Advanced Quantitative Methods

Difference-in-differences

Instructor: Gregory Eady
Office: 18.2.10
Office hours: Fridays 13-15

Today

- Difference-in-differences
- Exercise

Difference-in-differences setup

- There is often no opportunity for randomization—whether by a researcher or externally (an instrumental variable)
- Panel data, however, can help us estimate a causal effect
- How? In cases where some units are treated in a given time period and others are not
  - e.g. A new policy implemented
  - e.g. An event occurs in some place, but not others

Card and Krueger (1994): The classic DD setup

- There is debate among economists about whether increasing the minimum wage causes an increase in unemployment
- At the time, there is cross-sectional evidence that this is true
- But US states do not select a minimum wage at random, so cross-sectional regressions might be missing a bunch of unobserved confounders

Card and Krueger's (1994) solution

- Compare a treatment and control case over time
- New Jersey raised its minimum wage in April 1992
- Card and Krueger (1994) compare employment in New Jersey's fast food industry to that of neighboring Pennsylvania before and after the minimum wage increase
- Result: If anything, a *positive* effect on employment

Card and Krueger's (1994): Two-period diff-in-diff

- Two units:
  - Treatment case: New Jersey
  - Control case: Pennsylvania
- Four observations
  - Pre-treatment ($t = 0$)
    - Unemployment in New Jersey (untreated)
    - Unemployment in Pennsylvania (untreated)
  - Post-treatment ($t = 1$)
    - Unemployment in New Jersey (treated)
    - Unemployment in Pennsylvania (untreated)

Difference-in-differences: Compare the difference in employment in New Jersey at $t = 0$ and $t = 1$ to the difference in employment in Pennsylvania at $t = 0$ and $t = 1$

Why are we comparing the difference in employment between two states?

- We need a way to create some sort of counterfactual comparison for New Jersey (the treated state)
- We will assume that—if no new policy were implemented—*changes* in the number of employees for New Jersey and Pennsylvania would move *in parallel*
  - We'll call this the "parallel trends assumption"

Naive pre-post comparison ($11 - 7 = 4$?)

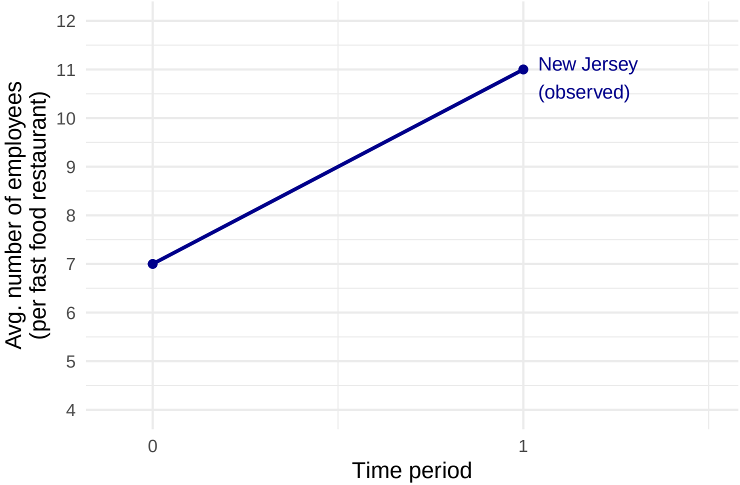

Naive pre-post comparison ($11 - 7 = 4$?)

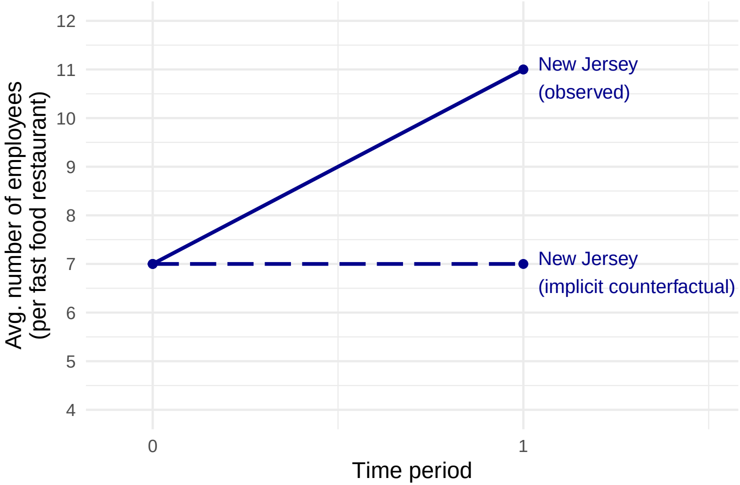

Pennsylvania as a control

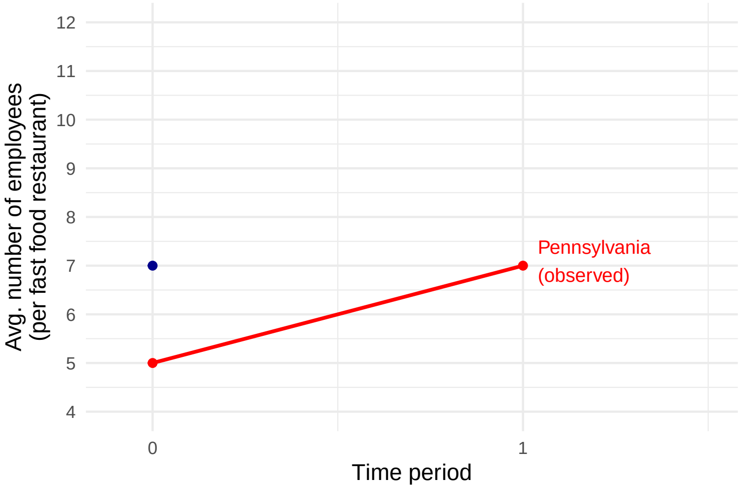

Pennsylvania versus “New Jersey”

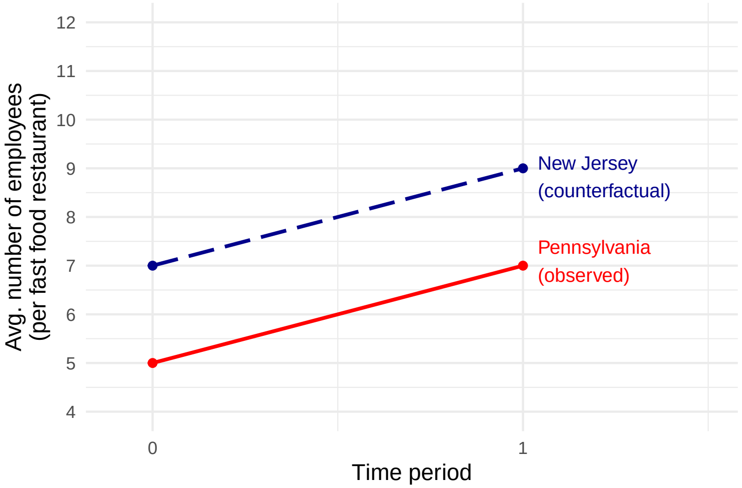

Pennsylvania versus New Jersey

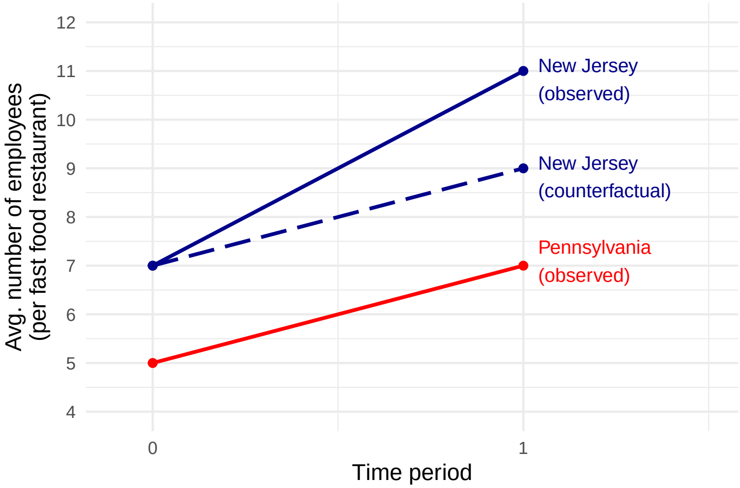

Difference-in-differences estimate

- Pennsylvania (**control**):
  - Before: 5
  - After: 7
  - $\textit{Difference}_{PA}$: $7 - 5 = 2$
- New Jersey (**treated**):
  - Before: 7
  - After: 11
  - $\textit{Difference}_{NJ}$: $11 - 7 = 4$
- The difference in these two differences? $4 - 2 = 2$
- This is the *difference-in-differences* estimate of the effect of a minimum wage increase on employment

A bit more formally

The difference-in-differences estimate is:

$$(\bar{Y}_{\text{NJ,t}=1} - \bar{Y}_{\text{NJ,t}=0}) - (\bar{Y}_{\text{PA,t}=1} - \bar{Y}_{\text{PA,t}=0}) \tag{1}$$

In potential outcomes notation:

Values of Y_i for treated unit (e.g. NJ) before and after:

$Y(1)_{i,\text{post}}|T = 1$ and $Y(1)_{i,\text{pre}}|T = 1$

Values of Y_i for the treated unit (e.g. NJ) before and after if it had *not* been treated:

$Y(0)_{i,\text{post}}|T = 1$ and $Y(0)_{i,\text{pre}}|T = 1$

A bit more formally

The difference-in-differences estimate in potential outcomes notation:

$$\overbrace{(Y(1)_{post}|T=1) - Y(1)_{pre}|T=1)}^{\substack{\text{Difference in New Jersey employment} \\ \text{if minimum wage policy change}}}$$

A bit more formally

The difference-in-differences estimate in potential outcomes notation:

$$\overbrace{(Y(1)_{post}|T=1) - Y(1)_{pre}|T=1)}^{\substack{\text{Difference in New Jersey employment} \\ \text{if minimum wage policy change}}} - \overbrace{(Y(0)_{post}|T=1) - Y(0)_{pre}|T=1)}^{\substack{\text{Difference in New Jersey employment} \\ \text{if } \textit{no} \text{ minimum wage policy change}}}$$

The big assumption that allows us to use an untreated unit as a control is the "parallel trends assumption"

$$\overbrace{(Y(0)_{post}|T=1)-Y(0)_{pre}|T=1)}^{\substack{\text{Difference in New Jersey employment}\\ \text{if } no \text{ minimum wage policy change}}}-$$

$$\overbrace{(Y(0)_{post}|T=0)-Y(0)_{pre}|T=0)}^{\substack{\text{Difference in Pennsylvania employment}\\ \text{if } no \text{ minimum wage policy change}}}$$

We need to assume that this is equal to zero (else the diff-in-diff estimate is biased)

Why might the parallel trends assumption be broken?

- Something else happens at the same time as the treatment that affects the groups differently
  - e.g. A big McDonald's ad campaign in New Jersey
- Other shocks or events
  - Macro- or micro-level economic forces affect Pennsylvania differently from New Jersey

Choice of control unit(s) thus might matter

- Card and Krueger (1994) recognized this:
  - New Jersey and Pennsylvania have similar economic composition
  - Same weather
  - Same region, so similar economic or other shocks
- We can test for violations of the parallel trends assumption using an "event study model" (next week's lecture)
- Are also new machine-learning based methods to create a counterfactual
  - Can automate selection of control comparison case with "synthetic control" methods (Adadie et al. 2003, 2010, 2015)

A small aside: Synthetic control

FIGURE 1 Trends in per Capita GDP: West Germany versus Rest of the OECD Sample

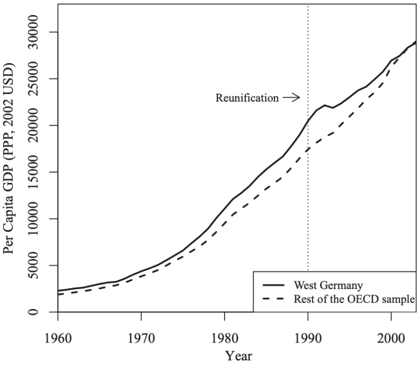

FIGURE 2 Trends in per Capita GDP: West Germany versus Synthetic West Germany

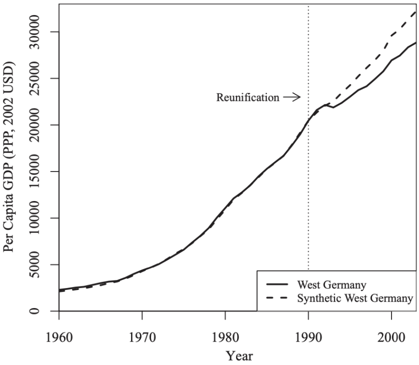

PA

Waking Up the Golden Dawn: Does Exposure to the Refugee Crisis Increase Support for Extreme-Right Parties?

Elias Dinas[1,2], Konstantinos Matakos[3], Dimitrios Xefteris[4] and Dominik Hangartner[5,6,7]

[1] *Department of Political and Social Sciences and RSCAS, European University Institute, San Domenico di Fiesole, I-50014, Italy*
[2] *Department of Politics and International Relations, Oxford University, Oxford, OX1 3UQ, UK*
[3] *Department of Political Economy, King's College London, London, WC2A 4PH, UK*
[4] *Department of Economics, University of Cyprus, Cyprus*
[5] *Center of Comparative and International Studies, ETH Zurich, 8092 Zurich, Switzerland*
[6] *Immigration Policy Lab, Stanford University and ETH Zurich, 8092 Zurich, Switzerland*
[7] *Department of Government, London School of Economics and Political Science, London, WC2A 2AE, UK.*
Email: d.hangartner@lse.ac.uk

Abstract

Does exposure to the refugee crisis fuel support for extreme-right parties? Despite heated debates about the political repercussions of the refugee crisis in Europe, there exists very little—and sometimes conflicting—evidence with which to assess the impact of a large influx of refugees on natives' political attitudes and behavior. We provide causal evidence from a natural experiment in Greece, where some Aegean islands close to the Turkish border experienced sudden and drastic increases in the number of Syrian refugees while other islands slightly farther away—but with otherwise similar institutional and socioeconomic characteristics—did not. Placebo tests suggest that precrisis trends in vote shares for exposed and nonexposed islands were virtually identical. This allows us to obtain unbiased estimates of the electoral consequences of the refugee crisis. Our study shows that among islands that faced a massive but transient inflow of refugees passing through just before the September 2015 election, vote shares for Golden Dawn, the most extreme-right party in Europe, moderately increased by 2 percentage points (a 44 percent increase at the average). The finding that mere exposure to the refugee crisis is sufficient to fuel support for extreme-right parties has important implications for our theoretical understanding of the drivers of antirefugee backlash.

Keywords: natural experiments, causal inference, instrumental variables, panel data

Political Analysis (2019)
vol. 27:244–254
DOI: 10.1017/pan.2018.48

Setup: Instrumental variables & diff-in-diff

The geography and timing of the refugee crisis in the Aegean Sea

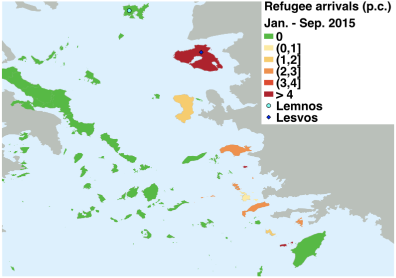

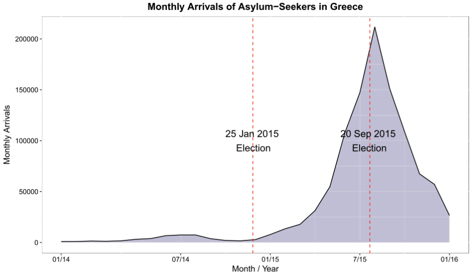

Figure 1. Panel A shows that islands close to the Turkish border received the most refugee arrivals per capita. Panel B shows the monthly number of asylum seekers arriving at Greek islands over the period from January 2014 to March 2016. During the study period, the first election took place just before the onset of the refugee crisis on January 25, 2015. A second election took place at the height of the refugee crisis on September 20, 2015.

What is the effect of exposure to refugees on far-right vote share?

- Multiple periods (4 elections) and multiple units (95 municipalities & 248 townships)
- Having multiple periods allows us to visually check for parallel trends

Multi-period, multi-unit diff-in-diff

Parallel trends at the municipal and township level

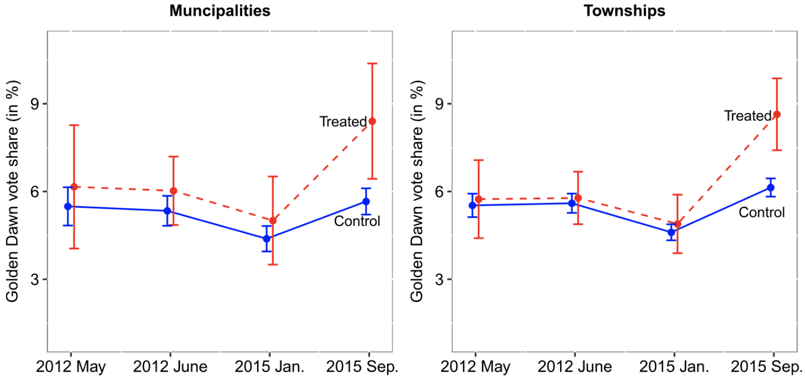

Figure 2. The analyses at the municipality (left panel) and township level (right panel) show that treated and control islands experience highly similar changes in support for GD prior to the refugee crisis, thereby strengthening our confidence in the parallel trend assumption. The blue connected line indicates the average vote share for GD in the municipalities (left panel) and townships (right panel) that received refugees. The red dashed line denotes the average GD vote share in municipalities and townships without refugee exposure.

Parallel trends

- We see *pre*-treatment parallel trends
- This is only an *indirect* test of the parallel trends assumption
- Parallel trends assumption concerns *post-treatment* trends
  - i.e. absent a treatment, we would counter-factually observe parallel trends in the post-treatment period
- Counter example: the average height of boys and girls evolves in parallel until about age 13 and then diverges
  - We should *not* conclude a causal effect of bar mitzvahs (when boys turn age 13) on a boy's height even if the heights of girls and boys is parallel before age 13

Calculating a difference-in-difference estimate

$$y_{it} = \alpha + \beta_1 \text{treatment}_i + \beta_2 \text{post}_{it} + \beta_3 (\text{treatment}_i \times \text{post}_{it}) + \epsilon_{it}$$

Calculating a difference-in-difference estimate

$$y_{it} = \alpha + \beta_1 \text{treatment}_i + \beta_2 \text{post}_{it} + \beta_3(\text{treatment}_i \times \text{post}_{it}) + \epsilon_{it}$$

- α: baseline pre-treatment vote share (y_{it}) for Golden Dawn among the control group

Calculating a difference-in-difference estimate

$$y_{it} = \alpha + \beta_1 \text{treatment}_i + \beta_2 \text{post}_{it} + \beta_3 (\text{treatment}_i \times \text{post}_{it}) + \epsilon_{it}$$

- α: baseline pre-treatment vote share (y_{it}) for Golden Dawn among the control group
- β_1: difference between treatment group and control group in the pre-treatment period

Calculating a difference-in-difference estimate

$$y_{it} = \alpha + \beta_1 \text{treatment}_i + \beta_2 \text{post}_{it} + \beta_3(\text{treatment}_i \times \text{post}_{it}) + \epsilon_{it}$$

- α: baseline pre-treatment vote share (y_{it}) for Golden Dawn among the control group
- β_1: difference between treatment group and control group in the pre-treatment period
- β_2: change in mean vote share for the control group between the pre-treatment and post-treatment period

Calculating a difference-in-difference estimate

$$y_{it} = \alpha + \beta_1 \text{treatment}_i + \beta_2 \text{post}_{it} + \beta_3 (\text{treatment}_i \times \text{post}_{it}) + \epsilon_{it}$$

- α: baseline pre-treatment vote share (y_{it}) for Golden Dawn among the control group
- β_1: difference between treatment group and control group in the pre-treatment period
- β_2: change in mean vote share for the control group between the pre-treatment and post-treatment period
- β_3: *difference* in the change in mean vote share for the treatment group relative to the control group between the pre-treatment and post-treatment period (i.e. our diff-in-diff estimate)

	Golden Dawn Vote Share (%)
Post × Treatment	2.218***
	(0.724)
Post	1.530***
	(0.215)
Treatment	0.286
	(0.512)
Intercept	4.605***
	(0.152)
Observations	498
R^2	0.161
Note:	*p<0.1; **p<0.05; ***p<0.01

Exposure to refugees caused a 2.2%-point increase in vote share for the Golden Dawn

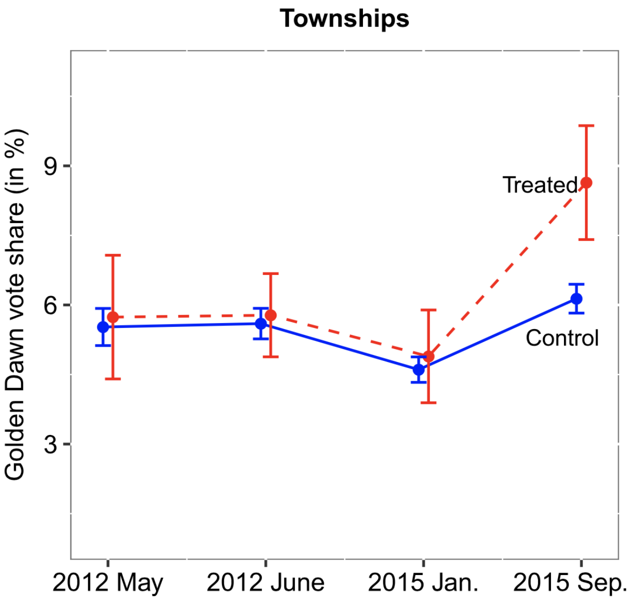
Townships
Golden Dawn vote share (in %)
9
6
3
Treated
Control
2012 May
2012 June
2015 Jan.
2015 Sep.

Two-period diff-in-diff is a special case

- Often we have many periods, so we can't run the simple regression as shown in the last slides
- Often treatment timing varies (e.g. a minimum wage increase is implemented in different states at different times)
- We thus need a *generalized* difference-in-differences model

Two-way fixed effects

- Before, we looked at fixed effects for panel data for a single unit (e.g. variation *within* how respondents answer a survey over time)
- But we can add in a time period fixed effect as well
- This accounts for changes up and down, on average, across all units in a given time period
- Without a treatment effect, all units should increase or decrease per time period more or less together (i.e. parallel trends)
- We can then test for whether units that receive a treatment increase or decrease in the outcome relative to the control units

Generalized difference-in-differences (a Two-Way Fixed Effects model)

$$y_{it} = \alpha_i + \omega_t + \beta \text{treatment}_{it} + \gamma X_{it} + \epsilon_{it}$$

- α_i: A unit-specific fixed effect
- ω_t: A time-specific fixed effect
- β: The treatment effect
- treatment_{it}: Takes the value 1 if unit i is treated in time period t, and the value 0 otherwise
- γ: Relationships between time-varying covariates, X_{it} and the outcome (i.e. controls)

Generalized difference-in-differences for Dinas et al. (2019)

$$\text{GD vote share}_{it} = \alpha_i + \omega_t + \beta \text{refugees}_{it} + \epsilon_{it}$$

- α_i: A unit-specific fixed effect capturing average levels of support for Golden Dawn in a given town
- ω_t: A time-specific fixed effect capturing average levels of support for Golden Dawn in a specific time period
- β: The treatment effect
- refugees_{it}: Takes the value 1 if town i received refugees in time period t, and the value 0 otherwise

Dinas et al. (2019) results (& lags as parallel trends check)

Table 1. Impact of refugee arrivals on GD vote share.

Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Outcome:	$GD_{(t)}$		$GD_{(t-1)}$		$GD_{(t)}$		$GD_{(t-1)}$		$GD_{(t)}$		$GD_{(t-1)}$	
Treatment:	Binary treatment		Binary treatment		Binary treatment		Binary treatment		Arrivals per capita		Arrivals per capita	
Unit:	Municipality		Municipality		Township		Township		Municipality		Municipality	
Exposure	2.079	2.112	−0.040	−0.055	2.272	2.193	0.093	0.127	0.604	0.600	−0.004	−0.033
	(0.351)	(0.674)	(0.392)	(0.713)	(0.263)	(0.455)	(0.262)	(0.439)	(0.178)	(0.264)	(0.119)	(0.262)
Unit FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Election FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Unit trends		✓		✓		✓		✓		✓		✓
N	380	380	285	285	992	992	744	744	379	379	284	284
Elections	4	4	3	3	4	4	3	3	4	4	3	3
Clusters	95	95	95	95	248	248	248	248	95	95	95	95

Note: Models 1–12 display ordinary least squares (OLS) regression coefficients with clustered standard errors in parentheses. Models 1–8 use a binary treatment indicator while models 9–12 use the number of refugee arrivals per capita. Models 1, 2, 5, 6, 9 and 10 show the effect on GD vote share. Models 3, 4, 7, 8, 11, and 12 use the GD vote share from the previous election as placebo outcome. All models control for election and unit of analysis (municipality or township) fixed effects. In addition, models 2, 4, 6, 8, 10 and 12 also include unit-specific linear time trends.

Note they also include "unit trends" for robustness. This means they fit the following model:

$$y_{it} = \alpha_i + \omega_t + \beta \text{treatment}_{it} + \lambda_i t + \epsilon_{it},$$

where the parameters λ_i denote a separate time trend for each unit i

Robustness and extensions

- Placebo tests: Whether treatment status is correlated with the value of the outcome as measured before the event occurs
  - e.g. Models 3-4, 7-8, 11-12 on the previous slide
- Unit-specific trends: Account for trends in the outcome for each unit
  - e.g. Models 2, 4, 6, 8, 10, 12 on the previous slide
- Heterogeneity in the effect of a treatment.
  - Triple differences / difference-in-difference-in-differences
- You can include time-varying control variables, if you think they are important
- If treatment assignment varies over time, need further adjustments
- "Event study" models allow you to observe the dynamics of a treatment (how the effect changes over time)

Effect of the 1975 Voting Rights Act on voting:

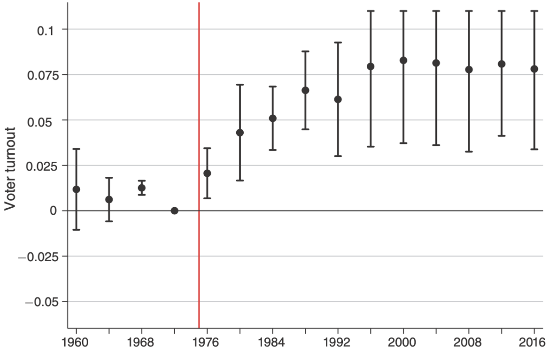

FIGURE 3. EFFECT ON VOTER TURNOUT

Notes: Data come from ICPSR and Dave Leip's Election Atlas. The graph shows DD coefficients and 95 percent confidence intervals from the estimation of equation (1). Observations are at the county-year level and weighted by the voting-eligible population. Standard errors are clustered at the state level. The red vertical line represents the passage of 1975 VRA. Full results are displayed in Table 2, column 1.

Currently a lot of new research on diff-in-diff

- Synthetic controls (various methods)
- Two-way fixed estimates are not as straightforward as they seem
  - https://andrewcbaker.netlify.app/2019/09/25/difference-in-differences-methodology/
  - https://blogs.worldbank.org/impactevaluations/what-are-we-estimating-when-we-estimate-difference-differences

Complete the exercise replicating the results from Dinas et al. (2019) from the data and R file on the course website

Exercise solutions

```
# Now we can estimate our simple two-period difference-in-differences model
# as in the regression equation from the lectures
model_classic_two_period <- lm(gdper ~ post * treated, data = C)
```

Exercise solutions

```
# If you had to calculate the difference-in-differences manually just using the
# function mean(), how would you calculate it?
# You need to calculate 4 quantities, take two differences and then the...
# difference in those differences.
# Your estimate from just using mean() should be the exact same thing as the
# regression model above gave you
#
# To give you a hint, this is how you would calculate the mean Golden Dawn
# vote share in the pre-treatment period among those townships that would have
# refugees.
(mean(C$gdper[C$treated == 1 & C$post == 1], na.rm = TRUE) -
mean(C$gdper[C$treated == 1 & C$post == 0], na.rm = TRUE)) -
(mean(C$gdper[C$treated == 0 & C$post == 1], na.rm = TRUE) -
mean(C$gdper[C$treated == 0 & C$post == 0], na.rm = TRUE))
```

Note how this is simply the difference in the Golden Dawn vote share between 2016 (post $==$ 1) and 2015 (post $==$ 0) among the treated... *minus* the difference in the Golden Dawn vote share between 2016 (post $==$ 1) and 2015 (post $==$ 0) among the control. i.e. the differences in these differences.

Exercise solutions

```
model_5 <- feols(gdper ~ towntr |      # Outcome regressed on treatment
                         town + year, # Unit and time fixed effects
                 cluster = ~ town,    # SE clusted on the unit
                 data = D)
summary(model_5)
```

Exercise solutions

```
model_6 <- feols(gdper ~ towntr |
                 town + year + town[year], # town FE, year FE, town x year trend
                 cluster = ~ town, # SEs clustered at the unit level
                 data = D)
summary(model_6)
```

Exercise solutions

```
model_6_alternative <- feols(gdper ~ towntr |
                             town + year + municipality[year],
                             cluster = ~ town, # SEs clustered at the unit level
                             data = D)
summary(model_6_alternative)
```