Principles of Empirical Analysis

Lecture 8: Difference-in-differences

Spring 2024
Miri Stryjan

Observational alternatives to experiments

1. **Selection on observables: treatment and control groups differ from each other only w.r.t. observable characteristics,** there is no problematic ***selection*** into treatment
   - Lecture 7 discussed such a situation
2. **Selection on unobservables: treatment and control groups differ from each other in unobservable characteristics**
   - Treatment and controls are observed before and after treatment – difference-in-differences (DID)
   - Selection mechanism is known – regression discontinuity designs (RDD)
   - Exogenous variable induces variation in treatment – instrumental variables (IV)

Outline

- **Basic idea of difference-in-difference (DID, DD, diff-in-diff) designs**
  - DID with two groups and two time periods
  - (Next time: More general case with many time periods)
- **Application:**
  - Card & Krueger (1994): classic paper on minimum wage

First: a note on pre-lecture assignments

- **Level of detail:** answer should show that you *understood* something
  - Not just repeat words straight from the material with no context.
  - A few sentences for each question can be a good rule of thumb.
- **Explanations:** especially if the question asks you to “explain” something, you need to include some actual explanation about the concept or issue asked about: what it is/why it is used/what it means…
- Idea: these short question prepare for both lectures and for writing answers in the exam.

DID

- **We have talked about the idea of using differences between groups to estimate causal effects**
  - We would like to find treatment and control groups who can be assumed to be similar in every way except receipt of treatment
  - Without randomization this is very difficult/implausible
- **A weaker assumption is that in the absence of treatment, the difference between treatment and control groups is constant over time (parallel or common trends)**
- **With this assumption we can relax the requirement that the treatm. and control groups are almost identical/as good as randomly assigned:**
  - use observations in treatment and control groups before and after the treatment to estimate a causal effect

DID

- **Idea:**
  - Pre-treatment difference between the groups is ‘normal’ difference
  - Post-treatment difference is ‘normal’ difference + causal effect of treatment
  - Difference-in-differences is the causal effect
- **DID relies heavily on common or parallel time trends, so visual inspection of the data is a very important part of any DID analysis**

First example: The Millennium Villages Project

This study does not use DID, but it is used to illustrate the idea.

- A joint project of the United Nations Development Program (UNDP), the Earth Institute at Columbia University, and an NGO called ”Millennium Promise”
- A large, expensive intervention at 15 sites in rural sub-Saharan Africa.
- Designed to show that ”people in the poorest regions of rural Africa can lift themselves out of extreme poverty in five years’ time” (MVP 2007).
- was launched, first in Sauri, Kenya in 2004 and then later at a number of sites across sub-Saharan Africa.

Thee MVP: a package of interventions

- **Video - Jeffrey Sachs and Angelina Jolie visit Sauri:**
  - https://www.youtube.com/watch?v=uUHf_kOUM74

Thee MVP: a package of interventions

The villages/villagers were offered several services:

- distribution of fertilizer
- school construction
- distribution of insecticide-treated bednets
- HIV testing
- Microfinance
- electric lines
- road construction
- water and irrigation

The intervention is in line with “big push” theory of economic development.

Simple difference: before/after

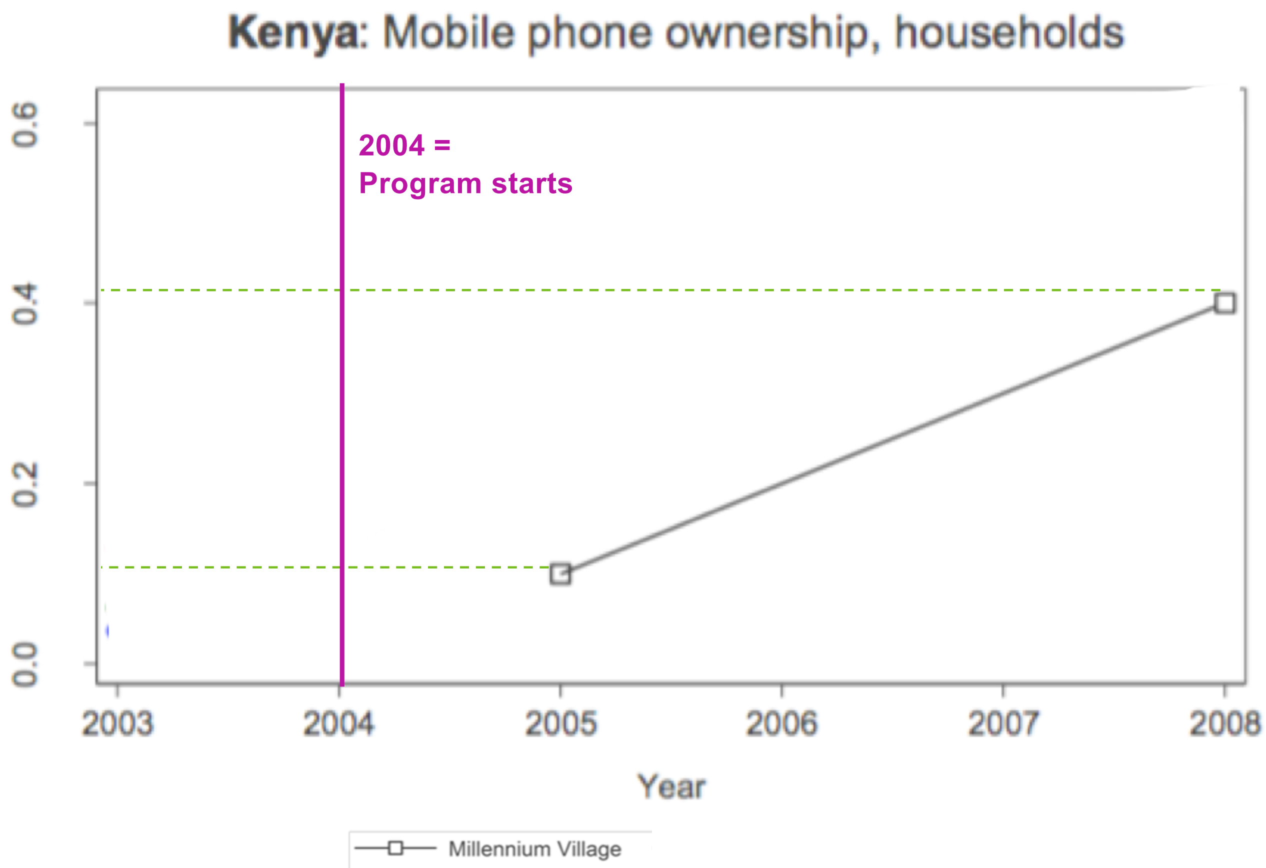

The share of mobile phone ownership increased substantially in the Millennium Village between 2005 and 2008. Does it mean the program had a positive effect?

Possible reasons for positive trend

Maybe the MVP led to better economic situation in the villages and thus to more mobile phones.

But there could also be alternative explanations for why mobile phones increased in the MV in this time-period.

How can we know which one it is when we have no control group?

Measuring *impact* of a program

– what is the counterfactual?

”Measuring the impact of the Project means asking this question: What happened at sites that received the Project’s package intervention, relative to what would have happened at those sites in the absence of the Project? ”In the absence of the Project” does not mean in the absence of any interventions whatsoever, it means what would have happened without that specific project.” (Clemens and Demombynes, 2010).

What do we need for evaluating the program impact

Goal: **measuring the effect of the program, T on some outcomes Y.**

We need to know the **counterfactual**: what would have happened in the absence of the program.

-> A control group that is *comparable** to the treated group!

How can we construct/find a control group with **observational data**?

With observational data – how could we construct a control group "ex post" (after the intervention)? Suggestions/ideas?

*Note: what we mean by "comparable" depends on the assumption of the specific estimation method used (DID, IV, RDD)

Adding possible “control groups” in 2008

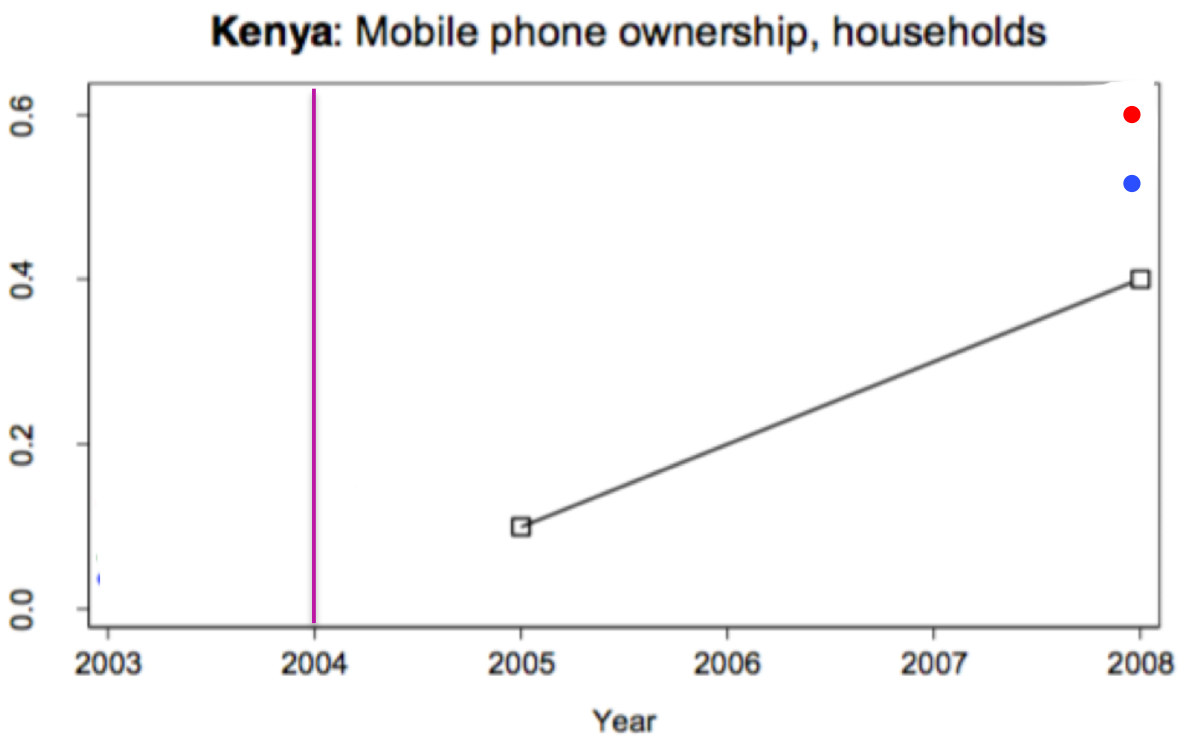

The share of mobile phone ownership in 2008 in two possible control regions (blue and red dots) is higher than in the MV. Does it mean the program had a negative effect?

Adding possible “control groups” for the same time period

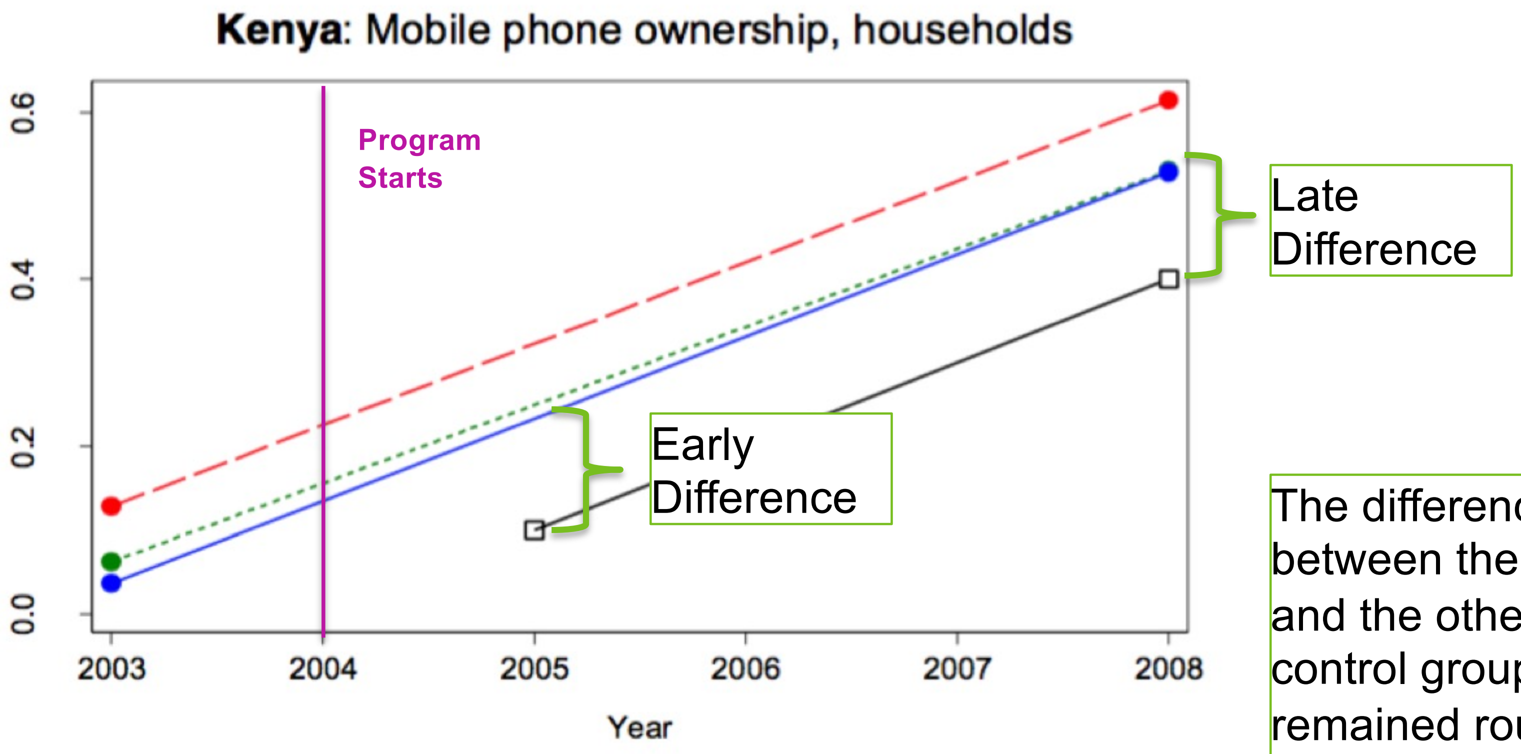

The difference between the MV and the other control groups remained roughly the same over the project period

Adding possible “control groups” for the same time period - idea

- Compare the changes at the sites to broader trends in the countries where the Millennium Villages are located.
- When doing this for Kenya, we see that the treated villages follow the same trend as the rest of Kenya
- But the rest of Kenya did not have the MVP.
- This suggests that mobile ownership would have risen at the MV sites with or without the project.
- **The counterfactual (what would have happened without treatment) would most likely have been a similar increase in mobile ownership as was found in similar villages in the rest of Kenya.**

Difference-in-differences – two groups, two time periods

DID: two groups, two time periods

- **The canonical DID design contains two time periods and two groups where the timing of treatment is the same for all treated**
  - Most current DID applications, however, use data from more than two time periods and often the treatment occurs at different times
  - There is currently a lively discussion going on regarding what to do when dealing with these more complicated designs!
- **The 2x2 design is still an excellent pedagogical point of departure**

Difference-in-differences graphically

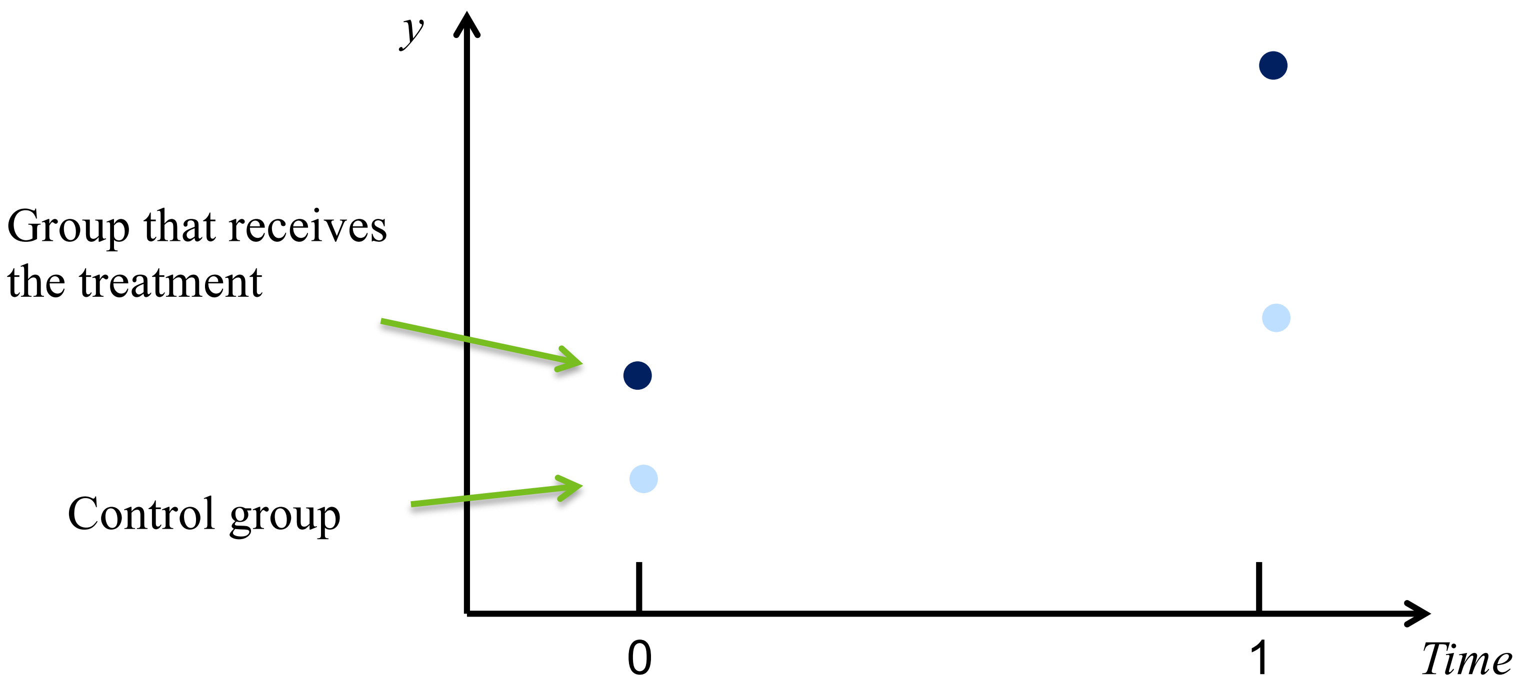

Difference-in-differences graphically

We are measuring the outcome of interest (y) in two time periods

The dots are means of the outcome for each group in each time period

The lines connecting the dots are just for visualization purposes

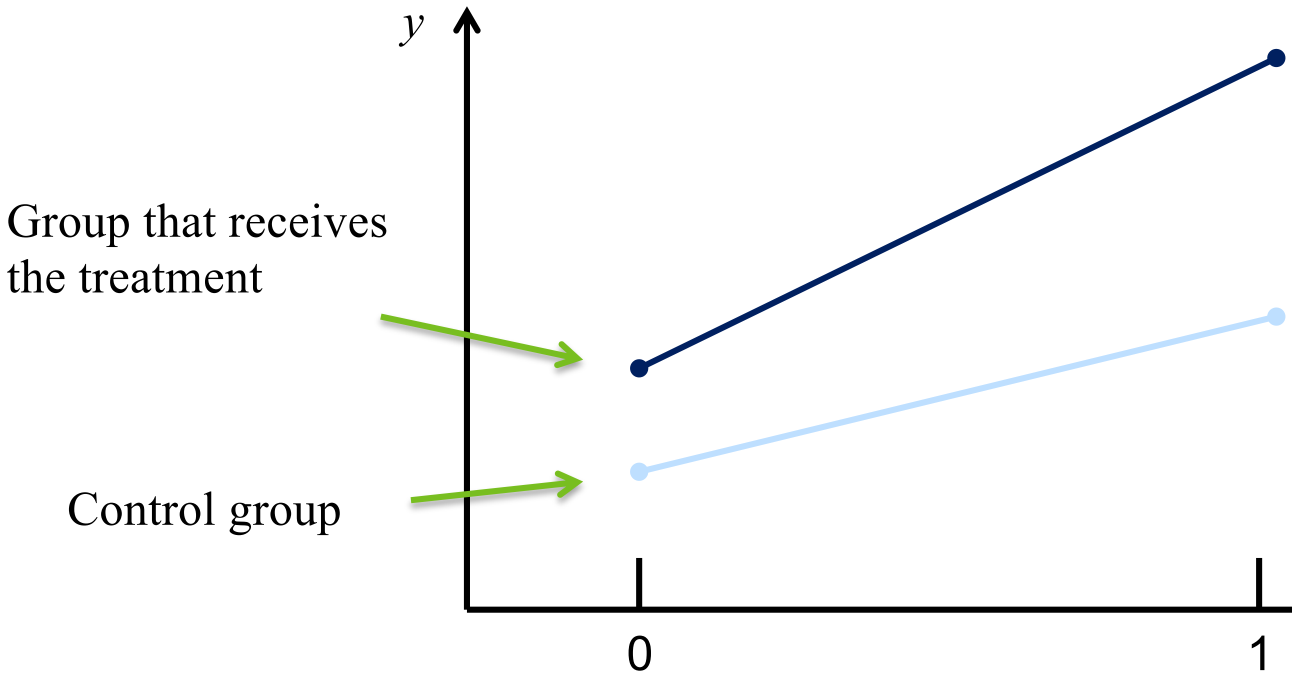

Difference-in-differences graphically

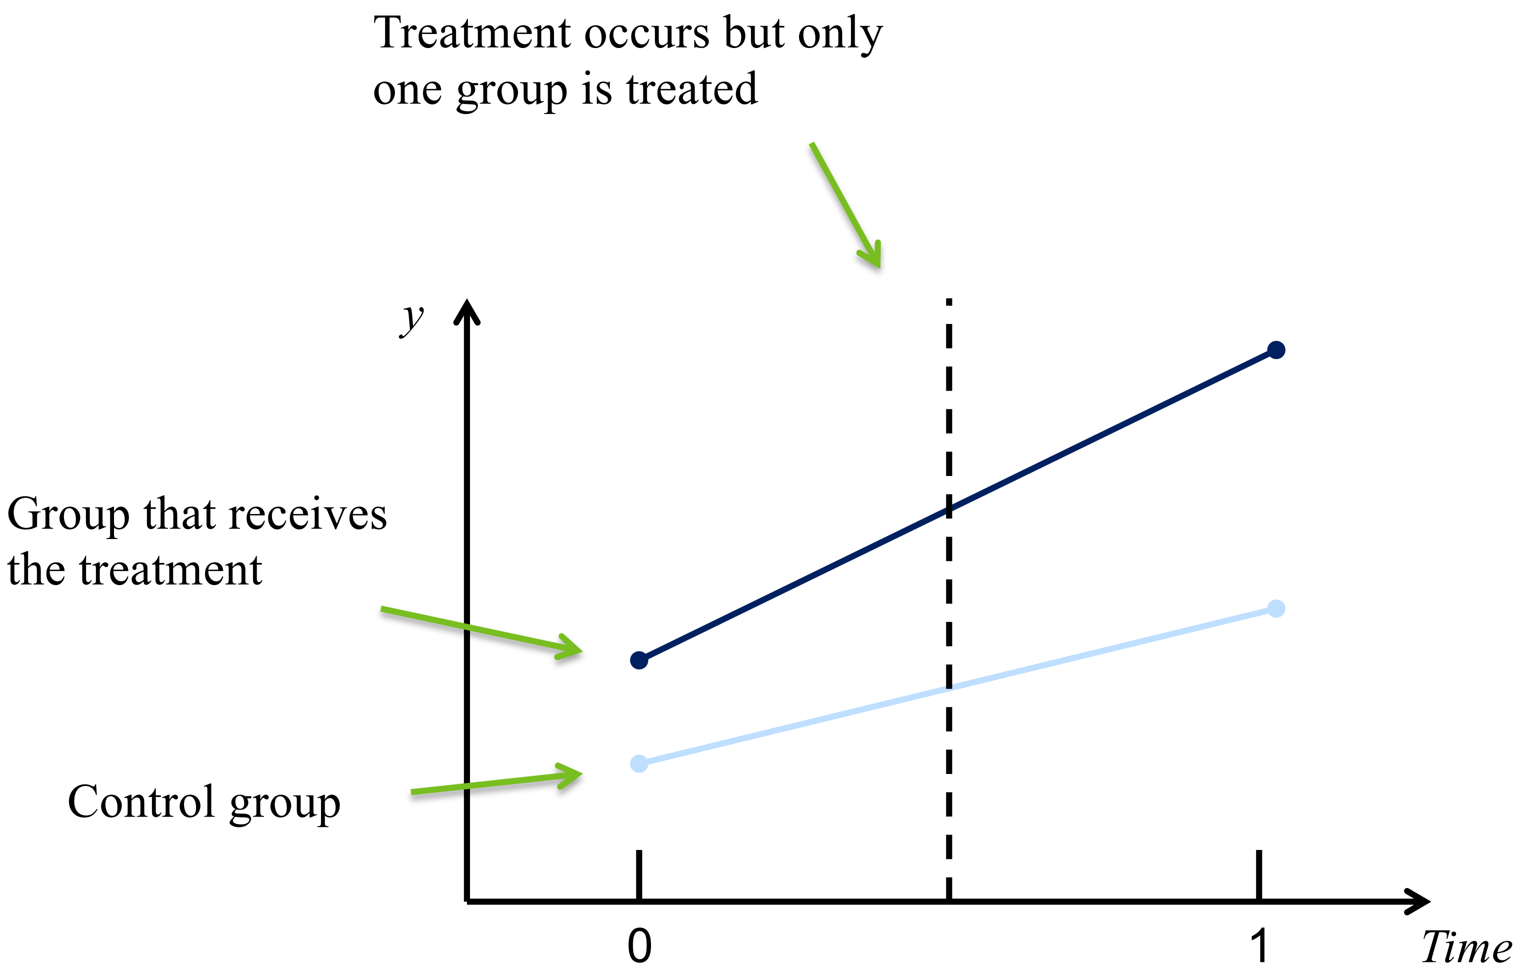

Difference-in-differences graphically

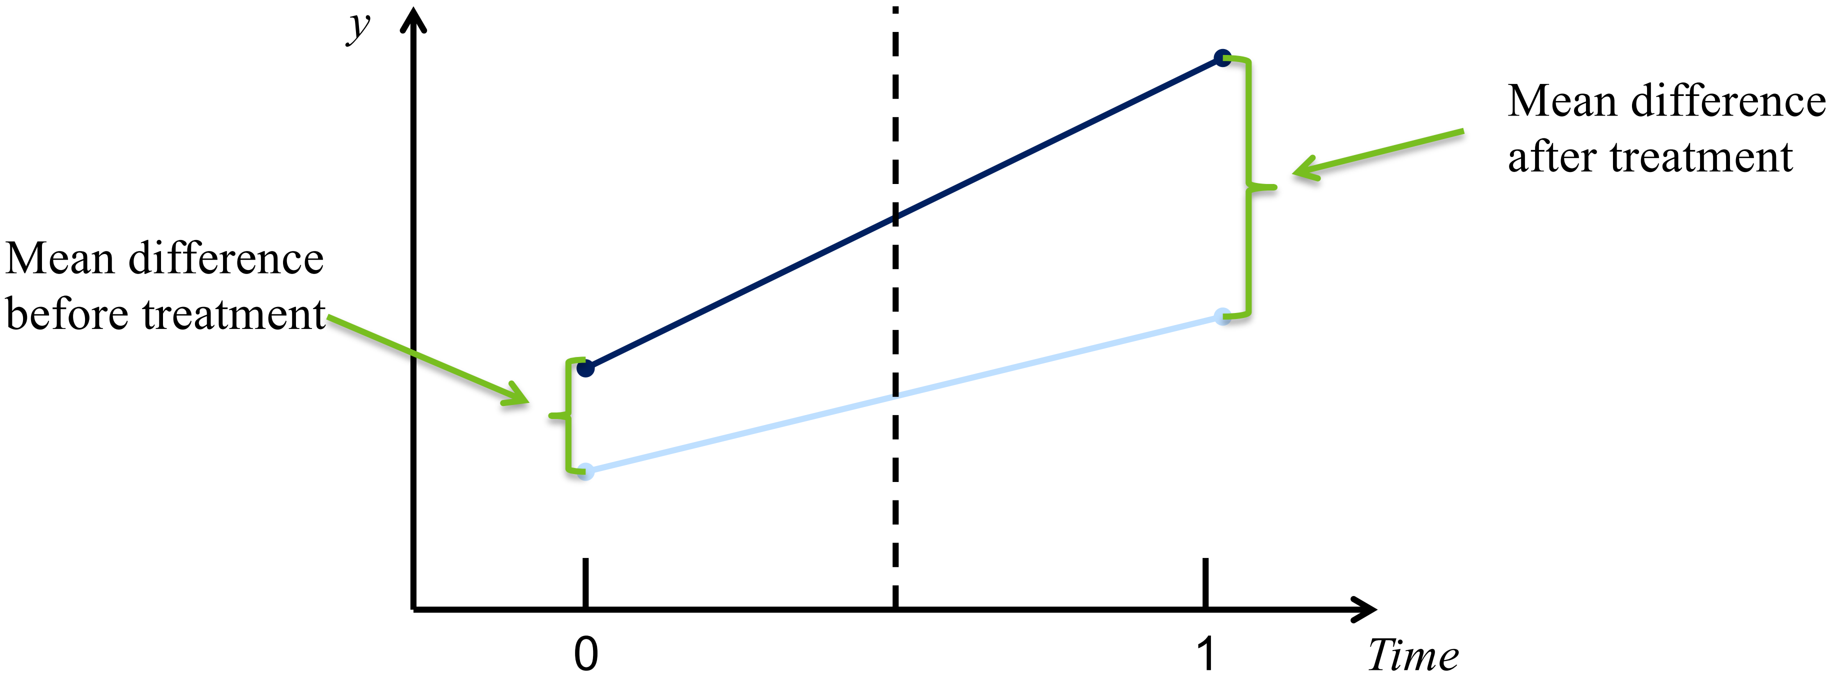

Difference-in-differences graphically

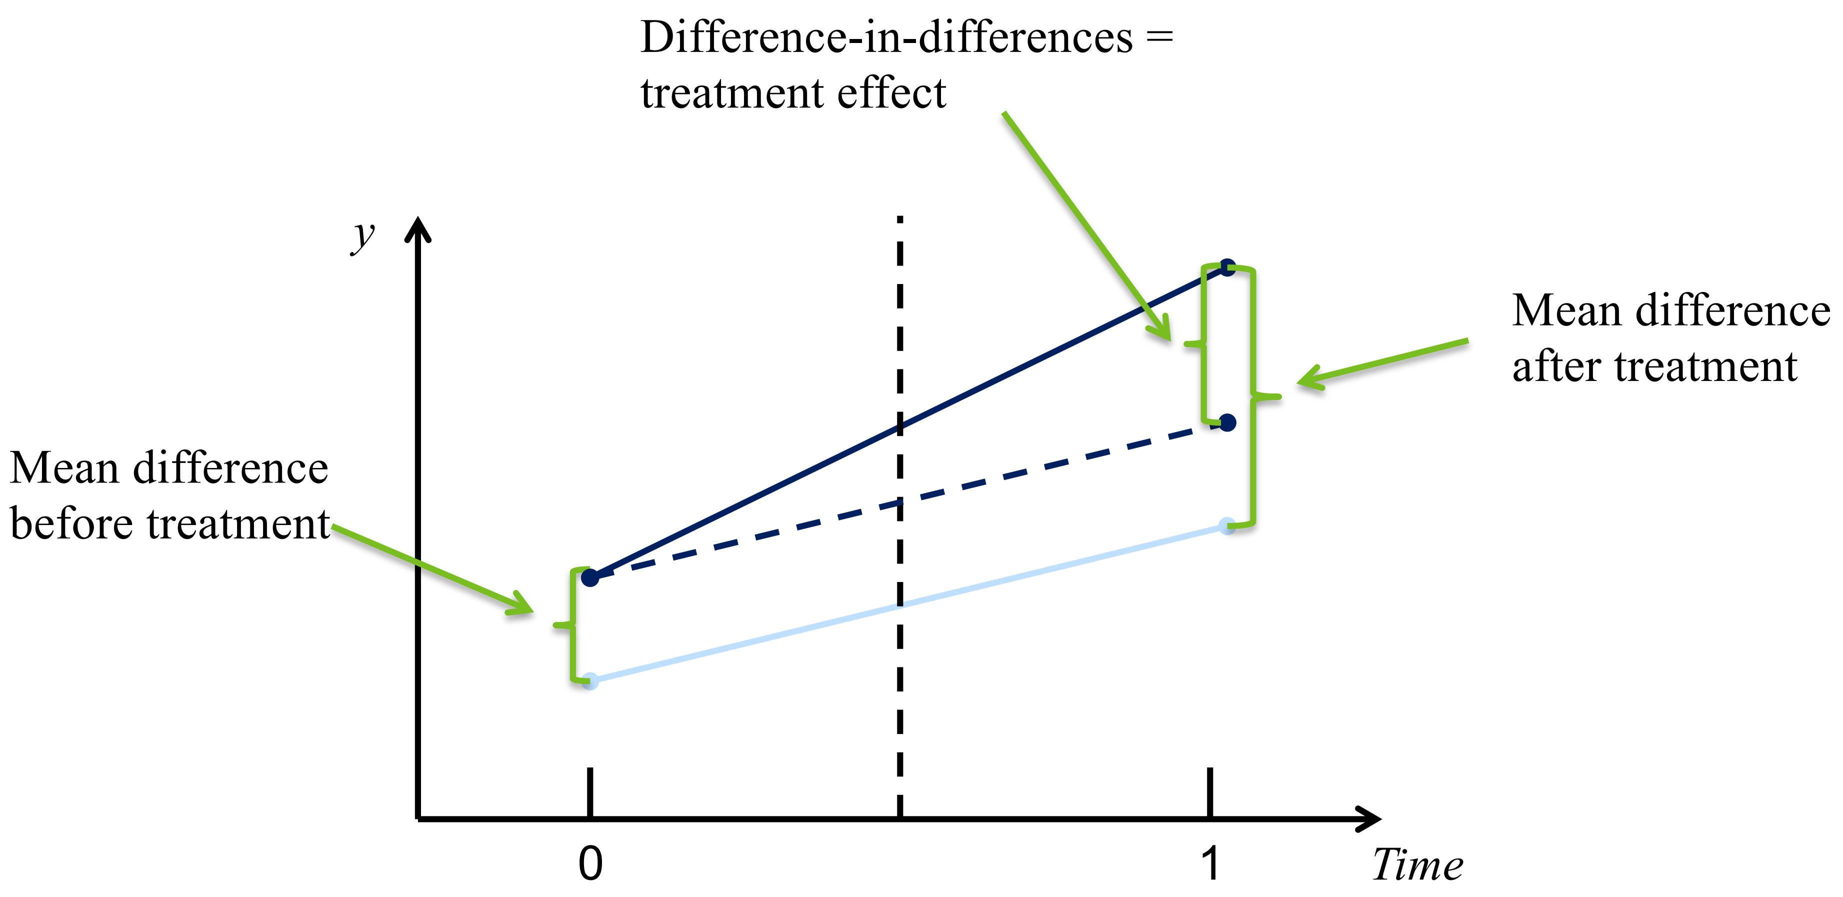

Difference-in-differences graphically

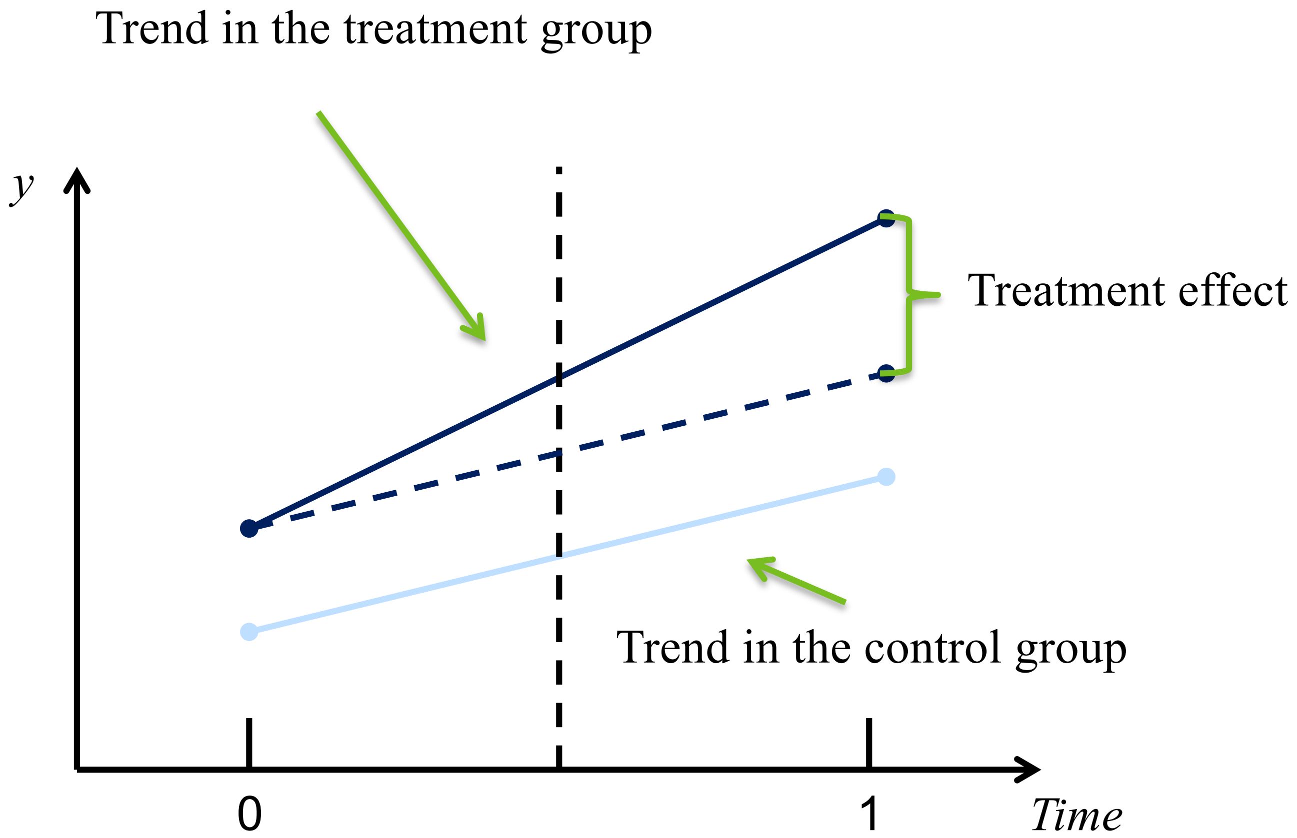

Difference-in-differences graphically

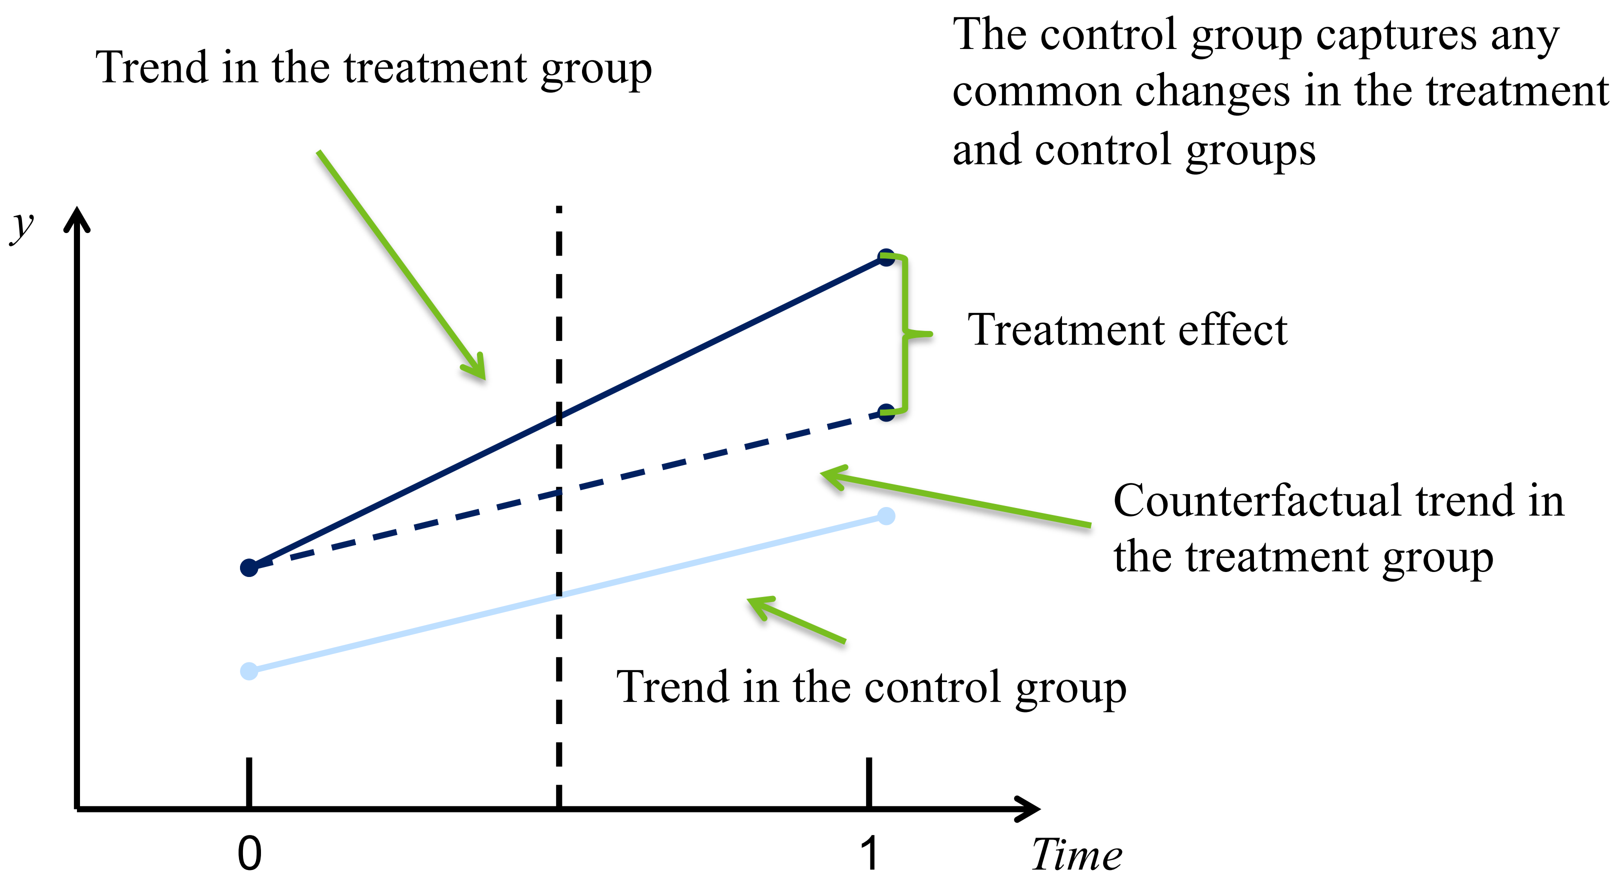

Difference-in-differences graphically

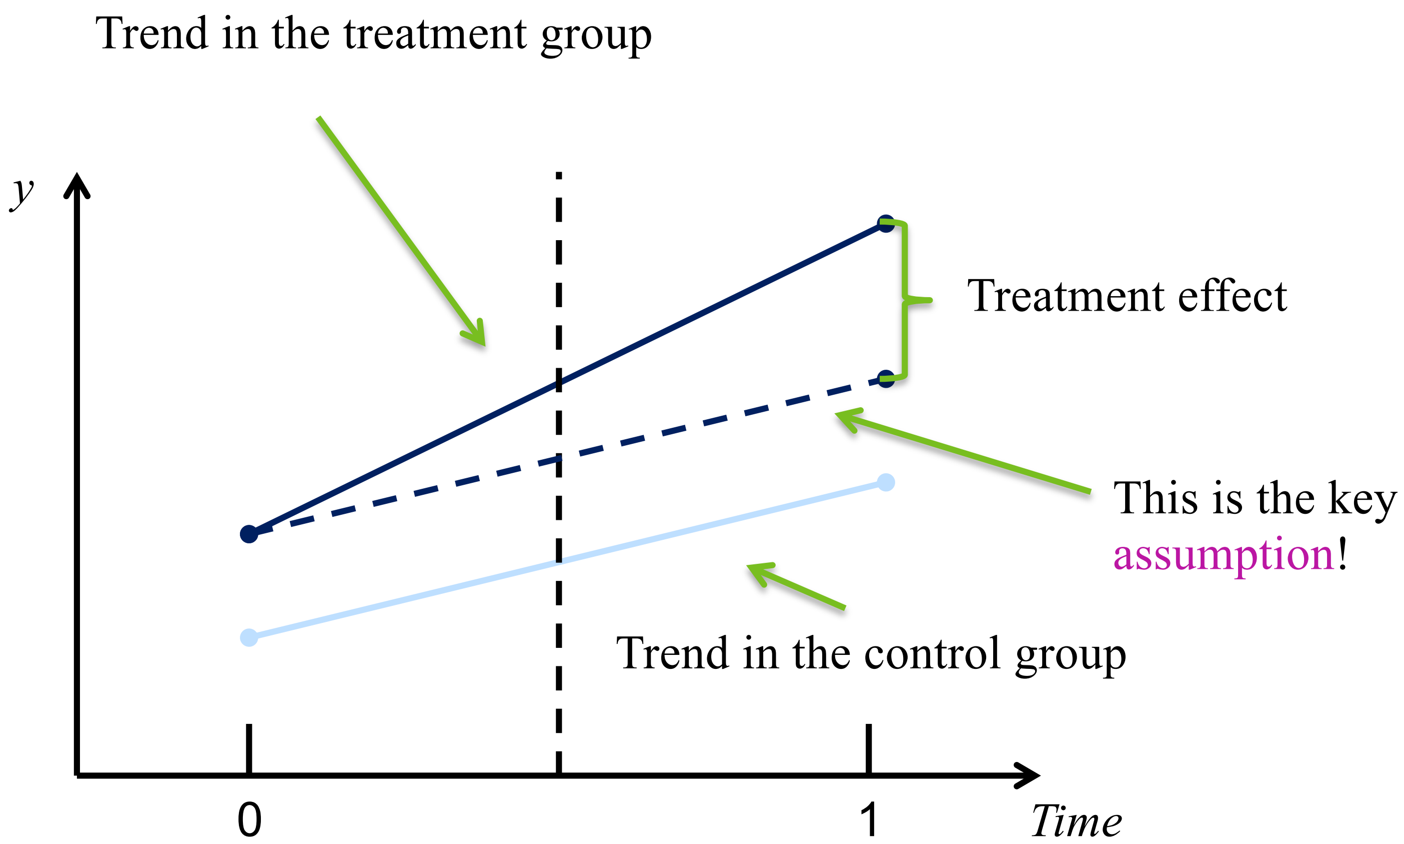

Example: New Jersey minimum wage increase

Treatment, T: higher minimum wage

Outcome, Y: Employment

Policy change: On April 1, 1992, NJ increased the state minimum wage from \$4.25 to \$5.05

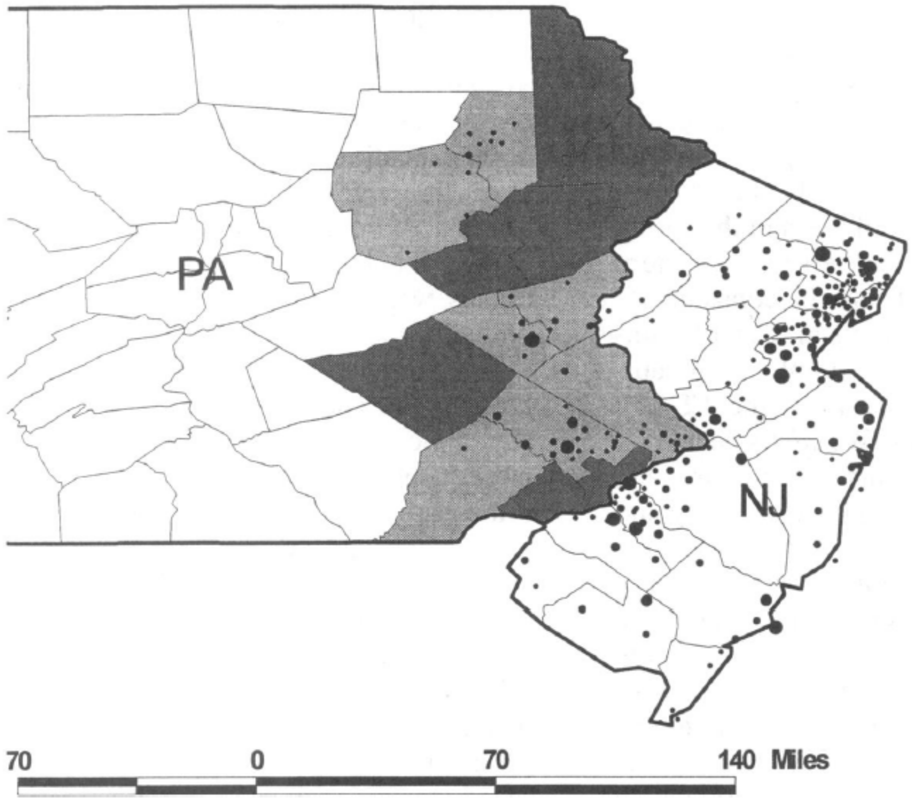

Card & Krueger (1994) wanted to measure how this change affected employment. Test the idea that min wage interferes with the demand and supply mechanism and leads to higher unemployment.

Possible evaluation strategy: compare employment in NJ in 1994 with employment in March 1992 (before/after). **Do you see any problems with this?**

Economy wide changes between 1992-94 could be a confounder.

Example: New Jersey minimum wage increase

Treatment, T: higher minimum wage

Outcome, Y: Employment

Enter: control group!

On April 1, 1992, NJ increased the state minimum wage from $4.25 to $5.05; **PA's minimum wage stayed at $4.25**

Card & Krueger (1994) surveyed about 400 fast food stores both in NJ and in PA before (February) and after (November) the minimum wage increase

Any common macroeconomic trends will be captured by using the control group in PA.

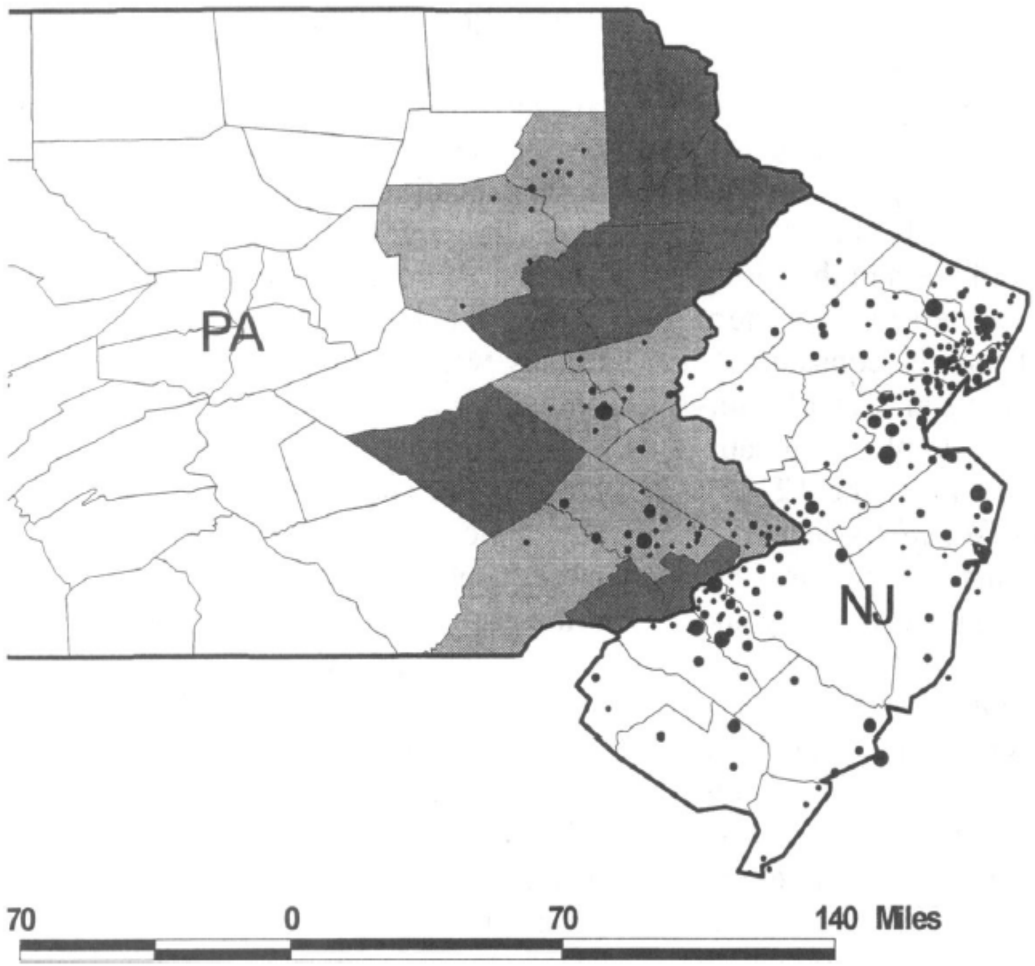

Card & Krueger (1994)

- **DID more formally**

  y_{ist}: employment at restaurant i, state s, time t

- **In DID, we need the following means**

$E[y_{ist}|s = NJ, t = Feb]$ **before**

$E[y_{ist}|s = NJ, t = Nov]$ **after**

$E[y_{ist}|s = NJ, t = Nov]$ –
$E[y_{ist}|s = NJ, t = Feb]$
= Diff 1: difference in employment in T

$E[y_{ist}|s = PA, t = Feb]$ **before**

$E[y_{ist}|s = PA, t = Nov]$ **after**

$E[y_{ist}|s = PA, t = Nov]$ –
$E[y_{ist}|s = PA, t = Feb]$
= Diff 2: difference in employment in C

Card & Krueger (1994)

- **In New Jersey:**

  $E[y_{ist}|s = NJ, t = Feb]$ = mean employment in February

  $E[y_{ist}|s = NJ, t = Nov]$ = mean employment in November

  $E[y_{ist}|s = NJ, t = Nov] - E[y_{ist}|s = NJ, t = Feb]$

  = Difference 1: difference in employment in NJ, the **treated** area

- **In Pennsylvania:**

  $E[y_{ist}|s = PA, t = Feb]$ = mean employment in February

  $E[y_{ist}|s = PA, t = Nov]$ = mean employment in November

  $E[y_{ist}|s = PA, t = Nov] - E[y_{ist}|s = PA, t = Feb]$

  = Difference 2: difference in employment in PA, the **control** area

Card & Krueger (1994)

- **The population DID is the treatment effect we are looking for: (Difference 1-Difference 2 from previous slide)**

$$\delta = (E[y_{ist}|s = NJ, t = Nov] - E[y_{ist}|s = NJ, t = Feb]) - (E[y_{ist}|s = PA, t = Nov] - E[y_{ist}|s = PA, t = Feb])$$

- **(The DID estimator is the sample analog):**

$$\hat{\delta} = \left(\bar{y}_{NJ,Nov} - \bar{y}_{NJ,Feb}\right) - \left(\bar{y}_{PA,Nov} - \bar{y}_{PA,Feb}\right)$$

Card & Krueger (1994)

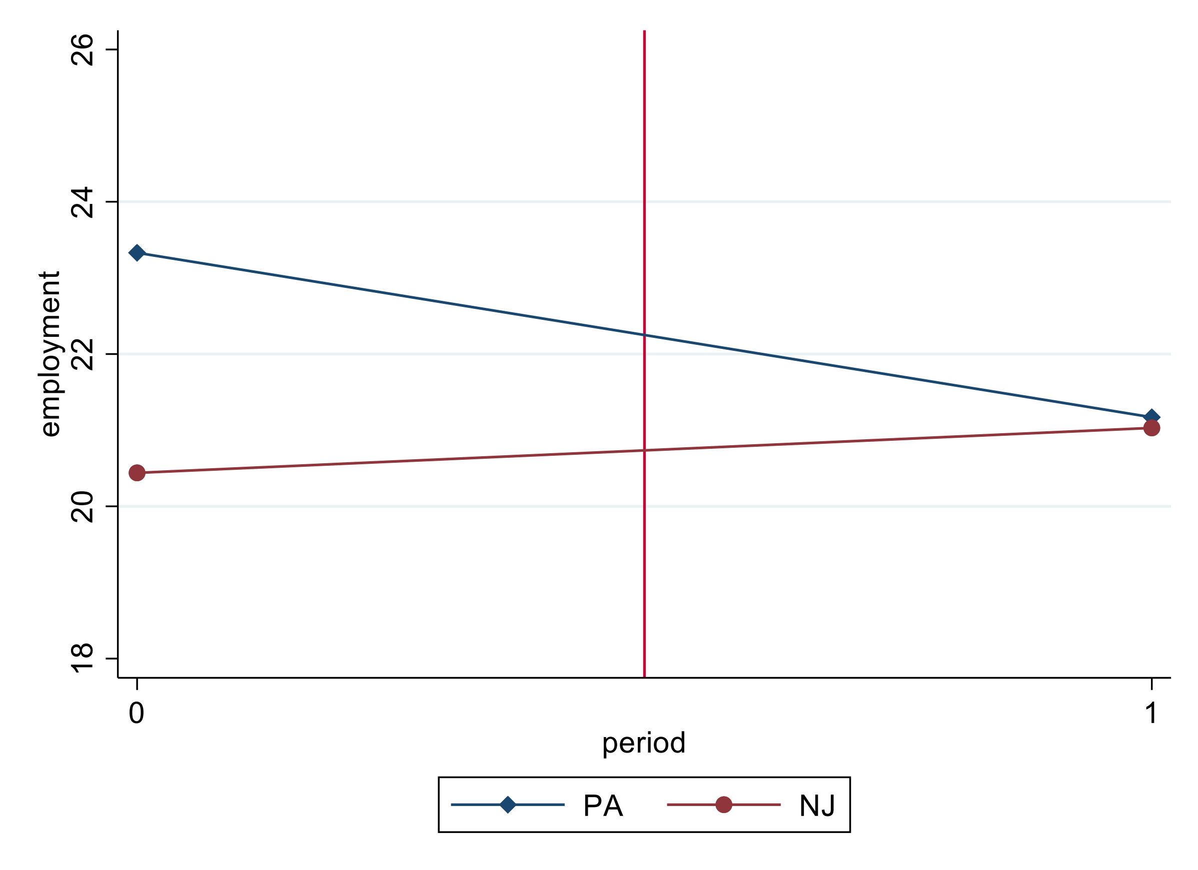

Variable	PA (i)	NJ (ii)	Difference, NJ – PA (iii)
	Stores by state		
1. FTE employment before, all available observations	23.33 (1.35)	20.44 (0.51)	–2.89 (1.44)
2. FTE employment after, all available observations	21.17 (0.94)	21.03 (0.52)	–0.14 (1.07)
3. Change in mean FTE employment	–2.16 (1.25)	0.59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

(21.03 – 20.44) – (21.17 – 23.33) = 2.76

- Surprisingly, if anything employment increased in New Jersey!

DID using regression

In the 2x2 case, the regression model would look like this

$$y_{it} = \alpha + \beta treated_i + \gamma after_t + \delta treated_i \cdot after_t + u_{it}$$

treated = 1 if observation is in the treatment group, 0 otherwise

after = 1 if observation is from the after period, 0 otherwise

*treated*after* = 1 for if observation is in the treatment group AND observed after the treatment

In econometrics jargon, *treated* and *after* are dummy variables and their product is called an interaction term

α is referred to as the intercept or the constant term

DID using regression

In Card & Krueger minimum wage study this would be

$$y_{ist} = \alpha + \beta NJ_s + \gamma Nov_t + \delta NJ_s \cdot Nov_t + u_{ist}$$

NJ = 1 if observation is in New Jersey the treatment group, 0 otherwise (regardless of the time period)

Nov = 1 if observation is from the after period, 0 otherwise (regardless of the state)

*NJ*Nov* = 1 for if observation is in New Jersey observed after the treatment

DID using regression* (extra, but useful for HW4)

In Card & Krueger minimum wage study this would be

$$y_{ist} = \alpha + \beta NJ_s + \gamma Nov_t + \delta NJ_s \cdot Nov_t + u_{ist}$$

- *NJ before:* $E[y_{ist} | NJ = 1, Nov = 0] = \alpha + \beta$
- *NJ after:* $E[y_{ist} | NJ = 1, Nov = 1] = \alpha + \beta + \gamma + \delta$
- *PA before:* $E[y_{ist} | NJ = 0, Nov = 0] = \alpha$
- *PA after:* $E[y_{ist} | NJ = 0, Nov = 1] = \alpha + \gamma$

- Assuming that $E[u_{ist} | NJ, Nov] = 0$

DID using regression* (extra, but useful for HW4)

$DID = (NJ\ after - NJ\ before) - (PA\ after - PA\ before)$

- $NJ\ after - NJ\ before = (\alpha + \beta + \gamma + \delta) - (\alpha + \beta) = \gamma + \delta$
- $PA\ after - PA\ before = (\alpha + \gamma) - \alpha = \gamma$

So, we have:

- $DID = (NJ\ after - NJ\ before) - (PA\ after - PA\ before) = \delta$

Estimating the regression model using OLS produces the DID estimate and standard errors which is very convenient

DID assumptions and testing them

Key assumption of DID

1. **Parallel trends:** The main assumption for any DID strategy is that the outcome in the treatment and control groups would follow the same time trend in the absence of treatment
   - Note: this does not mean that they must have the same mean (or level) of the outcome variable!!!

Testing the Key assumptions of DID

Parallel trends: The main assumption for any DID strategy is that the outcome in the treatment and control groups would **follow the same time trend in the absence of treatment**.
Since this counterfactual situation cannot be directly observed, support that this assumption is fulfilled is usually obtained by showing 2 things:

- **Parallel pre-trends:** Show that in the period before treatment, the two groups developed in a similar manner, i.e. followed a parallel trend. Better if we have observations from several points in time.
- **Common shocks:** Show that if other policies or changes coincided with the treatment period, these affected the treatment and control groups in the same way. Alternatively, convince reader that:
  - Nothing else (major) happens at the same time as the treatment takes place that would affect the control and treatment groups differently.

Showing support for parallel trends by looking at pre-treatment trends

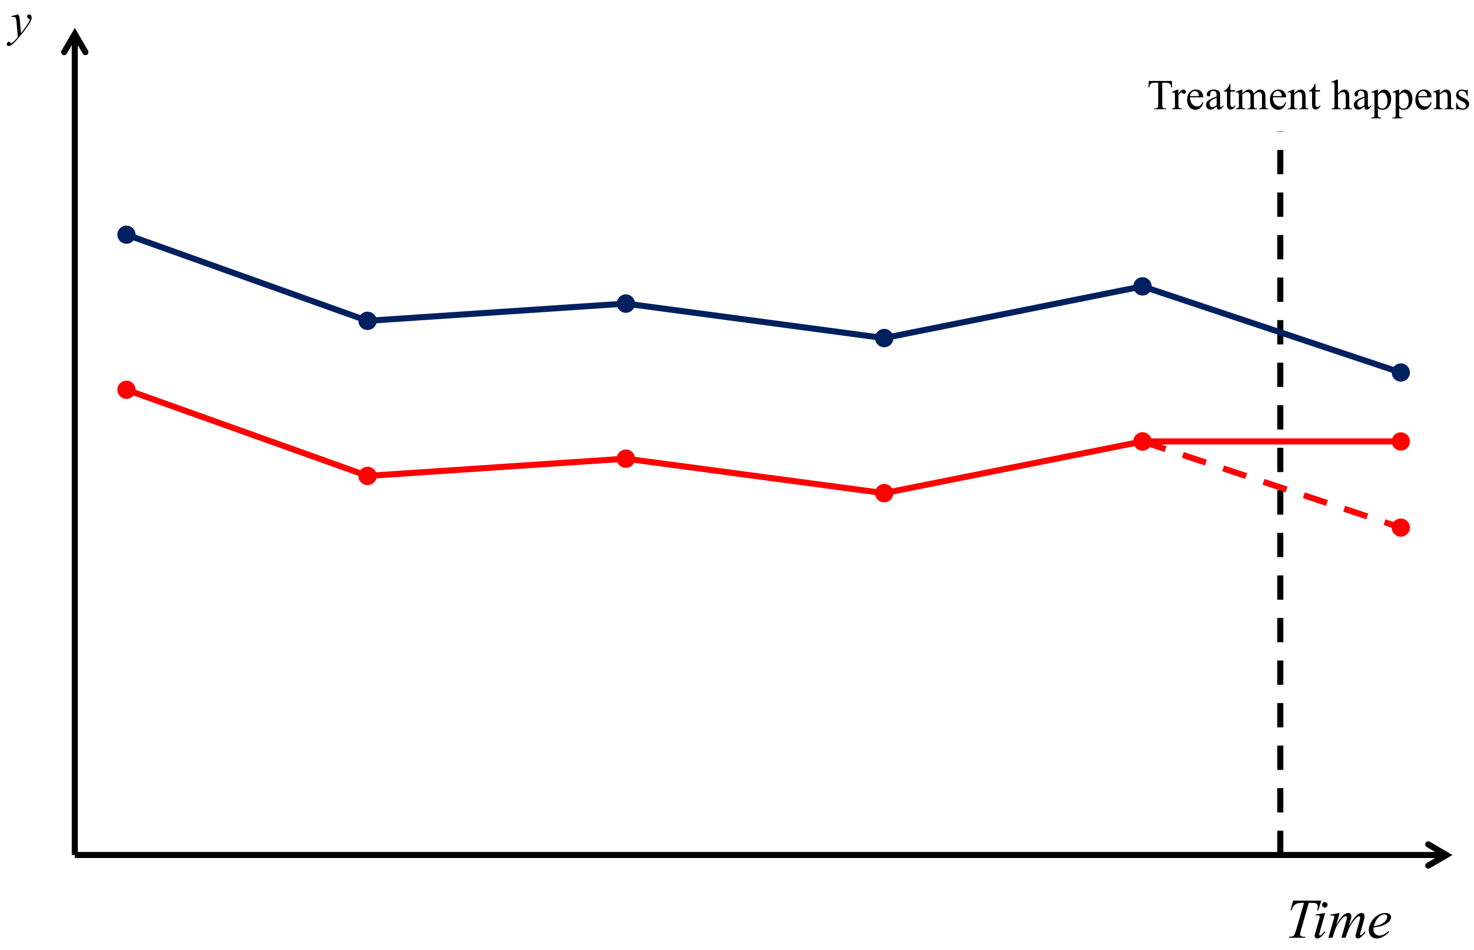

Showing support for no common shocks

- **Even if pre-trends are the same one still must worry about other policies or changes coinciding with the treatment**
  - Nothing else happens at the same time as the treatment takes place that would affect the control and treatment groups differently
- **It is very important for the researcher to be familiar with the institutional details of the reform/policy change to know:**
  - What were the macroeconomic events that took place during the studied period, and are there any worries that these affected T and C differently? E.g. a local recession in PA not affecting NJ?
  - Policies: Were there any other unemployment policies implemented in PA (the control group) during the studied period?

Additional assumptions*

Note – this is "extra" and not important in this course.

- **Another assumption is that there are no spillover effects of treatment or that group compositions do not change because of treatment (if using repeated cross-sections)**
  - In the minimum wage example, this would mean that New Jersey's minimum wage increase does not directly affect employment in Pennsylvania
- **Yet another, more technical assumption, is that** that **group composition does not change because of treatment** (i.e. people do not move disproportionally from the unaffected to the affected group, or vice versa, because of the treatment).

DID recap

- **Idea:**
  - Even if treated and control groups differ in baseline characteristics, we can use observations on treatment and control groups before and after the treatment to estimate a causal effect.
- **Assumptions:**
  - The potential outcomes (not observed) would have developed in a parallel manner for both groups in the absence of treatment.
    - *This assumption includes the “Common shocks” assumption: There can be no differential changes over time for the treated and control groups.*
- **Testing for design validity:**
  - Visualization and testing: are trends in outcomes parallel before treatment? Discuss: Is there anything else that could have happened to one group but not the other? (know your institutional setting!)