Event Studies

Matteo Paradisi
(EIEF and NBER)

Applied Micro - Lecture 6

Event Studies

- DID and its generalizations relied on a method to build Y_{i0} using info from a control group trend
- Event studies follow a similar logic
- They allow to study causal effects without a classical control group
- BUT, in some cases need a stronger identification assumption

Event Studies: Where do they come from?

- Originally, born to measure effects of an economic event on the value of firms
- Using financial market data, an event study measures the impact of a specific event on the value of a firm.
- Given rationality in the marketplace, effects of an event immediately reflected in security prices
- Measure impact using security prices observed over a short time period
- Avoid need of direct productivity measures, that may require months or years
- Used since Dolly (1933)!

Event Studies: Where do they come from?

- Finance Literature has worked with event studies in 50s and 60s
- Goal: try to "clean" measures of returns from economic fluctuations and confounding events
- Similar to the challenges that we discuss every day here!

Event Studies: Some Examples from the Past

How to build a classical event study:

- Define the event of interest, identify the event window
  - e.g. earnings announcement, the event window will include the day of the announcement
- Select outcome of interest: may be challenging in the case of stock prices
  - you want to avoid confounders
- Build a time series around the event time

Effects of Earnings Announcements (MacKinlay, 1997)

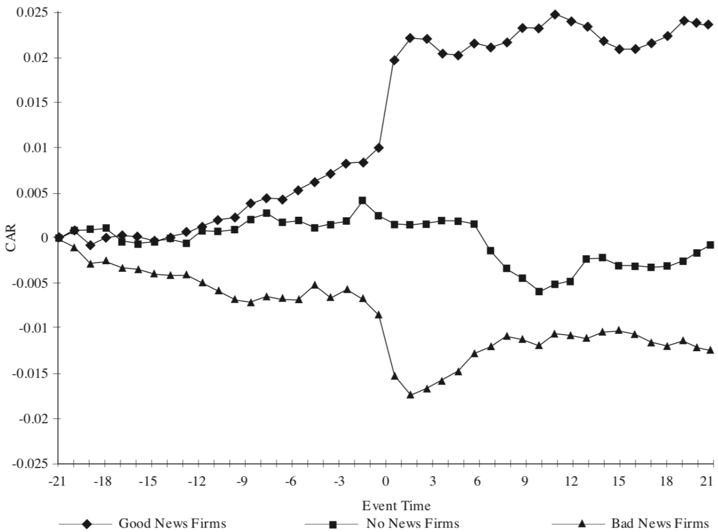

Figure 2a. Plot of cumulative abnormal return for earning announcements from event day -20 to event day 20. The abnormal return is calculated using the market model as the normal return measure.

More Recently: Kho et al. (2019)

- Kho et al. (2019) study the effect of internet access on test scores
- Exploit a program that introduced internet access in schools
- Schools enter the program in different years
- Start with simple before-after comparison in time-series

Effect of Internet Access on Math Scores

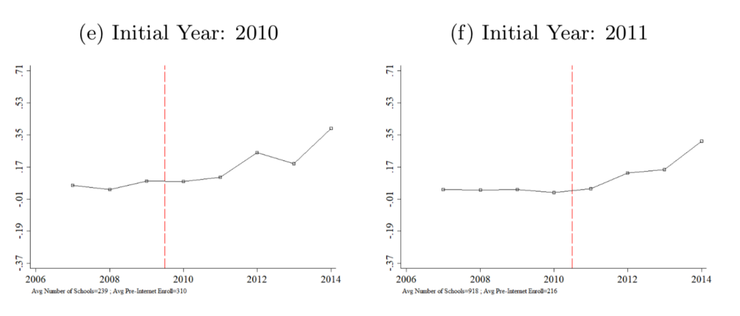

Identification

- How do we identify the causal effect here?
- Simply extrapolate the trend from the pre-event period and take deviation of post-event levels from trend
- What is the key identifying assumption?
- Assumption: trend would have been the same as the one observed before the treatment
- This assumption is obviously non-testable

Identification: Deviation From Pre-event Trend

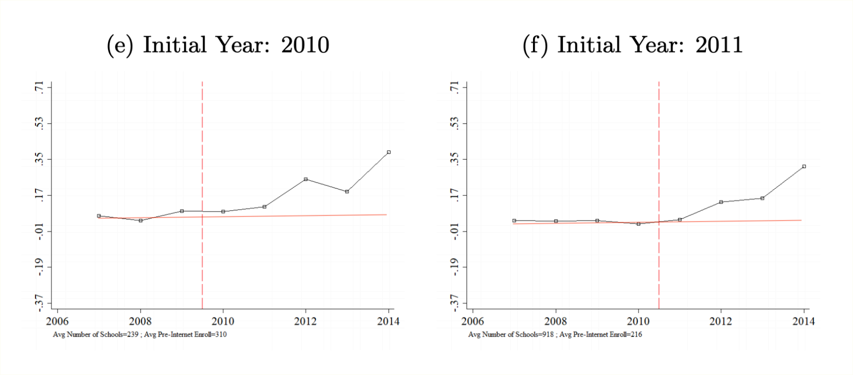

Difference from DID (with single event)

- In DID we observe a control group around the event time
- In the event study above, we don't
- Hence, the event study has a stonger assumption in this case
- We must assume that we can fit post-event trend using only pre-event info
- In DID we use control to gather info on post-event trend
- How can we achieve something similar in event study?

Generalizing the Event Study Design

Generalizing Event Study

- What if we have different units experiencing event at different times?
- Can we find a way to “pool” info on events happening at different times?
- If yes, can that allow us to relax our assumption relative to previous case?

Effects of the Economic Stimulus Payments

- Let's answer this question by studying the effects of the economic stimulus payment in 2008
- Broda and Parker (2014) study this question gathering data from the Nielsen's consumer panel
- The stimulus was paid at different moments in time
- Payments mailed to households during a nine-week period
- The week in which funds were disbursed depended on second-to-last digit of the Social Security number
- This number is effectively randomly assigned

Intuition Behind Identification

- Use those affected in the future as controls for those currently affected
- This makes the methodology more similar to a DID
- Key intuition: dichotomy between time of event and calendar time

Staggered Implementation of Cash Transfers

Figure 2: Economic stimulus payments in the NCP and as reported by the Treasury

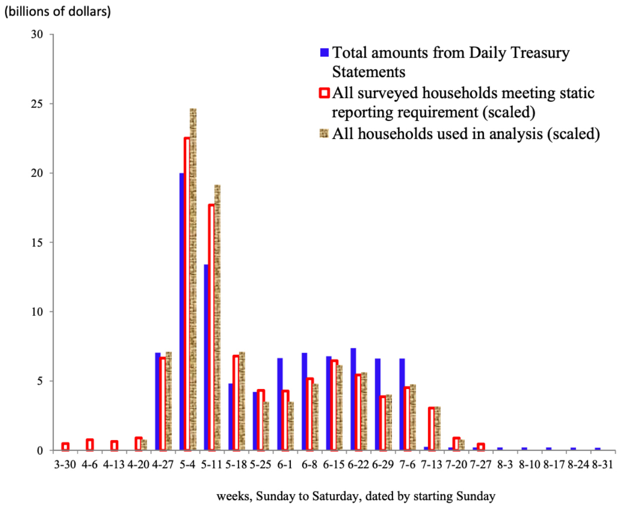

Econometric Specification

- Build time of event dummy variables
- Run the specification including these variables

$$C_{it} = \eta_i + \sum_{k=-\infty}^{+\infty} \beta_k I\left(K_{it} = k\right) + \tau_m \times \gamma_t + \varepsilon_{it}$$

- C is consumption
- τ_m is method of payment chosen, interacted with time effect γ_t
- β_ks are coefficients of interest

Econometric Specification: Discussion

$$C_{it} = \eta_i + \sum_{k=-\infty}^{+\infty} \beta_k I\left(K_{it} = k\right) + \tau_m \times \gamma_t + \varepsilon_{it}$$

- We must drop a β_k, normally drop β_{-1} and interpret in deviation from that period
- β_ks with $k < 0$ serve as placebo: we should not see any effect
- β_ks with $k > 0$ describe the effect of the cash transfer

The Effect of Cash Transfers

Table 3: The average household spending response by week

	Using all variation in time of receipt			Using only variation in timing within each method of receipt		
Regression Specification:	Dollars spent on indicator of ESP	Spending as pct of 2008Q1 spending on indicator of ESP	Dollars spent on average ESP/100	Dollars spent on indicator of ESP	Spending as pct of 2008Q1 spending on indicator of ESP	Dollars spent on average ESP/100
(Interpretation)	($ spent)	(% chg in spending)	(MPC, in %)	($ spent)	(% chg in spending)	(MPC, in %)
Week before	-0.1	(1.4)	-0.02	0.0	-0.1	0.03
	(1.8)	(1.5)	(0.21)	(2.0)	(1.7)	(0.24)
Contemporaneous week	14.0	9.9	1.55	12.9	10.7	1.47
	(2.1)	(1.8)	(0.24)	(2.3)	(2.0)	(0.27)
First week after	12.6	8.7	1.41	10.3	8.9	1.22
	(2.1)	(1.8)	(0.24)	(2.5)	(2.2)	(0.29)
Second week after	4.8	1.8	0.51	3.0	2.0	0.37
	(2.1)	(1.9)	(0.23)	(2.6)	(2.3)	(0.30)
Third week after	3.8	1.9	0.45	2.1	2.4	0.34
	(2.1)	(2.0)	(0.24)	(2.8)	(2.4)	(0.33)
Four week cumulative dollar increase or cumulative MPC	35.3		3.92	28.3		3.40
	(5.7)		(0.64)	(7.8)		(0.93)
Four week average percent increase in spending		5.6			6.0	
		(1.4)			(1.8)	

A More General Specification

- The previous specification did not control for year FE
- We can add these FE and write

$$y_{it} = \eta_i + \gamma_t + \sum_{k=-\infty}^{+\infty} \beta_k \mathbf{I}\left(K_{it} = k\right) + \varepsilon_{it}$$

- Obviously, you can always control for other individual specific trends, interacting observables with year FE
- We include all possible event time dummy variables, the model is "fully saturated"
- How do we identify β_ks here?

Graphical Intuition

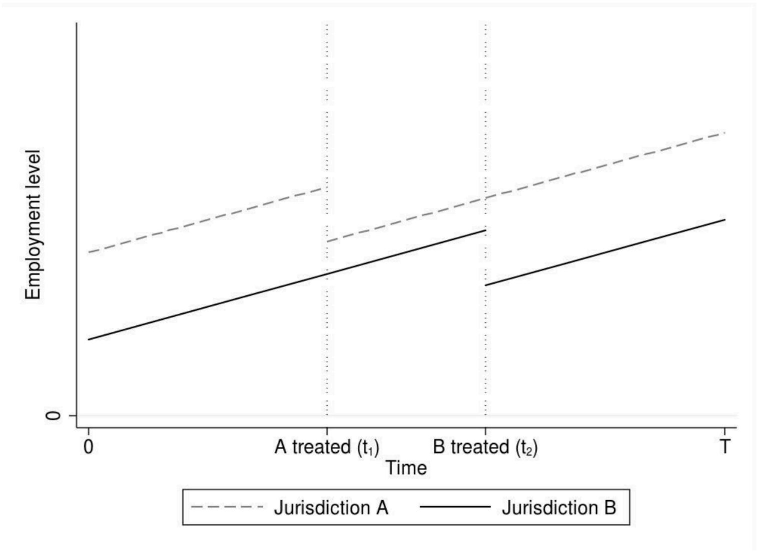

(a) Treatment effect discrete in levels

Intuition

- We can identify only β_ks in periods where we observe at least one "control"
- We use people who will eventually get treated as controls for those treated now
- How does this compare to DID?

In Practice: Kho et al. (2019)

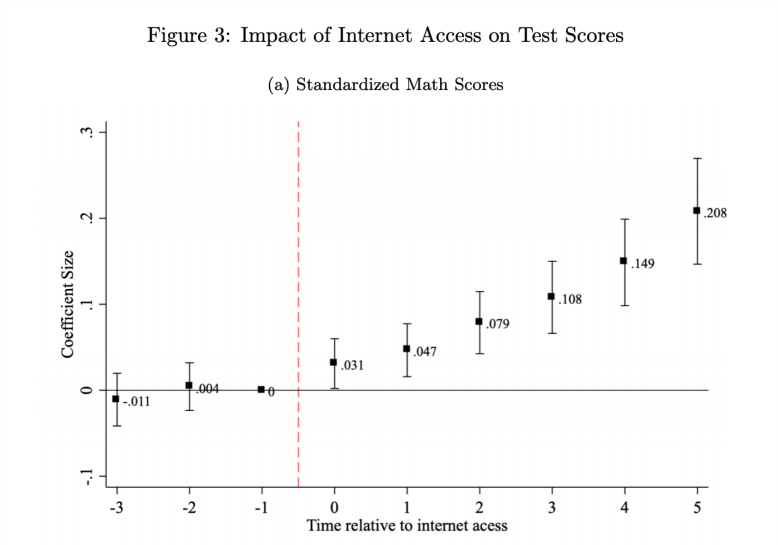

In Practice: Kho et al. (2019)

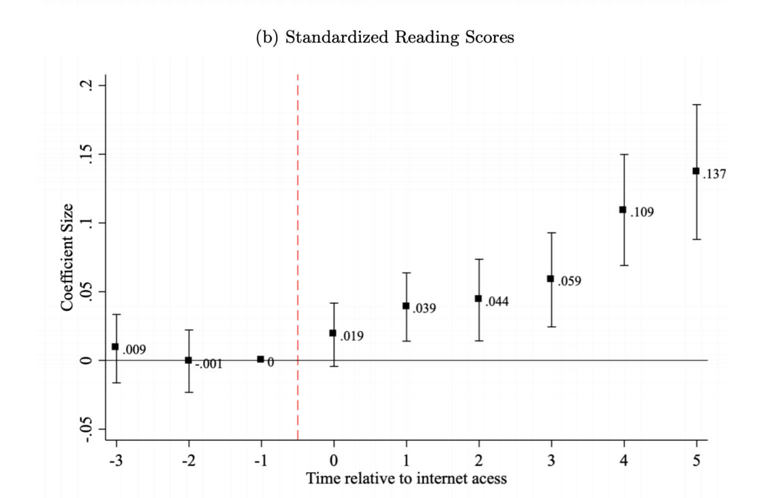

In Practice: Davis and von Wachter (2011)

Figure 5A: Average Annual Earnings Before and After Job Displacement Relative to Control Group Earnings, Men 50 or Younger with at Least 3 Years of Job Tenure

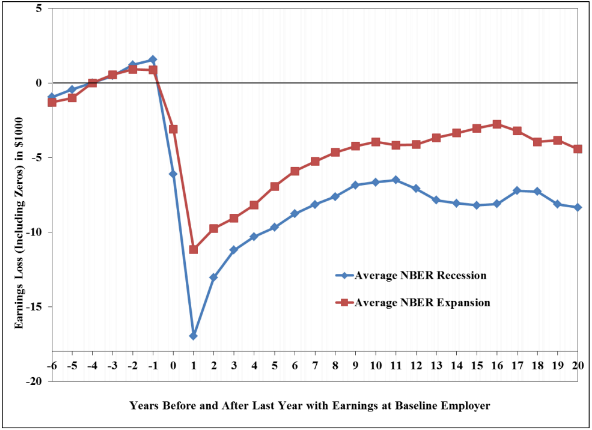

Adding Individual-Specific Time Trends

- You can extend the specification by adding individual-specific time trends
- The specification would become

$$y_{it} = \eta_i + \gamma_t + \sigma_i \times f(t) + \sum_{k=-\infty}^{+\infty} \beta_k I(K_{it} = k) + \varepsilon_{it}$$

- where $f(t)$ is a function of time, linear is one option
- Relaxes assumption: individuals can differ in trend, as long as individual-specific trend can capture this

Pitfalls of Event Studies

Event Studies: A Cautionary Tale

Some issues have been pointed out in event study implementation

1. Cannot identify β_ks in a fully saturated model with year FE (Borusyak and Jaravel, 2017)
2. If dependent variable has non-discrete change in growth, then event study has biased estimates (Meer and West, 2014)
3. A simplified specification that uses a single dummy to capture post-treatment effects implements a weird reweighting of short- and long-run treatment effects (Borusyak and Jaravel, 2017)
4. Individual-specific trends can create biases if we look at y in levels, but there is a non-discrete change in the growth of y (Meer and West, 2014)

β_ks in Fully Saturated Model with Year FE

- Borusyak and Jaravel (2017) point out that β_ks cannot be identified in

$$y_{it} = \eta_i + \gamma_t + \sum_{k=-\infty}^{+\infty} \beta_k I\left(K_{it} = k\right) + \varepsilon_{it}$$

- β_ks are identified up to a linear trend, underidentification problem!
- It means that data is equally fitted by β_k and $\tilde{\beta}_k = \beta_k + h \cdot k$
- Intuition: the time of event and calendar year go hand in hand "within individual"

β_ks in Fully Saturated Model with Year FE

- Let's see why!
- Notice that

$$
\begin{aligned}
y_{it} &= \eta_i + \gamma_t + \sum_{k=-\infty}^{+\infty} \beta_k I\left(K_{it} = k\right) + \varepsilon_{it} \\
&= \left(\eta_i + h \cdot E_i\right) + \left(\gamma_t - h \cdot t\right) + \sum_{k=-\infty}^{+\infty} \left(\beta_k + h \cdot k\right) I\left(K_{it} = k\right) + \varepsilon_{it}
\end{aligned}
$$

- This is because $t - E_i = K_{it}$

Simple Formal Intuition for Underidentification

- This problem can be thought as a problem of multicollinearity
- Take a specific example where

$$\eta_i = a + \eta E_i$$
$$\gamma_t = \gamma t$$
$$\beta_k = \beta K_{it}$$

- The specification becomes

$$\begin{aligned} y_{it} &= a + \eta E_i + \gamma t + \beta K_{it} + \varepsilon_{it} \\ &= a + \eta E_i + \gamma t + \beta t - \beta E_i + \varepsilon_{it} \end{aligned}$$

- In DID control group pins down the year FE!

A Similar Problem

- Similar to age-cohort-time problem in the regression

$$y_{it} = \underbrace{\eta_{E_i}}_{\text{Cohort FE}} + \underbrace{\gamma_t}_{\text{Time FE}} + \underbrace{\beta_{t-E_i}}_{\text{Age FE}} + \varepsilon_{it}$$

- Where E_i is date of birth

An Easy Solution

- Borusyak and Jaravel (2017) propose some solutions
- An easy solution is what many papers do in practice
- Bin some of the event years
- e.g. bin all $k < \underline{k}$ and $k > \bar{k}$ in two dummy variables
- Then "within dummy" you can identify the year FE

Discrete Change in y growth

- Suppose you are interested in growth of y
- Then issues may arise when different event times
- Change in growth creates a bias in control group
- Easy to visualize in a graph

Discrete Change in y growth

Bias in control above t_2

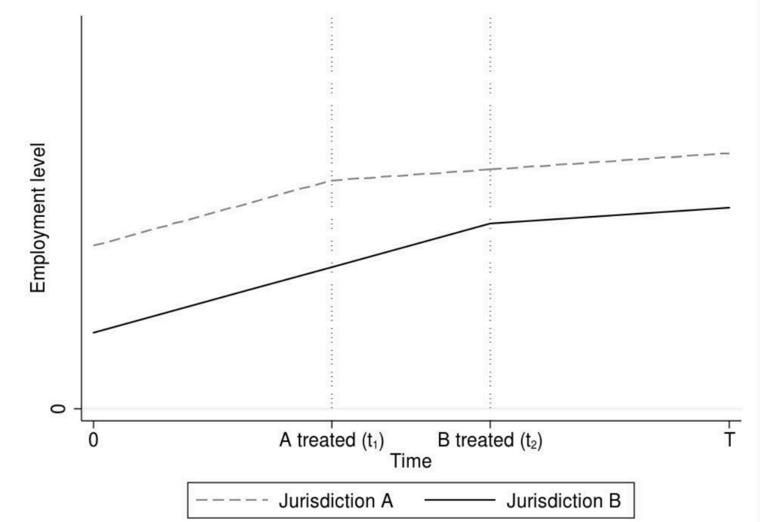

(b) Treatment effect discrete in growth

Discrete Change in y growth

- Cengiz et al. (2019) implement one possible solution
- They run a series of separate DID for every event time
- In each DID they have a control group build with those who are not affected at that time
- Then need a way to aggregate across event times (Abraham and Sun, 2019)

Problem With Simplified Specification

- Suppose we run the more parsimonious model

$$y_{it} = \eta_i + \gamma_t + \beta D_{it} + \varepsilon_{it}$$

- where $D_{it} = 1$ if $t > E_i$
- A single coefficient summarizes the treatment effect
- More parametric than generalized specification
- This specification generates some problems with heterogenous treatment effects over time
- Overweights short-run, underweights long-run

Problem With Simplified Specification

Table 1: Treatment Status in the Minimal Examples

Panel A

Period \ Group	A	B
0	0	0
1	0	1

Panel B

Period \ Group	A	B
0	0	0
1	0	1
2	1	1

- **Suppose treatment effect constant:** β
- **Panel A:**

$$(y_{B1} - y_{A1}) - (y_{B0} - y_{A0}) = (\eta_B + \beta - \eta_A) - (\eta_B - \eta_A) = \beta$$

- **Panel B:**

$$(y_{B1} - y_{A1}) - (y_{B2} - y_{A2}) = (\eta_B + \beta - \eta_A) - (\eta_B + \beta - \eta_A - \beta) = \beta$$

Problem With Simplified Specification

Table 1: Treatment Status in the Minimal Examples

Panel A

Period \ Group	A	B
0	0	0
1	0	1

Panel B

Period \ Group	A	B
0	0	0
1	0	1
2	1	1

- Suppose treatment effect not constant: β_S, β_L
- Panel A:

$$(y_{B1} - y_{A1}) - (y_{B0} - y_{A0}) = (\eta_B + \beta_S - \eta_A) - (\eta_B - \eta_A) = \beta_S$$

- Panel B:

$$\begin{aligned}(y_{B1} - y_{A1}) - (y_{B2} - y_{A2}) &= (\eta_B + \beta_S - \eta_A) - (\eta_B + \beta_L - \eta_A - \beta_S) \\ &= 2\beta_S - \beta_L\end{aligned}$$

Solution

- Estimate non-parametric model
- Take post-event coefficients and compute average
- Show average to summarize treatment effect in one number

Problem With Individual-Specific Trends

Trend may create bias in treatment group if change in growth

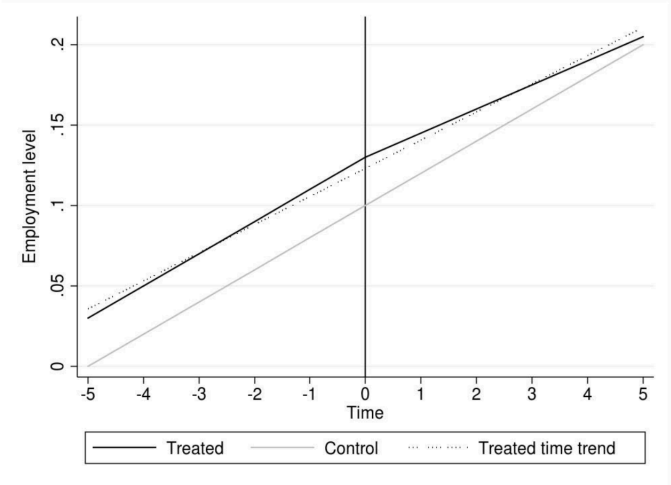

Figure 3: Simple example of disemployment effect in growth rate

Problem With Individual-Specific Trends

The problem arises when y is in levels

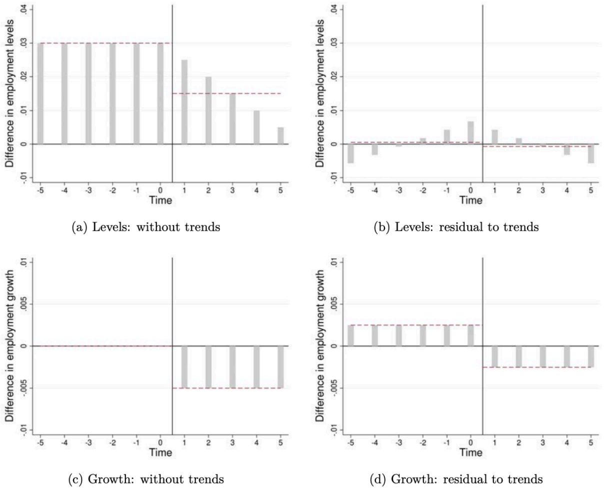

(a) Levels: without trends

(b) Levels: residual to trends

(c) Growth: without trends

(d) Growth: residual to trends

Figure 4: Example difference-in-differences without versus with jurisdiction time trends

Problem With Individual-Specific Trends

- Individual-specific trends relax identification assumption
- They may help your identification strategy
- Hence, there is a trade-off when including these trends

Possible Solution (Gruber Jensen, Kleven, 2019)

- Estimate individual-specific trends in the pre-period only for any group g

$$Y_{it}^{g} = \theta^{g} t$$

- Residualize dependent variable using estimated trend

$$\tilde{Y}_{it}^{g} = Y_{it}^{g} - \hat{\theta}^{g} t$$

- This "corrects" trends in post-period
- However, a standard error correction is needed!
- Probably bootstrap can work