE}} - \hat{\beta}_{\mathbf{RE}})'(\hat{\mathbf{V}}_{\mathbf{FE}} - \hat{\mathbf{V}}_{\mathbf{RE}})^{-1}(\hat{\beta}_{\mathbf{FE}} - \hat{\beta}_{\mathbf{RE}}) \tag{18}$$

That is, we're taking the difference between the two coefficient vectors, "weighted" by the difference between the VCV estimates; the result is distributed χ^2_k, where k is the number of *non-time-constant* variables in the model (minus the constant). The Stata command for this test is `hausman`.

Note several things about this test:

- Large values tend to reject the null that $\text{Cov}(\alpha_i, X_{it}) = 0$. However, there are plenty of other ways one can reject the null as well. Since this is a general specification test for errors correlated with the covariates, one can also reject the null if (e.g.) your model is mis- or under-specified. This means that, while random effects are probably OK if you fail to reject the null, they're not necessarily *not* OK if you can do so.

- Note as well that the test can only address variables that are time-varying (because its impossible to get estimates of $\hat{\beta}_{FE}$ for those that don't).

- Also, remember that the two estimators $\hat{\beta}_{FE}$ and $\hat{\beta}_{RE}$ converge as $T \to \infty$; so that the Hausman test will (correctly) tend to fail to reject the null with large T.

The Refugee Example, Once Again

Remember these results?

Variable	OLS	Fixed Effects	Between Effects	Random Effects
Constant	-0.32	-0.60	-0.30	-0.31
	(0.01)	(0.004)	(0.03)	(0.03)
Population Difference	-0.17	6.86	-0.25	0.09
	(0.22)	(2.55)	(0.53)	(0.52)
Distance	-0.13	(dropped)	-0.13	-0.13
	(0.005)		(0.01)	(0.01)
POLITY Difference	-0.0002	0.005	0.01	0.0008
	(0.016)	(0.022)	(0.05)	(0.0199)
War Difference	0.074	0.010	0.12	0.022
	(0.007)	(0.007)	(0.02)	(0.007)
$\hat{\rho}$	-	0.61	-	0.56

Note: $NT = 23618$ ($N = 2450$, $\bar{T} = 9.6$).

After estimating the random effects model, we can estimate these two tests. Typing the command `.xttest0` after the model estimates yields a test statistic of 28658.99 (again, $\sim \chi^2_1$), which gives a p-value that is off-the-chart significant. This ought not be so surprising, since our estimate of $\hat{\rho}$ was also very significant.

More interesting is the `hausman` results, which yield $\hat{W} = 30.9$, $\sim \chi^2_3$ for a p-value of < 0.001. This suggests that we can reject the null that random effects are (in some fashion) appropriate for the model we have here. Unless we care a lot about getting an estimate of the influence of the *Distance* variable, we might be better off going with fixed effects.

Other Practical Matters

There are practical as well as statistical reasons for choosing fixed or random effects models.

Panel vs. TSCS data

Recall the difference between panel (large, increasing N; small, fixed T) and TSCS data (smallish, fixed N; larger, increasing T).

- For TSCS data, when we really may actually care about the "names" of the units (e.g., nation-level OECD data), fixed effect can be a good idea. This is particularly true since we're (a) unlikely to have an increasing N, and (b) we are often interested in the influence / effects of variables conditional on there being some unit-specific effect. These are also more likely to be the situations in which we don't necessarily care about the effects of unit-level (i.e., non-time-varying) covariates.

- For panel data, we (a) don't care about the identities of the units (e.g., in surveys), and (b) often think in terms of N increasing while T remains fixed (and small). Here, random-effects are often a good idea, in part because it is here that we're more likely to care about unit-level covariate effects (e.g. demographic variables in survey research).

Covariate effects

One of the biggest reasons for using random effects rather than fixed is when we care, substantively, about the influence of variables that don't change over time. There was a big discussion in the IR literature a few years ago about fixed effects (we'll talk about that more later); one of the problems with fixed effects in that context is that they sweep out the influence of things like contiguity, which everyone knows is important to things in IR (conflict, trade, etc.).

In fact, fixed effects don't prevent one for *controlling for* contiguity, and even do a better job than a simple variable (since fixed-effects don't assume that the effect of contiguity is constant across all dyads the way a pooled analysis typically does). The problem, rather, is that they prevent us from saying anything *summarily* about contiguity effects (e.g., that "contiguous dyads are 552 percent more likely to engage in a dispute", etc.). So, the issue isn't whether you can *control* for the effects, but whether or not you want *estimates* of them.

Separating Between and Within Effects

It's also reasonable to estimate a model that is both a BE and an FE model. That is, if we omit the unit effects α_i, we can estimate:

$$Y_{it} = \alpha + \beta_B \bar{X}_i + \beta_W (X_{it} - \bar{X}_i) + u_{it} \tag{19}$$

This model separates the between and within unit effects, and estimates them separately. To do so, we simply generate a variable for $\bar{X}_i$, the unit-specific mean of X, and another for $X_{it} - \bar{X}_i$, the within-unit deviation around that mean. Once we've done this and estimated the model, it is straightforward to (e.g.) test the condition that $\hat{\beta}_B = \hat{\beta}_W$ for any or all of the independent variables in the model.

If we estimate this model for our refugee data, we get the following:

Variable	Estimate
Constant	-0.32
	(0.01)
Distance	-0.13
	(0.004)
Between (Mean) Population Difference	-0.22
	(0.22)
Within Population Difference	6.86
	(3.74)
Between (Mean) POLITY Difference	0.01
	(0.02)
Within POLITY Difference	0.005
	(0.032)
Between (Mean) War Difference	0.12
	(0.01)
Within War Difference	0.01
	(0.01)

Note: $NT = 23618$ ($N = 2450$, $\bar{T} = 9.6$).

Note that:

- The effect of *Distance* remains the same (no surprise there),
- The effects of the other covariates are the same as those in the respective (between- or fixed/within-effects) models in the table above, with the exception that
- The standard errors are slightly larger here (since there are more variables in the model, and so more collinearity and marginally fewer degrees of freedom).

Estimating Fixed and Random Effects Models using Stata

Not surprisingly, Stata can estimate all these models, and then some. To show this, let's take a closer look at the refugee data we've been using in our example.

Summary Statistics: xtdes and xtsum

The xtdes command is a special version of describe geared toward TSCS/panel data:

```
. xtdes

dirdyadID:  404411, 404420, ..., 651625                      n =       2450
    year:  1992, 1993, ..., 2001                             T =         10
           Delta(year) = 1; (2001-1992)+1 = 10
           (dirdyadID*year uniquely identifies each observation)

Distribution of T_i:   min      5%     25%       50%       75%     95%     max
                         1       9      10        10        10      10      10

     Freq.  Percent    Cum. |  Pattern
 ---------------------------+------------
      2256    92.08   92.08 |  1111111111
        96     3.92   96.00 |  .111111111
        96     3.92   99.92 |  11........
         2     0.08  100.00 |  .1........
 ---------------------------+------------
      2450   100.00         |  XXXXXXXXXX
```

This tells us that:

- There are 2450 units (defined by dirdyadID), and 10 time points (defined by year) in our data.
- The combination of dirdyadID and year uniquely identify every observation in our dataset.
- 2256 of the 2450 units have observations on all ten years; others are missing data, in various patterns.

`xtsum` summarizes the variability in our variables...

```
. xtsum ln_ref_flow pop_diff distance regimedif wardiff

Variable         |      Mean   Std. Dev.       Min        Max |    Observations
-----------------+--------------------------------------------+----------------
ln_ref~w overall | -.6011064    .925829  -.6931472   14.13433 |     N =   23618
         between |             .7517356  -.6931472   9.476154 |     n =    2450
         within  |             .5874289  -9.619386    10.8411 | T-bar =    9.64
                 |                                            |
pop_diff overall |         0   .0274139   -.117949    .117949 |     N =   23618
         between |             .0280993  -.1095172   .1095172 |     n =    2450
         within  |             .0015821  -.0088492   .0088492 | T-bar =    9.64
                 |                                            |
distance overall |  2.200497   1.259073          0      5.652 |     N =   23618
         between |             1.251927          0      5.652 |     n =    2450
         within  |                    0   2.200497   2.200497 | T-bar =    9.64
                 |                                            |
regime~f overall |         0    .381144         -1          1 |     N =   23618
         between |             .3338879      -.955       .955 |     n =    2450
         within  |              .181699      -1.18       1.18 | T-bar =    9.64
                 |                                            |
wardiff  overall |         0   .8806492         -4          4 |     N =   23618
         between |             .6827006       -2.3        2.3 |     n =    2450
         within  |             .5503981       -2.5        2.5 | T-bar =    9.64
```

These tell you the relative amount of within-unit and between-unit variation for each variable. So,

- *Distance* has no within-unit variation at all,
- Most of the variation in *Refugee Flows*, *Population Difference*, *POLITY Difference*, and *War Difference* occurs between units, rather than within – this is especially true for *Population Difference* (which makes sense, since population changes very slowly).

Model Estimation: `xtreg`

The command for estimating panel/TSCS models generally is `xtreg`:

- the `re` option (which is also the default, no-option option) estimates a random effects model,
- the `fe` option estimate a fixed (within) effects model,
- the `be` option estimates a between-effects model.