Difference-in-differences

Guidance and examples for using code repository to conduct GIEs.

Topics include:

- Difference-in-difference
- Pre-trends analysis
- Event study

AIDDATA
A Research Lab at William & Mary

Difference-in-Differences (DID)

- Quasi-experimental method of causal identification
- Construct counterfactual for treatment group using time trends of control group
- Treatment effect is the difference between the treatment and control groups in the difference in outcomes over time
- The treatment and control groups must have parallel trends in the outcome variable in the absence of treatment

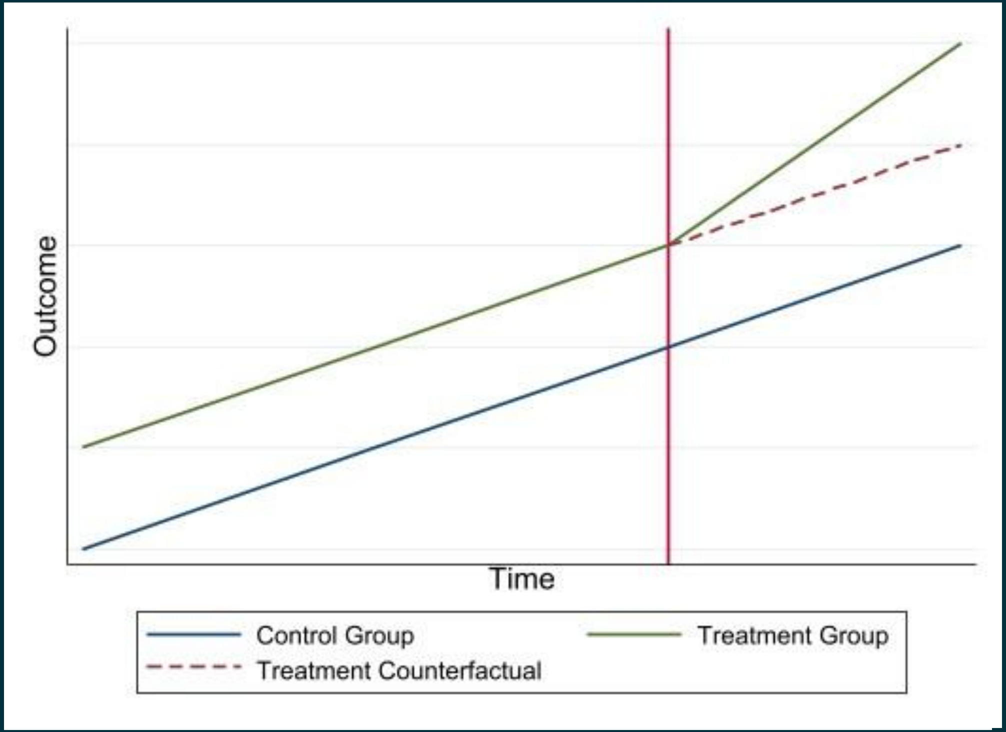

AIDDATA
A Research Lab at William & Mary

Single Treatment Time

- A canonical difference-in-difference has the following elements:
  - One treatment time
  - A treatment group
  - A control group
  - Observations before and after the treatment time
  - Repeated cross-section or panel data

Single Treatment Time

- The treatment effect is the difference between the treatment and control groups in the difference in outcomes over time

- $Y_{it} = \alpha + \beta_1{*}After_t + \beta_2{*}TreatmentGroup_i + \beta_3{*}After_t{*}TreatmentGroup_i +$

  $+ \delta{*}TimeFE_t + \theta{*}GeospatialFE_i + \epsilon_{it}$

  - $After_t$ is a binary variable indicating if observed after the time of treatment
  - $TreatmentGroup_i$ is a binary variable indicating if *i* is in the treatment group
  - β_3 is the treatment effect

Single Treatment Time

Code output

Example: Consider the effect of an irrigation project on child stunting, wasting, and anemia

AIDDATA
A Research Lab at William & Mary

- The code runs the difference-in-difference regression
- Outputs:
  - Log file with code output
  - Formatted regression table (Word, Excel, or LaTeX)

VARIABLES	(1) Height-for-Age
After	0.274***
	(0.0331)
Treatment Group	-0.0226
	(0.0220)
Treatment Group * After	0.811***
	(0.0369)
Observations	37,393
R-squared	0.082

Robust standard errors in parentheses
*** $p<0.01$, ** $p<0.05$, * $p<0.1$
SEs clustered two-way by cluster and wave.

  - Labeled coefficient plot of regression (jpg, png, svg, or pdf)

Difference-in-Difference (DID) Single Treatment Time

Code output

- The code runs the difference-in-difference regression
- Outputs:
  - Log file with code output
  - Formatted regression table (Word, Excel, or LaTeX)
  - Labeled coefficient plot of regression (jpg, png, svg, or pdf)

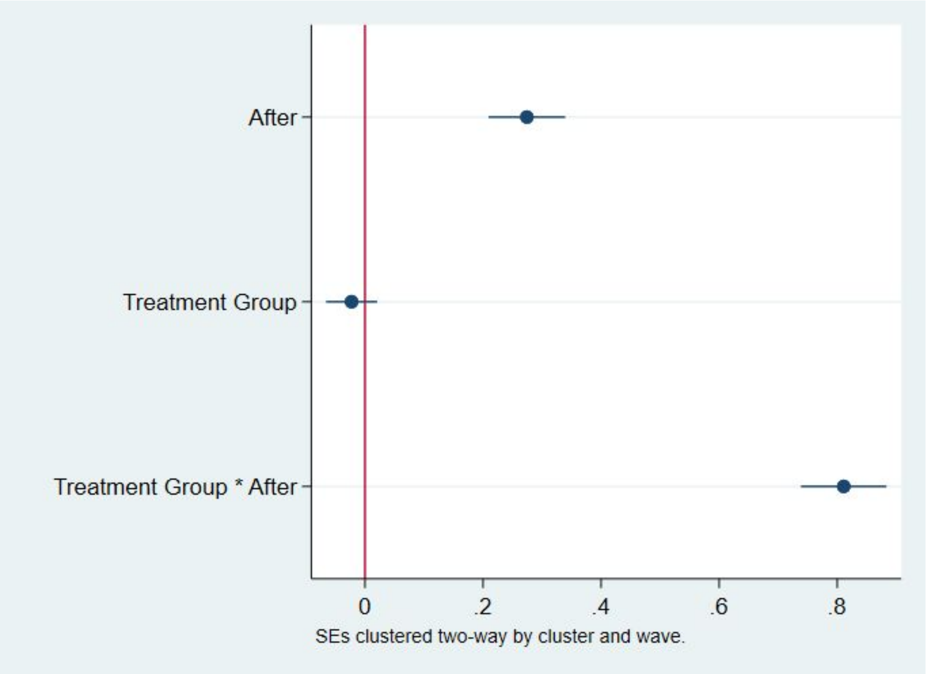

AIDDATA
A Research Lab at William & Mary

Pre-Trends

- Difference-in-Difference can only be used to estimate the treatment effect if the parallel trends assumption holds
  - The treatment and control group must have parallel time trends in the absence of treatment
- A pre-trends analysis tests this assumption
  - Considers if there was a parallel time trend between the two groups prior to the time of treatment

Difference-in-Difference (DID) Single Treatment Time Pre-Trends

Code output

- The code runs the pre-trend analysis
- Outputs:
    - Log file with code output
    - Labeled line graph of raw pre-trends (jpg, png, svg, or pdf)

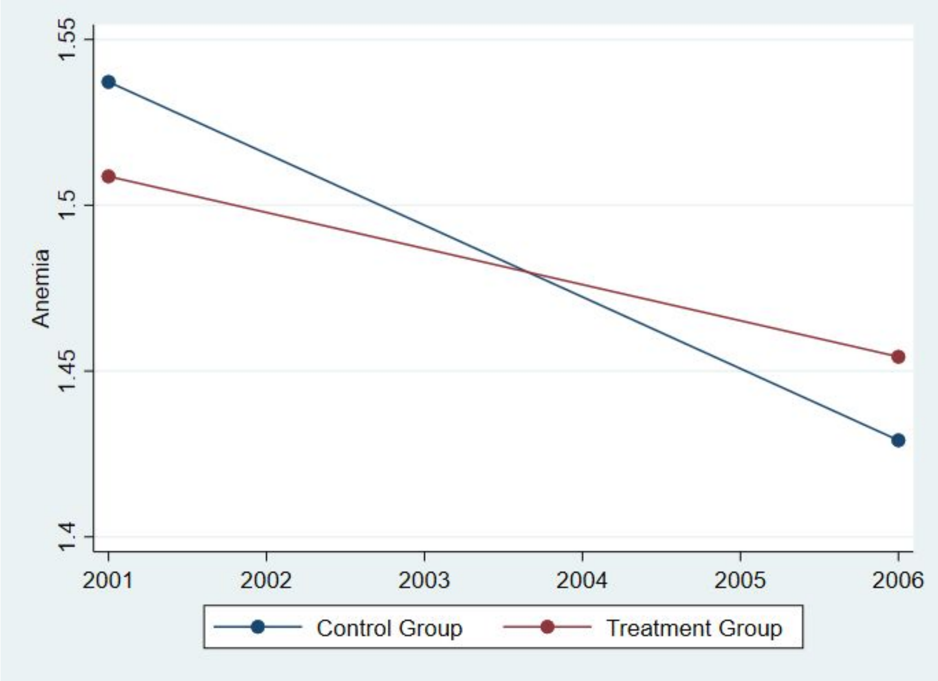

    - Labeled connected coefficient plot of pre-trends regression (jpg, png, svg, or pdf)

AIDDATA
A Research Lab at William & Mary

Difference-in-Difference (DID) Single Treatment Time Pre-Trends

Code output

- The code runs the pre-trend analysis
- Outputs:
  - Log file with code output
  - Labeled line graph of raw pre-trends (jpg, png, svg, or pdf)
  - Labeled connected coefficient plot of pre-trends regression (jpg, png, svg, or pdf)

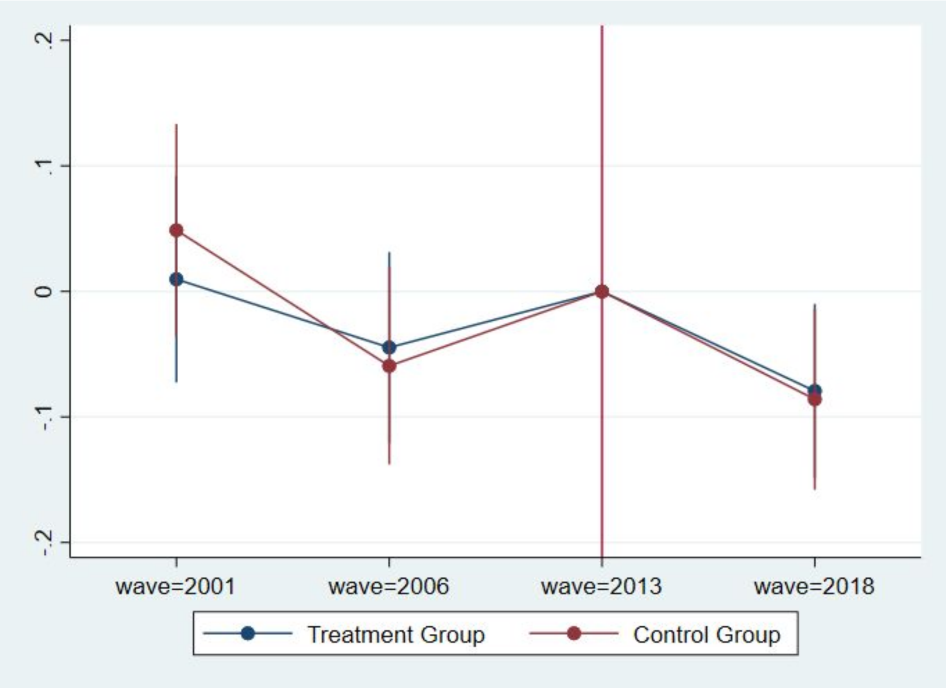

AIDDATA
A Research Lab at William & Mary

Multiple Treatment Times

- Multiple, staggered treatment times
- "Roll-out" designs

- Variation in treatment derives from only the treatment group (no true control group)

- In this approach, members of the treatment group serve as members of the control group, depending on when they are being observed

- Identifying assumption of single treatment time:
    - *Spatial* allocation of treatment not correlated with changes in outcomes

- Identifying assumption of multiple treatment times:
    - *Spatio-temporal* allocation of treatment not correlated with changes in outcomes

Multiple Treatment Times: Two-way Fixed Effects

- Time fixed effects and geospatial fixed effects are used to align treatment time

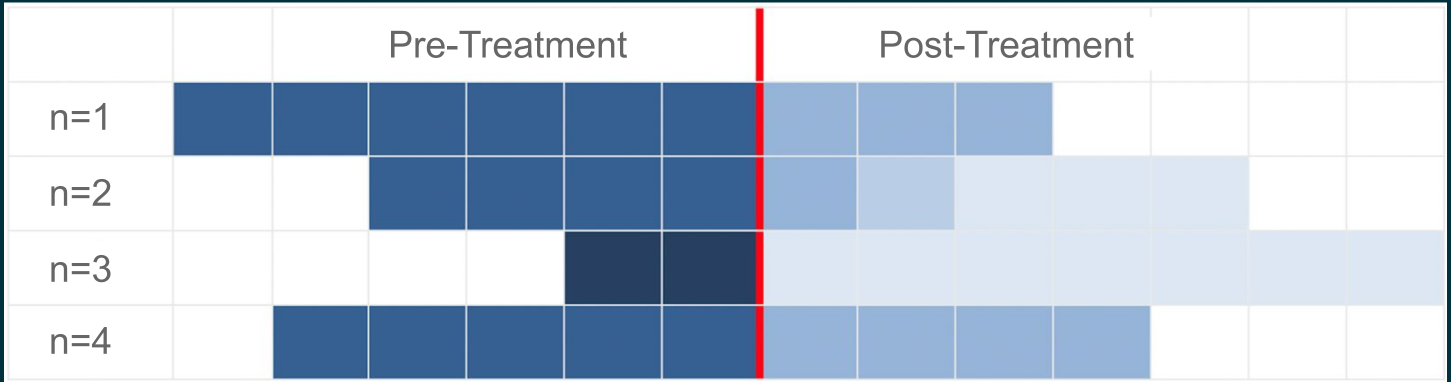

- $Y_{it} = \alpha + \square{*}Treated_{it} + \delta{*}TimeFE_t + \theta{*}GeospatialFE_i + \epsilon_{it}$
  - $Treated_i$ is a binary variable indicating if *i* was observed after it was treated
  - $\square$ is the treatment effect

AIDDATA
A Research Lab at William & Mary

Multiple Treatment Times

Code output

- The code runs the difference-in-difference regression
- Outputs:
  - Log file with code output
  - Formatted regression table (Word, Excel, or LaTeX)

VARIABLES	(1) NDVI
Treated	0.00580***
	(0.00111)
Observations	17,858
R-squared	0.743

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1
Includes grid-cell and region fixed effects. SEs two-way clustered by grid-cell and year.

  - Labeled coefficient plot of regression (jpg, png, svg, or pdf)

AIDDATA
A Research Lab at William & Mary

Recent literature on Multiple Treatment Times & TWFE

- The comparisons we "want":
  - Already-treated units vs. not-yet-treated units
- But standard TWFE model also includes other comparisons we don't typically think about, some of which may be problematic under some circumstances
  - Especially when the treatment effects vary over time
  - This is very common in many agriculture and environmental projects, when impacts build over time as people, markets, etc. respond gradually to interventions

Figure 1. Difference-in-Differences with Variation in Treatment Timing: Three Groups

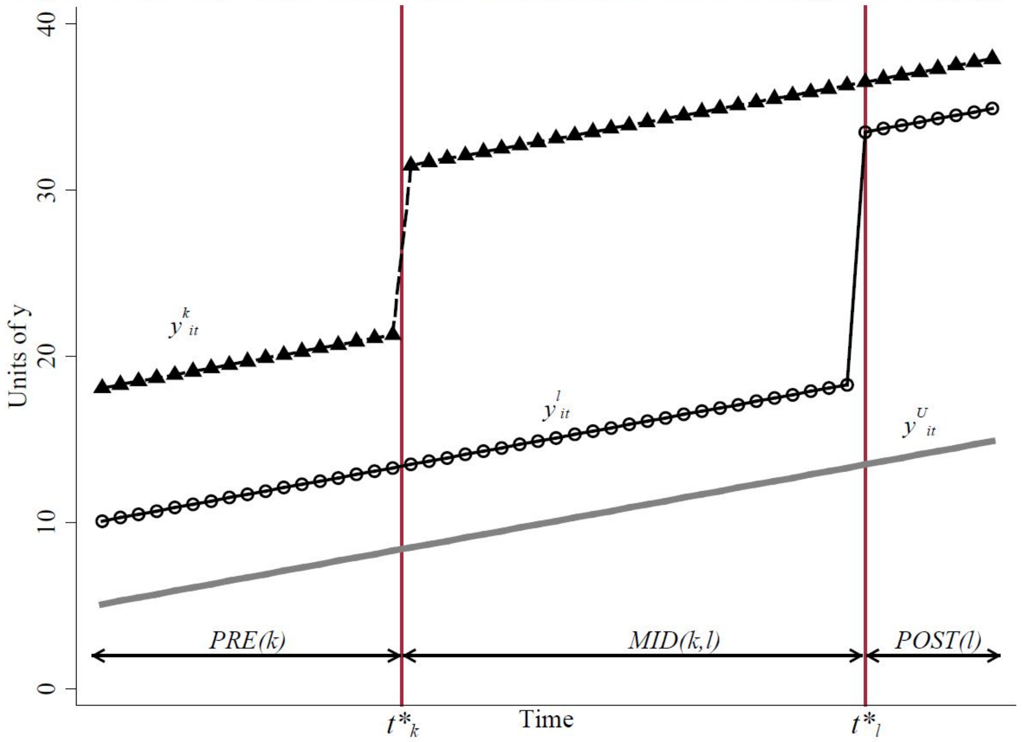

Notes: The figure plots outcomes in three groups: a control group, U, which is never treated; an early treatment group, E, which receives a binary treatment at $t_k^* = \frac{34}{100}T$; and a late treatment group, ℓ, which receives the binary treatment at $t_\ell^* = \frac{85}{100}T$. The x-axis notes the three sub-periods: the pre-period for group k, $[1, t_k^* - 1]$, denoted by $PRE(k)$; the middle period when group k is treated and group ℓ is not, $[t_k^*, t_\ell^* - 1]$, denoted by $MID(k, \ell)$; and the post-period for group ℓ, $[t_\ell^*, T]$, denoted by $POST(\ell)$. I set the treatment effect to 10 in group k and 15 in group ℓ.

Goodman-Bacon, Andrew. "Difference-in-differences with variation in treatment timing." *Journal of Econometrics* 225, no. 2 (2021): 254-277.

Figure 2. The Four Simple (2x2) Difference-in-Differences Estimates from the Three Group Case

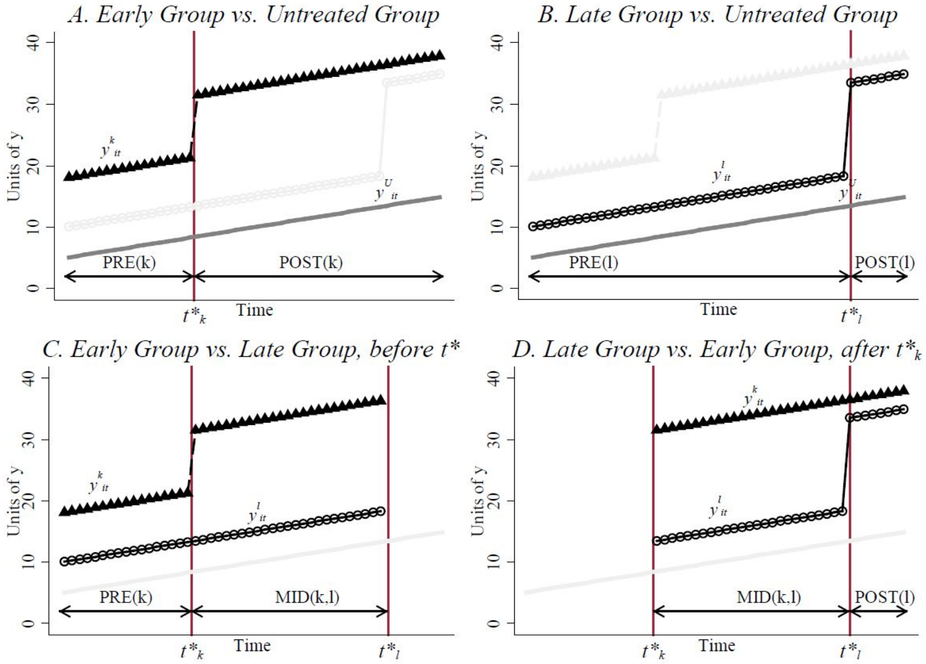

Notes: The figure plots the groups and time periods that generate the four simple 2x2 difference-in-difference estimates in the case with an early treatment group, a late treatment group, and an untreated group from Figure 1. Each panel plots the data structure for one 2x2 DD. Panel A compares early treated units to untreated units ($\hat{\beta}^{DD}_{kU}$); panel B compares late treated units to untreated units ($\hat{\beta}^{DD}_{\ell U}$); panel C compares early treated units to late treated units during the late group's pre-period ($\hat{\beta}^{DD,k}_{k\ell}$); panel D compares late treated units to early treated units during the early group's post-period ($\hat{\beta}^{DD,\ell}_{k\ell}$). The treatment times mean that $\overline{D}_k = 0.67$ and $\overline{D}_\ell = 0.16$, so with equal group sizes, the decomposition weights on the 2x2 estimate from each panel are 0.365 for panel A, 0.222 for panel B, 0.278 for panel C, and 0.135 for panel D.

Goodman-Bacon, Andrew. "Difference-in-differences with variation in treatment timing." *Journal of Econometrics* 225, no. 2 (2021): 254-277.

If no untreated group, we only have bottom two comparisons (C + D)

Figure 3. Difference-in-Differences Estimates with Variation in Timing Are Biased When Treatment Effects Vary Over Time

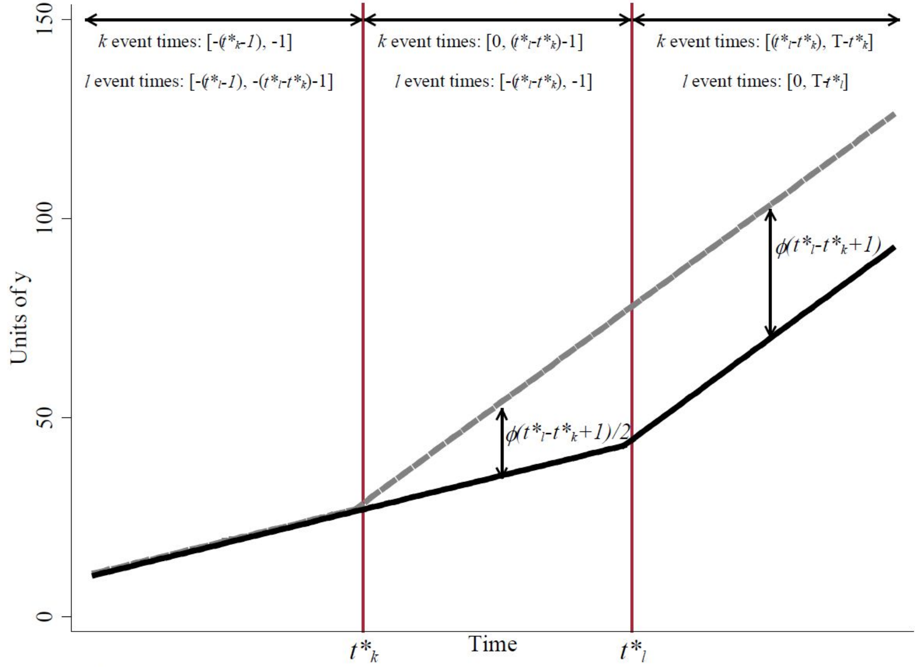

Notes: The figure plots a stylized example of a timing-only DD set up with a treatment effect that is a trend-break rather than a level shift (see Meer and West 2013). Following section II.A.ii, the trend-break effect equals $\phi \cdot (t - t^* + 1)$. The top of the figure notes which event-times lie in the $PRE(k)$, $MID(k,\ell)$, and $POST(\ell)$ periods for each unit. The figure also notes the average difference between groups in each of these periods. In the $MID(k,\ell)$ period, outcomes differ by $\frac{\phi}{2}(t^*_\ell - t^*_k + 1)$ on average. In the $POST(\ell)$ period, however, outcomes had already been growing in the early group for $t^*_\ell - t^*_k$ periods, and so they differ by $\phi(t^*_\ell - t^*_k + 1)$ on average. The 2x2 DD that compares the later group to the earlier group is biased and, in the linear trend-break case, weakly negative despite a positive and growing treatment effect.

Goodman-Bacon, Andrew. "Difference-in-differences with variation in treatment timing." *Journal of Econometrics* 225, no. 2 (2021): 254-277.

Problems arise when treatment effects vary over time!

Recent literature on Multiple Treatment Times & TWFE

- Fortunately, a slew of new estimators to "fix" this, each with different approach
- Goodman-Bacon (2021) offers a decomposition of sample across these windows to see how much of variation comes from each comparison
- Callaway & Sant'Anna (2021): use only not-yet-treated as comparisons
- Sun & Abraham (2021) use last-to-be-treated as comparisons
- Borusyak, Jaravel and Spiess (2021) create "imputation" estimator predicting counterfactuals from trends among not-yet-treated
- De Chaisemartin & d'Haultfoeuille (2020) use only not-yet-treated as comparisons but allow treatment to turn on and off
- Many others!

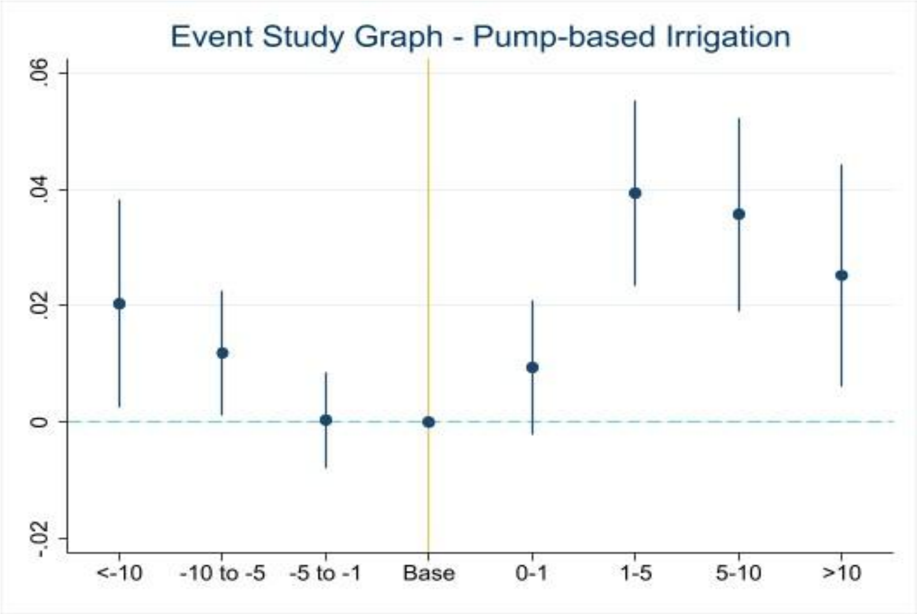

Standard TWFE

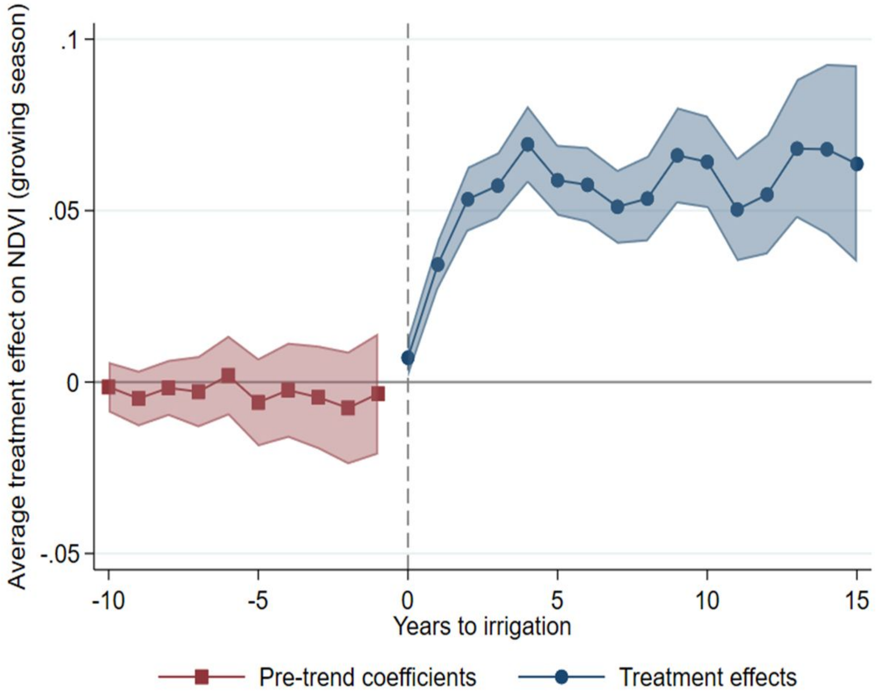

Borusyak et al (2021)

AIDDATA
A Research Lab at William & Mary

Main Takeaways

- "Cumulative" treatment effects very common

- Testing pre-trends either via
  - Window prior to *any* treatment, or
  - Use TWFE adjustments

- Other identification questions remain:
  - At what level is the spatiotemporal allocation as-good-as-random?
  - How to think about spillovers across units (especially since only-treated samples might be more clustered)

AIDDATA
A Research Lab at William & Mary