EC 320: An Intro to Difference-in-Differences (DiD)

David Hall
University of Oregon

August 3, 2026

Introduction

Recall from OLS that for our regression to be unbiased, we needed the regressors to be uncorrelated with the error. In fact, we derived that our estimated slope can be written as:

$$\hat{\beta}_1 \;=\; \beta_1 \;+\; \frac{Cov(X,u)}{Var(X)}.$$

This is the omitted variable bias term. If $Cov(X,u) \neq 0$, then $\hat{\beta}_1$ is biased because it contains part of the unobserved error term.

In practice this is extremely difficult to guarantee, because there are always unobserved forces affecting Y that we cannot observe or include in the regression. Up to now, your main tool has been to add more control variables. But we cannot measure or observe everything.

Difference-in-Differences (DiD) is a design that helps us handle this problem in a different way. The idea is that we observe a "treated" group and a "control" group that are similar in every way except for one thing: the treatment. And because that treatment is exactly the effect we want to learn about, if we compare how each group changes over time, we can difference out the part of the unobserved error that is fixed within each group across time. If that error component is the same in the pre and post periods for each group, then differencing removes it — eliminating the bias term from above through differencing rather than through adding controls.

The Specification

We run the following regression:

$$Y_{st} = \beta_0 + \beta_1 Treated_s + \beta_2 Post_t + \beta_3 (Treated_s \times Post_t) + \varepsilon_{st}$$

where:

- Treated = 1 if the group ever receives the treatment, 0 if not
- Post = 1 after the treatment begins, 0 before

The 2×2 Structure and Omitted Variable Bias

Now write each of the four group-time cells as the part we *want* (the betas) plus the part we *cannot observe* (the omitted variable bias). Importantly, we assume that this omitted variable bias term is **time-invariant** within each group.

	Pre (Post=0)	Post (Post=1)
Control Group (Treated=0)	$\beta_0 + \text{OVB}_c$	$\beta_0 + \beta_2 + \text{OVB}_c$
Treated Group (Treated=1)	$\beta_0 + \beta_1 + \text{OVB}_t$	$\beta_0 + \beta_1 + \beta_2 + \beta_3 + \text{OVB}_t$

Now compute the within-group before/after differences:

$$\text{Treated Post} - \text{Treated Pre} = (\beta_2 + \beta_3) + (\text{OVB}_t - \text{OVB}_t) = \beta_2 + \beta_3$$

$$\text{Control Post} - \text{Control Pre} = \beta_2 + (\text{OVB}_c - \text{OVB}_c) = \beta_2$$

Then difference those differences:

$$(\beta_2 + \beta_3) - \beta_2 = \beta_3.$$

The time-invariant omitted variable bias cancels out completely. This is why we are able to identify the causal effect of treatment through differencing rather than by observing every confounder directly.

Summary

Difference-in-Differences recovers the causal effect by subtracting away time-invariant unobserved differences between the treated and control groups. This is not about perfectly measuring every variable — it is about using the structure of the data to make the omitted variable bias drop out through differencing.

Parallel trends (which we will cover next) is what allows the control group's change to represent what would've happened to the treated group in the absence of treatment.