Applied Microeconometrics: Differences in Differences

Nicholas Giannakopoulos

University of Patras
Department of Economics

ngias@upatras.gr

December 17, 2019

Overview

What Is an Impact Evaluation?

"An impact evaluation assesses changes in the well-being of individuals, households, communities or firms that can be attributed to a particular project, program or policy. The central impact evaluation question is what would have happened to those receiving the intervention if they had not in fact received the program. Since we cannot observe this group both with and without the intervention, the key challenge is to develop a counterfactual which is a group as similar as possible (in observable and unobservable dimensions) to those receiving the intervention. This comparison allows for the establishment of definitive causality, attributing observed changes in welfare to the program, while removing confounding factors."

Gertler, P. J., Martinez, S., Premand, P., Rawlings, L. B., & Vermeersch, C. M. (2016). Impact evaluation in practice. The World Bank.

Why Evaluate?

- Programs and policies are typically designed to change outcomes, for example, to raise incomes, to improve learning, or to reduce illness.
- Whether or not these changes are actually achieved is a crucial public policy question but one that is not often examined.
- More commonly, program managers and policy makers focus on controlling and measuring the inputs and immediate outputs of a program—how much money is spent, how many textbooks are distributed—rather than on assessing whether programs have achieved their intended goals of improving well-being.
- Impact evaluations are part of a broader agenda of evidence-based policy making. This **growing global trend** is marked by a shift in focus from inputs to outcomes and results.
- Monitoring and evaluation are at the heart of evidence-based policy making.

Why Evaluate?

- An impact evaluation assesses the changes in the well-being of individuals that can be attributed to a particular project, program, or policy.
- The central challenge in carrying out effective impact evaluations is to identify the causal relationship between the project, program, or policy and the outcomes of interest.
- Impact evaluations generally estimate average impacts of a program on the welfare of beneficiaries.
- Examples:
  - did the introduction of a new curriculum raise test scores among students?
  - Did a water and sanitation program increase access to safe water and improve health outcomes?
  - Was a youth training program eff ective in fostering entrepreneurship and raising incomes?

What Is an Impact Evaluation?

Basic evaluation question: What is the impact or causal eff ect of a program on an outcome of interest?
Goal: measure causal impacts of policy on participants

- We choose A and as a result, B happened
- A is a policy or intervention
- B is an outcome of interest
- Examples:
  - We gave out insecticide-treated bednets, and fewer children under the age of 5 got sick with or died from malaria as a result
  - We distributed free lunches in elementary schools, and school attendance and/or academic performance went up as a result

Establishing Causality

Goal: measure causal impacts of policy on participants

- We want to be able to say B happened because of A
    - We need to rule out other possible causes of B
- If we can say this, then we can also say: if we did A again (in another place), we think that B would happen there as well

In an ideal world (research-wise), we could clone each program participant and observe the impacts of our program on their lives

 vs. 

Establishing Causality

In an ideal world (research-wise), we could clone each program participant and observe the impacts of our program on their lives

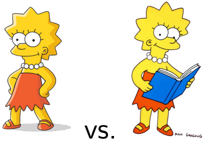

What is the impact of giving Lisa a book on her test score?

- Impact = Lisa's score with a book - Lisa's score without a book

In the real world, we either observe Lisa with a book or without

- We never observe the **counterfactual**

Establishing Causality

To measure the causal impact of giving Lisa a book on her test score, we need to find a comparison group that did not receive a book

 vs. 

Our estimate of the impact of the book is then the difference in test scores between the **treatment** group and the **comparison** group

- Impact = Lisa's score with a book - Bart's score without a book

However, finding a good comparison group is hard!

The Potential Outcomes Framework

Two potential outcomes for each individual, community, etc:

$$\text{Potential outcome} = \begin{cases} Y_{0i}, P_{0i} = 0 \\ Y_{1i}, P_{1i} = 1 \end{cases}$$

The problem: we only observe one of Y_{0i} and Y_{1i}

- Each individual either participates in the program or not
- The causal impact of program (P) on i is: $Y_{1i} - Y_{0i}$

We observe i's actual outcome:

$$Y_i = Y_{0i} + \underbrace{(Y_{1i} - Y_{0i})}_{\text{impact}} P_i$$

Defining the Counterfactual

To estimate the impact of a program, we need to know what would have happened to every participant i in the absence the program

- We call this the **counterfactual**

Of course, we can't actually clone our participants and see what happens to the clones if they don't participate in the program

- Instead, we estimate the counterfactual using a **comparison group**

The comparison group needs to:

- Look identical to the treatment group prior to the program
- Not be impacted by the program in anyway

You need a convincing comparison group!

Differences-in-Differences

- The individual-fixed effect approach requires repeated observations on the same individuals (units)...
- Often the variable of interest varies on more aggregate or group level, such as state or cohort
- Before and after: If we observe outcomes before and after treatment, we could use the treated before treatment as controls for the treated after treatment
- The problem of this comparison is that it can be contaminated by the effect of events other than the treatment that occurred between the two periods

Differences-in-Differences

- Policy change that affected only certain group at certain time
  - Suppose that only a fraction of the population is exposed to treatment. In such a case, we can use the group that never receives treatment to identify the temporal variation in outcomes that is not due to exposure to treatment. This is the basic idea of the DID
  - The fundamental identifying assumption is that the average changes in the two groups are the same in the absence of treatment:
    - No other simultaneous factors affecting the difference in outcomes between these groups
    - Parallel trends
  - Example: Suppose you are interested in the effect of minimum wages on employment (a classic and controversial question in labour economics)
  - In a competitive labour market, increases in the minimum wage would move us up a downward-sloping labour demand curve. Thus, employment would fall (Card & Krueger, 1994)

Differences-in-Differences

- Card & Krueger (1994) analyse the effect of a minimum wage increase in New Jersey using a differences-in-differences methodology.
- In February 1992 NJ increased the state minimum wage from \$4.25 to \$5.05. Pennsylvania's minimum wage stayed at \$4.25.

- They surveyed about 400 fast food stores both in NJ and in PA both before and after the minimum wage increase in NJ.

Differences-in-Differences

- DD is a version of fixed effects estimation. To see this more formally:

  Y_{1ist} : employment at restaurant i, state s, time t with a high w^{min}.

  Y_{0ist} : employment at restaurant i, state s, time t with a low w^{min}.
- In practice of course we only see one or the other.
- We then assume that:

$$E[Y_{0ist}|s,t] = \gamma_s + \lambda_t$$

- In the absence of a minimum wage change, employment is determined by the sum of a time-invariant state effect γ_s and a year effect λ_t that is common across states.
- Let D_{st} be a dummy for high-minimum wage states and periods.
- Assuming $E[Y_{1ist} - Y_{0ist}|s,t] = \delta$ is the treatment effect, observed employment can be written:

$$Y_{ist} = \gamma_s + \lambda_t + \delta D_{st} + \varepsilon_{ist}$$

Differences-in-Differences

- In New Jersey:
  - Employment in February is:
    $E[Y_{ist}|s = NJ, t = Feb] = \gamma_{NJ} + \lambda_{Feb}$
  - Employment in November is:
    $E[Y_{ist}|s = NJ, t = Nov] = \gamma_{NJ} + \lambda_{Nov} + \delta$
  - the difference between February and November is:
    $E[Y_{ist}|s = NJ, t = N] - E[Y_{ist}|s = NJ, t = F] = \lambda_N - \lambda_F + \delta$
- In Pennsylvania:
  - Employment in February is:
    $E[Y_{ist}|s = PA, t = Feb] = \gamma_{PA} + \lambda_{Feb}$
  - Employment in November is:
    $E[Y_{ist}|s = PA, t = Nov] = \gamma_{PA} + \lambda_{Nov}$
  - the difference between February and November is:
    $E[Y_{ist}|s = PA, t = Nov] - E[Y_{ist}|s = PA, t = Feb] = \lambda_{Nov} - \lambda_{Feb}$

Differences-in-Differences

- The differences-in-differences strategy amounts to comparing the change in employment in NJ to the change in employment in PA.
- The population differences-in-differences are:

$$E[Y_{ist}|s = NJ, t = N] - E[Y_{ist}|s = NJ, t = F]$$
$$-E[Y_{ist}|s = PA, t = Nov] - E[Y_{ist}|s = PA, t = Feb] = \delta$$

- This is estimated using the sample analog of the population means.

Differences-in-Differences

Wage distribution ("treatment intensity") before/after April 1992:
(note the changing y-axis scale)

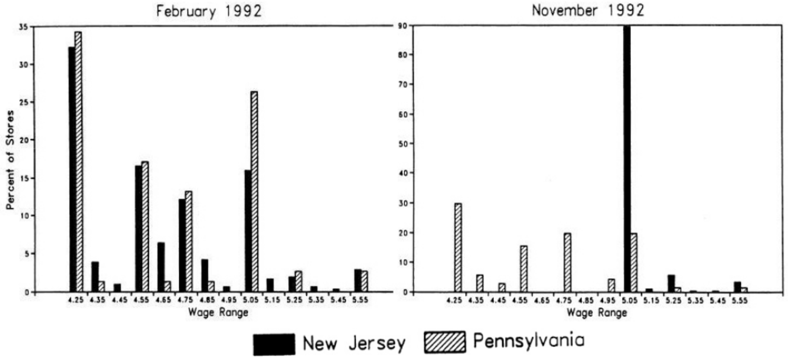

Differences-in-Differences

Variable	Stores by state		
	PA (i)	NJ (ii)	Difference, NJ − PA (iii)
1. FTE employment before, all available observations	23.33 (1.35)	20.44 (0.51)	−2.89 (1.44)
2. FTE employment after, all available observations	21.17 (0.94)	21.03 (0.52)	−0.14 (1.07)
3. Change in mean FTE employment	−2.16 (1.25)	0.59 (0.54)	2.76 (1.36)

Differences-in-Differences

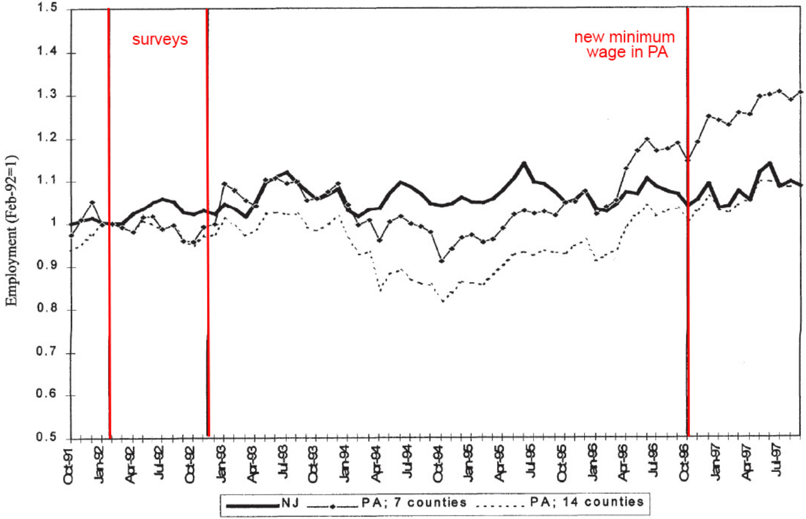

Differences-in-Differences

- We can estimate the differences-in-differences estimator in a regression framework.
- Advantages:
  - It is easy to calculate standard errors.
  - We can control for other variables which may reduce the residual variance (lead to smaller standard errors).
  - It is easy to include multiple periods.
  - We can study treatments with different treatment intensity. (e.g. varying increases in the minimum wage for different states).
- The typical regression model that we estimate is:

$$\text{Outcome}_{it} = \beta_1 + \beta_2 \text{ Treat}_i + \beta_3 \text{ Post}_t + \beta_4 \text{ (Treat * Post)}_{it} + \varepsilon$$

Treatment = a dummy if the observation is in the treatment group
Post = post treatment dummy

Differences-in-Differences

- In the Card & Krueger case the equivalent regression model would
  $Y_{ist} = \alpha + \gamma NJ_s + \lambda d_t + \delta(NJ_s * d_t) + \varepsilon_{ist}$
  - NJ is a dummy which is equal to 1 if the observation is from NJ.
  - d is a dummy which is equal to 1 if the observation is from Novemb (post).
- This equation takes the following values.
  - PA Pre: α
  - PA Post: $\alpha + \lambda$
  - NJ Pre: $\alpha + \gamma$
  - NJ Post: $\alpha + \gamma + \lambda + \delta$
- Differences-in-Differences estimate: (NJ Post - NJ Pre) - (PA Post PA Pre) $= \delta$

Differences-in-Differences

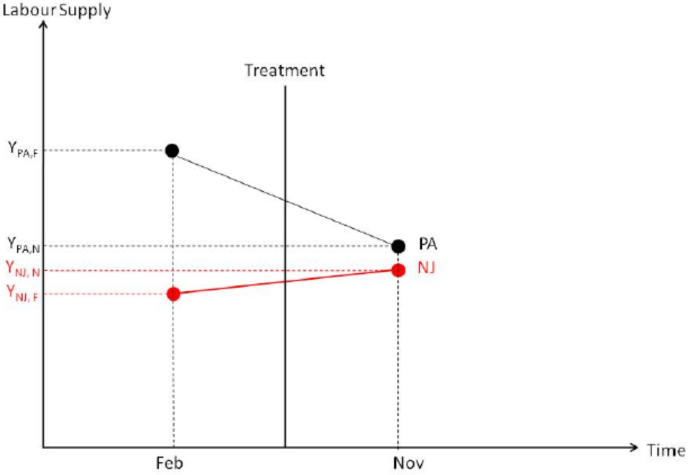

Differences-in-Differences

$$Y_{ist} = \alpha + \gamma NJ_s + \lambda d_t + \delta(NJ_s * d_t) + \varepsilon_{ist}$$

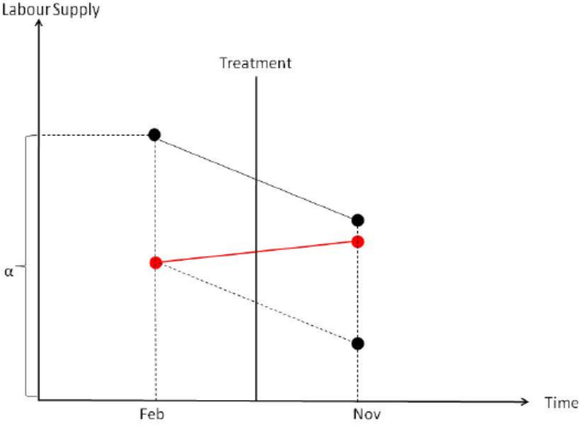

Differences-in-Differences

$$Y_{ist} = \alpha + \gamma NJ_s + \lambda d_t + \delta(NJ_s * d_t) + \varepsilon_{ist}$$

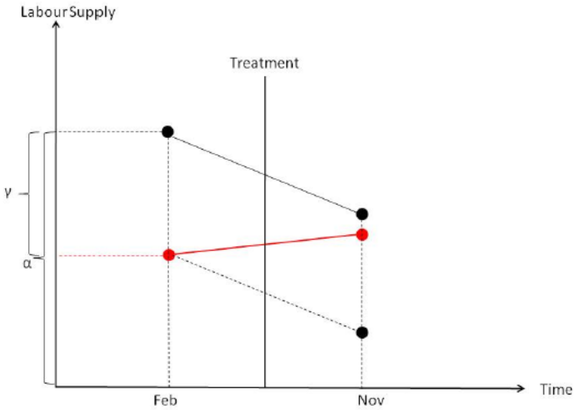

Differences-in-Differences

$$Y_{ist} = \alpha + \gamma NJ_s + \lambda d_t + \delta(NJ_s * d_t) + \varepsilon_{ist}$$

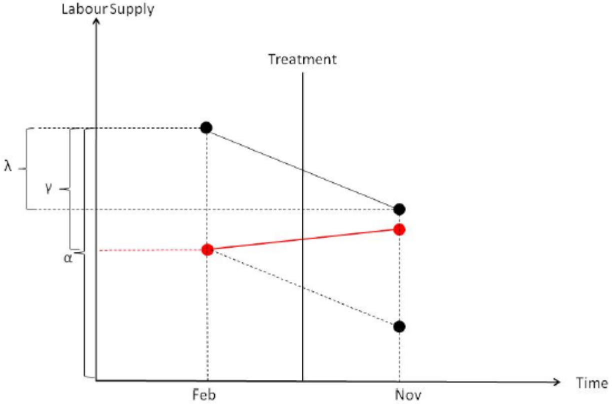

Differences-in-Differences

$$Y_{ist} = \alpha + \gamma NJ_s + \lambda d_t + \delta(NJ_s * d_t) + \varepsilon_{ist}$$

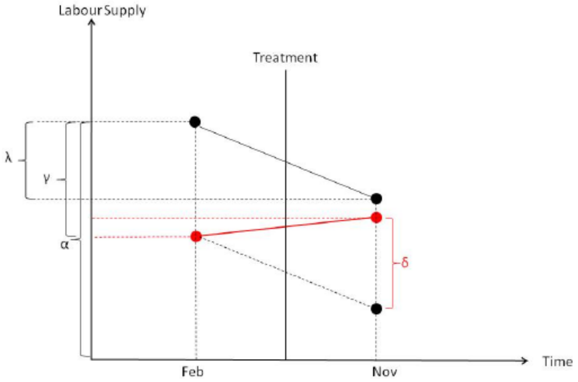

Differences-in-Differences

- The key assumption for any DD strategy is that the outcome in treatment and control group would follow the same time trend in the absence of the treatment.
- This does not mean that they have to have the same mean of the outcome!
- Common trend assumption is difficult to verify but one often uses pre-treatment data to show that the trends are the same.
- Even if pre-trends are the same one still has to worry about other policies changing at the same time.

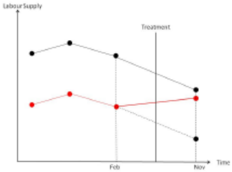

Differences-in-Differences

- Including leads into the DD model is an easy way to analyze pre-trends.
- Lags can be included to analyze whether the treatment effect chan over time after treatment.
- The estimated regression would be:

$$Y_{ist} = \gamma_s + \lambda_t + \sum_{\tau=-q}^{-1} \delta_\tau D_{s\tau} + \sum_{\tau=0}^{m} \delta_\tau D_{s\tau} + X_{ist} + \varepsilon_{ist}$$

  - treatment occurs in year 0.
  - includes q leads or anticipatory effects.
  - includes m leads or post treatment effects.