What's the Difference? Fixed Effects and Difference-in-Difference Approaches

Charlie Gibbons
Political Science 239

February 25, 2009

Outline

1. Fixed effects

2. Social returns to education
   - Demonstration of problem
   - Assumptions revisited
   - Solution

3. Diff-in-diff problems
   - Coverage
   - Achieving correct coverage

References

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan. 2004. "How Much Should We Trust Differences-in-Differences Estimates?" *Quarterly Journal of Economics*. 119(1): 249–275.

Gibbons, Charles E., Juan Carlos Suárez Serrato, and Mike Urbancic. 2009. "Broken or Fixed Effects?" Working paper.

An easier way?

We have seen many complicated ways of dealing with omitted variables—IV, control functions, and structural modeling (for selection on unobservables) to matching and RD (selection on observables).

An easier way?

We have seen many complicated ways of dealing with omitted variables—IV, control functions, and structural modeling (for selection on unobservables) to matching and RD (selection on observables).

Isn't there an easier way?

An easier way?

We have seen many complicated ways of dealing with omitted variables—IV, control functions, and structural modeling (for selection on unobservables) to matching and RD (selection on observables).

Isn't there an easier way?

Yes—make assumptions and get more data!

Fixed effects

Fixed effects attempt to overcome an omitted variables problem in panel data. Suppose that an outcome for individual i at time t Y_{it} is a function of observed covariates X_{it}, treatment status D_{it}, some unobservable factor A_i, and possibly some unobservable time factor L_t:

$$Y_{it} = \mu(X_{it}, D_{it}, A_i, L_t).$$

Fixed effects

What do we notice?

Fixed effects

What do we notice?

A_i is independent of t and L_t is independent of i.

Fixed effects

What do we notice?

A_i is independent of t and L_t is independent of i.

This implies that the individual-level unobserved variation does not change over time (though there may be a time trend in the data). Specifically we must assume that

$$\mathbb{E}(\epsilon_{it}|D_{it}, X_{it}, A_i, L_t) = 0.$$

Without the fixed effects, our estimates would be biased since the expectation of the error would not be 0.

Fixed effects

Additionally, there can be no selection into treatment, only an unobserved, time invariant effect on the outcome.

Fixed effects

Now we introduce linearity and assume that the true model is

$$Y_{it} = \alpha + D_{it}\beta + X_{it}\delta + A_i + L_t + \epsilon_{it}$$

Since A_i and L_t are not observed, we use a set of indicator variables for each $j \in \{1, \dots N\}$, $a_i = \mathbb{I}(i = j)$ and each $s \in \{1, \dots, T\}$, $l_t = \mathbb{I}(t = s)$ and estimate coefficients on these variables $\xi_i = A_i$ and $\lambda_t = L_t$:

$$Y_{it} = \alpha + D_{it}\beta + X_{it}\delta + a_i\xi_i + l_t\lambda_t + \epsilon_{it}.$$

Fixed effects

The individual fixed effects are estimated consistently if N stays fixed while T goes to infinity, but this means that the time fixed effects are not estimated consistently. This is because the number of parameters grows with the sample size—the incidental or nuisance parameters problem. But, we don't care about these effects *per se*—*i.e.*, we only care about β, which is estimated consistently as long as the sample size is growing.

Fixed effects

Suppose that we are estimating the effect of education on wages. Can we estimate the effect of sex, race, or religion as well?

Fixed effects

Suppose that we are estimating the effect of education on wages. Can we estimate the effect of sex, race, or religion as well?

No, we cannot because the fixed effect is time invariant, so too are fixed demographic features like sex and race. We cannot separate multiple fixed effects.

Fixed effects

Notice, too what we are estimating: the effect of education conditional on a set of covariates and on an individual effect. Since each individual is different, we cannot make outcome predictions for individuals outside the sample or even for those in the sample if the fixed effects are estimated imprecisely.

Fixed effects

We can also ignore the individual fixed effects in estimation by deviating observations from their means:

$$
\begin{aligned}
\bar{Y}_i &= \alpha + \bar{D}_i\beta + \bar{X}_i\delta + \bar{a}_i\xi_i + \bar{l}\lambda_t + \bar{\epsilon}_i \\
Y_{it} - \bar{Y}_i &= \alpha + (D_{it} - \bar{D}_i)\beta + (X_{it} - \bar{X}_i)\delta + (a_i - \bar{a}_i)\xi_i + (l_t - \bar{l})\lambda_t \\
&\quad + (\epsilon_{it} - \bar{\epsilon}_i) \\
Y_{it} - \bar{Y}_i &= \alpha + (D_{it} - \bar{D}_i)\beta + (X_{it} - \bar{X}_i)\delta + (l_t - \bar{l})\lambda_t + (\epsilon_{it} - \bar{\epsilon}_i)
\end{aligned}
$$

The first equation gives the *between-group estimator*, which is like a cross-section estimator, while the final line gives the *within-group estimator*. They differ in the variation they use for estimation and the efficiency of their standard errors (see, *e.g.*, Gibbons et al. 2009).

Fixed effects

We can also difference observations by year:

$$\Delta Y_{it} = \Delta D_{it}\beta + \Delta X_{it}\delta + \Delta L_t + \Delta\epsilon_{it}$$

Typically the X covariates (*e.g.*, sex, race) are time invariant and treatment, once on, remains on for the remainder of the sample. Then we can estimate

$$\beta = \mathbb{E}\left[\Delta Y_{it}^{\Delta D=1}\right] - \mathbb{E}\left[\Delta Y_{it}^{\Delta D=0}\right].$$

This is the *difference-in-difference* estimator.

Fixed effects

Notice that this parameter is identified *only when countries change from control to treatment status* (*i.e.*, $\Delta D_{it} = 1$).

Note that this holds for fixed effects generally—the estimate is identified based upon changes in the treatment status of individuals (recall that other "fixed" effects cannot be identified, so too would be treatment if it is unchanging).

But diff-in-diff allows for more flexibility in the unobserved individual component, so long as it is time invariant.

Both approaches assume that the time trend is the same across individuals and that the individual effect is the same across time.

Fixed effects

Let's say an author is looking at the effect of education on wages in a panel data set and uses individual and time fixed effects. Suppose that he reports the results using both wage and log wage as the outcomes. Which should you believe? Is there anything wrong with this?

Fixed effects

Let's say an author is looking at the effect of education on wages in a panel data set and uses individual and time fixed effects. Suppose that he reports the results using both wage and log wage as the outcomes. Which should you believe? Is there anything wrong with this?

Remember, the identifying assumption is that trends are constant across time and individuals. If the expectations of the log wage trends are constant, then the expectations for the wage level cannot be constant.

The problem

Suppose that we are looking at the social returns to education by examining the effect of education on incarceration for individuals across state-years and are ignoring treatment (*i.e.*, years of education) selection. We have the following model:

$$p_i = \alpha + e_i\beta + a_i\xi + l_i\lambda + s_i\rho + \epsilon_i$$

where p_i is whether or not the individual is in prison, e_i is education, a_i is age fixed effects in three year increments, l_i is year fixed effects, s_i is state fixed effects.

The problem

Let's estimate β using subsets of our data based upon ages of the included individuals:

Table: Effect of education on imprisonment

	Whites	Blacks
All ages	-0.078**	-0.298**
	(0.002)	(0.012)
60 or younger	-0.095**	-0.364**
	(0.002)	(0.014)
45 or younger	-0.139**	-0.547**
	(0.004)	(0.019)
30 or younger	-0.238**	-0.917**
	(0.007)	(0.028)

The problem

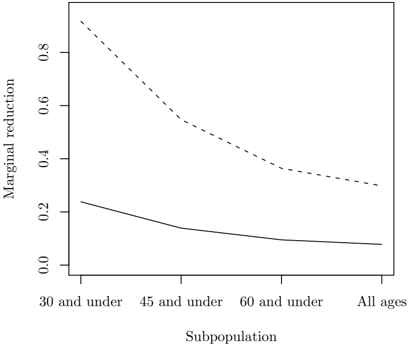

Figure: Reduction in criminal propensity due to a marginal increase in education for whites (solid) and blacks (dashed) across age subsamples

The problem

Notice that the estimate changes with the subpopulation under examination.

The problem

Notice that the estimate changes with the subpopulation under examination.

We expect fixed effects to remove or “control for” the influence of age on the effect of education but that does not happen here!

Assumptions revisited

Let's use a simpler model to look at the assumptions that we are making. Let's eliminate other covariates:

$$Y_i = \alpha + \beta e_i + \xi a_i + \epsilon_i.$$

Assumptions revisited

So what's the coefficient estimate of β?

$$\beta = \frac{\text{Cov}(e_i, Y_i)\text{Var}(a_i) - \text{Cov}(a_i, Y_i)\text{Cov}(e_i, a_i)}{\text{Var}(e_i)\text{Var}(a_i) - [\text{Cov}(e_i, a_i)]^2}.$$

Assumptions revisited

So what's the coefficient estimate of β?

$$\beta = \frac{\text{Cov}(e_i, Y_i)\text{Var}(a_i) - \text{Cov}(a_i, Y_i)\text{Cov}(e_i, a_i)}{\text{Var}(e_i)\text{Var}(a_i) - [\text{Cov}(e_i, a_i)]^2}.$$

Suppose that we add more individuals with a different age distribution. We don't want the composition of our sample to alter our estimate. So we need $\text{Cov}(e_i, A_i) = 0$.

Assumptions revisited

More generally, we need the fixed effects to be uncorrelated with the covariates. That is, the fixed effect must only be a level effect. Otherwise, there will still be an age effect embedded in our estimate.

Assumptions revisited

More generally, we need the fixed effects to be uncorrelated with the covariates. That is, the fixed effect must only be a level effect. Otherwise, there will still be an age effect embedded in our estimate.

How might we achieve this? Have an estimate of β for each fixed effect; *i.e.*, interact β with the fixed effects.

Solution

Now let's interact β and the fixed effects:

$$Y_i = \alpha + e_i\beta_1 + (e_i \times a_i)\beta_{A-1} + a_i\xi + l_i\lambda + s_i\rho + \epsilon_i$$

where A is the number of age groups.

Solution

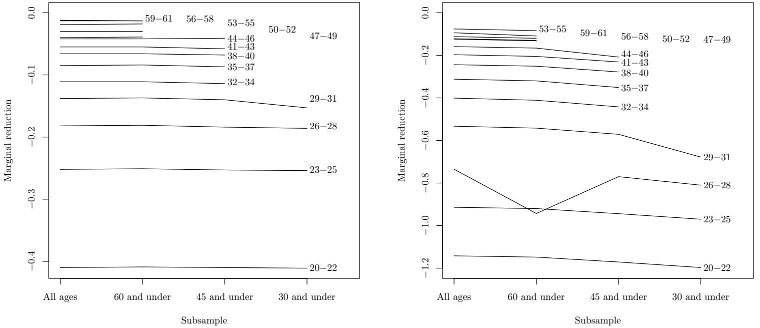

(a) Whites (b) Blacks

Figure: Marginal effect of education on crime by age across age subsamples

Solution

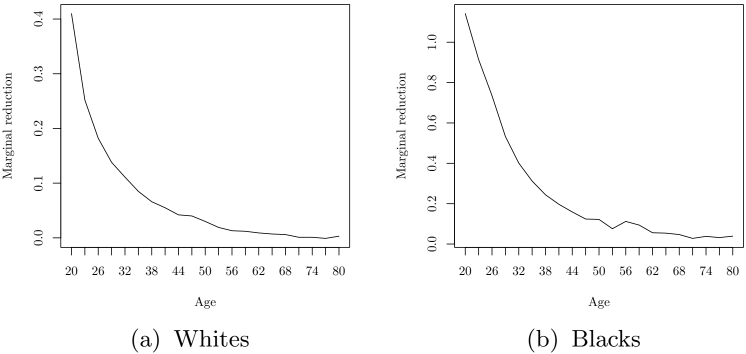

Figure: Reduction in criminal propensity due to a marginal increase in education across ages

Solution

Now the parameters are stable when we change age groups. But what should our baseline be?

Solution

Now the parameters are stable when we change age groups. But what should our baseline be?

We would like the estimates to be the same as if we ran the regression separately for each age group. Here, by definition, each estimate of β will not depend upon the underlying age distribution because each estimate is based upon a single age group. Using a single regression can reduce standard errors, hence we would like to estimate using one equation.

Solution

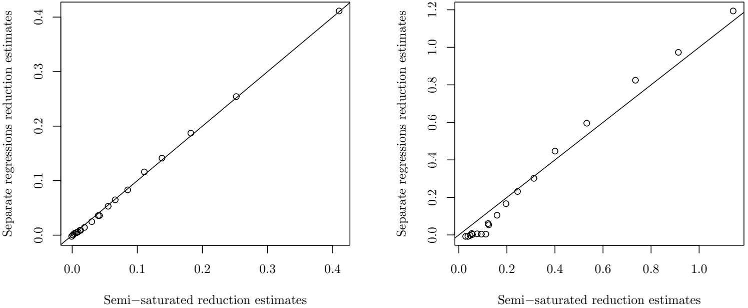

(a) Whites

(b) Blacks

Figure: Reduction in criminal propensity due to a marginal increase in education in semi-saturated and separate regressions

Solution

The semi-saturated interaction model yields almost the same estimates as separate regressions. Now our estimates are unaffected by changes in the age distribution of our subsamples.

Diff-in-diff problems

Bertrand et al. examine the standard errors given in most applications of diff-in-diff. They use a panel from the CPS and create random state laws and look for an effect of these made up laws using diff-in-diff.

Percentage of simulations that should yield a spurious, but statistically significant result if the standard errors are correct:

Diff-in-diff problems

Bertrand et al. examine the standard errors given in most applications of diff-in-diff. They use a panel from the CPS and create random state laws and look for an effect of these made up laws using diff-in-diff.

Percentage of simulations that should yield a spurious, but statistically significant result if the standard errors are correct: 5%.

Actual percentage:

Diff-in-diff problems

Bertrand et al. examine the standard errors given in most applications of diff-in-diff. They use a panel from the CPS and create random state laws and look for an effect of these made up laws using diff-in-diff.

Percentage of simulations that should yield a spurious, but statistically significant result if the standard errors are correct: 5%.

Actual percentage: 67.5%.

Diff-in-diff problems

Why?

Diff-in-diff problems

Why?

There is strong serial correlation in the errors—*e.g.*,

$$\mathbb{E}[\Delta\epsilon_{it}] = \mathbb{E}[\epsilon_{it} - \epsilon_{it-1}] \not\perp \mathbb{E}[\epsilon_{it-1} - \epsilon_{it-2}] = \mathbb{E}[\Delta\epsilon_{it-1}]$$

Achieving correct coverage

How do we fix this?

Achieving correct coverage

How do we fix this?

Parametric modeling of the autocorrelation is not particularly effective.

Achieving correct coverage

The best way is via a *block bootstrap*. Suppose you are observing N countries over time. To do a studentized block bootstrap:

1. Calculate the full sample naive OLS t-statistic: $t = \frac{\hat{\beta}}{\hat{\sigma}_{\beta}^{0.5}}$.
2. Draw N country mini-data sets (outcomes, treatment, fixed effects, and covariates for all observations across time for a given country) with replacement.
3. Run OLS and calculate β_b and $t_b = \frac{\beta_b - \hat{\beta}}{\hat{\sigma}_{b}^{0.5}}$.
4. Reject the null hypothesis of $\beta = 0$ at the 5% level if 95% of the $|t_b|$ are less than $|t|$.

Achieving correct coverage

Other solutions:

- Ignore the time series aspect by taking averages before and after treatment and creating a two-period panel.
- Assume no cross-sectional heteroskedasticity and a common autocorrelation process in all countries. Then allow for intra-state correlations via clusering in the covariance matrix.
- Allow for cross-sectional heteroskedasticity and varying autocorrelation processes using a “White-like” standard error formula.

The latter two do poorly in small N samples.