

Lecture 15: Difference-in-differences

Economics 326 — Introduction to Econometrics II

Vadim Marmer, UBC

Motivation

- In the previous lecture, we estimated treatment effects from **cross-sectional** data under the assumption of selection on observables: after controlling for covariates, treatment is as good as random.
- In many settings, this assumption is hard to justify. An alternative approach exploits **panel data** (repeated observations on the same units over time).
- The **difference-in-differences (DID)** method compares changes over time between a treatment group and a control group.

DID setup

- Two time periods: $t \in \{0, 1\}$ (before and after treatment).
- Two groups: $D_i \in \{0, 1\}$ (control and treatment).
- Treatment occurs between periods 0 and 1, and only the treatment group ($D_i = 1$) is affected.
- We observe Y_{it} : the outcome for individual i at time t .

DID regression model

- The DID regression is:

$$Y_{it} = \alpha + \delta \cdot t + \gamma D_i + \beta(t \cdot D_i) + U_{it},$$

where $E[U_{it} | D_i] = 0$.

- The regressors:
 - t : time indicator (0 = before, 1 = after).
 - D_i : group indicator (0 = control, 1 = treatment).
 - $t \cdot D_i$: interaction term, equals 1 only for the treatment group after treatment.

Interpreting the coefficients

•

$$Y_{it} = \alpha + \delta \cdot t + \gamma D_i + \beta(t \cdot D_i) + U_{it},$$

- Evaluate $E[Y_{it} | D_i]$ for each combination of t and D_i :

$$t = 0, D_i = 0: E[Y_{i0} | D_i = 0] = \alpha,$$

$$t = 0, D_i = 1: E[Y_{i0} | D_i = 1] = \alpha + \gamma,$$

$$t = 1, D_i = 0: E[Y_{i1} | D_i = 0] = \alpha + \delta,$$

$$t = 1, D_i = 1: E[Y_{i1} | D_i = 1] = \alpha + \delta + \gamma + \beta.$$

- Summarized as a 2x2 table:

	$D_i = 0$ (Control)	$D_i = 1$ (Treatment)
$t = 0$	α	$\alpha + \gamma$
$t = 1$	$\alpha + \delta$	$\alpha + \delta + \gamma + \beta$

- α : baseline expected outcome (control group, before treatment).
- γ : pre-existing **group difference** at baseline ($t = 0$).
- δ : **time effect** — the change in the control group from $t = 0$ to $t = 1$, capturing common trends.
- From the 2x2 table, the change over time for each group:
 - Treatment:** $E[Y_{i1} | D_i = 1] - E[Y_{i0} | D_i = 1] = \delta + \beta$.
 - Control:** $E[Y_{i1} | D_i = 0] - E[Y_{i0} | D_i = 0] = \delta$.
- Subtracting the control group's change from the treatment group's change:

$$(\delta + \beta) - \delta = \beta.$$

- The DID estimand as a double difference of conditional expectations:

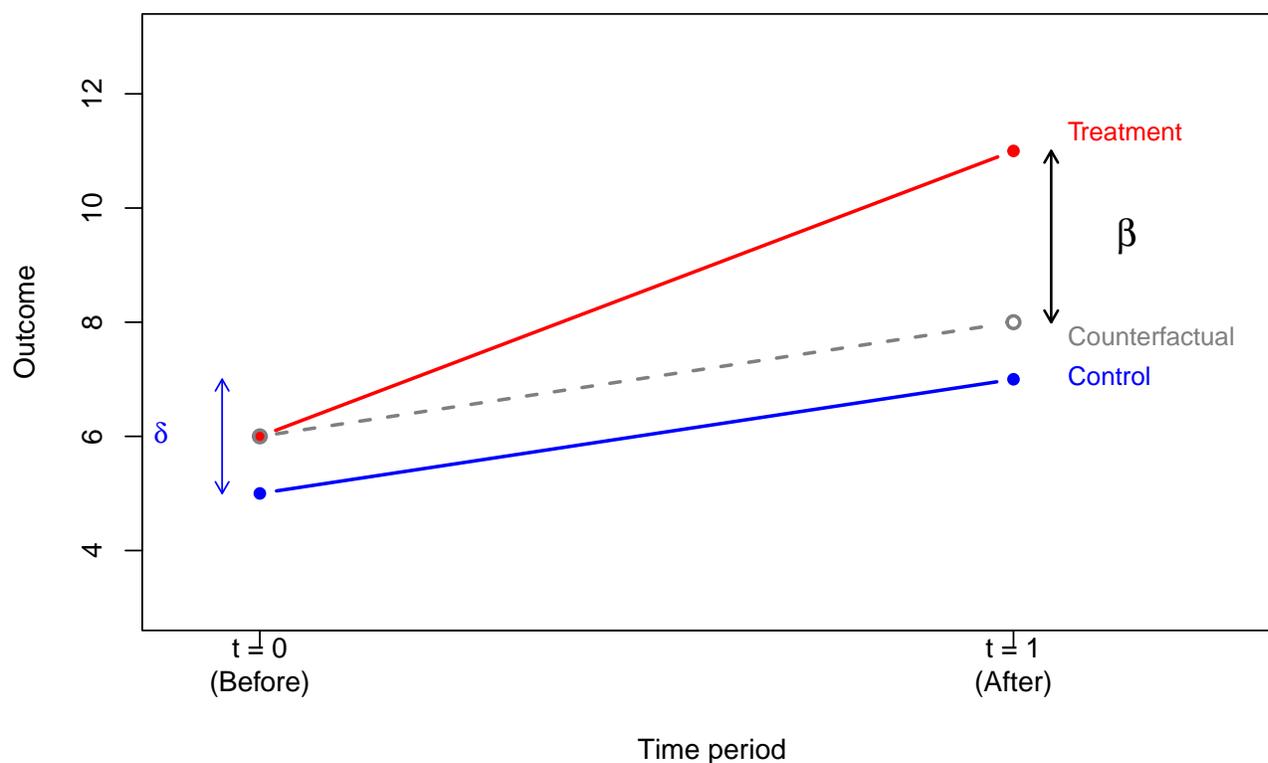
$$\beta = E[Y_{i1} - Y_{i0} | D_i = 1] - E[Y_{i1} - Y_{i0} | D_i = 0].$$

- By subtracting the control group's change, the common time trend δ cancels, isolating the treatment effect β .

DID diagram

- The classic DID diagram shows the control group, treatment group, and counterfactual:

Difference-in-differences



- We predict the counterfactual outcome for the treatment group at $t = 1$ (dashed gray line) by adding the control group's change δ to the treatment group's baseline $\alpha + \gamma$.

DID and potential outcomes

- To connect DID with the potential outcomes framework, define **panel potential outcomes**: $Y_{it}(d)$ is the outcome for individual i at time t if assigned to group $d \in \{0, 1\}$.
- The observed outcome is:

$$Y_{it} = D_i Y_{it}(1) + (1 - D_i) Y_{it}(0).$$

- What we observe for each group:

	Control ($D_i = 0$)	Treatment ($D_i = 1$)
$t = 0$	$Y_{i0}(0)$	$Y_{i0}(1)$
$t = 1$	$Y_{i1}(0)$	$Y_{i1}(1)$

- The treatment effect at time $t = 1$ for the treated group is:

$$\text{ATT} = \text{E}[Y_{i1}(1) - Y_{i1}(0) \mid D_i = 1].$$

The counterfactual $Y_{i1}(0)$ is unobserved for the treated group.

DID as a treatment effect

- $\beta = \text{E}[Y_{i1} - Y_{i0} \mid D_i = 1] - \text{E}[Y_{i1} - Y_{i0} \mid D_i = 0]$.
- Substituting observed outcomes with potential outcomes:

$$\beta = \text{E}[Y_{i1}(1) - Y_{i0}(1) \mid D_i = 1] - \text{E}[Y_{i1}(0) - Y_{i0}(0) \mid D_i = 0].$$

- To relate β to the ATT, add and subtract $\text{E}[Y_{i1}(0) \mid D_i = 1]$ and $\text{E}[Y_{i0}(0) \mid D_i = 1]$ inside the first expectation:

$$\begin{aligned} \beta &= \text{E}[Y_{i1}(1) - Y_{i0}(1) \mid D_i = 1] - \text{E}[Y_{i1}(0) - Y_{i0}(0) \mid D_i = 0] \\ &= \text{E} \left[Y_{i1}(1) - \underbrace{Y_{i1}(0) + Y_{i1}(0)}_{=0} - \underbrace{Y_{i0}(0) + Y_{i0}(0)}_{=0} - Y_{i0}(1) \mid D_i = 1 \right] \\ &\quad - \text{E}[Y_{i1}(0) - Y_{i0}(0) \mid D_i = 0]. \end{aligned}$$

- Rearranging and splitting the expectation:

$$\begin{aligned} \beta &= \underbrace{\text{E}[Y_{i1}(1) - Y_{i1}(0) \mid D_i = 1]}_{\text{ATT}} \\ &\quad + \underbrace{\text{E}[Y_{i1}(0) - Y_{i0}(0) \mid D_i = 1] - \text{E}[Y_{i1}(0) - Y_{i0}(0) \mid D_i = 0]}_{\text{difference in trends}} \\ &\quad + \underbrace{\text{E}[Y_{i0}(0) - Y_{i0}(1) \mid D_i = 1]}_{\text{anticipation effect}}. \end{aligned}$$

- For β to equal the ATT, the last two terms must be zero. This requires two assumptions.

Assumption 1: Parallel trends

- The “difference in trends” term equals zero if both groups would have experienced the same change over time in the absence of treatment:

$$E[Y_{i1}(0) - Y_{i0}(0) | D_i = 1] = E[Y_{i1}(0) - Y_{i0}(0) | D_i = 0].$$

- Under parallel trends, the decomposition reduces to:

$$\beta = \text{ATT} + \underbrace{E[Y_{i0}(0) - Y_{i0}(1) | D_i = 1]}_{\text{anticipation effect}}.$$

- The parallel trends assumption cannot be directly tested because $Y_{i1}(0)$ is unobserved for the treated group.
- If pre-treatment data for multiple periods exist, one can check whether trends were parallel before treatment.

Assumption 2: No anticipation

- The “anticipation effect” term equals zero if the outcome at $t = 0$ (before treatment) is not affected by future treatment assignment:

$$E[Y_{i0}(1) | D_i = 1] = E[Y_{i0}(0) | D_i = 1].$$

- Being assigned to the treatment group does not change pre-treatment outcomes in expectation.
- Under both parallel trends and no anticipation:

$$\beta = \text{ATT}.$$

Example: incinerator and house prices

- Kiel and McClain (1995) studied how the construction of a garbage incinerator affected nearby house prices in North Andover, Massachusetts. This is an example of an **event study**: a research design that estimates the causal effect of a specific event by comparing outcomes before and after it occurs.
- We use the `kielmc` dataset from the `wooldridge` package:

```
library(wooldridge)
data(kielmc)
# Show 2 observations from each of the 4 groups
rows <- c(
  head(which(kielmc$y81 == 0 & kielmc$nearinc == 0), 2),
  head(which(kielmc$y81 == 0 & kielmc$nearinc == 1), 2),
  head(which(kielmc$y81 == 1 & kielmc$nearinc == 0), 2),
  head(which(kielmc$y81 == 1 & kielmc$nearinc == 1), 2)
)
kielmc[rows, c("rprice", "y81", "nearinc", "age")]
```

	rprice	y81	nearinc	age
14	52000.00	0	0	32
15	49000.00	0	0	18
1	60000.00	0	1	48
2	40000.00	0	1	83
187	90245.77	1	0	1

```
188 46082.95 1 0 41
180 37634.41 1 1 81
181 39938.55 1 1 71
```

- `rprice`: house price in 1978 dollars (Y_{it}).
- `y81`: 1 if year is 1981 (after incinerator announced), 0 if 1978 ($t = 1$ if 1981, 0 if 1978).
- `nearinc`: 1 if house is near the incinerator site (D_i).

The 2×2 table of means

- Compute the four group means:

```
means <- tapply(kielmc$rprice, list(kielmc$y81, kielmc$nearinc), mean)
colnames(means) <- c("Far (nearinc=0)", "Near (nearinc=1)")
rownames(means) <- c("1978 (y81=0)", "1981 (y81=1)")
round(means, 2)
```

```
              Far (nearinc=0) Near (nearinc=1)
1978 (y81=0)      82517.23      63692.86
1981 (y81=1)     101307.51      70619.24
```

- Computing the DID by hand:

```
diff_near <- means[2, 2] - means[1, 2]
diff_far  <- means[2, 1] - means[1, 1]
DID <- diff_near - diff_far
cat("Change (near):", round(diff_near, 2), "\n")
```

```
Change (near): 6926.38
```

```
cat("Change (far): ", round(diff_far, 2), "\n")
```

```
Change (far): 18790.29
```

```
cat("DID:          ", round(DID, 2), "\n")
```

```
DID:          -11863.9
```

DID regression

- The DID regression, where $y81nrinc = y81 \times nearinc$ is the interaction term:

```
options(scipen = 999)
reg_did <- lm(rprice ~ y81 + nearinc + y81nrinc, data = kielmc)
round(summary(reg_did)$coefficients, 4)
```

```
              Estimate Std. Error t value Pr(>|t|)
(Intercept)  82517.23    2726.910  30.2603  0.0000
y81          18790.29    4050.065   4.6395  0.0000
nearinc      -18824.37    4875.322  -3.8612  0.0001
y81nrinc     -11863.90    7456.646  -1.5911  0.1126
```

- The coefficient on `y81nrinc` matches the DID computed from the 2×2 table.
- The estimated effect is negative (incinerator reduced nearby prices), but the p-value is around 0.11, so it is not statistically significant at the 5% level.

Assumptions in the incinerator example

- **Parallel trends**: without the incinerator, house prices near and far from the site would have followed the same trend over time.
- **No anticipation**: before the incinerator was announced, living near the future site did not affect house prices.

DID with covariates

- The basic DID estimator is unbiased only if the parallel trends assumption holds unconditionally. If houses near the incinerator site are systematically different from those farther away (e.g., older), and house age affects price trends, then near and far houses may follow different price trajectories even without the incinerator. This violates parallel trends and biases the DID estimate.
- Adding covariates addresses this: if parallel trends holds *conditional* on house characteristics, controlling for them removes the bias. It also reduces residual variance, improving precision.

```
reg_did_cov <- lm(rprice ~ y81 + nearinc + y81nrinc + age + I(age^2),
                 data = kielmc)
round(summary(reg_did_cov)$coefficients, 4)
```

	Estimate	Std. Error	t value	Pr(> t)
(Intercept)	89116.5354	2406.0511	37.0385	0.0000
y81	21321.0418	3443.6311	6.1914	0.0000
nearinc	9397.9359	4812.2218	1.9529	0.0517
y81nrinc	-21920.2700	6359.7454	-3.4467	0.0006
age	-1494.4240	131.8603	-11.3334	0.0000
I(age^2)	8.6913	0.8481	10.2476	0.0000

- After controlling for house age, the DID estimate becomes larger in magnitude and statistically significant. The change in the estimate suggests that the basic DID was biased: older houses near the incinerator appreciated differently than houses farther away, masking part of the incinerator's negative effect.