

Differences in Differences (DD)

Empirical Methods

Prof. Michael R. Roberts

Topic Overview

- Introduction
 - » Intuition and examples
 - » Experiments
 - » Single Difference Estimators
- DD
 - » What is it
 - » Identifying Assumptions
 - Graphical & Statistical Analysis
 - » Sensitivity Tests
 - » Potential Concerns & Limitations
 - » Extensions
- Alternative Perspective & Estimation Strategies
- References

DD Intuition

- DD is a **quasi-experimental** technique used to understand the effect of a sharp change in the economic environment or government policy.
 - » Examples
 - Card (1990) uses the Mariel Boatlift, which increased the Miami labor force by 7% between May and September of 1980, to understand the consequences of immigration on non-immigrant wages
 - Butler and Cornaggia (2008) use ethanol mandates from the EPA of 2005, which require the increased use of corn in fuel, to understand the effect of access to finance on productivity of farmers
- Used in conjunction with a **natural experiment** in which nature does the randomization for us
 - » Key: transparent exogenous source of variation that determine treatment assignment (e.g., policy changes, government randomization, etc.)

A Hypothetical Example

- Question: What is the effect of a decline in expected bankruptcy costs on corporate debt usage?
 - » Tax-bankruptcy cost theories of capital structure predict that debt usage and expected bankruptcy costs are inversely related
- Ideal but Impossible Experiment:
 - » Take a set of firms, reduce bankruptcy costs (e.g., streamline bankruptcy procedures) and measure debt usage
 - » “Rewind the clock,” take the same set of firms and measure their debt usage.
 - » Compare debt usage across two scenarios
- Desirable but Infeasible Experiment
 - » Take a set of firms and randomly select some fraction of firms to be subject to the new bankruptcy procedures.
 - » Compare debt usage across the two sets of firms

What was Good about these Experiments?

- The key to program evaluation is estimating the counterfactual: What would have happened had the treated not be treated?
- Therefore, quality of our evaluation is tied to how well we can estimate the counterfactual
 - » The Ideal but Impossible Experiment actually provides the counterfactual by “rewinding the clock.”
 - » The Desirable but Infeasible Experiment provides a good estimate of the counterfactual by the random assignment
 - There are no systematic differences between the treated and untreated groups that are related to the outcome of interest
 - Without random assignment (i.e., self-selection or manipulation) then we can't be sure if differences (or similarities) between the treated and untreated are due to the program or some other difference between the two groups.

The Natural Experiment

- At the end of 1991, Delaware passes a law that significantly streamlines bankruptcy proceedings to make litigation less costly and time-consuming.
 - » Must assume that this is a *random* (or outcome unrelated) event.
 - The event should not be a response to pre-existing differences between the treatment and control group (e.g., Delaware firms are much more likely to enter financial distress)
 - E.g., “Ashenfelter Dip” (Ashenfelter and Card (1985))
 - » Must understand what *caused* the event to occur
- This law change offers a potentially useful setting with which to test our hypothesized relation between bankruptcy costs and capital structure
 - » How do we empirically test the relation?

Cross-Sectional Difference After Treatment

- Let's compare the average leverage of firms registered in Delaware to that of firms registered elsewhere
 - » This can be accomplished via a cross-sectional regression

$$y_i = \beta_0 + \beta_1 I(\text{treat}_i) + \varepsilon_i$$

- » where y_i = Leverage for firm i in 1992, and $I(\text{treat}_i) = 1$ if firm is registered in Delaware
- Assuming $E(\varepsilon_i | I(\text{treat}_i)) = 0$:

$$E(y_i | I(\text{treat}_i) = 0) = \beta_0$$

$$E(y_i | I(\text{treat}_i) = 1) = \beta_0 + \beta_1$$

$$\Rightarrow E(y_i | I(\text{treat}_i) = 1) - E(y_i | I(\text{treat}_i) = 0) = \beta_1$$

- Our estimate is just the difference in average leverage in 1992 for the treatment group (Delaware firms) and control group (non-Delaware firms)

Cross-Sectional Difference After Treatment

Potential Concerns

- As usual, the concern lies with our assumption of $E(\varepsilon_i | I(treat_i)) = 0$
- What could threaten this assumption and, consequentially, the internal validity of our estimate?
 - » What if firms in Delaware are in more capital intensive industries relative to firms elsewhere?
 - Problem is that firms with more physical capital tend to be more levered so our assumption is violated because capital intensity is sitting in ε and it's correlated with treatment status
 - In other words, even if the law was never passed, we would expect firms in Delaware to have higher leverage than other firms because of differences in capital intensity

Cross-Sectional Difference After Treatment

Omitted Variables Bias

- This is just an omitted variables bias where the true relation is:

$$y_i = \beta_0 + \beta_1 I(\text{treat}_i) + \beta_2 CI_i + \eta_i$$

but we estimate

$$y_i = \beta_0 + \beta_1 I(\text{treat}_i) + \varepsilon_i$$

- Implies the OLS estimator of β_1 is:

$$\begin{aligned}\hat{\beta}_1 &= \frac{\text{Cov}(I(\text{treat}), y)}{\text{Var}(I(\text{treat}))} = \frac{\text{Cov}(I(\text{treat}), \beta_0 + \beta_1 I(\text{treat}) + \beta_2 CI + \eta)}{\text{Var}(I(\text{treat}))} \\ &= \beta_1 + \beta_2 \frac{\text{Cov}(I(\text{treat}), CI)}{\text{Var}(I(\text{treat}))} = \beta_1 + \beta_2 \gamma\end{aligned}$$

where γ is the slope coefficient from a regression of capital intensity CI on $I(\text{treat})$

- $\gamma > 0, \beta_2 > 0 \rightarrow$ OLS estimate biased up in favor of a treatment effect...not good.

Cross-Sectional Difference After Treatment

What to Do?

- One Solution:
 - » Insert a proxy for capital intensity (e.g., net plant, property, and equipment divided by assets) in the regression. But,
 - this is just a proxy and just one variable
 - » There is still some heterogeneity between treatment and control groups that is sitting in the error term and is correlated with our treatment indicator
 - A question of magnitude but we can't estimate this since error unobservable

Cross-Sectional Difference After Treatment

Direction of Bias

- Bias does not always have to work in favor of finding a result. For example,
 - » Imagine instead that more financially conservative CEOs tend to run firms in Delaware (perhaps because of an industry bias)
 - » Just replace *CI* on slide 9 with a measure of CEO conservatism (good luck measuring this thing) and we can see:
 - $\beta_2 < 0$ since more conservative CEOs use less debt
 - $\gamma > 0$ since more conservative CEOs are concentrated in Delaware
 - OLS estimate is biased down against finding a treatment effect
 - Therefore, any estimate is a lower bound

Cross-Sectional Difference After Treatment

Summing it Up

- What's common across these examples of threats to internal validity is that they correspond to unmeasured differences between the two groups
 - » Imagine a latent fixed effect across the two groups
 - This results from non-random assignment to treatment and control
 - » This will confound any comparison between the two groups

Time-Series Difference Within the Treatment Group

- Let's compare the average leverage of firms registered in Delaware in 1991 to that in 1992...avoids heterogeneous firm concern
 - » This can be accomplished via a two-period panel regression using *only* Delaware firms

$$y_{it} = \beta_0 + \beta_1 I(Post_{it}) + \varepsilon_{it}$$

where $I(Post) = 1$ if year = 1992 and 0 if year = 1991.

- Assuming $E(\varepsilon_{it} | I(Post_{it})) = 0$:

$$E(y_{it} | I(Post_{it}) = 0) = \beta_0$$

$$E(y_{it} | I(Post_{it}) = 1) = \beta_0 + \beta_1$$

$$\Rightarrow E(y_{it} | I(Post_{it}) = 1) - E(y_{it} | I(Post_{it}) = 0) = \beta_1$$

- Our estimate is just the difference in average leverage for Delaware firms in 1992 (the post-treatment era) and 1991 (the pre-treatment era)
 - » Consider the cross-sectional, first-difference regression

$$\Delta y_{it} = \beta_1 + \Delta \varepsilon_{it}$$

Time-Series Difference Within the Treatment Group

- Let's compare the average leverage of firms registered in Delaware in 1991 to that in 1992
 - » This can be accomplished via a two-period panel regression using *only* Delaware firms

$$y_{it} = \beta_0 + \beta_1 I(Post_{it}) + \varepsilon_{it}$$

where $I(Post) = 1$ if year = 1992 and 0 if year = 1991.

- » or, after differencing, a cross-sectional regression using *only* Delaware firms

$$\Delta y_i = \beta_1 + \Delta \varepsilon_i$$

Assuming $E(\varepsilon_{it} | I(Post_{it})) = 0$

$$E(y_{it} | I(Post_{it}) = 0) = \beta_0$$

$$E(y_{it} | I(Post_{it}) = 1) = \beta_0 + \beta_1$$

$$\Rightarrow E(y_{it} | I(Post_{it}) = 1) - E(y_{it} | I(Post_{it}) = 0) = \beta_1$$

Time-Series Difference Within the Treatment Group – Potential Concerns

- The primary concern here is other factors affecting leverage *over time*.
 - » E.g., Increase in the supply of credit due to financial innovation
 - » Leverage would likely have increased for firms even without the passage of the law
 - » Bias is positive → in favor of supporting a treatment effect
- This is just another form of omitted variables bias, as before
 - » We could insert control variables but difficult to measure all perfectly
- Also, bias can work in both directions
 - » E.g., 1992 may be a period of tight credit, which leads to a decline in debt usage and offsetting effect

DD Estimator

- It would be nice to combine the positive features of each single difference estimator
 - » Cross-sectional estimator avoided omitted common trends
 - » Time-series estimator avoided omitted cross-sectional differences
- The DD estimator does precisely that

$$y_{it} = \beta_0 + \beta_1 I(treat_{it}) + \beta_2 I(post_{it}) + \beta_3 I(treat_{it}) \times I(post_{it}) + \varepsilon_{it}$$

- » We need a full panel of firms consisting of Delaware ($I(treat) = 1$) and Non-Delaware ($I(treat) = 0$) registered firms observed before ($I(post) = 0$) and after ($I(post) = 1$) the passage of the law
- » β_3 is the parameter of interest (i.e., the DD estimator)

How is β_3 the DD?

- Compute the conditional expectations

$$E(y_{it} | I(treat_{it}) = 1, I(Post_{it}) = 1) = \beta_0 + \beta_1 + \beta_2 + \beta_3$$

$$E(y_{it} | I(treat_{it}) = 1, I(Post_{it}) = 0) = \beta_0 + \beta_1$$

$$E(y_{it} | I(treat_{it}) = 0, I(Post_{it}) = 1) = \beta_0 + \beta_2$$

$$E(y_{it} | I(treat_{it}) = 0, I(Post_{it}) = 0) = \beta_0$$

- Take difference over time in average leverage for control group and subtract from difference over time in average leverage for treatment group

$$\left\{ \begin{array}{l} E(y_{it} | I(treat_{it}) = 1, I(Post_{it}) = 1) - E(y_{it} | I(treat_{it}) = 1, I(Post_{it}) = 0) \\ \text{Average Leverage for Delaware firms in 1992 } (\beta_0 + \beta_1 + \beta_2 + \beta_3) \quad \text{Average Leverage for Delaware firms in 1991 } (\beta_0 + \beta_1) \end{array} \right\} -$$

Difference in Average Leverage for Delaware Firms ($\beta_2 + \beta_3$)

$$\left\{ \begin{array}{l} E(y_{it} | I(treat_{it}) = 0, I(Post_{it}) = 1) - E(y_{it} | I(treat_{it}) = 0, I(Post_{it}) = 0) \\ \text{Average Leverage for Non-Delaware firms in 1992 } (\beta_0 + \beta_2) \quad \text{Average Leverage for Non-Delaware firms in 1991 } (\beta_0) \end{array} \right\} = \beta_3$$

Difference in Average Leverage for Non-Delaware Firms (β_2)

Conditional Mean Estimates

- Easier to see what the parameters estimate by arranging things in a 2 x 2 matrix:

	<i>Post</i>	<i>Pre</i>	<i>Difference</i>
<i>Treatment</i>	$\beta_0 + \beta_1 + \beta_2 + \beta_3$	$\beta_0 + \beta_1$	$\beta_2 + \beta_3$
<i>Control</i>	$\beta_0 + \beta_2$	β_0	β_2
<i>Difference</i>	$\beta_1 + \beta_3$	β_1	β_3

- Note:
 - » Implies alternative route to DD estimator: subtract difference in average leverage for treatment and control in 1991 from the difference in average leverage for treatment and control in 1992
 - » β_1 is the average leverage of treatment firms in the pre-treatment era *relative to* (i.e. minus) the average leverage of control firms in the pre-treatment era
 - » β_2 is the average leverage of control firms in the post-treatment era *relative to* (i.e. minus) the average leverage of control firms in the pre-treatment era

Re-Visiting the Single Difference Estimators

- Cross-sectional difference:

$$E(y | I(\text{treat}) = 1, I(\text{post}) = 1) - E(y | I(\text{treat}) = 0, I(\text{post}) = 1) = \beta_1 + \beta_3$$

- » Unbiased if β_1 is equal to zero \rightarrow no permanent difference between the treatment and control groups

- Time-series difference:

$$E(y | I(\text{treat}) = 1, I(\text{post}) = 1) - E(y | I(\text{treat}) = 1, I(\text{post}) = 0) = \beta_2 + \beta_3$$

- » Unbiased if β_2 is equal to zero \rightarrow no common trend over the pre- and post-treatment eras

- DD estimator avoids these two threats by differencing away any permanent differences between the groups *and* any common trend affecting both groups

Key Assumption Behind DD

- In the absence of treatment, the average change in the response variable would have been the same for both the treatment and control groups.
 - » a.k.a. the **parallel trends** assumption since it requires that the trend in the outcome variable for both treatment and control groups during the pre-treatment era are similar.
- What this assumptions does *not* mean
 - » It does *not* mean that the there is *no* trend in the outcome variable during the pre-treatment era (just the same trend across groups)
 - » It does *not* require that the *level* of the outcome variable for the two groups be same in the pre-treatment era
 - This can create other issues concerning function form...more on this later

Testing the Key Assumption Behind DD

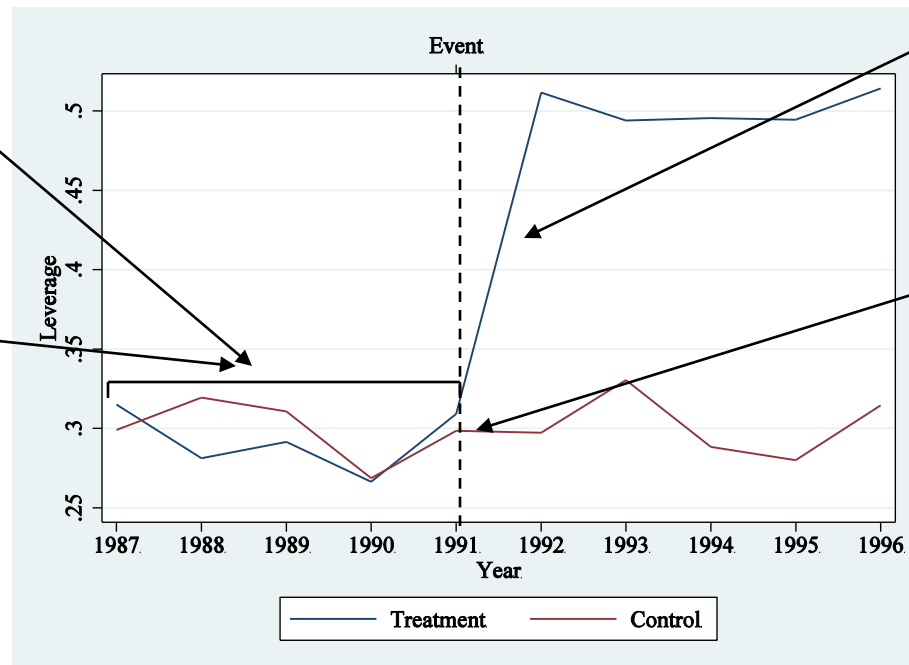
- Strictly speaking, untestable
 - » It's just saying that $E(\varepsilon | I(\text{treat}) \times I(\text{post})) = 0$ and we don't observe ε
- But, we can inspect pre-treatment era growth rates in outcome variables.
 - » By definition, this requires more than two periods of data

Graphical Analysis

- A “nice” DD picture of our bankruptcy law experiment

Similar leverage growth (~0) in pre-treatment era

Bonus: Similar leverage level (~30%) in pre-treatment era



Sharp change in leverage for treatment group *right after* event

No *change* in leverage after event for control group

Statistical Analysis

- We could also perform a t-test of the difference in average leverage growth rates across the treatment and control groups during the pre-treatment era.
 - » This should turn out insignificant (statistically *and* economically) if the parallel trends assumption is valid
 - Statistical insignificance is *not* enough...could be low power
- You should be skeptical of any DD that doesn't show you a picture or an explicit test of this assumption.
 - » Of course, this requires sufficient data on pre-treatment outcomes

Level Regression Estimates

- Remember:
 - » β_1 = Avg. leverage of treatment group in pre-treatment era less avg. leverage of control group in pre-treatment era
 - » β_2 = Avg. leverage of control group in post-treatment era less avg. leverage of control group in pre-treatment era
 - » β_3 = DD estimate
 - » β_0 = Avg. leverage of control group in the pre-treatment era
- Note:
 - » Single difference estimators will give you basically same result since $\beta_1 = \beta_2 = 0$ (i.e., no time-invariant differences between groups or common trends across groups)
 - » Back out conditional means via linear combinations of parameters (e.g., avg. leverage of treatment group in pre-treatment era = $\beta_0 + \beta_1 = 29\%$)
 - » Reconcile estimates with figure

$\beta_1: I(treat)$	-0.007 (-0.673)
$\beta_2: I(post)$	0.003 (0.296)
$\beta_3: I(treat) \times I(post)$	0.206** (14.939)
$\beta_0: Intercept$	0.299** (43.403)
Observations	5,000
Fraction Treatment	0.498
Adjusted R ²	0.115

Sensitivity Tests I

- Redo the DD analysis on pre-event years
 - » The estimated treatment effect should be zero
- Multiple control groups
 - » E.g., define control groups to be firms in adjacent states
- For the these previous two sensitivity tests, we could do a triple diff: Difference-in-difference-in-differences
 - » DD-in-dif estimated difference between DD of interest and the other DD. But,
 - If the other DD is zero, doing a triple dif will only work to inflate the standard errors.
 - If the other DD is *not* zero, you have to wonder whether the original DD is unbiased
- Compare treatment and control groups along various covariates to see if they are similar, at least along *observable* dimensions
 - » Can also include (time-varying) covariates in the regression specification
 - My want to include higher order polynomial terms to relax linearity assumption

Sensitivity Tests II

- Make sure change is concentrated around the event
 - » Identification is coming from exogenous event
 - Moving away from the event allows other factors to creep in
- Make sure other outcome variables that should be unaffected by the event are actually unaffected
- Control by systematic variation
 - » See if treatment and control groups respond to other factors similarly (Rosenbaum (1987))
- Multiple treatment groups
 - » E.g., If Delaware implemented the changes in three stages, one could define separate control groups corresponding to each stage
 - » E.g., could separate out levered from unlevered firms in order to focus on firms most likely to be affected by the change.
- Treatment Reversal
 - » If the new law is repealed, do we see a reversal of the effect? (e.g., Leary (2008))
- Use robust (to heteroskedasticity *and*, more importantly, dependence) SEs
 - » Bertrand, Duflo, and Mulainathan (2004), Petersen (2007), Donald and Lang (2007)

Potential Concerns

- Heterogeneous Responses

- » Ideally control group shows no response to the treatment
- » Otherwise we need different responses by the treatment and control groups
 - But, if intensity of responses is different, estimated effect may be meaningless. E.g., Imagine the following model

$$y_{it} = \mu_i + \alpha_t + \beta_i x_{it}$$

where x_{it} increases from 0 to x^T for treatment subjects and 0 to x^C for control subjects. The DD estimate is:

$$\beta = [y_1^T - y_0^T] [y_1^C - y_0^C] = \beta_T x^T - \beta_C x^C$$

If

$$\beta_C = 2\beta_T \text{ and } x^T = 2x^C$$

then the DD estimate is zero, despite a potentially large effect. (See Feldstein (1995))

Additional Concerns and Limitations

- Fixed effects make sense only if response to event is immediate.
 - » If delayed than the model may require lagged dependent variable
 - Blundell and Bond (1998) GMM is one way to estimate this **dynamic linear panel data model** appropriately
- As with all natural experiments and quasi-experimental methods, extrapolation is tricky
 - » Take care in extrapolating the results
 - » Internal validity is often inversely related to external validity

Extensions

Fixed Effects

- What if we have more than two groups (treatment and control) and more than two periods (Pre and Post)?
 - » Use group (α) and time-period (γ) fixed effects:
$$y_{igt} = \alpha_g + \gamma_t + \beta I_{gt} + \delta X_{igt} + \varepsilon_{igt}$$
 - » Identification of β comes from group-specific changes over time.
 - Take care to account-for within group dependence in SEs (Bertrand, Duflo, and Mulainathan (2004), Petersen (2007), Donald and Lang (2007))
 - Visually inspect the effect and parallel trends assumption for a few treatment groups if possible
 - » What would this general specification look like in our hypothetical corporate finance setting if we have data from 1987 to 1996?
 - 9 Year dummies
 - 1 treatment indicator
 - 5 interaction terms (treatment indicator x 5 post-treatment yearly dummies)
 - Can capture heterogeneity in the effect over time

Matching Perspective

- The DD estimator is simply:

$$E(y_{ia} - y_{ib} \mid I(\text{treat}) = 1) - E(y_{ia} - y_{ib} \mid I(\text{treat}) = 0)$$

where a = post-treatment, b = pre-treatment

- So, we could match the treatment subjects and the control subjects based on *pre-treatment* characteristics
- Estimate the conditional expectations using sample analogues (i.e., conditional averages)
- Compute the appropriate SEs and voila!
 - » E.g., Lemmon and Roberts (2007) for a simple example and Heckman, Ichimura, and Todd (1997) for a more complex example

Summary

- DD can be a simple but powerful empirical strategy
- Relies crucially on:
 - » exogeneity *and* sharpness of the event/treatment
 - » Comparability of the treatment and control groups
 - Parallel trends assumption
- As with many empirical papers, a picture is worth a thousand words...
- Like other quasi-experimental methods, should be accompanied by a battery of robustness tests ensuring internal validity

References I

- Ashenfelter, Orley, and David Card, 1985, Using the Longitudinal Structure of Earnings to Estimate the Effect of Training Programs, *Review of Economics and Statistics* 68, 648-660
- Bertrand, Marianne, Esther Duflo, Sendhil Mullainathan, How Much Should We Trust Difference-in-Differences Estimates? *Quarterly Journal of Economics* , 249-275
- Blundell, Richard, and Stephen Bond, 1998, Initial Conditions and Moment Restrictions in Dynamic Panel Data Models, *Journal of Econometrics* 87, 115-143
- Butler, Alexander, and Jess Cornaggia, 2008, Does Access to Finance Improve Productivity? Evidence from a Natural Experiment, *Working Paper*, University of Texas
- Card, David, 1990, The Impact of the Mariel Boatlift on the Miami Labor Market, *Industrial and Labor Relations Review* 43, 245-257
- Donald, S., and Kevin Lang, 2007, Inference with Difference-in-Differences and Other Panel Data, *Review of Economics and Statistics* 89, 221-233
- Duflo, Esther, *Empirical Methods Handout*

References II

- Gruber, Jonathan, The Incidence of Mandated Maternity Benefits, *American Economic Review* 84, 622-641
- Heckman, James, H. Ichimura, and Petra Todd, 1997, Matching as an Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Program, *Review of Economic Studies* 64, 605-654
- Leary, Mark T., 2008, Bank Loan Supply, Lender Choice, and Corporate Capital Structure, forthcoming *Journal of Finance*
- Lemmon, Michael, and Michael R. Roberts, 2007, The Response of Corporate Financing and Investment to Changes in the Supply of Credit, *Working Paper*, University of Pennsylvania
- Meyer, Bruce D., 1995, Natural and Quasi-Experiments in Economics, *Journal of Business and Economic Statistics* 13, 151-161
- Petersen, Mitchell, 2007, Estimating Standard Errors in Finance Panel Datasets: Comparing Approaches, forthcoming *Review of Financial Studies*
- Rosenbaum, Paul, 1987, Observational Studies, in *Series in Statistics*, Heidelberg, New York.