Causal inference week 6: Differences-and-Differences

Andy Eggers

Oxford DPIR

HT 2019

Overview

Strategies for estimating effects of treatments so far:

- Randomize treatment and take the DIGM
- Identify and control for confounding variables such that the CIA holds
- Identify an instrumental variable and use two-stage-least-squares to estimate average treatment effect for compliers
- Identify a situation in which the treatment depends on a cutoff

Today: using observations at more than one point in time.

As with IV, works even if there is (certain types of) confounding.

Madrid train bombing, 11 March 2004

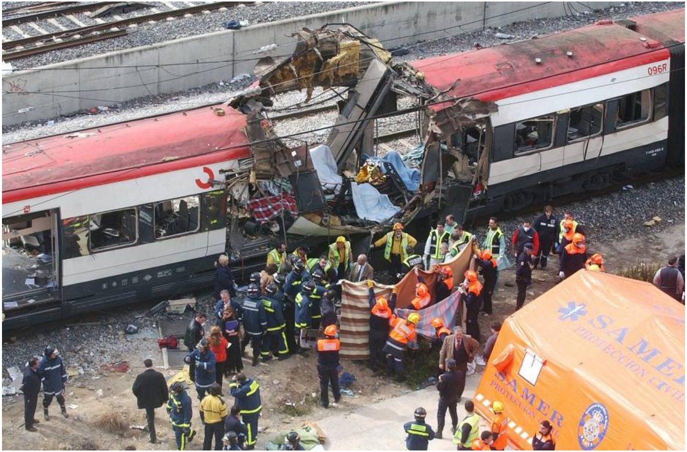

Question: How did the M11 attack affect the election three days later?

Possible research designs

How could you use

- polls
- post-election surveys (which asked e.g. “Did the terrorist attack of March 11th in Madrid influence your vote?”) (see Bali 2007, *Electoral Studies*)

to estimate the effect of the attacks on relative support for the Conservatives vs Socialists?

Two differences you could estimate

Montalvo (2011) points out that Spanish nationals living abroad voted *before* the bombing.

Two differences you could estimate

Montalvo (2011) points out that Spanish nationals living abroad voted *before* the bombing.

What about estimating the effect of the attacks by

- comparing the results in 2004 for resident and non-resident voters? (individual or province-level, with some covariates) **(cross-sectional)**
- comparing the results for non-resident voters in 2004 and 2000? (individual or province-level, with some covariates). **(before-and-after, "time-series")**

Using both: diff-in-diff

Diff-in-diff idea: what about comparing the before-and-after (diff) for residents and non-residents (in diff)?

Using both: diff-in-diff

Diff-in-diff idea: what about comparing the before-and-after (diff) for residents and non-residents (in diff)?

Measure the difference between non-resident and resident voters in 2004, but then subtract this same difference measured in 2000 → difference-in-differences.

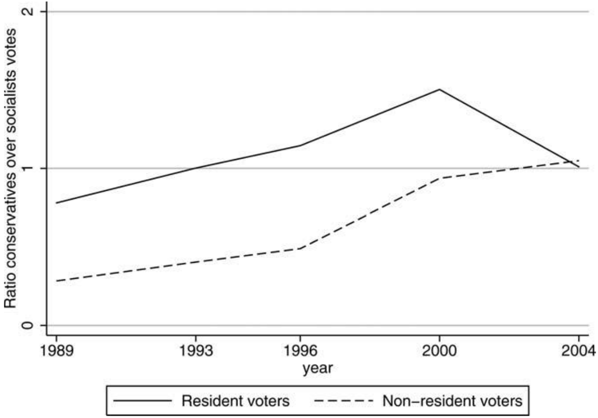

Scope of application

Simple case (today): binary treatment, applied at one point in time (but not to everyone)

More general case (next week): general treatment, applied in any pattern

Commonalities: ;

- multiple observations over time, with treatment varying within group or unit over time
- estimation via a regression that controls for time period and group or unit **(fixed effects)**
- CIA relies on **no time-variant confounders**: all omitted variables must be constant over time

Notation for time periods

Up to now:

- Potential outcomes: Y_{0i}, Y_{1i}
- Definition linking them: $\tau_i \equiv Y_{1i} - Y_{0i}$

With two time periods:

- Potential outcomes: $Y_{0i,t}$, $Y_{1i,t}$
- Definitions linking them:

		time period t: 0	time period t: 1
treatment condition d	0	$Y_{0i,0}$	$Y_{0i,1} = Y_{0i,0} + \lambda_i$
	1	$Y_{1i,0} = Y_{0i,0} + \tau_{i,0}$	$Y_{1i,1} = Y_{0i,0} + \lambda_i + \tau_{i,1}$

NB: This is notation, not an assumption.

Notation for groups

Suppose units belong to one of two groups, T and C, with neither exposed to treatment in period 0 and group T exposed to treatment in period 1. Let g_i denote i's group.

For example,

$$E[Y_{1i,1} \mid g_i = T]$$

is the average potential outcome under treatment in time period 1 for units in group T.

Two groups, two time periods

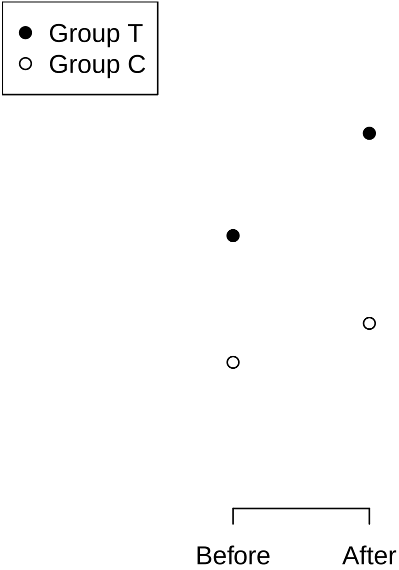

Adding notation

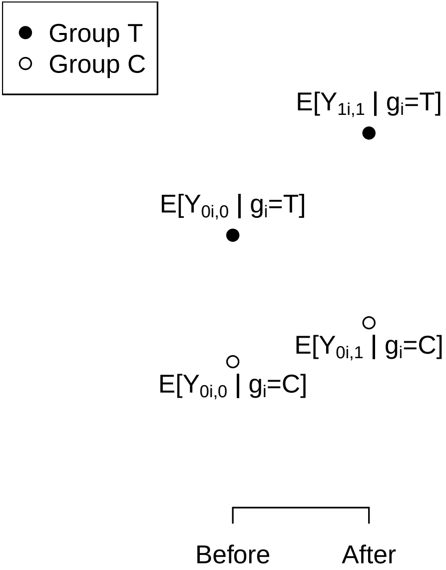

Before-and-after in group *C*

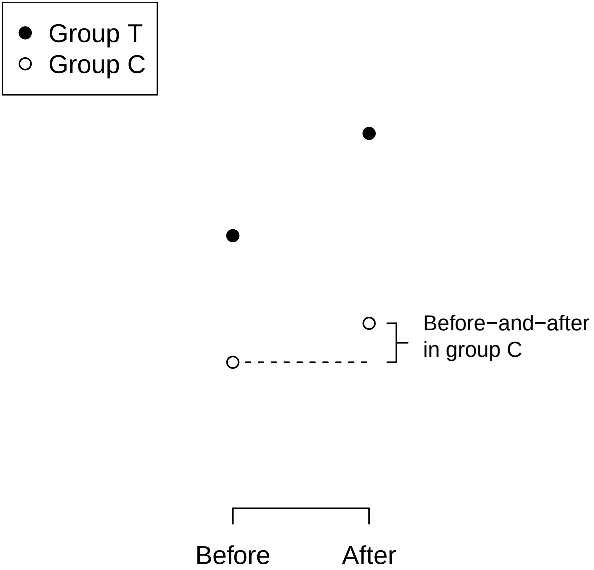

Before-and-after in group C

After-minus-before in group C is

$$E[Y_{0i,1} \mid g_i = C] - E[Y_{0i,0} \mid g_i = C]$$

We use the definitions above to restate in terms of the time trend:

$$\begin{aligned}
&= E[Y_{0i,0} + \lambda_i \mid g_i = C] - E[Y_{0i,0} \mid g_i = C] \\
&= E[\lambda_i \mid g_i = C] + E[Y_{0i,0} \mid g_i = C] - E[Y_{0i,0} \mid g_i = C] \\
&= E[\lambda_i \mid g_i = C] \\
&= \text{Time trend in group C}
\end{aligned}$$

Before-and-after in group *T*

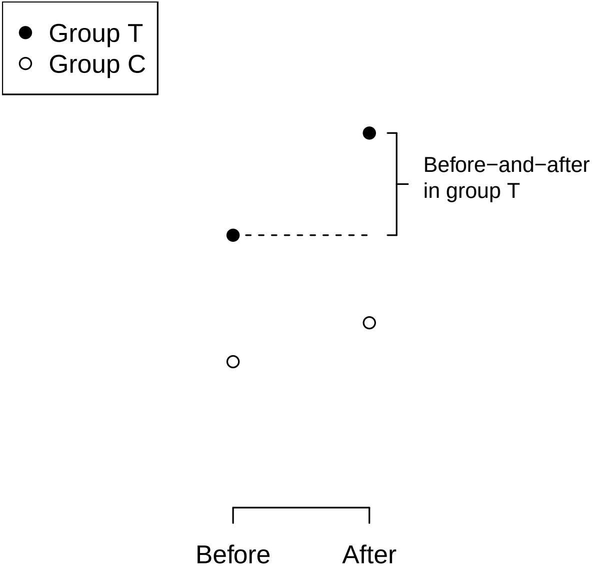

Before-and-after in group T

After-minus-before in group T is

$$E[Y_{1i,1} \mid g_i = T] - E[Y_{0i,0} \mid g_i = T]$$

We use the definitions above to restate in terms of time trend and ATE:

$$\begin{aligned}
&= E[Y_{0i,0} + \lambda_i + \tau_{i,1} \mid g_i = T] - E[Y_{0i,0} \mid g_i = T] \\
&= E[\lambda_i \mid g_i = T] + E[\tau_{i,1} \mid g_i = T] + E[Y_{0i,0} \mid g_i = T] - E[Y_{0i,0} \mid g_i = T] \\
&= E[\lambda_i \mid g_i = T] + E[\tau_{i,1} \mid g_i = T] \\
&= \text{Time trend in group T} + \text{ATE in group T}
\end{aligned}$$

Before-and-after in both groups

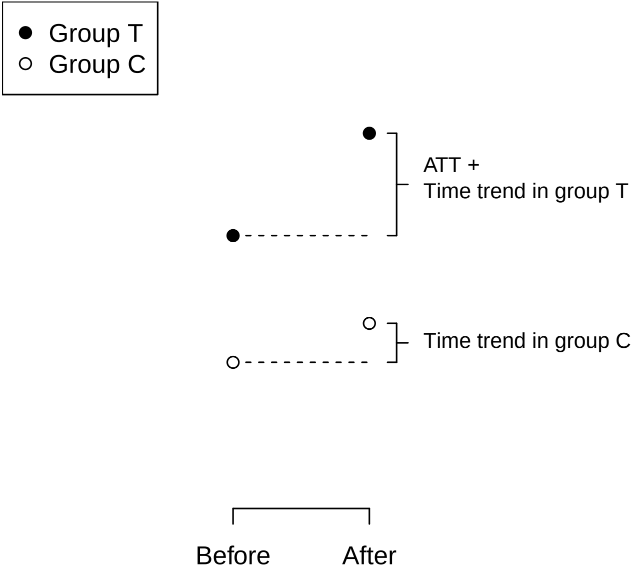

Common trend?

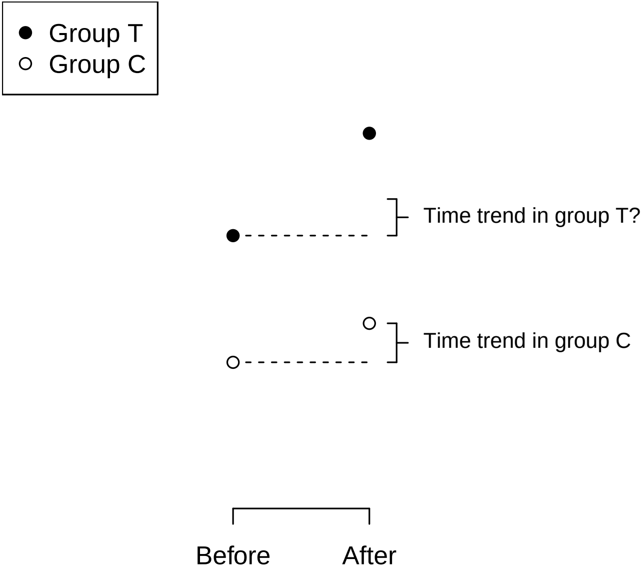

ATT given common trend assumption

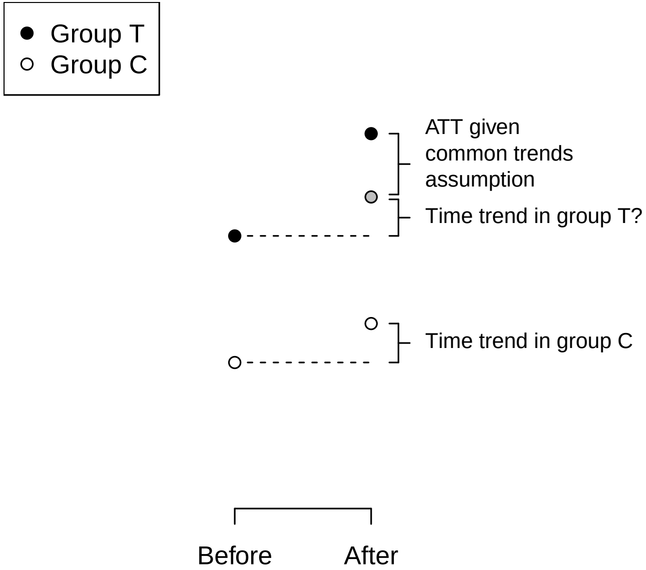

DIGM before

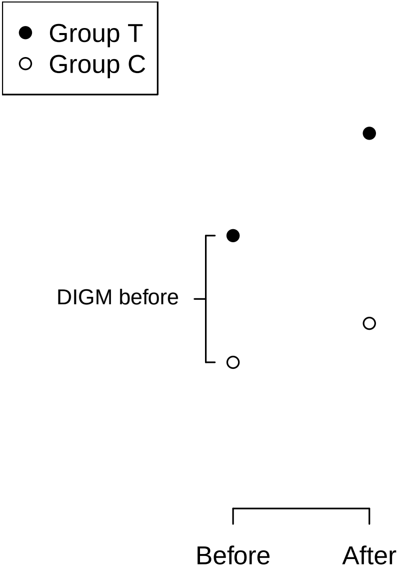

DIGM before

DIGM in the pre-treatment period is

$$E[Y_{0i,0} \mid g_i = T] - E[Y_{0i,0} \mid g_i = C]$$

By definition, this is selection bias.

DIGM after

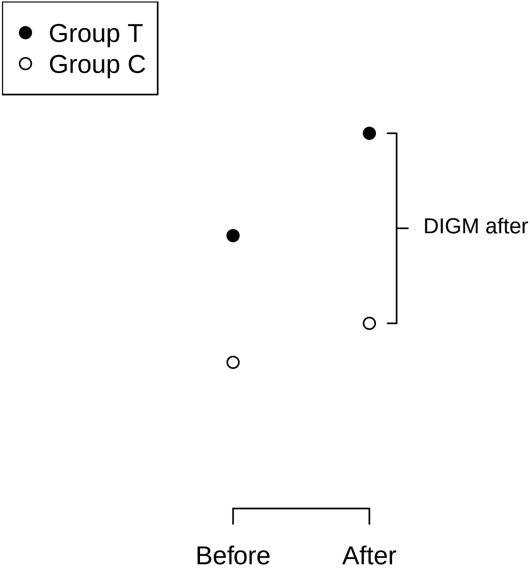

DIGM after

The DIGM at time 1 is

$$E[Y_{1i,1} \mid g_i = T] - E[Y_{0i,1} \mid g_i = C]$$

We use the definitions above to restate in terms of time trend, selection bias, and ATE:

$$\begin{aligned}
&= E[Y_{0i,0} + \lambda_i + \tau_{i,1} \mid g_i = T] \quad - \quad E[Y_{0i,0} + \lambda_i \mid g_i = C] \\
&= E[Y_{0i,0} \mid g_i = T] + E[\lambda_i \mid g_i = T] + E[\tau_{i,1} \mid g_i = T] - E[Y_{0i,0} \mid g_i = C] - E[\lambda_i \mid g_i = C] \\
&= E[Y_{0i,0} \mid g_i = T] - E[Y_{0i,0} \mid g_i = C] + E[\lambda_i \mid g_i = T] - E[\lambda_i \mid g_i = C] + E[\tau_{i,1} \mid g_i = T] \\
&= \text{Selection bias} + \text{Time trend in group T} - \text{Time trend in group C} + \text{ATE in group T}
\end{aligned}$$

Both DIGMs

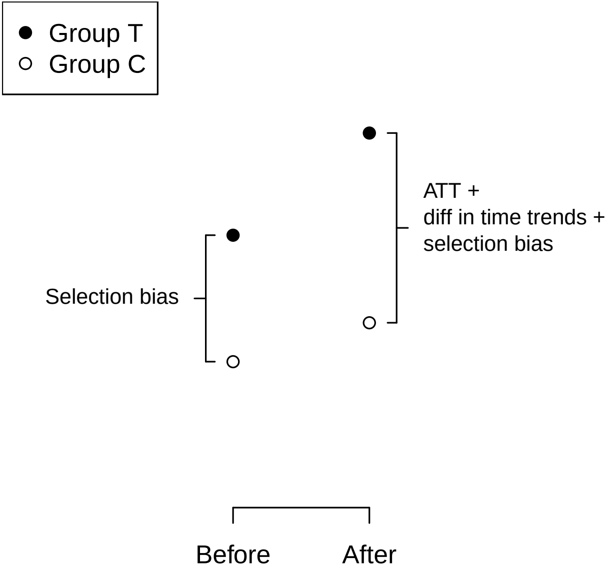

Group T
Group C
ATT +
diff in time trends +
selection bias
Selection bias
Before
After

ATT given common trends assumption

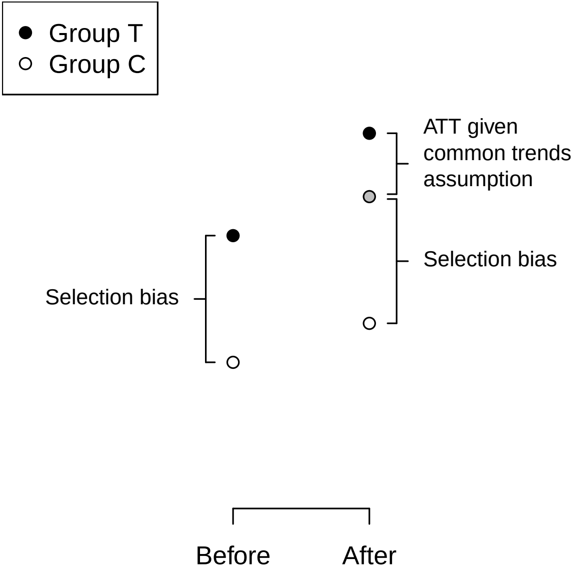

Can the common trends assumption be tested?

No. But common trends in several pre-treatment periods is suggestive:

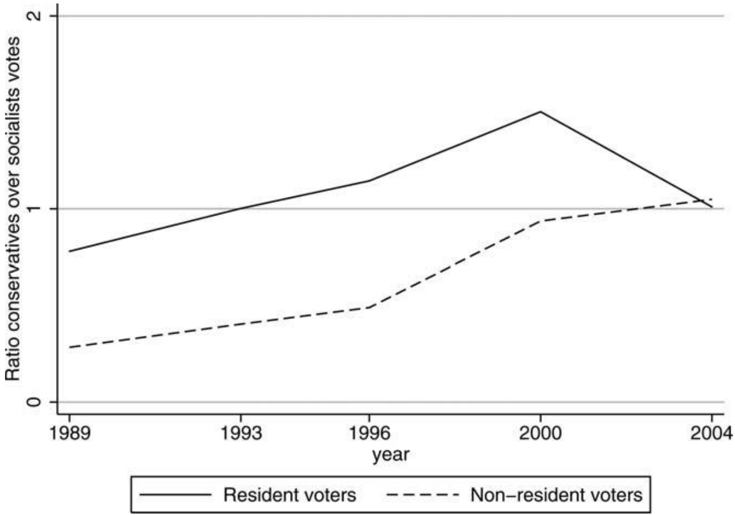

Dinas et al (2018) on political impact of refugees

- **Question**: Did the influx of refugees in Greece increase support for the right-wing Golden Dawn party in 2015?
- **Treatment**: Large number of refugees arriving in locality
- **Outcome**: Golden Dawn vote share in locality

To consider:

- What about a cross-sectional approach? What covariates might help?
- How might an IV approach help?
- How can we use variation over time in a diff-in-diff?

Dinas et al on the Golden Dawn (2)

Islands that received lots of refugees may vote differently even without the refugee influx.

Maybe that difference is constant over time.

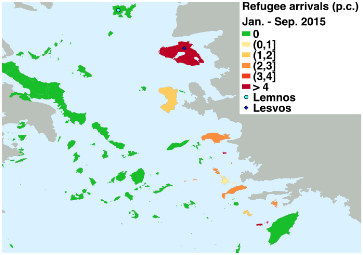

Common trends assumption: if they had not received refugees, islands that did receive refugee would have seen the same **change** in support for Golden Dawn as other islands.

To consider: are these other islands really *untreated*?

Dinas et al on the Golden Dawn (3)

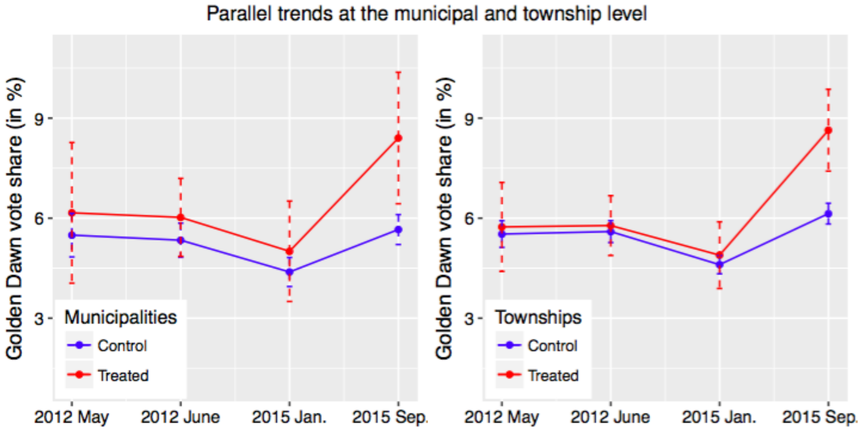

Diff-in-diff: implementation: method 1

Method 1: group-period interactions

- data structure: two rows for each municipality (elections of Jan. 2015, Sept. 2015)
- `evertr`: 1 for municipalities that received refugees
- `post`: 1 for election after the influx
- `gdper`: support for Golden Dawn

municipality	evertr	post	gdper
Αίγινας	0	0	6.363300
Αίγινας	0	1	7.617789
Αγίου Βασιλείου	0	0	2.714932
Αγίου Βασιλείου	0	1	3.694069
Αγίου Ευστρατίου	0	0	4.878048
Αγίου Ευστρατίου	0	1	5.988024
Αγίου Νικολάου	0	0	3.159049
Αγίου Νικολάου	0	1	4.604597
Αγαθονησίου	1	0	3.278688
Αγαθονησίου	1	1	5.000000
Αγκιστρίου	0	0	6.129032
Αγκιστρίου	0	1	9.981852
Αλοννήσου	0	0	5.727377
Αλοννήσου	0	1	5.976096

Use command

```
lm(gdper ~ evertr*post)
```

to run regression:

$$\text{gdper}_{mt} = \beta_0 + \beta_1 \text{evertr}_m + \beta_2 \text{post}_t + \beta_3 \text{evertr}_m \times \text{post}_t$$

Interpretation of coefficients using method 1

$$\text{gdper}_{mt} = \beta_0 + \beta_1 \text{evertr}_m + \beta_2 \text{post}_t + \beta_3 \text{evertr}_m \times \text{post}_t$$

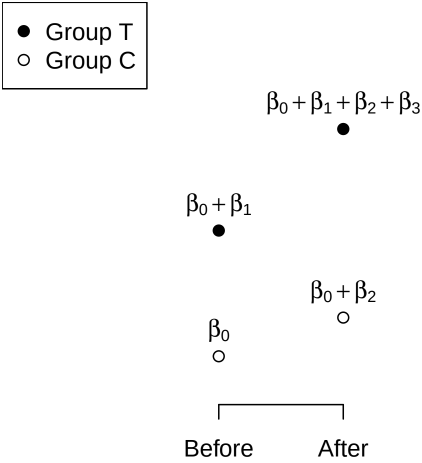

Interpretation of coefficients using method 1

$$\text{gdper}_{mt} = \beta_0 + \beta_1 \text{evertr}_m + \beta_2 \text{post}_t + \beta_3 \text{evertr}_m \times \text{post}_t$$

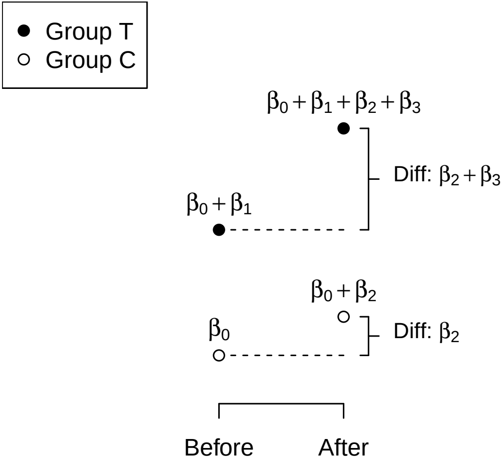

Interpretation of coefficients using method 1

$$\text{gdper}_{mt} = \beta_0 + \beta_1 \text{evertr}_m + \beta_2 \text{post}_t + \beta_3 \text{evertr}_m \times \text{post}_t$$

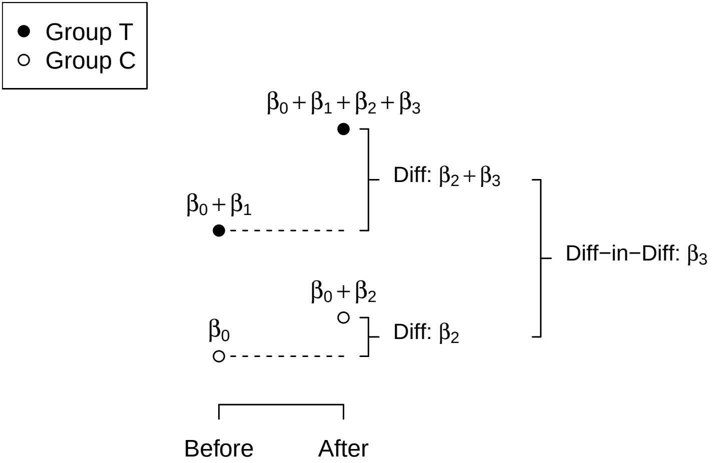

Diff-in-diff implementation: method 1

Method 1: group-period interactions

Regression output:

```
Call:
lm(formula = gdper ~ evertr * post, data = d[!is.na(d$muni) &
    d$election %in% c("pre3", "post"), ])

Residuals:
    Min      1Q  Median      3Q     Max
-5.6730 -1.6899 -0.2142  1.3753  9.1088

Coefficients:
            Estimate Std. Error t value Pr(>|t|)
(Intercept)   4.3810     0.2440  17.954  < 2e-16 ***
evertr        0.6257     0.6866   0.911 0.363315
post          1.2921     0.3451   3.744 0.000241 ***
evertr:post   2.1052     0.9710   2.168 0.031413 *
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 2.223 on 186 degrees of freedom
Multiple R-squared:  0.1769,	Adjusted R-squared:  0.1637
F-statistic: 13.33 on 3 and 186 DF,  p-value: 6.435e-08
```

Diff-in-diff implementation: method 2

Method 2: group & time dummies and treatment indicator

- data structure: four rows for each municipality (elections of May 2012, June 2012, Jan. 2015, Sept. 2015)
- `evertr`: 1 for municipalities that received refugees
- `election`: date of election (factor)
- `treatment`: 1 if `evertr` = 1 and Sept. 2015
- `gdper`: support for Golden Dawn

municipality	evertr	election	treatment	gdper
Αίγινας	0	May12	0	7.9822884
Αίγινας	0	June12	0	7.2771678
Αίγινας	0	Jan15	0	6.3633003
Αίγινας	0	Sept15	0	7.6177893
Αγίου Βασιλείου	0	May12	0	2.5829175
Αγίου Βασιλείου	0	June12	0	4.2843981
Αγίου Βασιλείου	0	Jan15	0	2.7149322
Αγίου Βασιλείου	0	Sept15	0	3.6940687
Αγίου Ευστρατίου	0	May12	0	4.9549551
Αγίου Ευστρατίου	0	June12	0	4.7619047
Αγίου Ευστρατίου	0	Jan15	0	4.8780484
Αγίου Ευστρατίου	0	Sept15	0	5.9880238
Αγίου Νικολάου	0	May12	0	2.8652139
Αγίου Νικολάου	0	June12	0	3.0493212
Αγίου Νικολάου	0	Jan15	0	3.1590488
Αγίου Νικολάου	0	Sept15	0	4.6045966
Αγαθονησίου	1	May12	0	3.5714288
Αγαθονησίου	1	June12	0	4.6875000
Αγαθονησίου	1	Jan15	0	3.2786884
Αγαθονησίου	1	Sept15	1	5.0000000

Use command

```
lm(gdper ~ as.factor(election) + evertr + treatment -1)
```

to run regression:

$$\text{gdper}_{mt} = \alpha_t + \beta_1 \text{evertr}_m + \beta_2 \text{treatment}_{mt}$$

Diff-in-diff implementation: method 2

Method 2: group & time dummies and treatment indicator

Regression output:

```
Call:
lm(formula = gdper ~ evertr + as.factor(election) + treatment -
    1, data = d[!is.na(d$muni), ])

Residuals:
    Min      1Q  Median      3Q     Max
-5.6730 -1.8094 -0.3837  1.2926 13.3359

Coefficients:
                          Estimate Std. Error t value Pr(>|t|)
evertr                      0.6521     0.4488   1.453   0.1471
as.factor(election)Sept15   5.6730     0.2763  20.534   <2e-16 ***
as.factor(election)Jan15    4.3776     0.2644  16.558   <2e-16 ***
as.factor(election)June12   5.3529     0.2644  20.247   <2e-16 ***
as.factor(election)May12    5.5027     0.2644  20.813   <2e-16 ***
treatment                   2.0788     0.8976   2.316   0.0211 *
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Residual standard error: 2.517 on 374 degrees of freedom
Multiple R-squared:  0.8253,  Adjusted R-squared:  0.8225
F-statistic: 294.4 on 6 and 374 DF,  p-value: < 2.2e-16
```

Diff-in-diff implementation: method 3

Method 3: unit & time dummies and treatment indicator

We have controlled for group differences with a group dummy.

What about using *municipality* dummies instead?

municipality	evertr	election	treatment	gdper
Αίγινας	0	May12	0	7.9822884
Αίγινας	0	June12	0	7.2771678
Αίγινας	0	Jan15	0	6.3633003
Αίγινας	0	Sept15	0	7.6177893
Αγίου Βασιλείου	0	May12	0	2.5829175
Αγίου Βασιλείου	0	June12	0	4.2843981
Αγίου Βασιλείου	0	Jan15	0	2.7149322
Αγίου Βασιλείου	0	Sept15	0	3.6940687
Αγίου Ευστρατίου	0	May12	0	4.9549551
Αγίου Ευστρατίου	0	June12	0	4.7619047
Αγίου Ευστρατίου	0	Jan15	0	4.8780484
Αγίου Ευστρατίου	0	Sept15	0	5.9880238
Αγίου Νικολάου	0	May12	0	2.8652139
Αγίου Νικολάου	0	June12	0	3.0493212
Αγίου Νικολάου	0	Jan15	0	3.1590488
Αγίου Νικολάου	0	Sept15	0	4.6045966
Αγαθονησίου	1	May12	0	3.5714288
Αγαθονησίου	1	June12	0	4.6875000
Αγαθονησίου	1	Jan15	0	3.2786884
Αγαθονησίου	1	Sept15	1	5.0000000

Use command

```
lm(gdper ~ treatment + as.factor(election) + as.factor(muni) -1)
```

to run regression

$$\text{gdper}_{mt} = \beta_1 \text{treatment}_{mt} + \alpha_t + \gamma_m$$

Diff-in-diff implementation: method 3

Method 3: unit & time dummies and a treatment indicator

Regression output:

```
Call:
lm(formula = gdper ~ treatment + as.factor(election) + as.factor(muni) -
    1, data = d[use, ])

Residuals:
    Min      1Q  Median      3Q     Max
-4.5855 -0.5236 -0.0003  0.4404  6.9990

Coefficients:
                                   Estimate Std. Error t value Pr(>|t|)
treatment                            2.0788     0.3948   5.265 2.79e-07 ***
as.factor(election)Sept15            7.7566     0.5635  13.764  < 2e-16 ***
as.factor(election)Jan15             6.4612     0.5624  11.488  < 2e-16 ***
as.factor(election)June12            7.4365     0.5624  13.222  < 2e-16 ***
as.factor(election)May12             7.5862     0.5624  13.489  < 2e-16 ***
as.factor(muni)Αγίου Βασιλείου      -3.9911     0.7829  -5.098 6.33e-07 ***
as.factor(muni)Αγίου Ευστρατίου     -2.1644     0.7829  -2.765 0.006078 **
as.factor(muni)Αγίου Νικολάου       -3.8906     0.7829  -4.969 1.17e-06 ***
as.factor(muni)Αγαθονησίου          -3.6954     0.7891  -4.683 4.41e-06 ***
as.factor(muni)Αγκιστρίου            4.2533     0.7829   5.433 1.20e-07 ***
as.factor(muni)Αλοννήσου            -2.1973     0.7829  -2.807 0.005357 **
as.factor(muni)Αμαρίου              -4.5633     0.7829  -5.828 1.53e-08 ***
```

[result clipped]

Diff-in-diff implementation: group dummy or unit dummies?

Unit dummies produce lower standard errors, so why not always use them instead of **group dummies**?

Basic diff-in-diff can be done in two kinds of data:

- panel data: same units at several points in time
- repeated cross-section: may not be same units

Cannot use unit dummies with repeated cross-section.

Problem with repeated observations

Above we got lower standard errors by using more periods:

- using elections of Jan 2015 and Sept 2015: 0.97
- adding elections of May 2012 and June 2012: 0.90

Where does this stop? What if Greece had more elections – still okay to use all of them?

Assumptions for standard errors

What does the standard error mean?

How could you tell from a simulation if it were correct?

Basic assumptions behind OLS standard errors:

- Variance of regression errors independent of X (homoskedastic)
- Regression errors independent each other (uncorrelated across units)

Second assumption likely to be met in DiD case?

Addressing correlations among errors

Common assumption is that regression errors are independent except within clusters → **cluster-robust standard errors.**

See `estimatr` or `lfe` packages (in `R`).

In Stata, see `cluster()`.