Causal inference week 7: Panel diff-in-diff

Andy Eggers

Oxford DPIR

HT 2019

Introduction

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
- Use observations at more than one point in time

Today: Generalizing the diff-in-diff.

Simplest diff-in-diff

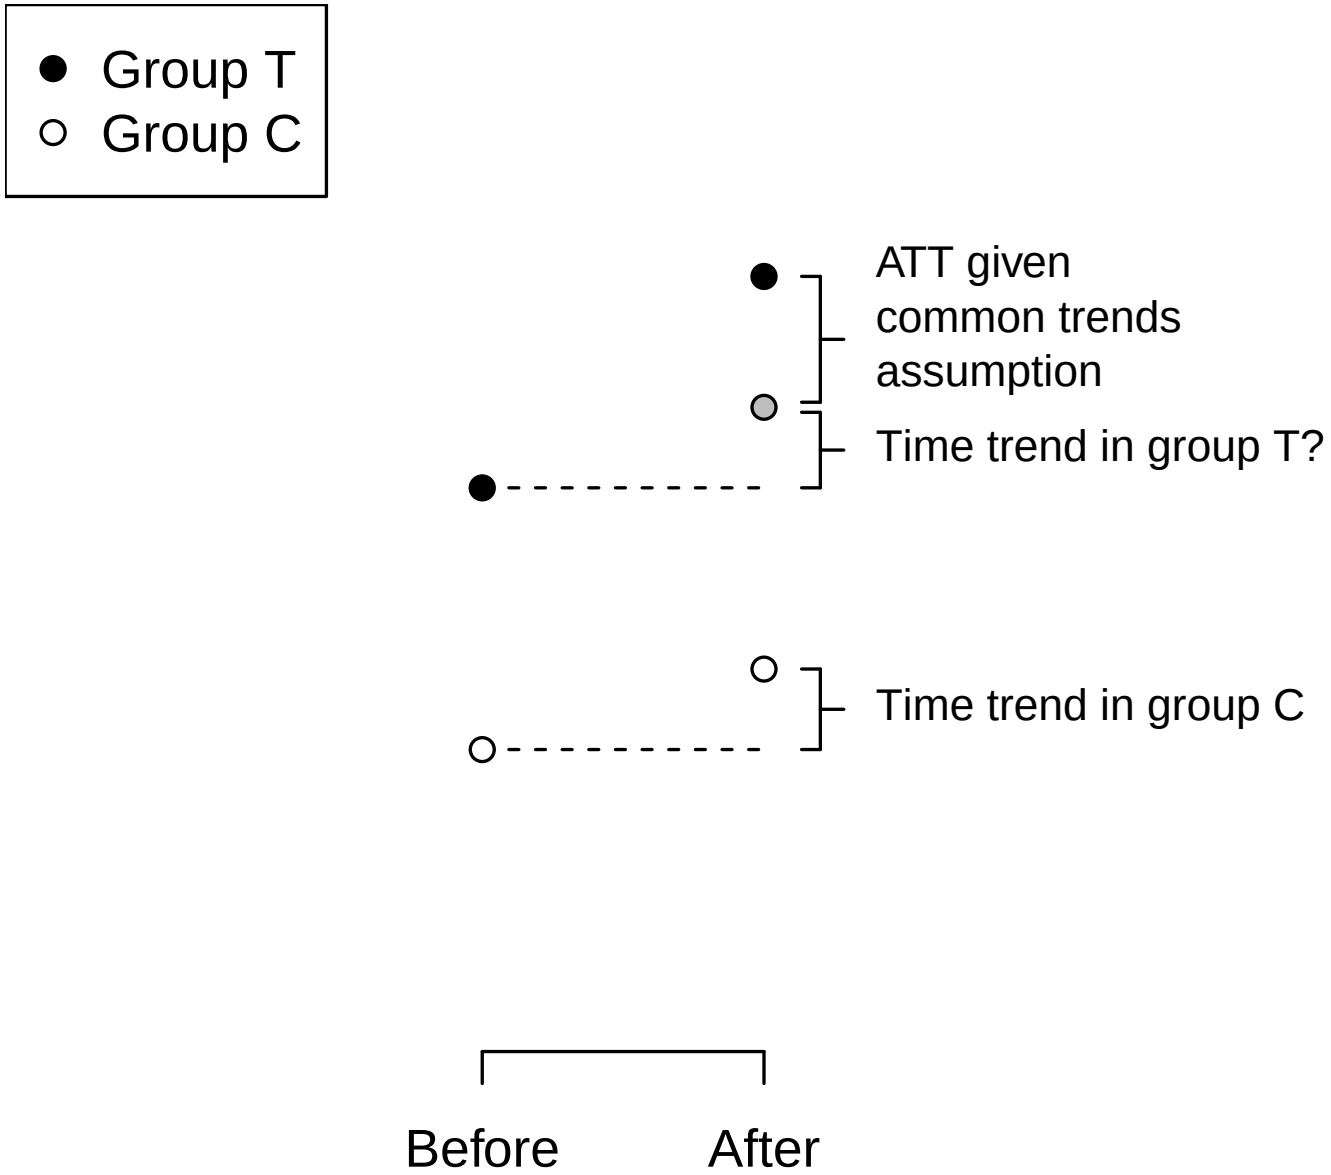

Dinas et al on the Golden Dawn

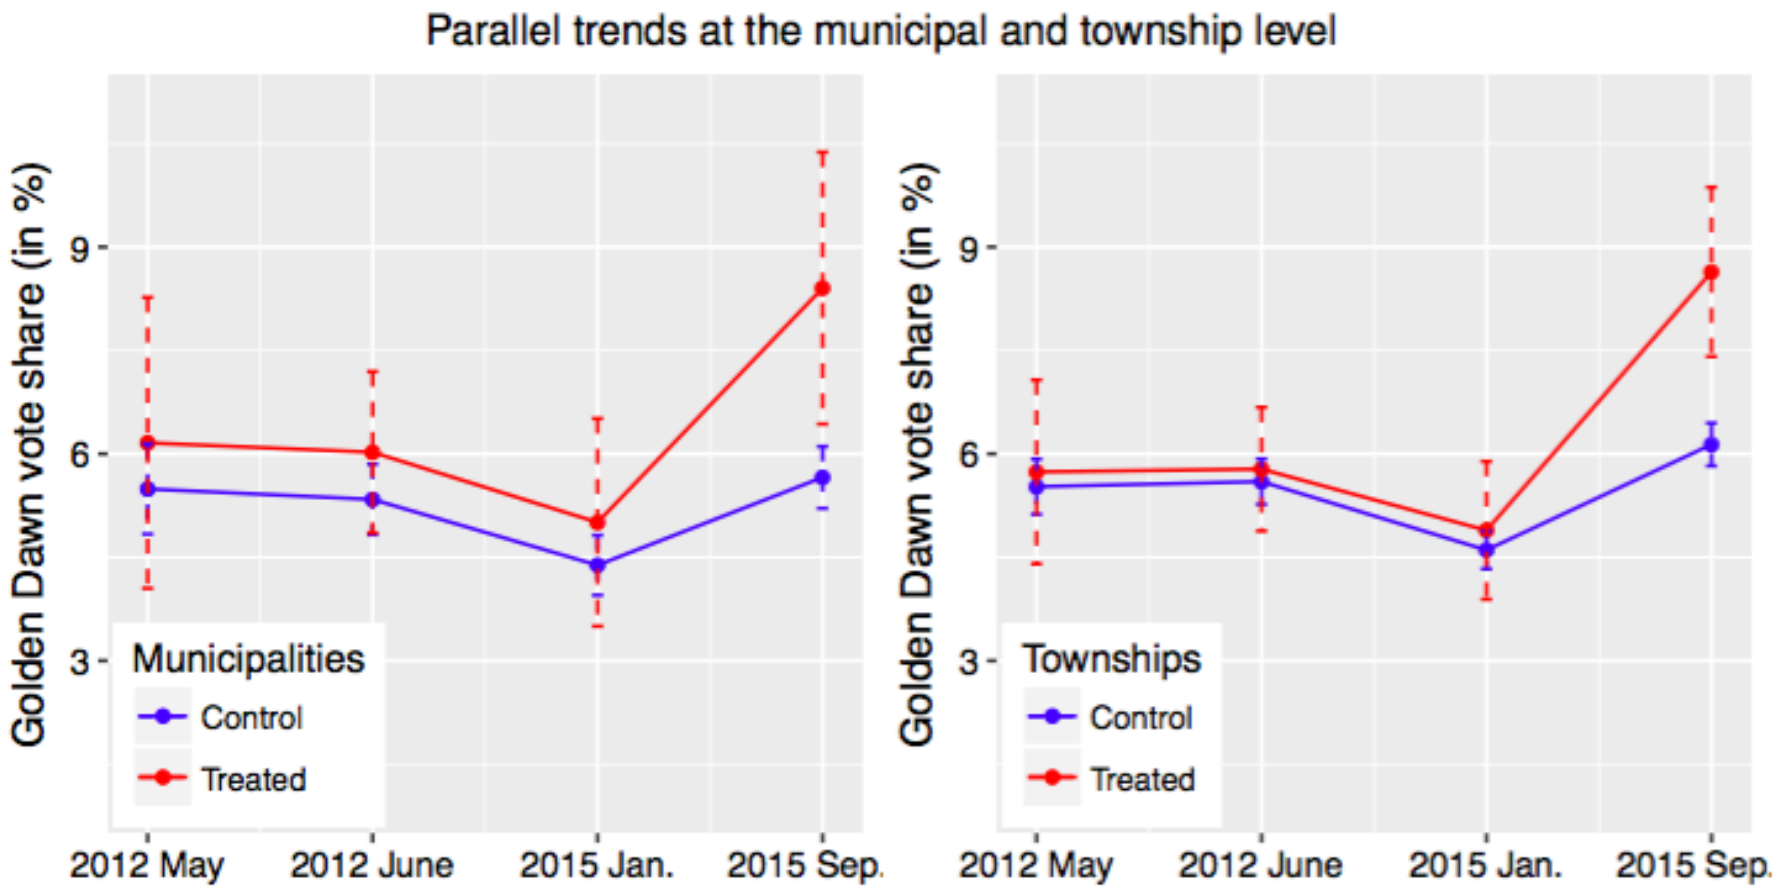

Diff-in-diff with unit and time period dummies

Given **panel data**, you can run

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

to estimate coefficients of regression

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

which **MM** would write as

$$\text{gdper}_{mt} = \beta_1 \text{treatment}_{mt} + \sum_{j=1}^{T} \alpha_j \text{Election}_{jt} + \sum_{k=1}^{M} \gamma_k \text{Municipality}_k.$$

Regression output (truncated):

```
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 ***
```

Panel diff-in-diff: main idea

Given a simple diff-in-diff in panel data, we can run this regression:

$$Y_{it} = \beta_1 \text{treatment}_{it} + \alpha_t + \gamma_i$$

But in panel data we can run this regression for **any** type of treatment applied in **any** pattern.

Under what assumptions is β_1 an unbiased estimator of the ATT?
Two ways of putting it:

- **parallel trends**: time trends unrelated to treatment received; i.e., if treatment did not vary, treated and untreated units would follow common trends
- **no time-varying confounders**: any omitted variables related to treatment must be fixed over time

“English Bacon”: research question

Does the UK government favor politically-aligned local councils when distributing targeted grants?

Consider assessing this with cross-sectional data (Ward & John, 1999).

- What covariates would you need?
- What about IV?
- What about RDD?

"English Bacon": overview

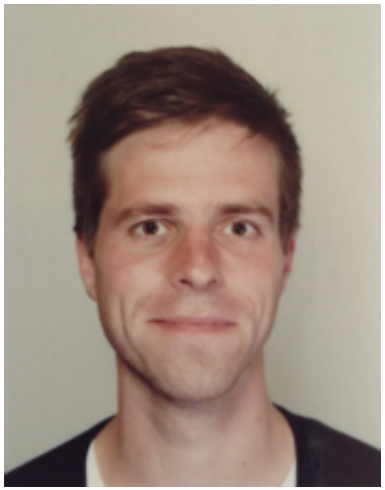

Alex Fouirnaies

Hande Mutlu-Eren

- Assemble panel data for 1992-2012 with
  - partisan composition of local councils
  - grants allocated (per capita)
- Define treatment Copartisan_{it} as: council i's majority and PM are copartisans in year t
- Regress grants on (lagged) treatment and
  - council dummies (council fixed effects)
  - year dummies (year fixed effects)
  - council-year interactions (council-specific linear time trends)
- Test for larger effects before elections, in swing councils, etc. (more next week on treatment effect heterogeneity)

"English Bacon": basic regression (no unit-specific time trends)

We might expect grants at t to depend on $\text{Copartisan}_{i,t-1}$.

We estimate

$$\text{LogOfGrantsPerCapita}_{it} = \beta_1 \text{Copartisan}_{i,t-1} + \alpha_t + \gamma_i$$

with this syntax

```
lm(lngrants ~ treatment_lag1 + as.factor(year) + as.factor(council) )
```

to estimate effect of alignment k years ago on grants now.

Regression output (truncated)

```
> summary(lm(lngrants ~ treatment_lag1 + as.factor(year) + as.factor(councilnumber), data = d[use,]))

Call:
lm(formula = lngrants ~ treatment_lag1 + as.factor(year) + as.factor(councilnumber),
    data = d[use, ])

Residuals:
    Min      1Q  Median      3Q     Max
-3.2089 -0.2575  0.0148  0.2418  4.3960

Coefficients:
                          Estimate Std. Error t value Pr(>|t|)
(Intercept)               0.4260079  0.0998185   4.268 2.00e-05 ***
treatment_lag1            0.1269693  0.0149935   8.468  < 2e-16 ***
as.factor(year)1993       0.1383544  0.0318947   4.338 1.46e-05 ***
as.factor(year)1994       0.1507899  0.0317807   4.745 2.14e-06 ***
as.factor(year)1995       0.0719591  0.0321439   2.239 0.025214 *
as.factor(year)1996       0.0982419  0.0319827   3.072 0.002138 **
as.factor(year)1997       0.0837433  0.0321113   2.608 0.009132 **
as.factor(year)1998       0.0833194  0.0318026   2.620 0.008818 **
as.factor(year)1999       0.1550595  0.0317998   4.876 1.11e-06 ***
as.factor(year)2000       0.2804133  0.0317496   8.832  < 2e-16 ***
as.factor(year)2001       0.4673901  0.0315067  14.835  < 2e-16 ***
as.factor(year)2002       0.6083286  0.0312453  19.469  < 2e-16 ***
as.factor(year)2003       1.1727693  0.0309422  37.902  < 2e-16 ***
as.factor(year)2004       1.3882406  0.0311179  44.612  < 2e-16 ***
as.factor(year)2005       1.5416901  0.0311378  49.512  < 2e-16 ***
as.factor(year)2006       2.1168448  0.0310975  68.071  < 2e-16 ***
as.factor(year)2007       2.2289501  0.0313889  71.011  < 2e-16 ***
as.factor(year)2008       2.2081314  0.0313613  70.409  < 2e-16 ***
as.factor(year)2009       2.3290924  0.0322764  72.161  < 2e-16 ***
as.factor(year)2010       2.3613410  0.0322684  73.178  < 2e-16 ***
as.factor(councilnumber)2 0.4049491  0.1336061   3.031 0.002448 **
as.factor(councilnumber)3 0.2732103  0.1393311   1.961 0.049940 *
```

Effect of partisan alignment at $t - k$ on log grants

For lags of $k = 0, 1, \ldots, 6$ years:

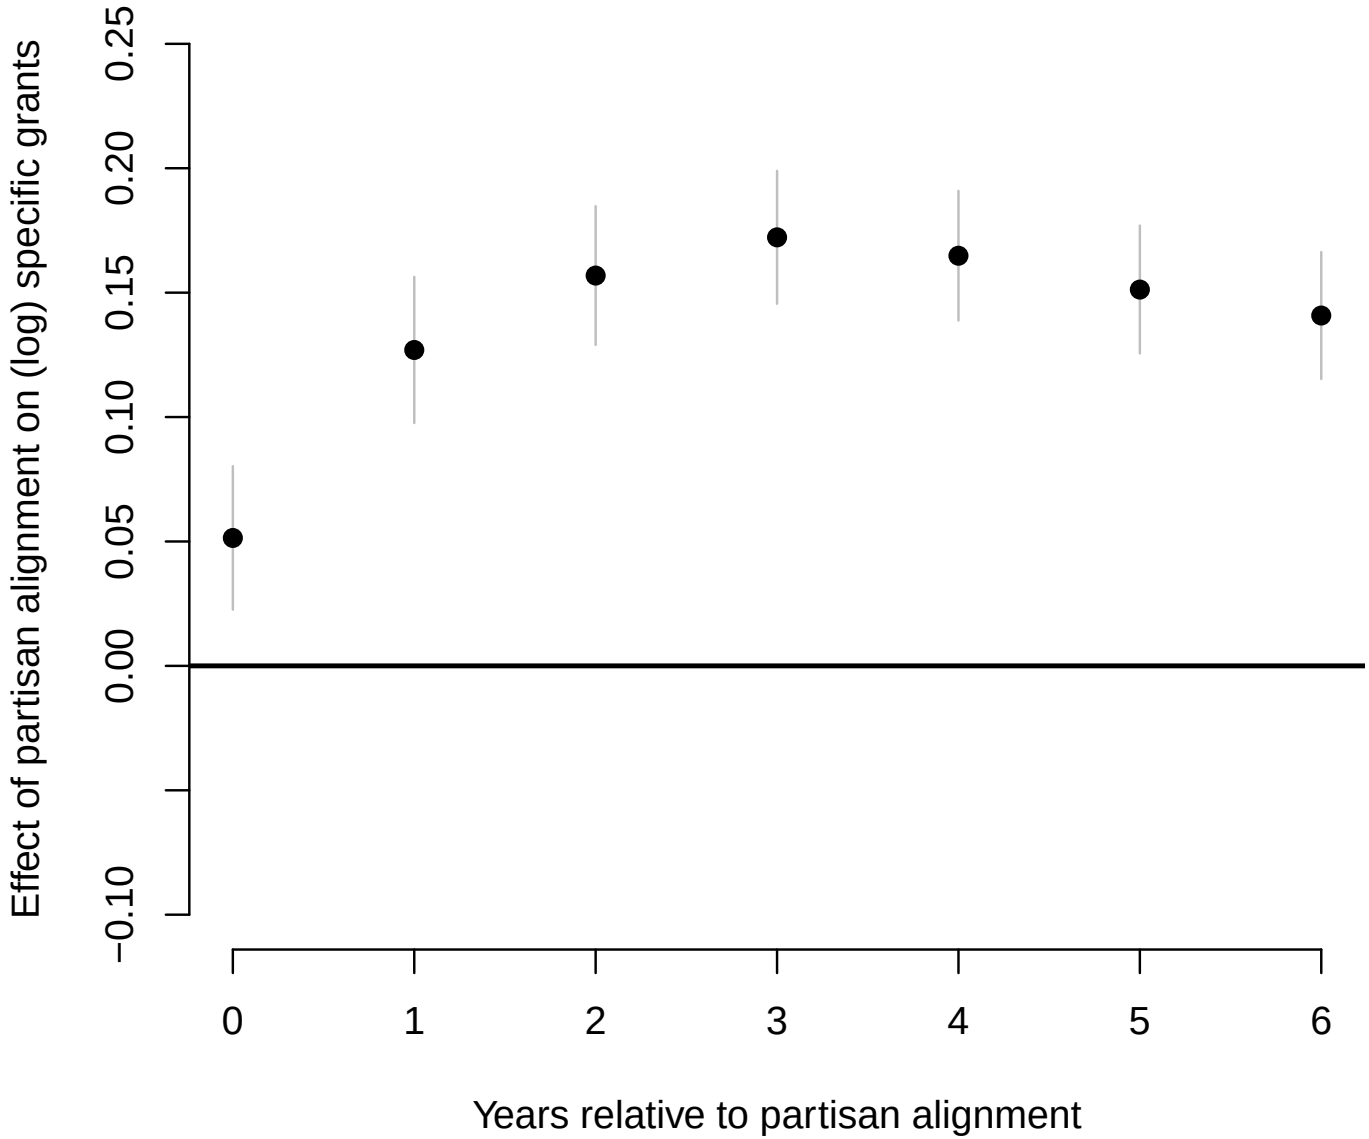

What could explain this finding?

Recall: regression equation was

$$\text{LogOfGrantsPerCapita}_{it} = \beta_1 \text{Copartisan}_{i,t-k} + \alpha_t + \gamma_i$$

Could we find positive β_1 because

- rural councils get fewer per-capita grants and tend to be Conservative; mostly Labour governments in 1992-2012?
- Labour governments gave more grants when they were in government, and there are more Labour councils in the data?

What else could explain it?

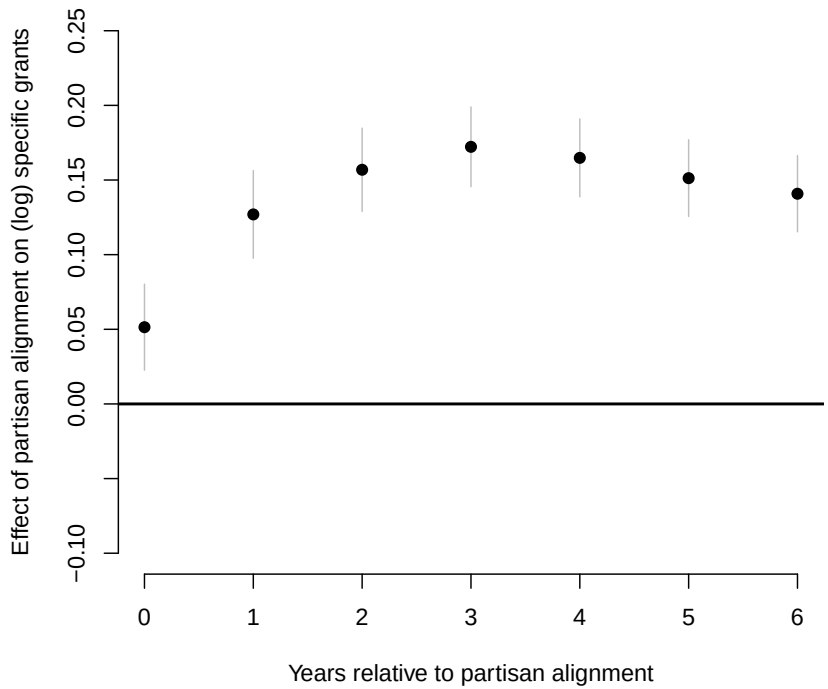

Explaining panel DiD findings

Suppose the **data generating process (DGP)** is

$$Y_{it} = \beta_1 D_{it} + \eta \mathbf{X}_t + \zeta \mathbf{U}_i + \psi \mathbf{V}_{it} + \omega_{it}$$

where

- $\mathbf{X}_t$ are time-specific variables that affect outcomes for all units the same way (e.g. budget for targeted grants),
- $\mathbf{U}_i$ are unit-specific variables that are constant over time (e.g. urban/rural character, presence of Roman ruins),
- $\mathbf{V}_{it}$ are variables that may vary within units over time (e.g. presence of ambitious council member, local economic situation), and
- ω_{it} is random noise.

In panel-DiD analysis where we estimate $Y_{it} = \beta_1 D_{it} + \alpha_t + \gamma_i + \epsilon_{it}$,

- time dummies (α_t) control for all $\mathbf{X}_t$
- unit dummies (γ_i) control for all $\mathbf{U}_i$

so the only possible confounders are $\mathbf{V}_{it}$.

What could explain this finding? (2)

Recall: regression equation was

$$\text{LogOfGrantsPerCapita}_{it} = \beta_1 \text{Copartisan}_{i,t-k} + \alpha_t + \gamma_i$$

What confounders might vary with treatment over time within units?

- Labour councils had growing needs, Conservative councils shrinking needs?
- Labour councillors improving?
- others?

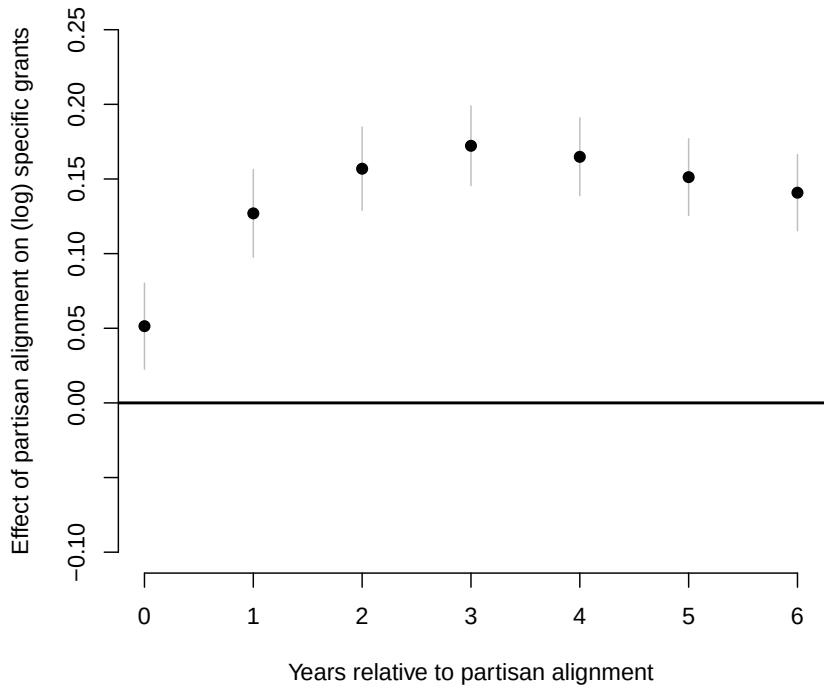

Relaxing the parallel trends assumption

Regression equation was

$$\text{LogOfGrantsPerCapita}_{it} = \beta_1 \text{Copartisan}_{i,t-k} + \alpha_t + \gamma_i$$

but consider adding unit-specific linear time trends:

$$\text{LogOfGrantsPerCapita}_{it} = \beta_1 \text{Copartisan}_{i,t-k} + \alpha_t + \gamma_i + \gamma_i t$$

where t is the year. To implement (needs at least 3 years):

```
lm(lngrants ~ treatment_k + as.factor(year) + as.factor(council)*year )
```

(Could add $\texttt{year}^2$ or $\sqrt{\text{year}}$ or $\ln(\text{year})$ to make time trends non-linear.)

Effect over time, w. unit specific time trends

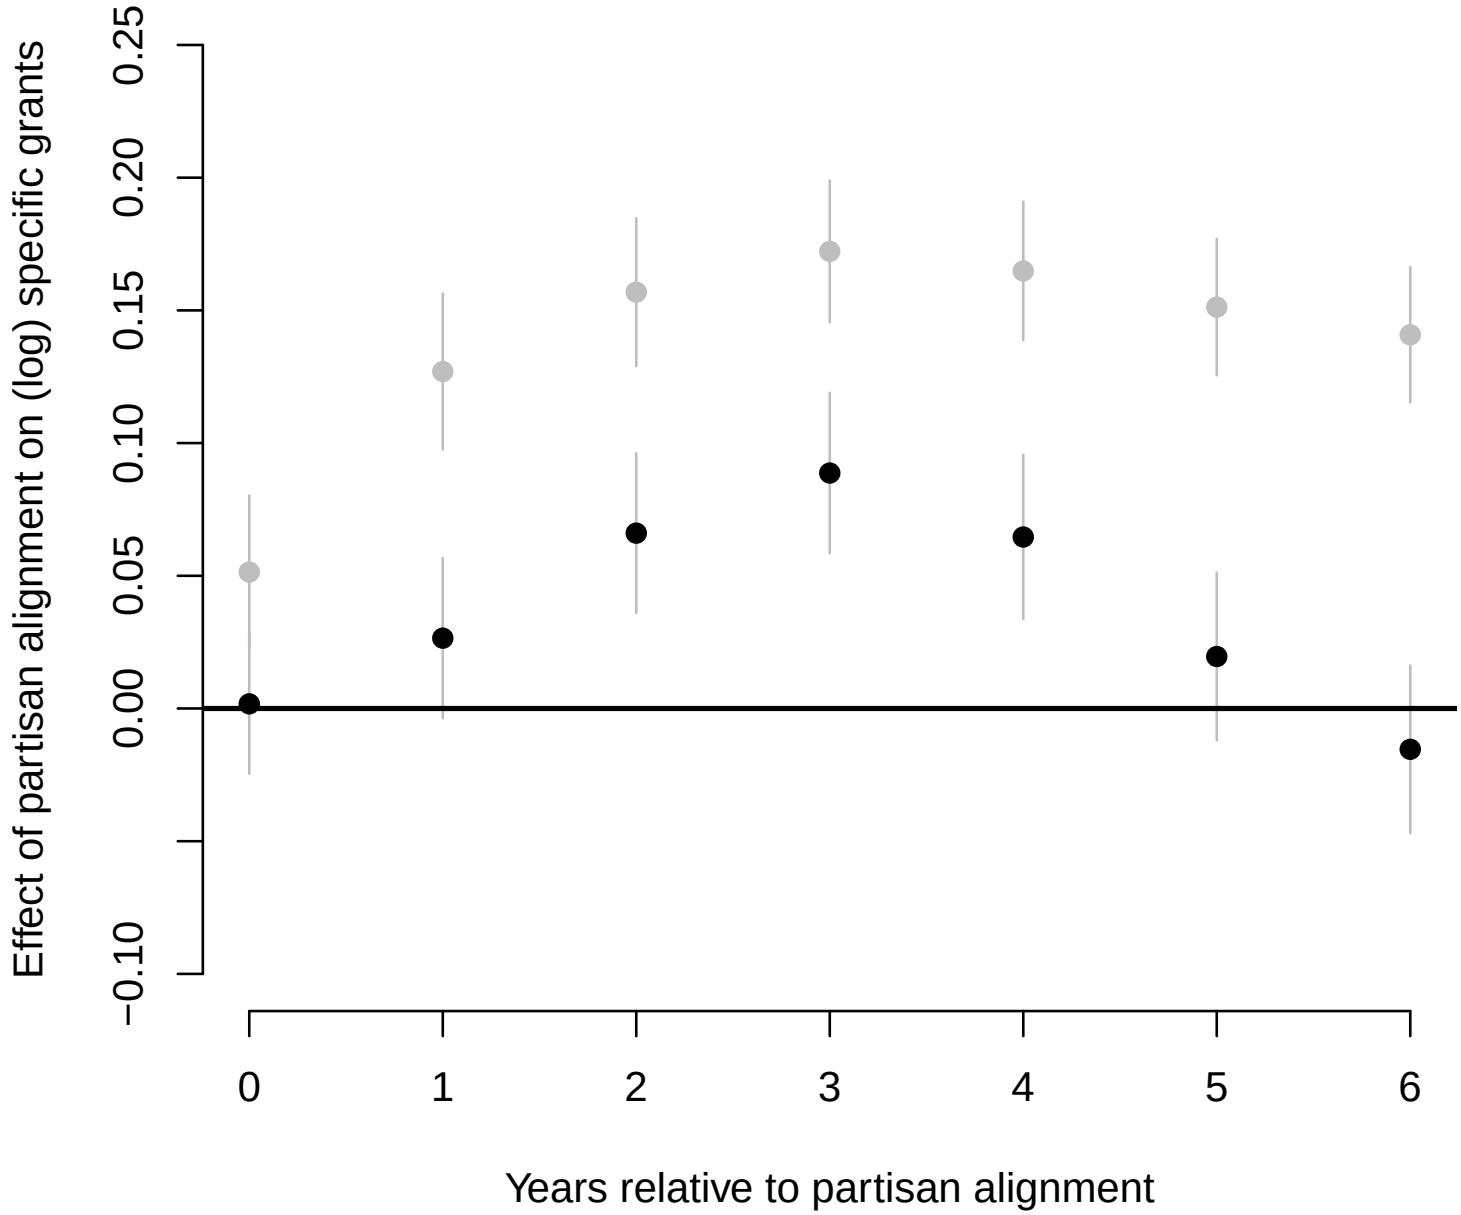

Testing assumptions in panel DiD

Unfortunately, no test as simple and transparent as the parallel trends plot.

The alternative:

$$\text{LogOfGrantsPerCapita}_{it} = \sum_{k=0}^{5} \beta_k \text{Copartisan}_{i,t-k} + \sum_{k=1}^{3} \theta_k \text{Copartisan}_{i,t+k} + \alpha_t + \gamma_i + \gamma_i t$$

i.e. include **lags** and **leads** of treatment in one regression.

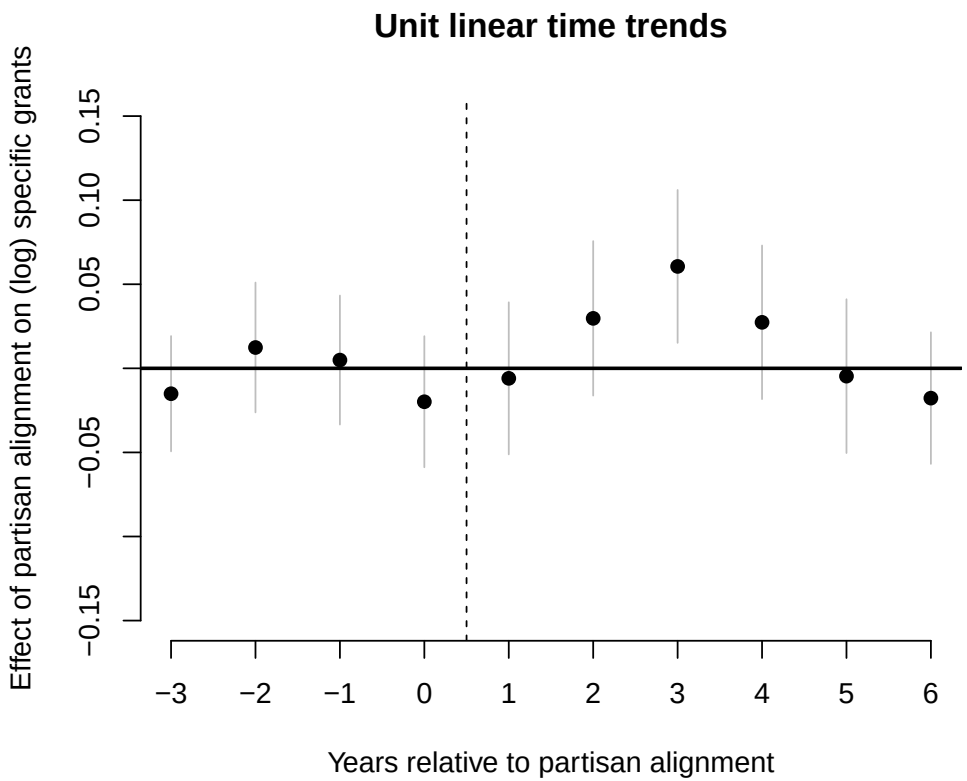

Kuziemko & Werker: effect of UNSC seat on US aid

Kuziemko and Werker on effects of UNSC seat

Kuziemko and Werker (2006), "How Much Is a Seat on the Security Council Worth? Foreign Aid and Bribery at the United Nations".

Question: What is the effect of a non-permanent seat on the UNSC affect aid from the US and UN?

Consider running this cross-sectional regression:

$$\text{AidFromUS}_i = \beta_0 + \beta_1 \text{UNSCseat}_i + \beta_2 \text{GDPperCapita}_i + \epsilon_i$$

- Would you expect β_1 to be positive or negative?
- What assumption is necessary to interpret that coefficient causally?
- Why might this assumption be violated?

Empirical Strategy

A positive association between foreign aid and council memberships would hardly be conclusive evidence of the vote-for-aid deals that we have hypothesized. Any omitted variable at the country level associated with both a country's propensity to serve on the council and its ability to extract aid from donor nations would lead to biased coefficients, almost certainly in the positive direction. Thus our basic empirical strategy is to look within countries across time and measure how their aid receipts changed as a function of their Security Council status. This

receipts changed as a function of their Security Council status. This estimation can be captured by the following equation, using a logarithmic specification following Alesina and Dollar (2000):

$$\ln(\text{Aid}_{irt}) = \alpha + \beta \times \text{SCMember}_{it} + \gamma \times \boldsymbol{X}_{it} + \boldsymbol{W}_{rt} + \boldsymbol{\eta}_t + \boldsymbol{\mu}_i + \epsilon_{1it}, \quad (1)$$

where i indexes countries, r indexes regions, t indexes years, SCMember is a dummy variable coded as one if country i is serving on the Security Council in year t, $\boldsymbol{X}$ is a vector of time-varying political and economic controls for each country, $\boldsymbol{W}$ is a regional quartic time trend,[2] $\boldsymbol{\eta}$ is a vector of year fixed effects, and $\boldsymbol{\mu}$ is a vector of country fixed effects. In the results that follow, we set Aid to equal either U.S. foreign aid or U.N. development aid.

TABLE 2
ECONOMIC AND MILITARY AID FROM THE UNITED STATES, 1946–2001
Dependent Variable: ln(Total Aid and Loans from U.S., $1996)

	(1)	(2)	(3)	(4)	(5)	(6)
SCMember	1.527	.466				
	[.379]***	[.239]*				
On SC, unimportant year			−.086	.03	.337	
			[.436]	[.407]	[.423]	
On SC, somewhat important year			.432	.474	.478	
			[.282]	[.294]	[.256]*	
On SC, important year			.99	.993	.741	
			[.440]**	[.455]**	[.397]*	
War occurring (> 1,000 deaths)				.007	−.058	−.051
				[.535]	[.624]	[.624]
Polity 2 score				.101	.037	.038
				[.034]***	[.028]	[.028]
ln(GDP per capita, $1996)					−.993	−1.009
					[.887]	[.888]
One year before election to SC						−.045
						[.204]
Year of election to SC						.42
						[.213]*
First year of serving on SC						.44
						[.321]
Second year of serving on SC						.715
						[.260]***
First year after finishing SC term						.202
						[.363]
Second year after finishing SC term						.15
						[.331]
Country and year fixed effects	No	Yes	Yes	Yes	Yes	Yes
Region quartics	No	Yes	Yes	Yes	Yes	Yes
Observations	5,425	5,425	5,425	4,902	3,616	3,616
R^2	0	.64	.64	.61	.6	.6

Using covariates in panel DiD analysis

receipts changed as a function of their Security Council status. This estimation can be captured by the following equation, using a logarithmic specification following Alesina and Dollar (2000):

$$\ln(\text{Aid}_{irt}) = \alpha + \beta \times \text{SCMember}_{it} + \gamma \times \boldsymbol{X}_{it} + \boldsymbol{W}_{rt} + \boldsymbol{\eta}_t + \boldsymbol{\mu}_i + \epsilon_{1it}, \quad (1)$$

where i indexes countries, r indexes regions, t indexes years, SCMember is a dummy variable coded as one if country i is serving on the Security Council in year t, $\boldsymbol{X}$ is a vector of time-varying political and economic controls for each country, $\boldsymbol{W}$ is a regional quartic time trend,[2] $\boldsymbol{\eta}$ is a vector of year fixed effects, and $\boldsymbol{\mu}$ is a vector of country fixed effects. In the results that follow, we set Aid to equal either U.S. foreign aid or U.N. development aid.

Using covariates in panel DiD analysis

TABLE 2
ECONOMIC AND MILITARY AID FROM THE UNITED STATES, 1946–2001
Dependent Variable: ln(Total Aid and Loans from U.S., $1996)

	(1)	(2)	(3)	(4)	(5)	(6)
SCMember	1.527	.466				
	[.379]***	[.239]*				
On SC, unimportant year			−.086	.03	.337	
			[.436]	[.407]	[.423]	
On SC, somewhat important year			.432	.474	.478	
			[.282]	[.294]	[.256]*	
On SC, important year			.99	.993	.741	
			[.440]**	[.455]**	[.397]*	
War occurring (> 1,000 deaths)				.007	−.058	−.051
				[.535]	[.624]	[.624]
Polity 2 score				.101	.037	.038
				[.034]***	[.028]	[.028]
ln(GDP per capita, $1996)					−.993	−1.009
					[.887]	[.888]
One year before election to SC						−.045
						[.204]
Year of election to SC						.42
						[.213]*
First year of serving on SC						.44
						[.321]
Second year of serving on SC						.715
						[.260]***
First year after finishing SC term						.202
						[.363]
Second year after finishing SC term						.15
						[.331]
Country and year fixed effects	No	Yes	Yes	Yes	Yes	Yes
Region quartics	No	Yes	Yes	Yes	Yes	Yes
Observations	5,425	5,425	5,425	4,902	3,616	3,616
R^2	0	.64	.64	.61	.6	.6

Levitt on effects of campaign spending

Levitt (1994), "Using Repeat Challengers to Estimate the Effect of Campaign Spending on Election Outcomes in the U.S. House".

Question: What is the effect of campaign spending on election outcomes?

Consider running this cross-sectional regression:

$$\text{DemVoteShare}_i = \beta_0 + \beta_1(\text{DemSpend}_i - \text{RepSpend}_i) + \beta_2 \text{DemPresVoteShare}_i + \epsilon_i$$

- Would you expect β_1 to be positive or negative?
- What assumption is necessary to interpret that coefficient causally?
- Why might this assumption be violated?

Levitt on effects of campaign spending

Levitt (1994), "Using Repeat Challengers to Estimate the Effect of Campaign Spending on Election Outcomes in the U.S. House".

Question: What is the effect of campaign spending on election outcomes?

Consider running this panel regression:

$$\text{DemVoteShare}_{it} = \beta_0 + \beta_1(\text{DemSpend}_{it} - \text{RepSpend}_{it}) + \alpha_t + \gamma_i + \epsilon_i$$

where γ_i is a dummy for each **candidate pair**.

- Would you expect β_1 to be positive or negative?
- What assumption is necessary to interpret that coefficient causally?
- Why might this assumption be violated?

First differences approach

Suppose again the **data generating process (DGP)** is

$$Y_{it} = \beta_1 D_{it} + \alpha \mathbf{X}_t + \gamma \mathbf{U}_i + \psi \mathbf{V}_{it} + \omega_{it}.$$

We estimated β_1 via regression with unit and time-period dummies.

First differences approach: Generate first difference of each variable, e.g.

$$\Delta Y_{it} = Y_{it} - Y_{i,t-1}$$

and then estimate

$$\Delta Y_{it} = \beta_1 \Delta D_{it} + \alpha_t,$$

i.e. regress differenced outcome on differenced treatment and year dummies (could add unit dummies for unit-specific linear time trends).

Generally gives similar results; **same** results if only two periods.

Ansell on effect of house prices on welfare attitudes

Ansell (2014), "The political economy of ownership: housing markets and the welfare state"

Question: How does variation in house prices affect homeowners? preferences regarding redistribution?

Consider running this cross-sectional regression:

$$\text{SupportForRedistribution}_i = \beta_0 + \beta_1 \text{PriceOfHome}_i + \beta_2 \text{Income}_i + \beta_3 \text{Age}_i + \epsilon_i.$$

- Would you expect β_1 to be positive or negative?
- What assumption is necessary to interpret β_1 causally?
- Why might this assumption be violated?

Ansell on effect of house prices on welfare attitudes (2)

Ansell (2014), "The political economy of ownership: housing markets and the welfare state"

Question: How does variation in house prices affect homeowners? preferences regarding redistribution?

Consider running this panel regression:

$$\text{SupportForRedistribution}_{it} = \beta_1 \text{PriceOfHome}_{it} + \alpha_t + \gamma_i$$

or (Ansell's actual basic specification – first differences)

$$\Delta\text{SupportForRedistribution}_{it} = \beta_1 \Delta\text{PriceOfHome}_{it} + \alpha_t$$

- Would you expect β_1 to be positive or negative?
- What assumption is necessary to interpret β_1 causally?
- Why might this assumption be violated?

Ansell's control strategy

Ansell (2014) controls for changes in

- home ownership
- household income
- party ID
- retired status

and controls for (i.e. allows time trends to vary by)

- age
- gender
- race

Adler on the “Waitrose effect”

Adler (2017 MPhil dissertation), “The other Waitrose effect”

Question: How does gentrification affect renters?

Consider running this cross-sectional regression:

$$\text{EvictionRate}_i = \beta_0 + \beta_1 \text{WaitroseNearby}_i + \beta_2 \text{UnemploymentRate}_i + \beta_3 \text{CrimeRate}_i + \epsilon_i.$$

- Would you expect β_1 to be positive or negative?
- What assumption is necessary to interpret β_1 causally?
- Why might this assumption be violated?

Adler on the “Waitrose effect” (2)

Adler (2017 MPhil dissertation), “The other Waitrose effect”

Question: How does gentrification affect renters?

Consider running this panel regression:

$$\text{EvictionRate}_{it} = \beta_1 \text{WaitroseNearby}_{it} + \alpha_t + \gamma_i$$

- Would you expect β_1 to be positive or negative?
- What assumption is necessary to interpret β_1 causally?
- Why might this assumption be violated?