

Mgmt 469

Fixed Effects Models

Suppose you want to learn the effect of price on the demand for back massages. You have the following data from four Midwest locations:

Table 1: A Single Cross-section of Data

Location	Year	Price	Per capita Quantity
Chicago	2003	\$75	2.0
Peoria	2003	\$50	1.0
Milwaukee	2003	\$60	1.5
Madison	2003	\$55	0.8

This is *cross-section* data – data from several locations at a single point in time.

If you eyeball this data, you will see that *across* the four cities, price and quantity are positively correlated. In fact, if you regress per capita quantity on price, you will obtain a coefficient on price of 0.45, suggesting that each \$1 increase in price is associated with a .45 increase in per capita massages! You highly doubt that the demand curve for massages is upward sloping so you think a bit more about the data. You wonder if, perhaps, the quality of massages is higher in Chicago. Or perhaps “sophisticated urbanites” have a higher demand for massages. This could boost the demand for massages and simultaneously drive up the price. In other words, the regression probably suffers from endogeneity bias in the form of omitted variable bias (because quality and “sophistication” are omitted.)

You really can't say very much about the role of price with this cross-section data.

Now suppose you had the following cross-section/time series data (data from multiple cities and multiple years):

Table 2: A Cross-section/Time Series

<i>Location</i>	<i>Year</i>	<i>Price</i>	<i>Per capita Quantity</i>
Chicago	2003	\$75	2.0
Chicago	2004	\$85	1.8
Peoria	2003	\$50	1.0
Peoria	2004	\$48	1.1
Milwaukee	2003	\$60	1.5
Milwaukee	2004	\$65	1.4
Madison	2003	\$55	0.8
Madison	2004	\$60	0.7

If you eyeball this data, you will see that *within* each of the four cities, price and quantity are inversely correlated, as you would expect if demand is downward sloping. By obtaining multiple observations about each city and looking at the effect of price within each city, we have removed the pernicious effect of omitted variable bias. This is the intuition behind fixed effects regression.

If you examine the sales data (per capita quantity) in Table 2, you will observe that there is considerable variation (“action”) from one row to the next. This variation/action comes in two “flavors”

- *Intercity (across city) variation*: variation in the average quantity from one city to the next
- *Intracity (within city) variation*: variation within each city over time.

The single cross-section of data (Table 1) offered only intercity (across) variation. We now know that regressions relying on intercity variation are problematic due to potential omitted variable bias. The solution is to focus on intracity (within) variation.

We must make one key assumption if we are to believe we have really removed omitted variable bias by focusing on within variation. We must assume that there are no changes in quality or demand over time within each city that we cannot control for. For example, suppose prices went up in Chicago between 2003 and 2004 because quality increased over that time period. If we cannot adequately control for quality in our regression, then we have an omitted variable bias, because omitted quality is correlated with both price and quantity.¹

Using econometrics parlance, we must assume the following:

Identifying assumption: Unobservable factors that might simultaneously affect the LHS and RHS of the regression are time-invariant.

If we accept this identifying assumption – and it is usually an innocuous assumption – then we have a powerful tool for removing omitted variable bias. This tool is known as fixed effects regression, and it exploits within-group variation over time. Across-group variation is not used to estimate the regression coefficients, because this variation might reflect omitted variable bias.

The theory behind fixed effects regressions

Examining the data in Table 2, it is *as if* there were four “before and after” experiments. In other words, there are sales and price data before and after prices change in each of four cities. This concept of “before and after” offers some insight into the estimation of fixed effects models.

Suppose for now that we believe the demand curve has the same shape in each of the four cities. This is a parsimonious assumption and requires that we estimate the average of the four “before-and-after” effects. One way to do this is to compute the changes in prices and the

¹ Be sure you can confirm the direction of the bias and see me before class if you cannot.

changes in quantities in each city (ΔP and ΔQ) and then regress ΔQ on ΔP . (The appropriate data from our massage example are shown in the last two columns of Table 3.) Many research papers estimate such *difference equations*.

Table 3: Data for Difference Equation Estimation

<i>Location</i>	<i>Year</i>	<i>Price</i>	<i>Per capita Quantity</i>	ΔP	ΔQ
Chicago	2003	\$75	2.0		
Chicago	2004	\$85	1.8	10	-0.2
Peoria	2003	\$50	1.0		
Peoria	2004	\$48	1.1	-2	0.1
Milwaukee	2003	\$60	1.5		
Milwaukee	2004	\$65	1.4	5	-0.1
Madison	2003	\$55	0.8		
Madison	2004	\$60	0.7	5	-0.1

If we have lots of years of data, things get messier. We could, in principle, compute all of the differences (i.e., 2004 versus 2003, 2005 versus 2004, etc.) and then run a single regression, but there is an easier way. Instead of thinking of each year's observation in terms of how much it differs from the prior year for the same city, let's think about how much each observation differs from the average for that city. With two key variables, price and quantity, we will be concerned with the following:

- Variations in prices around the mean price for each city
- Variations in quantities around the mean quantity for each city.

We want to know whether variations in the quantities (around their means) are related to variations in prices (around their means). With our relatively small massage data set, we can do this by hand.

Table 4 reports the mean prices and quantities of massages in each of the four cities, where P_{ct}^* and Q_{ct}^* denote these means.

Table 4: Cross-section/Time Series with Mean Within-Group Values

Location	Year	Price	P_{ct}^*	Quantity	Q_{ct}^*
Chicago	2003	75	80	2.0	1.9
Chicago	2004	85	80	1.8	1.9
Peoria	2003	50	49	1.0	1.05
Peoria	2004	48	49	1.1	1.05
Milwaukee	2003	60	62.5	1.5	1.45
Milwaukee	2004	65	62.5	1.4	1.45
Madison	2003	55	57.5	0.8	0.75
Madison	2004	60	57.5	0.7	0.75

Table 5 reports the variation in prices and quantities around their means.

Table 5: Cross-section/Time Series with Variation around the Means

Location	Year	Price	P_{ct}^*	$Price - P_{ct}^*$	Quantity	Q_{ct}^*	$Quantity - Q_{ct}^*$
Chicago	2003	75	80	-5.0	2.0	1.9	0.1
Chicago	2004	85	80	5.0	1.8	1.9	-0.1
Peoria	2003	50	49	-1.0	1.0	1.05	.05
Peoria	2004	48	49	1.0	1.1	1.05	-.05
Milwaukee	2003	60	62.5	-2.5	1.5	1.45	.05
Milwaukee	2004	65	62.5	2.5	1.4	1.45	-.05
Madison	2003	55	57.5	-2.5	0.8	0.75	.05
Madison	2004	60	57.5	2.5	0.7	0.75	-.05

The final step is to regress $Quantity - Q_{ct}^*$ on $Price - P_{ct}^*$. Note that by subtracting the means, we have restricted all of the action in the regression to within-city action. Thus, we have eliminated the key source of omitted variable bias, namely, unobservable across-city differences in quality, sophistication, or whatever!

Fixed effects regression in practice

It turns out that there is a simple way to do this in practice:

- 1) Start with your original price and quantity data.
- 2) Regress quantity on price, but include dummy variables for the cities (remembering to omit one city).
- 3) Add a control for the time trend if you think such a trend might be important.

In Stata, you could run:

regress quantity price peoriadummy milwaukeedummy madisondummy 2004dummy

I ran this regression and obtained the following results:

Variable	Coefficient	What it signifies
Intercept	3.91	This is the intercept of the demand curve in the omitted city (Chicago).
price	-.025	Each \$1 increase in price causes per capita demand to fall by 0.025.
peoriadummy	-1.64	This is the Peoria intercept relative to the Chicago intercept.
milwaukeedummy	-1.72	This is the Milwaukee intercept relative to the Chicago intercept.
madisondummy	-0.89	This is the Madison intercept relative to the Chicago intercept.
2004dum	.039	Per capita sales are higher in 2004 than in 2003 by .039

This is known as a “*fixed effects*” regression because it holds constant (fixes) the average effects of each city.

There is a shortcut in Stata that eliminates the need to create all the dummy variables.

Suppose that our variable names are **quantity**, **price**, **city** and **year**. If we type:

xi: regress quantity price i.city i.year

Stata will automatically create dummies for all but one of the city categories as well as for the year category and then run the fixed effects regression. You can do this procedure with any

number of years of data, provided there are at least two observations per city. (Cities with only one observation will drop out of the regression.)

Behind the scenes of fixed effect regressions

By including fixed effects (group dummies), you are controlling for the average differences across cities in any observable *or unobservable* predictors, such as differences in quality, sophistication, etc. The fixed effect coefficients soak up all the across-group action. What is left over is the within-group action, which is what you want. You have greatly reduced the threat of omitted variable bias.

Because fixed effects models rely on within-group action, you need repeated observations for each group, and a reasonable amount of variation of your key X variables within each group. The more action the better, of course. (When you work with the full yogurt data set, be sure to examine the data to see if you have sufficient intra-group action.)

One potentially significant limitation of fixed effects models is that you cannot assess the effect of variables that have little within-group variation. For example, if you want to know the effect of spectator sports attendance on the demand for massages, you might not be able to use a fixed effects model, because sports attendance within a city does not vary very much from one year to the next. If it is crucial that you learn the effect of a variable that does not show much within-group variation, then you will have to forego fixed effects estimation. But this exposes you to potential omitted variable bias. Unfortunately, there is no easy solution to this dilemma.

The importance of fixed effects regression

Fixed effects regressions are very important because data often fall into categories such as industries, states, families, etc. When you have data that fall into such categories, you will normally want to control for characteristics of those categories that might affect the LHS variable. Unfortunately, you can never be certain that you have all the relevant control variables, so if you estimate a plain vanilla OLS model, you will have to worry about unobservable factors that are correlated with the variables that you included in the regression. Omitted variable bias would result. If you believe that these unobservable factors are time-invariant, then fixed effects regression will eliminate omitted variable bias.

In some cases, you might believe that your set of control variables is sufficiently rich that any unobservables are part of the regression noise, and therefore omitted variable bias is nonexistent. But you can never be certain about unobservables because, well, they are unobservable! So fixed effects models are a nice precaution even if you think you might not have a problem with omitted variable bias.

Of course, if the unobservables are not time-invariant – if they move up and down over time within categories in a way that is correlated with the variables included in the regression – then you still have omitted variable bias. You may never be able to fully rule out this possibility. There are other, more sophisticated solutions that we will discuss later in the quarter.

Advice on using fixed effects

- 1) If you are concerned about omitted factors that may be correlated with key predictors at the group level, then you should try to estimate a fixed effects model.
- 2) Include a dummy variable for each group, remembering to omit one of them. (Use the **xi:** feature in Stata.)
- 3) The coefficient on each key predictor tells you the average effect of that predictor; i.e., the common slope averaged across the groups.
- 4) You can use a partial-F (Chow) test to determine if the groups have different intercepts.

Different slopes for different folks?

In the basic fixed effects model, the effect of each predictor variable (i.e., the slope) is assumed to be identical across all the groups, and the regression merely reports the average within-group effect. What if you believe the slopes differ across all groups? In the extreme, you could estimate a different regression for each group. This will generate a different slope for each predictor variable in each market, which can quickly get out of hand. A more parsimonious solution is to estimate a single fixed effects regression but include slope dummies (i.e. interactions between the group dummies and the key predictor variables) for key predictors and use a Chow test to see if the slopes really are different. We will do this for the yogurt project.

An application

I will walk you through many applications of fixed effect models. Here is one to ponder. There have recently been high profile studies of the relationship between staffing in hospitals and patient outcomes. These studies use traditional OLS regression where the unit of observation is the patient, the dependent variable is some outcome measure like mortality (a dummy variable that equals 1 if the patient died in the hospital) and the key predictor is staffing (e.g., nurses per patient). These studies do not use fixed effects, and invariably show that hospitals with more staff have better patient outcomes. These results have had enormous policy implications.

Unfortunately, these studies may suffer from omitted variable bias. For example, the key unobservable variable might be the severity of patients' illnesses, which is notoriously difficult to control for with available data. Severity is likely to be correlated with both mortality and staffing, so that the coefficient on staffing will be biased. (What is the direction of the bias?)

If you ran a *hospital fixed-effects* model, you would include hospital dummies in the regression that would control for observable and unobservable differences in severity (and all other factors) across hospitals. This would greatly reduce potential omitted variable bias. None of the current research in this field has done so, perhaps because there is not enough intra-hospital variation in staffing to allow for fixed-effects estimation.

Even a fixed effects model would not completely eliminate potential omitted variable bias. You must also hope that any changes over time in unobservable patient severity *within each hospital* are uncorrelated with changes over time in staffing. This might not be such a good assumption. As hospitals experience increases in severity, they may increase staffing. If so, then unobservable severity within a hospital is correlated with staffing and the omitted variable bias is still present.