Panel Data II: Difference-in-Differences

Francisco Villamil

Applied Quantitative Methods II

MA in Social Sciences, Spring 2026

Today's goals

- Understand Difference-in-Differences as an extension of TWFE

Today's goals

- Understand Difference-in-Differences as an extension of TWFE
- Learn the 2×2 DiD setup and the parallel trends assumption

Today's goals

- Understand Difference-in-Differences as an extension of TWFE
- Learn the 2$\times$2 DiD setup and the parallel trends assumption
- Interpret the DiD regression coefficient

Today's goals

- Understand Difference-in-Differences as an extension of TWFE
- Learn the $2{\times}2$ DiD setup and the parallel trends assumption
- Interpret the DiD regression coefficient
- Use event studies to visualize dynamic effects and test pre-trends

Today's goals

- Understand Difference-in-Differences as an extension of TWFE
- Learn the 2×2 DiD setup and the parallel trends assumption
- Interpret the DiD regression coefficient
- Use event studies to visualize dynamic effects and test pre-trends
- Understand the staggered DiD problem and modern solutions

Roadmap

Recap: what fixed effects does

$$y_{it} = \alpha_i + \gamma_t + \beta x_{it} + \varepsilon_{it}$$

- α_i: absorbs all **time-invariant** unit differences

Recap: what fixed effects does

$$y_{it} = \alpha_i + \gamma_t + \beta x_{it} + \varepsilon_{it}$$

- α_i: absorbs all **time-invariant** unit differences
- γ_t: absorbs all **unit-invariant** time shocks

Recap: what fixed effects does

$$y_{it} = \alpha_i + \gamma_t + \beta x_{it} + \varepsilon_{it}$$

- α_i: absorbs all **time-invariant** unit differences
- γ_t: absorbs all **unit-invariant** time shocks
- $\hat{\beta}$: identified from within-unit variation, net of common trends

Recap: what fixed effects does

$$y_{it} = \alpha_i + \gamma_t + \beta x_{it} + \varepsilon_{it}$$

- α_i: absorbs all **time-invariant** unit differences
- γ_t: absorbs all **unit-invariant** time shocks
- $\hat{\beta}$: identified from within-unit variation, net of common trends
- But TWFE treats x_{it} as a continuous variable that varies continuously over time

Recap: what fixed effects does

$$y_{it} = \alpha_i + \gamma_t + \beta x_{it} + \varepsilon_{it}$$

- α_i: absorbs all **time-invariant** unit differences
- γ_t: absorbs all **unit-invariant** time shocks
- $\hat{\beta}$: identified from within-unit variation, net of common trends

- But TWFE treats x_{it} as a continuous variable that varies continuously over time
- What if the variation comes from a **specific, discrete intervention**?

A new question: policy interventions

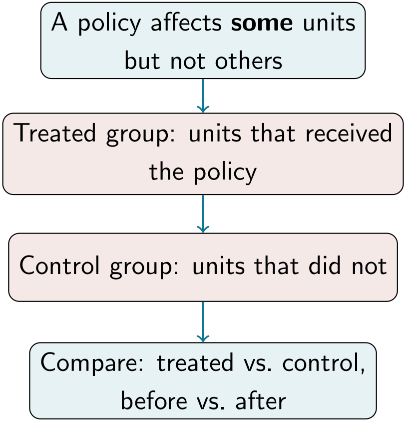

Running example: Card & Krueger (1994)

- **Question**: Does raising the minimum wage reduce employment?

Running example: Card & Krueger (1994)

- **Question**: Does raising the minimum wage reduce employment?
- **Policy**: NJ raised the minimum wage on April 1, 1992 — from $4.25 to $5.05 per hour

Running example: Card & Krueger (1994)

- **Question**: Does raising the minimum wage reduce employment?
- **Policy**: NJ raised the minimum wage on April 1, 1992 — from \$4.25 to \$5.05 per hour
- **Design**:

Running example: Card & Krueger (1994)

- **Question**: Does raising the minimum wage reduce employment?
- **Policy**: NJ raised the minimum wage on April 1, 1992 — from \$4.25 to \$5.05 per hour
- **Design**:
  - $\rightarrow$ **Treatment**: fast-food restaurants in New Jersey (NJ)

Running example: Card & Krueger (1994)

- **Question**: Does raising the minimum wage reduce employment?
- **Policy**: NJ raised the minimum wage on April 1, 1992 — from \$4.25 to \$5.05 per hour
- **Design**:
  - → **Treatment**: fast-food restaurants in New Jersey (NJ)
  - → **Control**: fast-food restaurants in Pennsylvania (PA)

Running example: Card & Krueger (1994)

- **Question**: Does raising the minimum wage reduce employment?
- **Policy**: NJ raised the minimum wage on April 1, 1992 — from \$4.25 to \$5.05 per hour
- **Design**:
  - → **Treatment**: fast-food restaurants in New Jersey (NJ)
  - → **Control**: fast-food restaurants in Pennsylvania (PA)
  - → **Before**: Feb–Mar 1992; **After**: Nov–Dec 1992

Running example: Card & Krueger (1994)

- **Question**: Does raising the minimum wage reduce employment?
- **Policy**: NJ raised the minimum wage on April 1, 1992 — from \$4.25 to \$5.05 per hour
- **Design**:
  - → **Treatment**: fast-food restaurants in New Jersey (NJ)
  - → **Control**: fast-food restaurants in Pennsylvania (PA)
  - → **Before**: Feb–Mar 1992; **After**: Nov–Dec 1992
- Pennsylvania did not change its minimum wage

The 2×2 DiD table

	Before	**After**
Control (PA)	$\bar{y}_{C,pre}$	$\bar{y}_{C,post}$
Treated (NJ)	$\bar{y}_{T,pre}$	$\bar{y}_{T,post}$

$$\hat{\delta}_{\text{DiD}} = (\bar{y}_{T,post} - \bar{y}_{T,pre}) - (\bar{y}_{C,post} - \bar{y}_{C,pre})$$

- Control group change: what would have happened **without** treatment

The 2×2 DiD table

	Before	**After**
Control (PA)	$\bar{y}_{C,pre}$	$\bar{y}_{C,post}$
Treated (NJ)	$\bar{y}_{T,pre}$	$\bar{y}_{T,post}$

$$\hat{\delta}_{\text{DiD}} = (\bar{y}_{T,post} - \bar{y}_{T,pre}) - (\bar{y}_{C,post} - \bar{y}_{C,pre})$$

- Control group change: what would have happened **without** treatment
- DiD subtracts this "counterfactual trend" from the treated group change

Roadmap

The DiD formula

$$\hat{\delta} = \underbrace{(\bar{y}_{T,post} - \bar{y}_{T,pre})}_{\text{treated group change}} - \underbrace{(\bar{y}_{C,post} - \bar{y}_{C,pre})}_{\text{control group change}}$$

- **Treated change**: NJ employment went from 20.4 to 21.0 FTEs per restaurant (+0.59)

The DiD formula

$$\hat{\delta} = \underbrace{(\bar{y}_{T,post} - \bar{y}_{T,pre})}_{\text{treated group change}} - \underbrace{(\bar{y}_{C,post} - \bar{y}_{C,pre})}_{\text{control group change}}$$

- **Treated change**: NJ employment went from 20.4 to 21.0 FTEs per restaurant ($+0.59$)
- **Control change**: PA employment went from 23.3 to 21.2 FTEs per restaurant (-2.17)

The DiD formula

$$\hat{\delta} = \underbrace{(\bar{y}_{T,post} - \bar{y}_{T,pre})}_{\text{treated group change}} - \underbrace{(\bar{y}_{C,post} - \bar{y}_{C,pre})}_{\text{control group change}}$$

- **Treated change**: NJ employment went from 20.4 to 21.0 FTEs per restaurant (+0.59)
- **Control change**: PA employment went from 23.3 to 21.2 FTEs per restaurant (−2.17)
- $\hat{\delta} = 0.59 - (-2.17) = +2.76$

The DiD formula

$$\hat{\delta} = \underbrace{(\bar{y}_{T,post} - \bar{y}_{T,pre})}_{\text{treated group change}} - \underbrace{(\bar{y}_{C,post} - \bar{y}_{C,pre})}_{\text{control group change}}$$

- **Treated change**: NJ employment went from 20.4 to 21.0 FTEs per restaurant ($+0.59$)
- **Control change**: PA employment went from 23.3 to 21.2 FTEs per restaurant (-2.17)

- $\hat{\delta} = 0.59 - (-2.17) = +2.76$

- Interpretation: the minimum wage **raised** NJ employment by $\approx$2.76 FTEs relative to PA

The regression formulation

$$y_{it} = \alpha + \beta_1 \underbrace{Post_t}_{\text{time}} + \beta_2 \underbrace{Treat_i}_{\text{group}} + \delta \underbrace{(Post_t \times Treat_i)}_{\text{interaction = DiD}} + \varepsilon_{it}$$

- $Post_t = 1$ if observation is in the post-treatment period

The regression formulation

$$y_{it} = \alpha + \beta_1 \underbrace{Post_t}_{\text{time}} + \beta_2 \underbrace{Treat_i}_{\text{group}} + \delta \underbrace{(Post_t \times Treat_i)}_{\text{interaction} = \text{DiD}} + \varepsilon_{it}$$

- $Post_t = 1$ if observation is in the post-treatment period
- $Treat_i = 1$ if unit is in the treatment group

The regression formulation

$$y_{it} = \alpha + \beta_1 \underbrace{Post_t}_{\text{time}} + \beta_2 \underbrace{Treat_i}_{\text{group}} + \delta \underbrace{(Post_t \times Treat_i)}_{\text{interaction = DiD}} + \varepsilon_{it}$$

- $Post_t = 1$ if observation is in the post-treatment period
- $Treat_i = 1$ if unit is in the treatment group
- $\delta =$ **DiD coefficient of interest**: causal effect of treatment

The regression formulation

$$y_{it} = \alpha + \beta_1 \underbrace{Post_t}_{\text{time}} + \beta_2 \underbrace{Treat_i}_{\text{group}} + \delta \underbrace{(Post_t \times Treat_i)}_{\text{interaction} = \text{DiD}} + \varepsilon_{it}$$

- $Post_t = 1$ if observation is in the post-treatment period
- $Treat_i = 1$ if unit is in the treatment group
- $\delta =$ **DiD coefficient of interest**: causal effect of treatment
- α: control group mean, pre-period

The regression formulation

$$y_{it} = \alpha + \beta_1 \underbrace{Post_t}_{\text{time}} + \beta_2 \underbrace{Treat_i}_{\text{group}} + \delta \underbrace{(Post_t \times Treat_i)}_{\text{interaction = DiD}} + \varepsilon_{it}$$

- $Post_t = 1$ if observation is in the post-treatment period
- $Treat_i = 1$ if unit is in the treatment group
- δ = **DiD coefficient of interest**: causal effect of treatment

- α: control group mean, pre-period
- β_1: time trend (common to both groups)

The regression formulation

$$y_{it} = \alpha + \beta_1 \underbrace{Post_t}_{\text{time}} + \beta_2 \underbrace{Treat_i}_{\text{group}}$$
$$+ \delta \underbrace{(Post_t \times Treat_i)}_{\text{interaction = DiD}} + \varepsilon_{it}$$

- $Post_t = 1$ if observation is in the post-treatment period
- $Treat_i = 1$ if unit is in the treatment group
- $\delta =$ **DiD coefficient of interest**: causal effect of treatment

- α: control group mean, pre-period
- β_1: time trend (common to both groups)
- β_2: pre-period difference between groups

Regression: recovering the 2×2 cells

	Before ($Post = 0$)	**After** ($Post = 1$)
Control ($Treat = 0$)	α	$\alpha + \beta_1$
Treated ($Treat = 1$)	$\alpha + \beta_2$	$\alpha + \beta_1 + \beta_2 + \delta$

- DiD from table: $((\alpha + \beta_1 + \beta_2 + \delta) - (\alpha + \beta_2)) - ((\alpha + \beta_1) - \alpha)$

Regression: recovering the 2×2 cells

	Before ($Post = 0$)	**After** ($Post = 1$)
Control ($Treat = 0$)	α	$\alpha + \beta_1$
Treated ($Treat = 1$)	$\alpha + \beta_2$	$\alpha + \beta_1 + \beta_2 + \delta$

- DiD from table: $((\alpha + \beta_1 + \beta_2 + \delta) - (\alpha + \beta_2)) - ((\alpha + \beta_1) - \alpha)$
- $= (\beta_1 + \delta) - \beta_1 = \delta$ ✓

Regression: recovering the 2×2 cells

	Before $(\mathit{Post} = 0)$	**After** $(\mathit{Post} = 1)$
Control $(\mathit{Treat} = 0)$	α	$\alpha + \beta_1$
Treated $(\mathit{Treat} = 1)$	$\alpha + \beta_2$	$\alpha + \beta_1 + \beta_2 + \delta$

- DiD from table: $((\alpha + \beta_1 + \beta_2 + \delta) - (\alpha + \beta_2)) - ((\alpha + \beta_1) - \alpha)$
- $= (\beta_1 + \delta) - \beta_1 = \delta$ ✓
- The regression recovers $\hat{\delta}$ automatically via OLS

DiD in R with `feols`

- Basic 2×2 DiD (two groups, two periods):

DiD in R with `feols`

- Basic 2×2 DiD (two groups, two periods):

```
feols(employment ~ post * treat,
```

DiD in R with `feols`

- Basic 2×2 DiD (two groups, two periods):

```
feols(employment ~ post * treat,
  data = df, cluster = ~state)
```

DiD in R with `feols`

- Basic 2×2 DiD (two groups, two periods):

```
feols(employment ~ post * treat,
  data = df, cluster = ~state)
```

- With multiple units and periods, use TWFE:

DiD in R with `feols`

- Basic 2×2 DiD (two groups, two periods):

```
feols(employment ~ post * treat,
  data = df, cluster = ~state)
```

- With multiple units and periods, use TWFE:

```
feols(employment ~ treated | unit + time,
```

DiD in R with `feols`

- Basic 2×2 DiD (two groups, two periods):

```
feols(employment ~ post * treat,
  data = df, cluster = ~state)
```

- With multiple units and periods, use TWFE:

```
feols(employment ~ treated | unit + time,
  data = df, cluster = ~unit)
```

DiD in R with `feols`

- Basic 2×2 DiD (two groups, two periods):

```
feols(employment ~ post * treat,
  data = df, cluster = ~state)
```

- With multiple units and periods, use TWFE:

```
feols(employment ~ treated | unit + time,
  data = df, cluster = ~unit)
```

- `treated` $= Post_t \times Treat_i$: a single binary indicator

DiD in R with `feols`

- Basic 2×2 DiD (two groups, two periods):

```
feols(employment ~ post * treat,
  data = df, cluster = ~state)
```

- With multiple units and periods, use TWFE:

```
feols(employment ~ treated | unit + time,
  data = df, cluster = ~unit)
```

- `treated` $= Post_t \times Treat_i$: a single binary indicator
- Unit and time FE replace the $Treat_i$ and $Post_t$ dummies

DiD in R with `feols`

- Basic 2×2 DiD (two groups, two periods):

```
feols(employment ~ post * treat,
  data = df, cluster = ~state)
```

- With multiple units and periods, use TWFE:

```
feols(employment ~ treated | unit + time,
  data = df, cluster = ~unit)
```

- `treated` $= Post_t \times Treat_i$: a single binary indicator
- Unit and time FE replace the $Treat_i$ and $Post_t$ dummies
- Always cluster standard errors at the **treatment level**

Roadmap

The key assumption: parallel trends

Parallel Trends Assumption

In the absence of treatment, the treated and control groups would have followed the **same time trend**.

- What this allows: use the control group's post-period change as the **counterfactual trend** for the treated group

The key assumption: parallel trends

Parallel Trends Assumption

In the absence of treatment, the treated and control groups would have followed the **same time trend**.

- What this allows: use the control group's post-period change as the **counterfactual trend** for the treated group
- What this does **not** require:

The key assumption: parallel trends

Parallel Trends Assumption

In the absence of treatment, the treated and control groups would have followed the **same time trend**.

- What this allows: use the control group's post-period change as the **counterfactual trend** for the treated group
- What this does **not** require:
  - $\rightarrow$ The two groups need not have the same baseline level

The key assumption: parallel trends

Parallel Trends Assumption

In the absence of treatment, the treated and control groups would have followed the **same time trend**.

- What this allows: use the control group's post-period change as the **counterfactual trend** for the treated group
- What this does **not** require:
  - → The two groups need not have the same baseline level
  - → Pre-period outcomes can differ — just not the trends

The key assumption: parallel trends

Parallel Trends Assumption

In the absence of treatment, the treated and control groups would have followed the **same time trend**.

- What this allows: use the control group's post-period change as the **counterfactual trend** for the treated group
- What this does **not** require:
  - → The two groups need not have the same baseline level
  - → Pre-period outcomes can differ — just not the trends
- This is an **assumption**, not a testable fact
  (we cannot observe the counterfactual)

Parallel trends: visualization

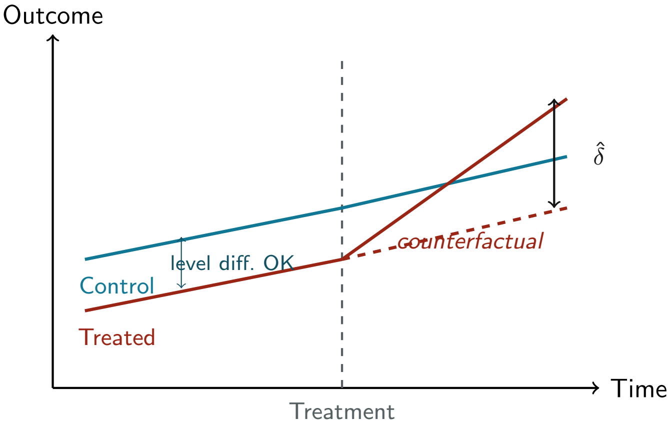

What can violate parallel trends?

- **Selection into treatment correlated with trends**

What can violate parallel trends?

- **Selection into treatment correlated with trends**
  - $\rightarrow$ Units that adopt the policy were already on different trajectories

What can violate parallel trends?

- **Selection into treatment correlated with trends**
  - → Units that adopt the policy were already on different trajectories
  - → Example: states with rising employment pre-trend more likely to adopt minimum wage

What can violate parallel trends?

- **Selection into treatment correlated with trends**
  - → Units that adopt the policy were already on different trajectories
  - → Example: states with rising employment pre-trend more likely to adopt minimum wage
- **Simultaneous confounders**

What can violate parallel trends?

- **Selection into treatment correlated with trends**
  - → Units that adopt the policy were already on different trajectories
  - → Example: states with rising employment pre-trend more likely to adopt minimum wage
- **Simultaneous confounders**
  - → Another event coincides with the treatment

What can violate parallel trends?

- **Selection into treatment correlated with trends**
  - → Units that adopt the policy were already on different trajectories
  - → Example: states with rising employment pre-trend more likely to adopt minimum wage
- **Simultaneous confounders**
  - → Another event coincides with the treatment
  - → Example: NJ also changed another labor market policy in 1992

What can violate parallel trends?

- **Selection into treatment correlated with trends**
  - → Units that adopt the policy were already on different trajectories
  - → Example: states with rising employment pre-trend more likely to adopt minimum wage
- **Simultaneous confounders**
  - → Another event coincides with the treatment
  - → Example: NJ also changed another labor market policy in 1992
- **Anticipation effects**

What can violate parallel trends?

- **Selection into treatment correlated with trends**
  - → Units that adopt the policy were already on different trajectories
  - → Example: states with rising employment pre-trend more likely to adopt minimum wage
- **Simultaneous confounders**
  - → Another event coincides with the treatment
  - → Example: NJ also changed another labor market policy in 1992
- **Anticipation effects**
  - → Treated units change behavior before the official treatment date

What can violate parallel trends?

- **Selection into treatment correlated with trends**
  - → Units that adopt the policy were already on different trajectories
  - → Example: states with rising employment pre-trend more likely to adopt minimum wage
- **Simultaneous confounders**
  - → Another event coincides with the treatment
  - → Example: NJ also changed another labor market policy in 1992
- **Anticipation effects**
  - → Treated units change behavior before the official treatment date
  - → Example: firms start hiring/firing when the wage increase is announced

The parallel trends assumption says treatment would not have changed the outcome trajectory.

When is this plausible?

What makes NJ and PA a good comparison?

Roadmap

Beyond the single post-period

- The 2×2 DiD uses one pre and one post period

Beyond the single post-period

- The 2×2 DiD uses one pre and one post period
- With panel data spanning many periods, we can do **more**:

Beyond the single post-period

- The 2$\times$2 DiD uses one pre and one post period
- With panel data spanning many periods, we can do **more**:
  - $\rightarrow$ Estimate treatment effects at each time point relative to treatment

Beyond the single post-period

- The 2$\times$2 DiD uses one pre and one post period
- With panel data spanning many periods, we can do **more**:
  - $\rightarrow$ Estimate treatment effects at each time point relative to treatment
  - $\rightarrow$ Trace out the **dynamic path** of the effect

Beyond the single post-period

- The 2×2 DiD uses one pre and one post period
- With panel data spanning many periods, we can do **more**:
  - → Estimate treatment effects at each time point relative to treatment
  - → Trace out the **dynamic path** of the effect
- Two purposes:

Beyond the single post-period

- The 2$\times$2 DiD uses one pre and one post period
- With panel data spanning many periods, we can do **more**:
  - $\rightarrow$ Estimate treatment effects at each time point relative to treatment
  - $\rightarrow$ Trace out the **dynamic path** of the effect
- Two purposes:
  - $\rightarrow$ **Pre-trend test**: were pre-treatment trends actually parallel?

Beyond the single post-period

- The 2×2 DiD uses one pre and one post period
- With panel data spanning many periods, we can do **more**:
  - → Estimate treatment effects at each time point relative to treatment
  - → Trace out the **dynamic path** of the effect
- Two purposes:
  - → **Pre-trend test**: were pre-treatment trends actually parallel?
  - → **Dynamic effects**: does the effect grow, fade, or remain constant?

Beyond the single post-period

- The 2×2 DiD uses one pre and one post period
- With panel data spanning many periods, we can do **more**:
  - → Estimate treatment effects at each time point relative to treatment
  - → Trace out the **dynamic path** of the effect
- Two purposes:
  - → **Pre-trend test**: were pre-treatment trends actually parallel?
  - → **Dynamic effects**: does the effect grow, fade, or remain constant?
- This is the **event study** design

Event study regression

$$y_{it} = \alpha_i + \gamma_t + \sum_{k \neq -1} \delta_k \cdot \mathbf{1}[t - T_i^* = k] + \varepsilon_{it}$$

- T_i^*: treatment date for unit i

Event study regression

$$y_{it} = \alpha_i + \gamma_t + \sum_{k \neq -1} \delta_k \cdot \mathbf{1}[t - T_i^* = k] + \varepsilon_{it}$$

- T_i^*: treatment date for unit i
- $k = t - T_i^*$: time relative to treatment ($k < 0$: pre, $k \geq 0$: post)

Event study regression

$$y_{it} = \alpha_i + \gamma_t + \sum_{k \neq -1} \delta_k \cdot \mathbf{1}[t - T_i^* = k] + \varepsilon_{it}$$

- T_i^*: treatment date for unit i
- $k = t - T_i^*$: time relative to treatment ($k < 0$: pre, $k \geq 0$: post)
- $k = -1$: **omitted reference period** (normalization)

Event study regression

$$y_{it} = \alpha_i + \gamma_t + \sum_{k \neq -1} \delta_k \cdot \mathbf{1}[t - T_i^* = k] + \varepsilon_{it}$$

- T_i^*: treatment date for unit i
- $k = t - T_i^*$: time relative to treatment ($k < 0$: pre, $k \geq 0$: post)
- $k = -1$: **omitted reference period** (normalization)

- **Pre-treatment** ($k < 0$): $\delta_k \approx 0$ if parallel trends holds

Event study regression

$$y_{it} = \alpha_i + \gamma_t + \sum_{k \neq -1} \delta_k \cdot \mathbf{1}[t - T_i^* = k] + \varepsilon_{it}$$

- T_i^*: treatment date for unit i
- $k = t - T_i^*$: time relative to treatment ($k < 0$: pre, $k \geq 0$: post)
- $k = -1$: **omitted reference period** (normalization)

- **Pre-treatment** ($k < 0$): $\delta_k \approx 0$ if parallel trends holds
- **Post-treatment** ($k \geq 0$): trace out dynamic causal effects

Event study regression

$$y_{it} = \alpha_i + \gamma_t + \sum_{k \neq -1} \delta_k \cdot \mathbf{1}[t - T_i^* = k] + \varepsilon_{it}$$

- T_i^*: treatment date for unit i
- $k = t - T_i^*$: time relative to treatment ($k < 0$: pre, $k \geq 0$: post)
- $k = -1$: **omitted reference period** (normalization)

- **Pre-treatment** ($k < 0$): $\delta_k \approx 0$ if parallel trends holds
- **Post-treatment** ($k \geq 0$): trace out dynamic causal effects

- Unit and time FE absorb levels; δ_k captures relative changes

Event study: coefficient plot

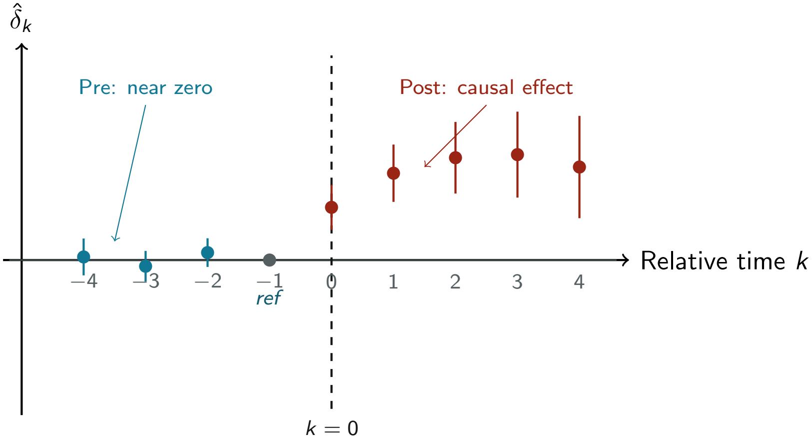

Event study in R

- Using `feols()` with `i()` for interaction syntax:

Event study in R

- Using `feols()` with `i()` for interaction syntax:

```
feols(y ~ i(rel_time, treated, ref = -1)
```

Event study in R

- Using `feols()` with `i()` for interaction syntax:

```
feols(y ~ i(rel_time, treated, ref = -1)
  | unit + year, data = df, cluster = ~unit)
```

Event study in R

- Using `feols()` with `i()` for interaction syntax:

```
feols(y ~ i(rel_time, treated, ref = -1)
  | unit + year, data = df, cluster = ~unit)
```

- `rel_time`: variable with $k = t - T_i^*$ (relative time to treatment)

Event study in R

- Using `feols()` with `i()` for interaction syntax:

```
feols(y ~ i(rel_time, treated, ref = -1)
  | unit + year, data = df, cluster = ~unit)
```

- `rel_time`: variable with $k = t - T_i^*$ (relative time to treatment)
- `treated`: binary indicator for treated units

Event study in R

- Using `feols()` with `i()` for interaction syntax:

```
feols(y ~ i(rel_time, treated, ref = -1)
   | unit + year, data = df, cluster = ~unit)
```

- `rel_time`: variable with $k = t - T_i^*$ (relative time to treatment)
- `treated`: binary indicator for treated units
- `ref = -1`: omit $k = -1$ as reference period

Event study in R

- Using `feols()` with `i()` for interaction syntax:

```
feols(y ~ i(rel_time, treated, ref = -1)
   | unit + year, data = df, cluster = ~unit)
```

- `rel_time`: variable with $k = t - T_i^*$ (relative time to treatment)
- `treated`: binary indicator for treated units
- `ref = -1`: omit $k = -1$ as reference period

- Plot with `iplot()` (coefficient plot with CIs):

Event study in R

- Using `feols()` with `i()` for interaction syntax:

```
feols(y ~ i(rel_time, treated, ref = -1)
   | unit + year, data = df, cluster = ~unit)
```

- `rel_time`: variable with $k = t - T_i^*$ (relative time to treatment)
- `treated`: binary indicator for treated units
- `ref = -1`: omit $k = -1$ as reference period

- Plot with `iplot()` (coefficient plot with CIs):

```
iplot(m_es, main = "Event Study")
```

Event study in R

- Using `feols()` with `i()` for interaction syntax:

```
feols(y ~ i(rel_time, treated, ref = -1)
   | unit + year, data = df, cluster = ~unit)
```

- `rel_time`: variable with $k = t - T_i^*$ (relative time to treatment)
- `treated`: binary indicator for treated units
- `ref = -1`: omit $k = -1$ as reference period

- Plot with `iplot()` (coefficient plot with CIs):

```
iplot(m_es, main = "Event Study")
```

- `fixest` makes event studies very easy to run and plot

Reading an event study plot

- **Pre-trend test**: are pre-treatment $\hat{\delta}_k$ jointly near zero?

Reading an event study plot

- **Pre-trend test**: are pre-treatment $\hat{\delta}_k$ jointly near zero?
  - $\rightarrow$ If yes: supports parallel trends (not proof, but reassuring)

Reading an event study plot

- **Pre-trend test**: are pre-treatment $\hat{\delta}_k$ jointly near zero?
  - $\rightarrow$ If yes: supports parallel trends (not proof, but reassuring)
  - $\rightarrow$ If no: parallel trends likely violated — DiD is not credible

Reading an event study plot

- **Pre-trend test**: are pre-treatment $\hat{\delta}_k$ jointly near zero?
  - $\rightarrow$ If yes: supports parallel trends (not proof, but reassuring)
  - $\rightarrow$ If no: parallel trends likely violated — DiD is not credible
- **Dynamic effects**: what do post-treatment $\hat{\delta}_k$ look like?

Reading an event study plot

- **Pre-trend test**: are pre-treatment $\hat{\delta}_k$ jointly near zero?
  - $\rightarrow$ If yes: supports parallel trends (not proof, but reassuring)
  - $\rightarrow$ If no: parallel trends likely violated — DiD is not credible
- **Dynamic effects**: what do post-treatment $\hat{\delta}_k$ look like?
  - $\rightarrow$ Immediate jump at $k = 0$: effect appears right away

Reading an event study plot

- **Pre-trend test**: are pre-treatment $\hat{\delta}_k$ jointly near zero?
  - $\rightarrow$ If yes: supports parallel trends (not proof, but reassuring)
  - $\rightarrow$ If no: parallel trends likely violated — DiD is not credible
- **Dynamic effects**: what do post-treatment $\hat{\delta}_k$ look like?
  - $\rightarrow$ Immediate jump at $k = 0$: effect appears right away
  - $\rightarrow$ Gradual build-up: effect grows over time (lagged adjustment)

Reading an event study plot

- **Pre-trend test**: are pre-treatment $\hat{\delta}_k$ jointly near zero?
  - $\rightarrow$ If yes: supports parallel trends (not proof, but reassuring)
  - $\rightarrow$ If no: parallel trends likely violated — DiD is not credible
- **Dynamic effects**: what do post-treatment $\hat{\delta}_k$ look like?
  - $\rightarrow$ Immediate jump at $k = 0$: effect appears right away
  - $\rightarrow$ Gradual build-up: effect grows over time (lagged adjustment)
  - $\rightarrow$ Decay: initial effect fades (people adapt, spillovers, attrition)

Reading an event study plot

- **Pre-trend test**: are pre-treatment $\hat{\delta}_k$ jointly near zero?
  - $\rightarrow$ If yes: supports parallel trends (not proof, but reassuring)
  - $\rightarrow$ If no: parallel trends likely violated — DiD is not credible
- **Dynamic effects**: what do post-treatment $\hat{\delta}_k$ look like?
  - $\rightarrow$ Immediate jump at $k = 0$: effect appears right away
  - $\rightarrow$ Gradual build-up: effect grows over time (lagged adjustment)
  - $\rightarrow$ Decay: initial effect fades (people adapt, spillovers, attrition)
- **Symmetry**: sometimes useful to check behavior far pre-treatment

Roadmap

Staggered treatment adoption

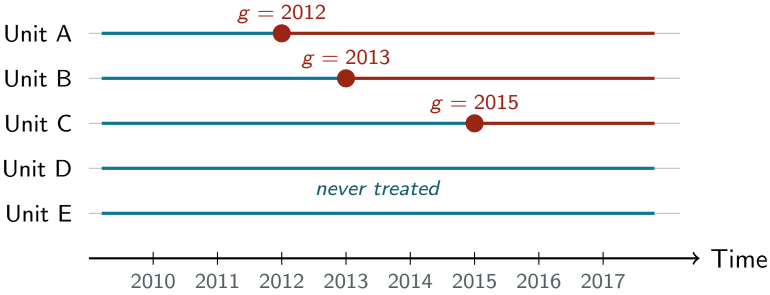

The problem with naive TWFE in staggered settings

- With staggered timing, TWFE uses **all available 2$\times$2 comparisons**

The problem with naive TWFE in staggered settings

- With staggered timing, TWFE uses **all available 2$\times$2 comparisons**
- Including: already-treated units as the **control group** for later-treated units

The problem with naive TWFE in staggered settings

- With staggered timing, TWFE uses **all available 2×2 comparisons**
- Including: already-treated units as the **control group** for later-treated units
  - → Unit adopted in 2012 becomes "control" for unit adopted in 2015

The problem with naive TWFE in staggered settings

- With staggered timing, TWFE uses **all available 2×2 comparisons**
- Including: already-treated units as the **control group** for later-treated units
  - → Unit adopted in 2012 becomes "control" for unit adopted in 2015
- This is the "**forbidden comparison**" problem

The problem with naive TWFE in staggered settings

- With staggered timing, TWFE uses **all available 2×2 comparisons**
- Including: already-treated units as the **control group** for later-treated units
  - → Unit adopted in 2012 becomes "control" for unit adopted in 2015
- This is the "**forbidden comparison**" problem
  - → Already-treated units are not a clean control

The problem with naive TWFE in staggered settings

- With staggered timing, TWFE uses **all available 2×2 comparisons**
- Including: already-treated units as the **control group** for later-treated units
  - → Unit adopted in 2012 becomes "control" for unit adopted in 2015
- This is the "**forbidden comparison**" problem
  - → Already-treated units are not a clean control
  - → If treatment effects vary over time, their included effect contaminates the estimate

The problem with naive TWFE in staggered settings

- With staggered timing, TWFE uses **all available 2×2 comparisons**
- Including: already-treated units as the **control group** for later-treated units
  - → Unit adopted in 2012 becomes "control" for unit adopted in 2015
- This is the "**forbidden comparison**" problem
  - → Already-treated units are not a clean control
  - → If treatment effects vary over time, their included effect contaminates the estimate
- Result: TWFE can produce a **weighted average with negative weights**

The problem with naive TWFE in staggered settings

- With staggered timing, TWFE uses **all available 2$\times$2 comparisons**
- Including: already-treated units as the **control group** for later-treated units
  - $\rightarrow$ Unit adopted in 2012 becomes "control" for unit adopted in 2015
- This is the "**forbidden comparison**" problem
  - $\rightarrow$ Already-treated units are not a clean control
  - $\rightarrow$ If treatment effects vary over time, their included effect contaminates the estimate
- Result: TWFE can produce a **weighted average with negative weights**
  - $\rightarrow$ Estimate may not correspond to any valid ATT

Negative weights: the intuition

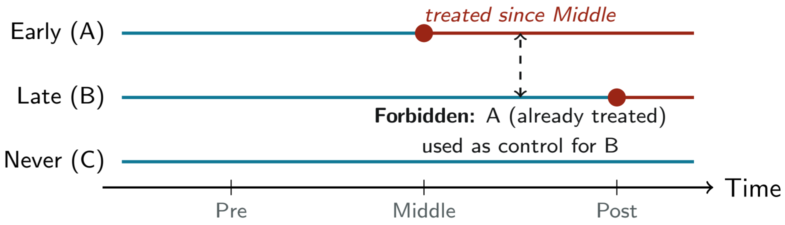

Modern alternatives: Callaway-Sant'Anna (2021)

- Key idea: estimate **group-time ATTs** $ATT(g,t)$ separately

Modern alternatives: Callaway-Sant'Anna (2021)

- Key idea: estimate **group-time ATTs** $ATT(g, t)$ separately
  - $\rightarrow$ g: treatment cohort (when the unit was first treated)

Modern alternatives: Callaway-Sant'Anna (2021)

- Key idea: estimate **group-time ATTs** $ATT(g, t)$ separately
  - $\rightarrow$ g: treatment cohort (when the unit was first treated)
  - $\rightarrow$ t: calendar time

Modern alternatives: Callaway-Sant'Anna (2021)

- Key idea: estimate **group-time ATTs** $ATT(g, t)$ separately
  - → g: treatment cohort (when the unit was first treated)
  - → t: calendar time
  - → Compare cohort g to clean controls (never-treated or not-yet-treated)

Modern alternatives: Callaway-Sant'Anna (2021)

- Key idea: estimate **group-time ATTs** $ATT(g,t)$ separately
  - → g: treatment cohort (when the unit was first treated)
  - → t: calendar time
  - → Compare cohort g to clean controls (never-treated or not-yet-treated)
- Then **aggregate** $ATT(g,t)$ as desired:

Modern alternatives: Callaway-Sant'Anna (2021)

- Key idea: estimate **group-time ATTs** $ATT(g, t)$ separately
  - → g: treatment cohort (when the unit was first treated)
  - → t: calendar time
  - → Compare cohort g to clean controls (never-treated or not-yet-treated)
- Then **aggregate** $ATT(g, t)$ as desired:
  - → Overall average ATT

Modern alternatives: Callaway-Sant'Anna (2021)

- Key idea: estimate **group-time ATTs** $ATT(g, t)$ separately
  - → g: treatment cohort (when the unit was first treated)
  - → t: calendar time
  - → Compare cohort g to clean controls (never-treated or not-yet-treated)
- Then **aggregate** $ATT(g, t)$ as desired:
  - → Overall average ATT
  - → ATT by event-time (dynamic/event study)

Modern alternatives: Callaway-Sant'Anna (2021)

- Key idea: estimate **group-time ATTs** $ATT(g, t)$ separately
  - → g: treatment cohort (when the unit was first treated)
  - → t: calendar time
  - → Compare cohort g to clean controls (never-treated or not-yet-treated)
- Then **aggregate** $ATT(g, t)$ as desired:
  - → Overall average ATT
  - → ATT by event-time (dynamic/event study)
  - → ATT by calendar time

Modern alternatives: Callaway-Sant'Anna (2021)

- Key idea: estimate **group-time ATTs** $ATT(g, t)$ separately
  - $\rightarrow$ g: treatment cohort (when the unit was first treated)
  - $\rightarrow$ t: calendar time
  - $\rightarrow$ Compare cohort g to clean controls (never-treated or not-yet-treated)
- Then **aggregate** $ATT(g, t)$ as desired:
  - $\rightarrow$ Overall average ATT
  - $\rightarrow$ ATT by event-time (dynamic/event study)
  - $\rightarrow$ ATT by calendar time
- Avoids forbidden comparisons entirely

Modern alternatives: Callaway-Sant'Anna (2021)

- Key idea: estimate **group-time ATTs** $ATT(g, t)$ separately
  - $\rightarrow$ g: treatment cohort (when the unit was first treated)
  - $\rightarrow$ t: calendar time
  - $\rightarrow$ Compare cohort g to clean controls (never-treated or not-yet-treated)
- Then **aggregate** $ATT(g, t)$ as desired:
  - $\rightarrow$ Overall average ATT
  - $\rightarrow$ ATT by event-time (dynamic/event study)
  - $\rightarrow$ ATT by calendar time
- Avoids forbidden comparisons entirely
- Works even with **heterogeneous treatment effects**

Modern alternatives: ETWFE (Wooldridge 2021)

- **Extended TWFE**: stays within the regression framework

Modern alternatives: ETWFE (Wooldridge 2021)

- **Extended TWFE**: stays within the regression framework
- Key idea: interact treatment with cohort dummies and time dummies

Modern alternatives: ETWFE (Wooldridge 2021)

- **Extended TWFE**: stays within the regression framework
- Key idea: interact treatment with cohort dummies and time dummies
  - $\rightarrow$ Each cohort gets its own treatment effect profile

Modern alternatives: ETWFE (Wooldridge 2021)

- **Extended TWFE**: stays within the regression framework
- Key idea: interact treatment with cohort dummies and time dummies
  - $\rightarrow$ Each cohort gets its own treatment effect profile
  - $\rightarrow$ Avoids the averaging-over-cohorts problem of naive TWFE

Modern alternatives: ETWFE (Wooldridge 2021)

- **Extended TWFE**: stays within the regression framework
- Key idea: interact treatment with cohort dummies and time dummies
  - → Each cohort gets its own treatment effect profile
  - → Avoids the averaging-over-cohorts problem of naive TWFE
- Estimating a richer model, not a simpler one

Modern alternatives: ETWFE (Wooldridge 2021)

- **Extended TWFE**: stays within the regression framework
- Key idea: interact treatment with cohort dummies and time dummies
  - → Each cohort gets its own treatment effect profile
  - → Avoids the averaging-over-cohorts problem of naive TWFE
- Estimating a richer model, not a simpler one
- Same identifying assumptions as Callaway-Sant'Anna

Modern alternatives: ETWFE (Wooldridge 2021)

- **Extended TWFE**: stays within the regression framework
- Key idea: interact treatment with cohort dummies and time dummies
  - → Each cohort gets its own treatment effect profile
  - → Avoids the averaging-over-cohorts problem of naive TWFE
- Estimating a richer model, not a simpler one
- Same identifying assumptions as Callaway-Sant'Anna
- Easier to implement in standard regression software

Modern alternatives: ETWFE (Wooldridge 2021)

- **Extended TWFE**: stays within the regression framework
- Key idea: interact treatment with cohort dummies and time dummies
  - → Each cohort gets its own treatment effect profile
  - → Avoids the averaging-over-cohorts problem of naive TWFE
- Estimating a richer model, not a simpler one
- Same identifying assumptions as Callaway-Sant'Anna
- Easier to implement in standard regression software
- In R: `etwfe` package by Grant McDermott

In R: `did` and `etwfe` packages

- **Callaway-Sant'Anna** via `did` package:

In R: `did` and `etwfe` packages

- **Callaway-Sant'Anna** via `did` package:

```
library(did)
```

In R: `did` and `etwfe` packages

- **Callaway-Sant'Anna** via `did` package:

```
library(did)
cs = att_gt(yname = "y", gname = "g", idname = "id",
```

In R: `did` and `etwfe` packages

- **Callaway-Sant'Anna** via `did` package:

```
library(did)
cs = att_gt(yname = "y", gname = "g", idname = "id",
    tname = "year", data = df)
```

In R: `did` and `etwfe` packages

- **Callaway-Sant'Anna** via `did` package:

```
library(did)
cs = att_gt(yname = "y", gname = "g", idname = "id",
    tname = "year", data = df)
aggte(cs, type = "simple") # overall ATT
```

In R: `did` and `etwfe` packages

- **Callaway-Sant'Anna** via `did` package:

```
library(did)
cs = att_gt(yname = "y", gname = "g", idname = "id",
    tname = "year", data = df)
aggte(cs, type = "simple") # overall ATT
aggte(cs, type = "dynamic") # event-study ATT
```

In R: `did` and `etwfe` packages

- **Callaway-Sant'Anna** via `did` package:

```
library(did)
cs = att_gt(yname = "y", gname = "g", idname = "id",
    tname = "year", data = df)
aggte(cs, type = "simple") # overall ATT
aggte(cs, type = "dynamic") # event-study ATT
```

- **ETWFE** via `etwfe` package:

In R: `did` and `etwfe` packages

- **Callaway-Sant'Anna** via `did` package:

```
library(did)
cs = att_gt(yname = "y", gname = "g", idname = "id",
    tname = "year", data = df)
aggte(cs, type = "simple") # overall ATT
aggte(cs, type = "dynamic") # event-study ATT
```

- **ETWFE** via `etwfe` package:

```
library(etwfe)
```

In R: `did` and `etwfe` packages

- **Callaway-Sant'Anna** via `did` package:

```
library(did)
cs = att_gt(yname = "y", gname = "g", idname = "id",
    tname = "year", data = df)
aggte(cs, type = "simple") # overall ATT
aggte(cs, type = "dynamic") # event-study ATT
```

- **ETWFE** via `etwfe` package:

```
library(etwfe)
m = etwfe(y ~ 1, tvar = year, gvar = g,
```

In R: did and etwfe packages

- **Callaway-Sant'Anna** via did package:

```
library(did)
cs = att_gt(yname = "y", gname = "g", idname = "id",
    tname = "year", data = df)
aggte(cs, type = "simple") # overall ATT
aggte(cs, type = "dynamic") # event-study ATT
```

- **ETWFE** via etwfe package:

```
library(etwfe)
m = etwfe(y ~ 1, tvar = year, gvar = g,
    data = df, vcov = ~id)
```

In R: `did` and `etwfe` packages

- **Callaway-Sant'Anna** via `did` package:

```
library(did)
cs = att_gt(yname = "y", gname = "g", idname = "id",
    tname = "year", data = df)
aggte(cs, type = "simple") # overall ATT
aggte(cs, type = "dynamic") # event-study ATT
```

- **ETWFE** via `etwfe` package:

```
library(etwfe)
m = etwfe(y ~ 1, tvar = year, gvar = g,
    data = df, vcov = ~id)
emfx(m) # marginal effects = ATT
```

Naive TWFE vs. modern estimators

	Naive TWFE	**Callaway-Sant'Anna**	**ETWFE**
Clean control groups	No	Yes	Yes
Heterogeneous effects	Biased	Handles	Handles
Negative weights	Possible	None	None
Pre-trend test	Via event study	Built-in	Via `emfx`
Implementation	`feols`	`did`	`etwfe`

- With simultaneous treatment (all treated at same time): TWFE is fine

Naive TWFE vs. modern estimators

	Naive TWFE	**Callaway-Sant'Anna**	**ETWFE**
Clean control groups	No	Yes	Yes
Heterogeneous effects	Biased	Handles	Handles
Negative weights	Possible	None	None
Pre-trend test	Via event study	Built-in	Via `emfx`
Implementation	`feols`	`did`	`etwfe`

- With simultaneous treatment (all treated at same time): TWFE is fine
- With staggered adoption and homogeneous effects: TWFE is approximately fine

Naive TWFE vs. modern estimators

	Naive TWFE	**Callaway-Sant'Anna**	**ETWFE**
Clean control groups	No	Yes	Yes
Heterogeneous effects	Biased	Handles	Handles
Negative weights	Possible	None	None
Pre-trend test	Via event study	Built-in	Via `emfx`
Implementation	`feols`	`did`	`etwfe`

- With simultaneous treatment (all treated at same time): TWFE is fine
- With staggered adoption and homogeneous effects: TWFE is approximately fine
- With staggered adoption and **heterogeneous effects**: use CS or ETWFE

Roadmap

DiD: putting it all together

Parallel trends + treatment/control design $\Rightarrow$ DiD

$\downarrow$

Event study: test pre-trends + trace dynamic effects

$\downarrow$

Staggered adoption: naive TWFE can be biased

$\downarrow$

Modern fix: Callaway-Sant'Anna or ETWFE

Key takeaways

- **DiD**: compare treated vs. control, before vs. after

Key takeaways

- **DiD**: compare treated vs. control, before vs. after
  - $\rightarrow$ Subtract the control trend as the counterfactual

Key takeaways

- **DiD**: compare treated vs. control, before vs. after
  - → Subtract the control trend as the counterfactual
- **Parallel trends**: the key identification assumption

Key takeaways

- **DiD**: compare treated vs. control, before vs. after
  - $\rightarrow$ Subtract the control trend as the counterfactual
- **Parallel trends**: the key identification assumption
  - $\rightarrow$ Untestable in the post-period; probe with pre-trend tests

Key takeaways

- **DiD**: compare treated vs. control, before vs. after
  - $\rightarrow$ Subtract the control trend as the counterfactual
- **Parallel trends**: the key identification assumption
  - $\rightarrow$ Untestable in the post-period; probe with pre-trend tests
- **Event studies**: leads and lags around treatment

Key takeaways

- **DiD**: compare treated vs. control, before vs. after
  - $\rightarrow$ Subtract the control trend as the counterfactual
- **Parallel trends**: the key identification assumption
  - $\rightarrow$ Untestable in the post-period; probe with pre-trend tests
- **Event studies**: leads and lags around treatment
  - $\rightarrow$ Pre-treatment $\hat{\delta}_k \approx 0$: supports parallel trends

Key takeaways

- **DiD**: compare treated vs. control, before vs. after
  - $\rightarrow$ Subtract the control trend as the counterfactual
- **Parallel trends**: the key identification assumption
  - $\rightarrow$ Untestable in the post-period; probe with pre-trend tests
- **Event studies**: leads and lags around treatment
  - $\rightarrow$ Pre-treatment $\hat{\delta}_k \approx 0$: supports parallel trends
- **Staggered DiD**: naive TWFE can fail with heterogeneous effects

Key takeaways

- **DiD**: compare treated vs. control, before vs. after
  - $\rightarrow$ Subtract the control trend as the counterfactual
- **Parallel trends**: the key identification assumption
  - $\rightarrow$ Untestable in the post-period; probe with pre-trend tests
- **Event studies**: leads and lags around treatment
  - $\rightarrow$ Pre-treatment $\hat{\delta}_k \approx 0$: supports parallel trends
- **Staggered DiD**: naive TWFE can fail with heterogeneous effects
  - $\rightarrow$ Use Callaway-Sant'Anna (`did`) or ETWFE (`etwfe`)

Key takeaways

- **DiD**: compare treated vs. control, before vs. after
  - $\rightarrow$ Subtract the control trend as the counterfactual
- **Parallel trends**: the key identification assumption
  - $\rightarrow$ Untestable in the post-period; probe with pre-trend tests
- **Event studies**: leads and lags around treatment
  - $\rightarrow$ Pre-treatment $\hat{\delta}_k \approx 0$: supports parallel trends
- **Staggered DiD**: naive TWFE can fail with heterogeneous effects
  - $\rightarrow$ Use Callaway-Sant'Anna (`did`) or ETWFE (`etwfe`)
- Always **cluster SEs** at the treatment-assignment level

For next session

- Complete Assignment 6 (DiD application)
- Read the assigned paper (DiD design)
- Next session: Spatial Data (I)
  - $\rightarrow$ Spatial data structures and visualization
  - $\rightarrow$ Spatial autocorrelation
  - $\rightarrow$ Spatial regression models

Questions?