Panel Data Analysis

Lecture 2: From randomized controlled trial to two-way fixed effects

5 May, 2015

Prof. Andrew Eggers

Last week

Diff-in-diff: Binary treatment applied at a point in time to a subset of the units in the dataset

Parallel trends assumption

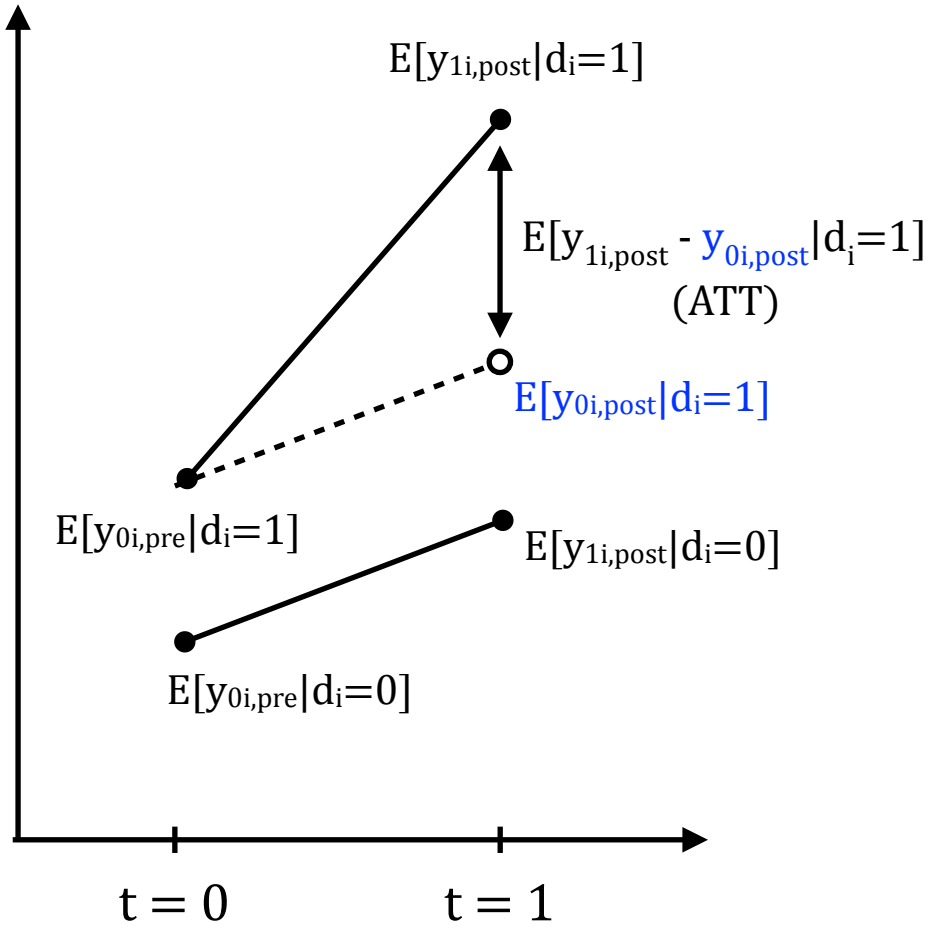

This week

Generalizing and expanding

— other kinds of grouped data (not necessarily time)

— general pattern of treatment application (not necessarily at **same point** in time)

— generalized treatment (not necessarily binary)

Back to RCT

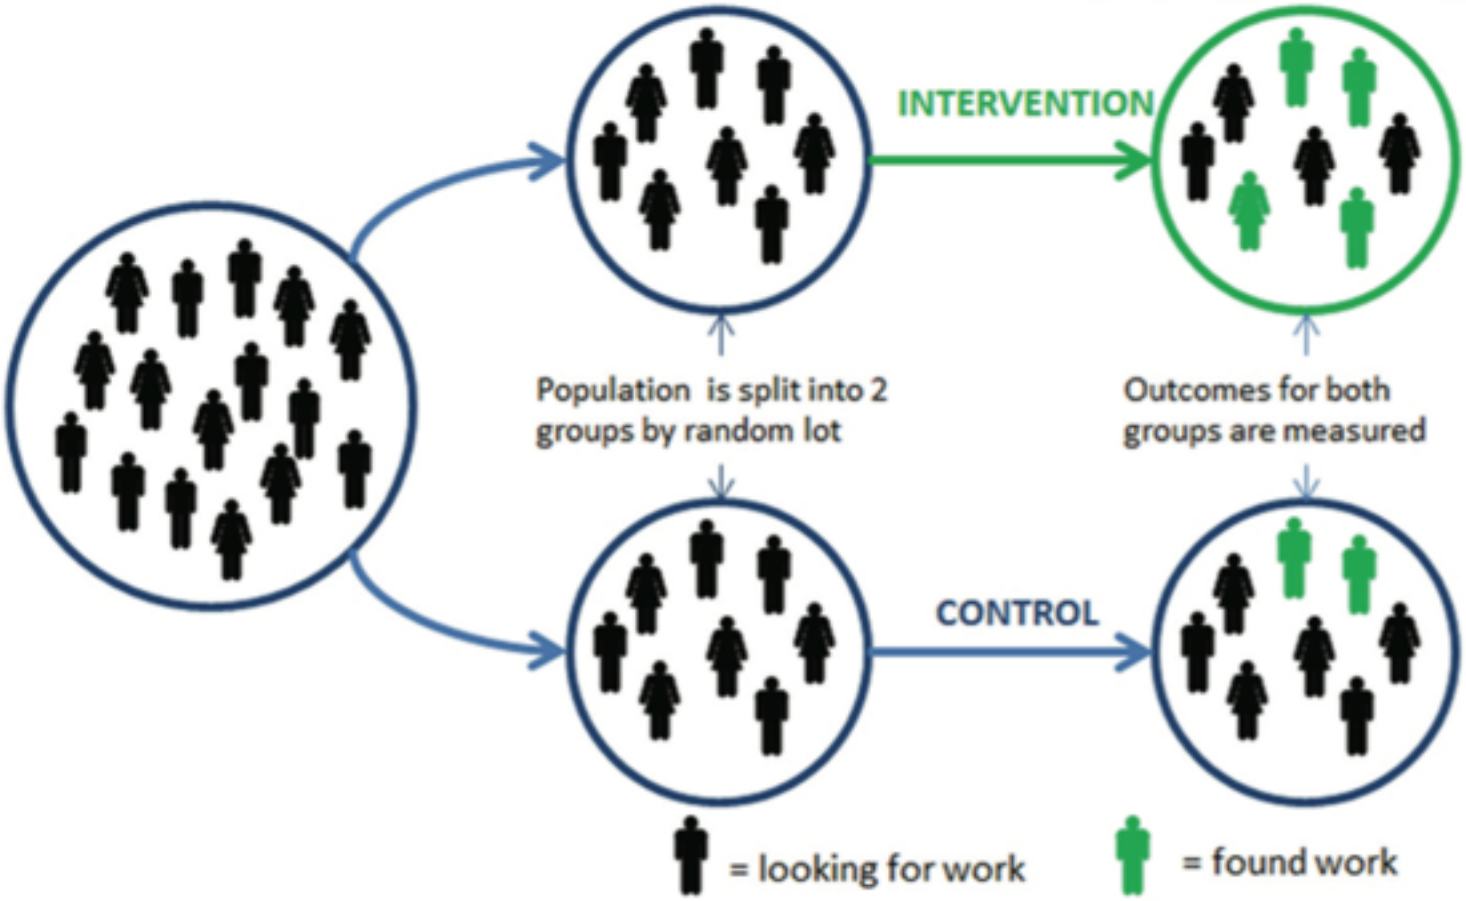

Figure 1. The basic design of a randomised controlled trial (RCT), illustrated with a test of a new 'back to work' programme.

More simulations!

Last time: from randomized controlled trial (RCT), we added a "baseline" measure: i.e. pre-treatment outcome for all units, with possible time trend.

This time: we start with an RCT in grouped data (e.g. pairs of twins participating in a drug trial; municipalities in districts participating in a field experiment) and add a second group (e.g. time periods).

Simulation 1: random assignment in grouped data

(1) Generate data according to

$$x_j \sim N(0,1)$$

$$y_{0i} \sim N(x_{j(i)},1)$$

$$y_{1i} \sim N(x_{j(i)} + \tau,1)$$

$$\tau = 1$$

where i indexes units, j indexes groups, and j(i) indicates the group of unit i

(2) Assign treatment (d) randomly

(3) Estimate ATT (effect of d on y) by

(3a) **Difference-in-means**: average difference in observed y between treated and control units

i.e. $E[y_i|d_i = 1]$ - $E[y_i|d_i = 0]$

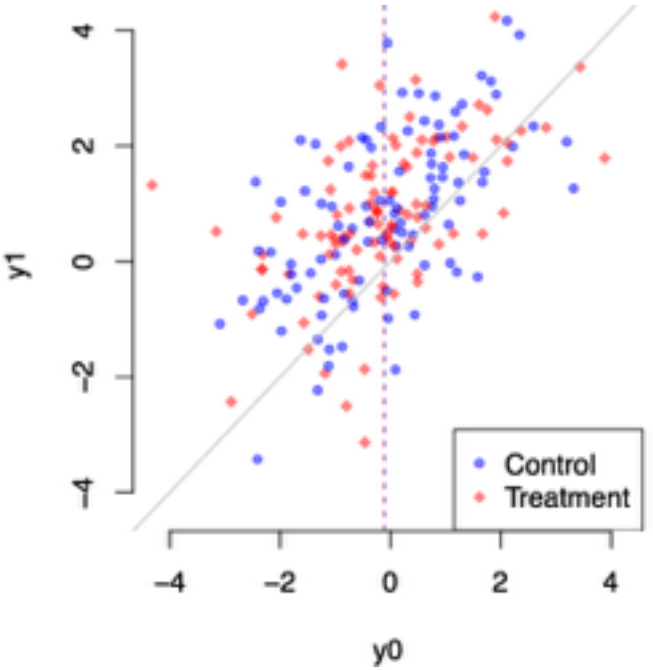

Is the unconfoundedness assumption met in this case?

(3b) **Controlling for x_j**: Regression of observed y_i on d_i and $x_{j(i)}$

i.e. $y_i = \alpha + \beta_1 d_i + \beta_2 x_{j(i)} + \varepsilon_i$

(3c) **LSDV:** Regression of observed y_i on d_i and indicator for each j

i.e. $y_i = \alpha_{j(i)} + \beta_1 d_i + \varepsilon_i$

Simulation 1 (random assignment in grouped data): distribution of estimates across replications

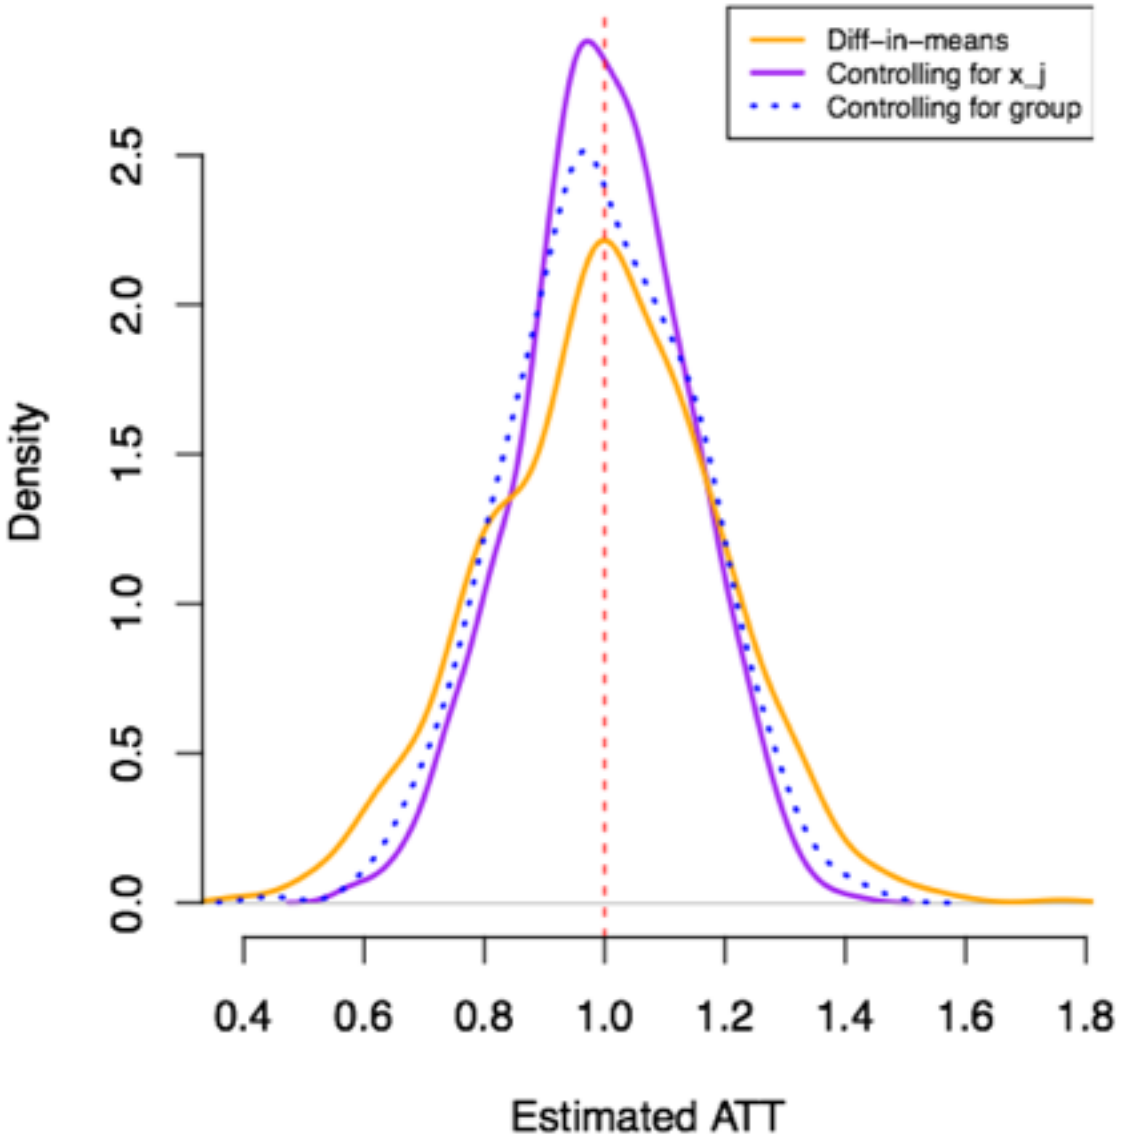

(The more units per group, the more similar the densities for the two control strategies.)

Simulation 2: non-random assignment in grouped data

(1) Generate data according to

$$x_j \sim N(0,1)$$

$$y_{0i} \sim N(x_{j(i)},1)$$

$$y_{1i} \sim N(x_{j(i)} + \tau,1)$$

$$\tau = 1$$

where i indexes units, j indexes groups, and j(i) indicates the group of unit i

(2) Assign treatment (d) as function of x_j:

$$Pr(d_i=1) = 1/(1 + \exp(-x_j))$$

(3) Estimate ATT (effect of d on y) by

(3a) **Difference-in-means**: average difference in observed y between treated and control units

i.e. $E[y_i|d_i = 1] - E[y_i|d_i = 0]$

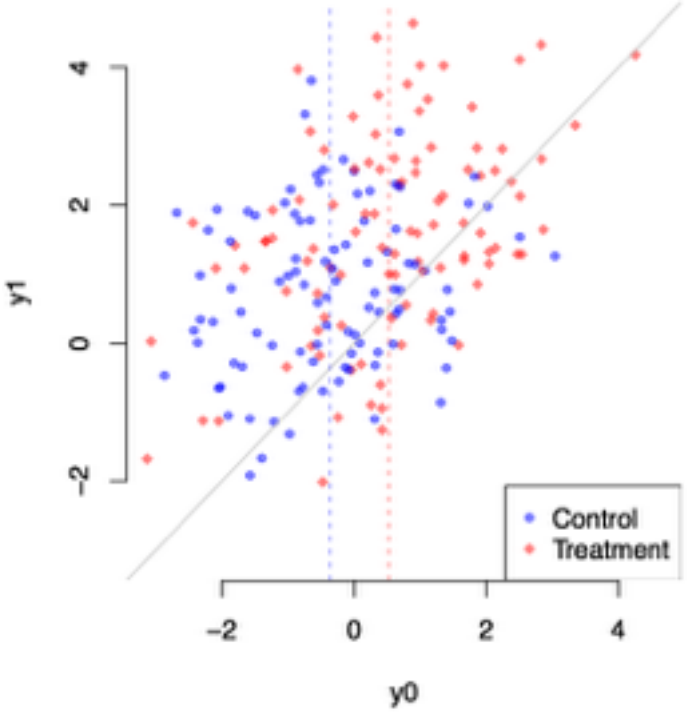

Is the unconfoundedness assumption met in this case?

(3b) **Controlling for x_j**: Regression of observed y_i on d_i and $x_{j(i)}$

i.e. $y_i = \alpha + \beta_1 d_i + \beta_2 x_{j(i)} + \varepsilon_i$

(3c) **LSDV:** Regression of observed y_i on d_i and indicator for each j

i.e. $y_i = \alpha_{j(i)} + \beta_1 d_i + \varepsilon_i$

Simulation 2 (non-random assignment in grouped data): distribution of estimates across replications

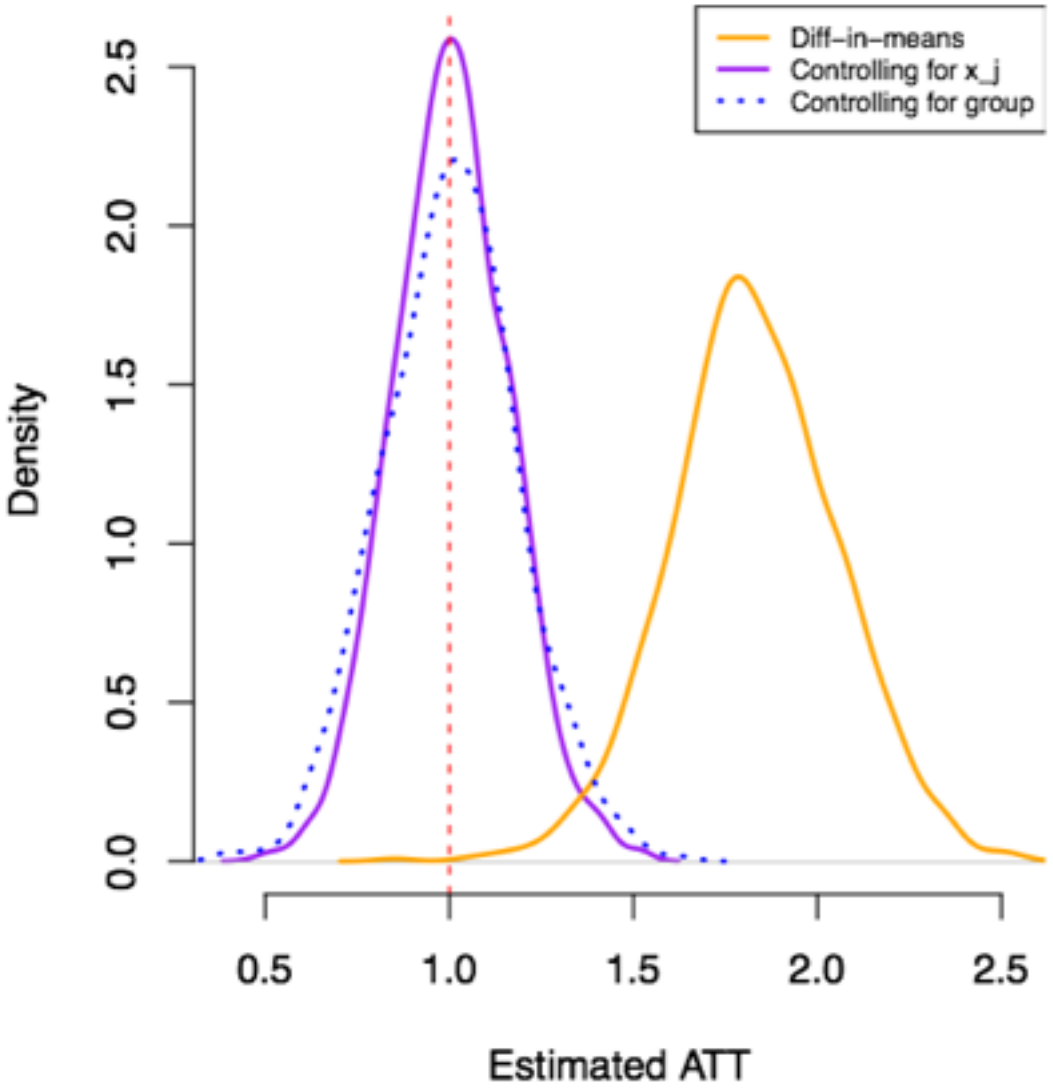

(The more units per group, the more similar the densities for the two control strategies.)

The point: Unobserved confounder? OK if it's shared by an identifiable group (e.g. district, school, judge)

LSDV regression and deviation from means: why do they give the same result?

LSDV regression: including a dummy for each group. (What was this in the Elbe example? Snow example?)

"**Deviation from means**" regression:

Define $\tilde{y}_i = y_i - \overline{y}_{j(i)}$, i.e. i's deviation from group mean. Similar for $\tilde{d}_i$.

Then estimate $\tilde{y}_i = \alpha + \beta_1 \tilde{d}_i + \tilde{\varepsilon}_i$

Important fact: They yield the same estimate for β_1. Why?

LSDV regression and deviation from means: why do they give the same result? (2)

Key idea from MHE's regression anatomy:

Given regression formula $y = \beta_0 + \beta_1 d + \gamma_1 z_1 + \gamma_2 z_2 + \gamma_3 z_3 + \ldots + \varepsilon$

Define $\tilde{d}$ as the residuals from a regression of d on all the z's.

Define $\tilde{y}$ as the residuals from a regression of y on all the z's.

Then regress $\tilde{y}$ on $\tilde{d}$:

$$\tilde{y} = \tilde{\beta}_0 + \tilde{\beta}_1 \tilde{d} + \tilde{\varepsilon}$$

Here's the exciting part: $\beta_1 = \tilde{\beta}_1$

Connection: If the z's are the group dummies, $\tilde{y} = \tilde{\beta}_0 + \tilde{\beta}_1 \tilde{d} + \tilde{\varepsilon}$ is the deviations-from-means regression.

Partial regression plot

No matter how complex your analysis (fixed effects, clustering, etc), when your treatment is continuous you should show the conditional bivariate relationship (y on x, controlling for z_1, z_2, z_3, etc) in a **partial regression plot:**

y-axis: residuals from regression of y on controls

x-axis: residuals from regression of treatment on controls

Example later.

More intuition of fixed effects

Fixed effects regression addresses group-specific unobservable confounders. How? Think about what would happen if you **observed** and **included** these confounders:

- Consider LSDV version: $y_i = \sum_{j=1}^{J} \alpha_j + \beta_1 d_i + \epsilon_i$

  If you included group-specific unobservable confounders, they would drop out due to multicollinearity with group dummies.

- Consider deviation from means version: $\tilde{y_i} = \beta_1 \tilde{d_i} + \tilde{\epsilon_i}$

  If you included group-specific unobservable confounders (before calculating deviations from means), they would drop out because they are constant within groups.

Applying one-way fixed effects

Key features:

- Data are organized into groups (e.g. individuals within households, households within municipalities, municipalities within districts, countries within regions, etc).
- Treatment varies within groups.
- Unconfoundedness/"selection on observables" may not hold in general: Treated and control units would be different even in absence of treatment
- But unconfoundedness holds *within groups*: Treated and control units *in the same group* are comparable

Does this apply to your case?

Adding a second group dimension

In simulations 1 and 2, each unit belonged to 1 group

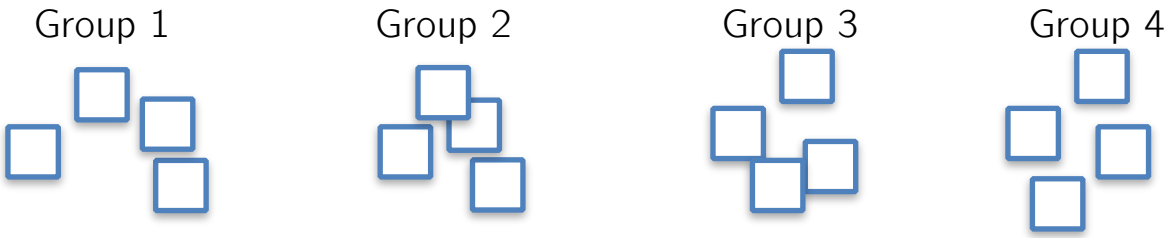

We can easily think about situations where each belongs to 2 or more groups

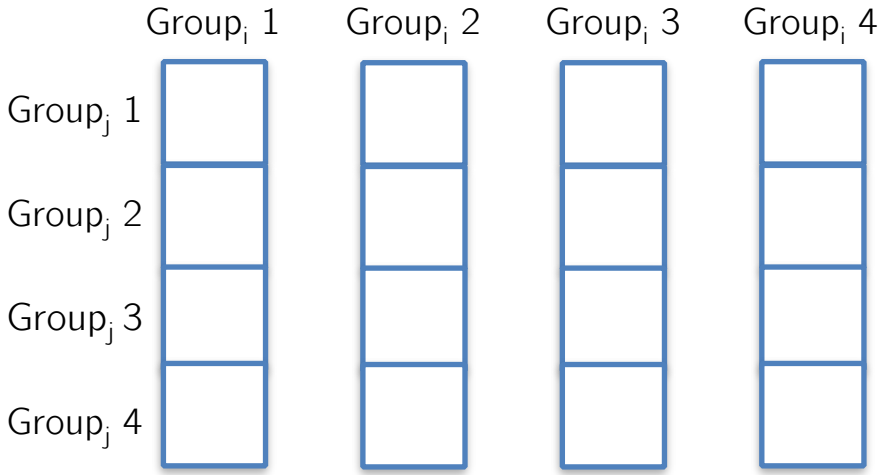

Most common: geographic unit and time period.

But also: class and ethnicity, education level and gender, etc.

Treatment assignment on the basis of two group dimensions

Suppose there is an unobserved confounder (i.e. something related to treatment that affects potential outcomes) that varies by district and year.

Suppose the relationship between the confounder and the district and year is additive (**not** interactive), i.e.

A. the confounder is higher for some **districts** than others

B. the confounder is higher for some **years** than others

- but **NOT** true that confounder is higher for some **district-year combinations** than others, once we account for A and B

Then we get the right answer through LSDV with dummies for district and year, or regression with district & year fixed effects.

Simulation 3: random assignment in twice-grouped data

(1) Generate data according to

$$x_j \sim N(0,1)$$

$$\lambda_t \sim N(0,1)$$

$$y_{0i} \sim N(x_{j(i)}+\lambda_{t(i)},1)$$

$$y_{1i} \sim N(x_{j(i)}+\lambda_{t(i)}+ \tau,1)$$

$$\tau = 1$$

where i indexes units, j indexes group 1, t indexes group t, and j(i) and t(i) indicate the group and time of unit i

(2) Assign treatment (d) randomly

(3) Estimate ATT (effect of d on y) by

(3a) **Difference-in-means**: average difference in observed y between treated and control units

i.e. $E[y_i|d_i = 1] - E[y_i|d_i = 0]$

(3b) **Controlling for x_j and λ_t**: Regression of observed y_i on d_i, $x_{j(i)}$, and $\lambda_{t(i)}$

i.e. $y_i = \alpha + \beta_1 d_i + \beta_2 x_{j(i)} + \beta_3 \lambda_{t(i)} + \varepsilon_i$

(3c) **Controlling for group 1:** Regression of observed y_i on d_i and indicator for each j

i.e. $y_i = \alpha_{j(i)} + \beta_1 d_i + \varepsilon_i$

(3d) **Controlling for groups 1 and 2:** Regression of observed y_i on d_i and indicator for each j and t

i.e. $y_i = \alpha_{j(i)} + \alpha_{t(i)} + \beta_1 d_i + \varepsilon_i$

Simulation 3 (random assignment in twice-grouped data): distribution of estimates across replications

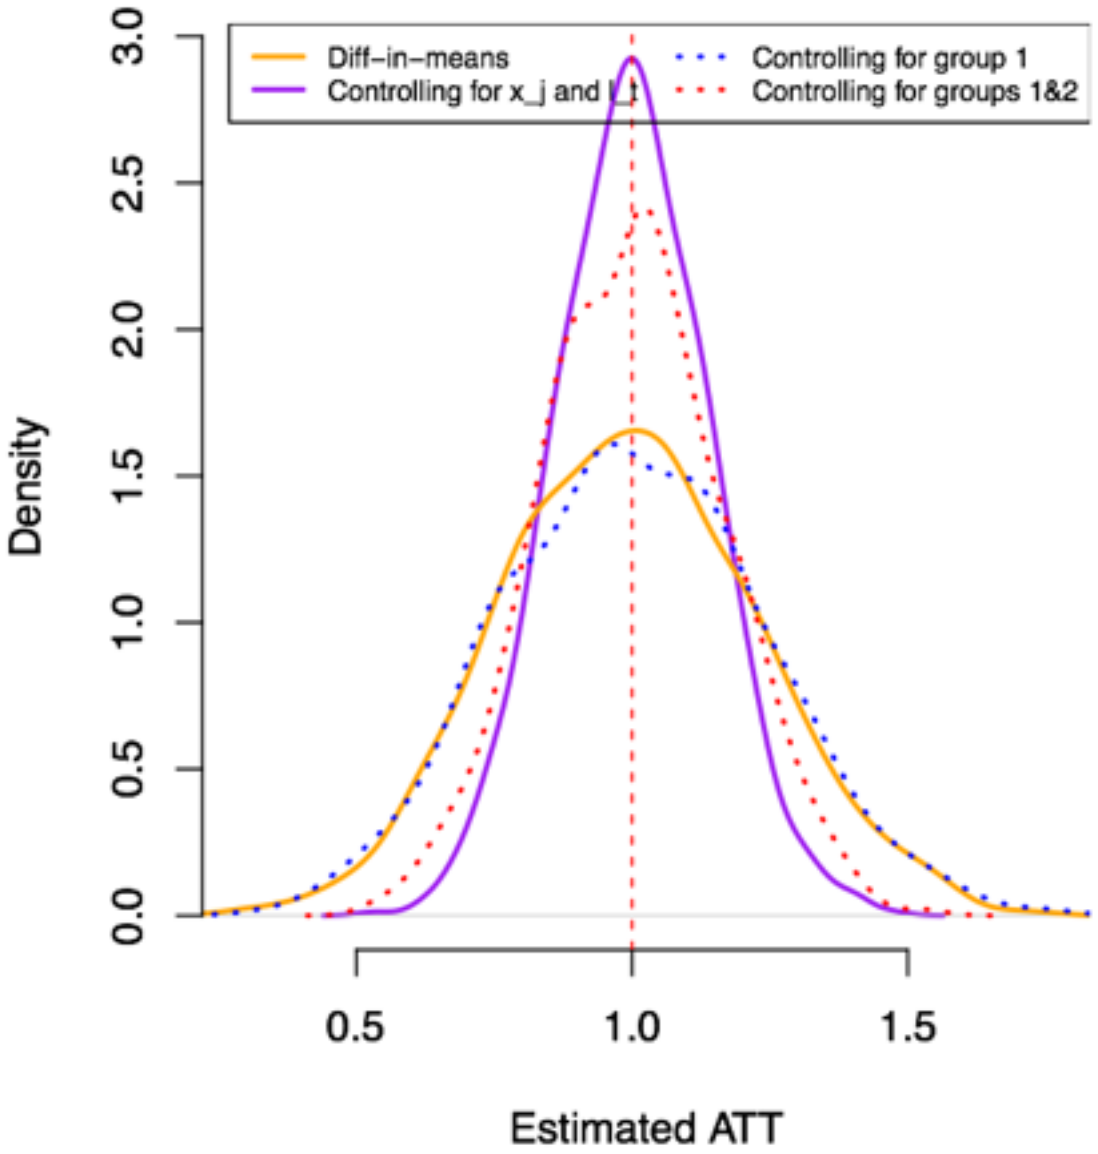

Simulation 3: random assignment in twice-grouped data

(1) Generate data according to

$$x_j \sim N(0,1)$$

$$\lambda_t \sim N(0,1)$$

$$y_{0i} \sim N(x_{j(i)}+\lambda_{t(i)},1)$$

$$y_{1i} \sim N(x_{j(i)}+\lambda_{t(i)}+ \tau,1)$$

$$\tau = 1$$

where i indexes units, j indexes group 1, t indexes group t, and j(i) and t(i) indicate the group and time of unit i

(2) Assign treatment (d) as function of x_j and λ_t:

$$Pr(d_i=1) = 1/(1 + \exp(-x_j-\lambda_t))$$

(3) Estimate ATT (effect of d on y) by

(3a) **Difference-in-means**: average difference in observed y between treated and control units

i.e. $E[y_i|d_i = 1] - E[y_i|d_i = 0]$

(3b) **Controlling for x_j and λ_t**: Regression of observed y_i on d_i, $x_{j(i)}$, and $\lambda_{t(i)}$

i.e. $y_i = \alpha + \beta_1 d_i + \beta_2 x_{j(i)} + \beta_3 \lambda_{t(i)} + \varepsilon_i$

(3c) **Controlling for group 1:** Regression of observed y_i on d_i and indicator for each j

i.e. $y_i = \alpha_{j(i)} + \beta_1 d_i + \varepsilon_i$

(3d) **Controlling for groups 1 and 2:** Regression of observed y_i on d_i and indicator for each j and t

i.e. $y_i = \alpha_{j(i)} + \alpha_{t(i)} + \beta_1 d_i + \varepsilon_i$

Simulation 4 (non-random assignment in twice grouped data): distribution of estimates across replications

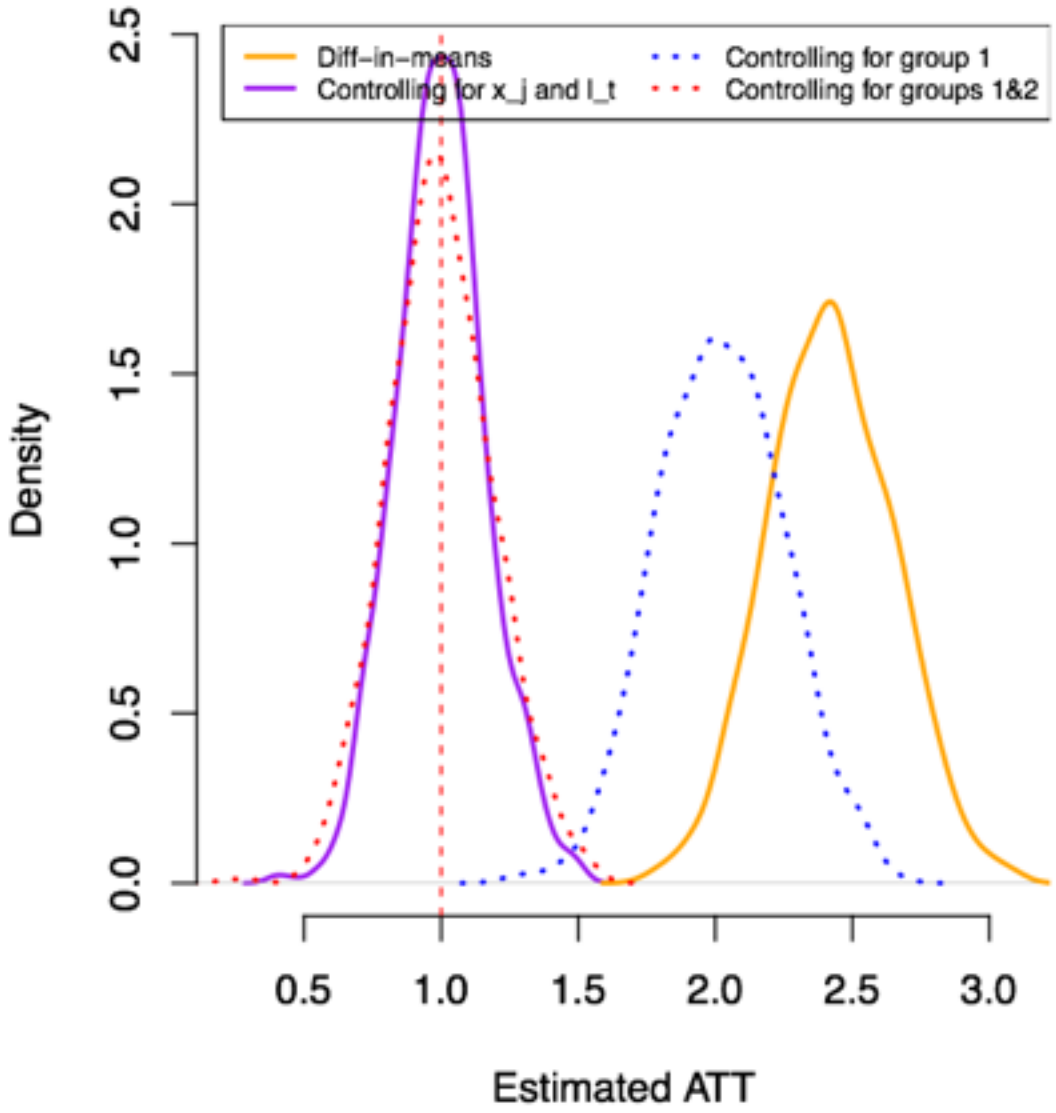

Toward the conditional independence assumption (CIA)

Recall: for binary treatment, the difference in means

$$E[y_{1i}|d_i=1] - E[y_{0i}|d_i=1]$$

can be rewritten as

$$\underbrace{E[y_{1i}|d_i=1] - E[y_{0i}|d_i=1]}_{\text{ATT}} + \underbrace{E[y_{0i}|d_i=1] - E[y_{0i}|d_i=0]}_{\text{Selection bias}}$$

Sometimes the independence assumption does not hold:

$$E[y_{0i}|d_i=1] \neq E[y_{0i}|d_i=0]$$

But if the **conditional independence assumption (CIA)** does hold

$$E[y_{0i}|X_i,d_i=1] = E[y_{0i}|X_i,d_i=0],$$

then we have

$$\underbrace{E[y_{1i}|X_i,d_i=1] - E[y_{0i}|X_i,d_i=0]}_{\text{Diff-in-means controlling for X}} = \underbrace{E[y_{1i}|X_i,d_i=1] - E[y_{0i}|X_i,d_i=1]}_{\text{ATT}}$$

CIA and panel data

From the perspective of causal inference, panel data offers intriguing X variables that make the CIA plausible.

For example, support for SPD in flooded and non-flooded districts would be different in absence of flooding, but maybe not once we control for past support for SPD and the time trend.

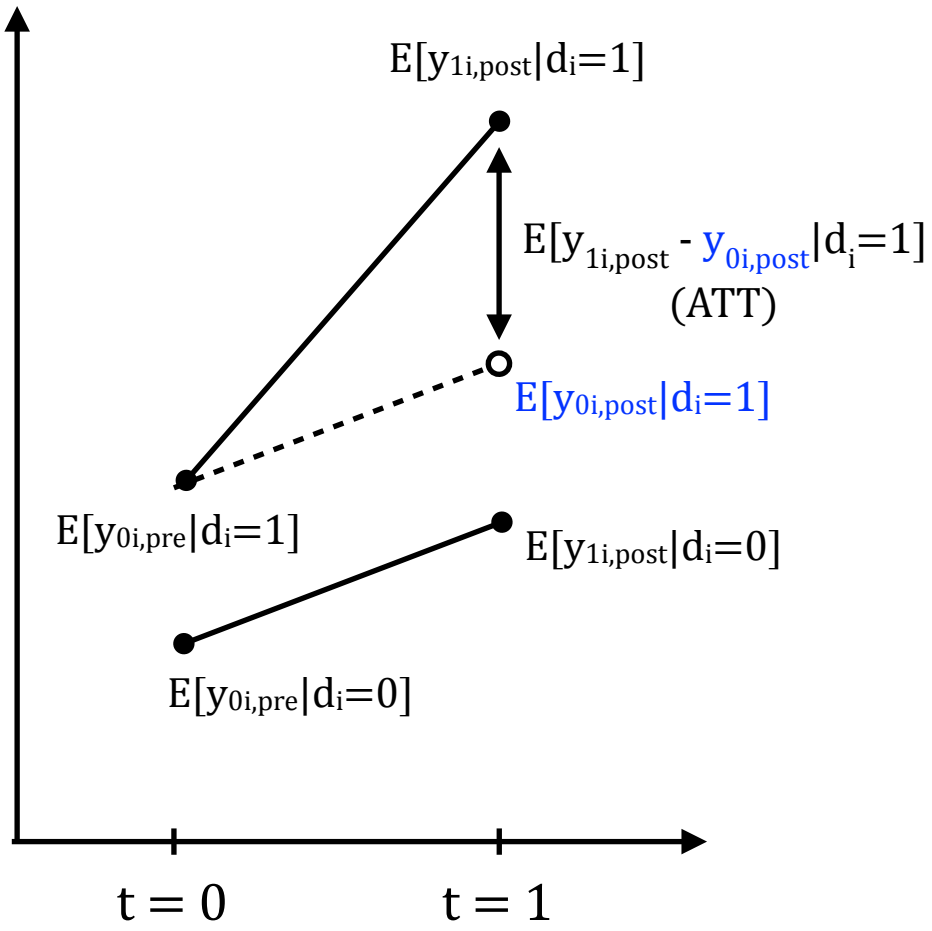

In diff-in-diff, CIA amounts to: **parallel trends assumption**.

Generally in panel data (with unit and time fixed effects), CIA amounts to: **no time-varying confounders within units.**

CIA beyond binary treatment

Recall: for binary treatment, the difference in means

$$E[y_{1i}|d_i=1] - E[y_{0i}|d_i=0]$$

can be rewritten as

$$\underbrace{E[y_{1i}|d_i=1] - E[y_{0i}|d_i=1]}_{\text{ATT}} + \underbrace{E[y_{0i}|d_i=1] - E[y_{0i}|d_i=0]}_{\text{Selection bias}}$$

Sometimes the independence assumption does not hold:

$$E[y_{0i}|d_i=1] \neq E[y_{0i}|d_i=0]$$

But if the **conditional independence assumption (CIA)** does hold

$$E[y_{0i}|X_i,d_i=1] = E[y_{0i}|X_i,d_i=0],$$

then we have

$$\underbrace{E[y_{1i}|X_i,d_i=1] - E[y_{0i}|X_i,d_i=0]}_{\text{Diff-in-means controlling for X}} = \underbrace{E[y_{1i}|X_i,d_i=1] - E[y_{0i}|X_i,d_i=1]}_{\text{ATT}}$$

Going beyond binary treatment

CIA extends beyond binary treatment:

Conditional on X, the units that get a particular value of treatment have the same potential outcomes for ALL values of treatment as the units that get another particular value of treatment. (See MHE 54-59).

Panel fixed effects regression can be seen as a particular set of controls that might make the CIA believable.

Big picture: the CIA is everywhere

Whenever we causally interpret a regression, we make a conditional independence assumption (CIA); what varies is what is in the conditioning set and how convincing it is.

- In RCT, CIA conditions on nothing.
- In generic cross-section regression (e.g. trade and democracy), CIA that conditions on a large set of confounders may be convincing.
- In grouped cross-section, CIA that conditions on group (i.e. one-way fixed effects) may be convincing.
- In typical panel fixed effects regression, CIA that conditions on group and time period (i.e. panel fixed effects) may be convincing.
- If CIA is more plausible in panel than cross-section, it is because time-invariant unit-level confounders are important.

You should always state your CIA and provide conditions in which it might be violated.

Example study: Fouirnaies, "How Do Campaign Spending Limits Affect Electoral Competition? Evidence from Great Britain 1885-2010"

In UK from 1885, constituency spending limits are a **deterministic** function of population and borough/county classification.

$$limit_{it} = a_t + b_t\, population_{it} + c_t\, county_{it} + d_t\, county_{it} \times population_{it}$$

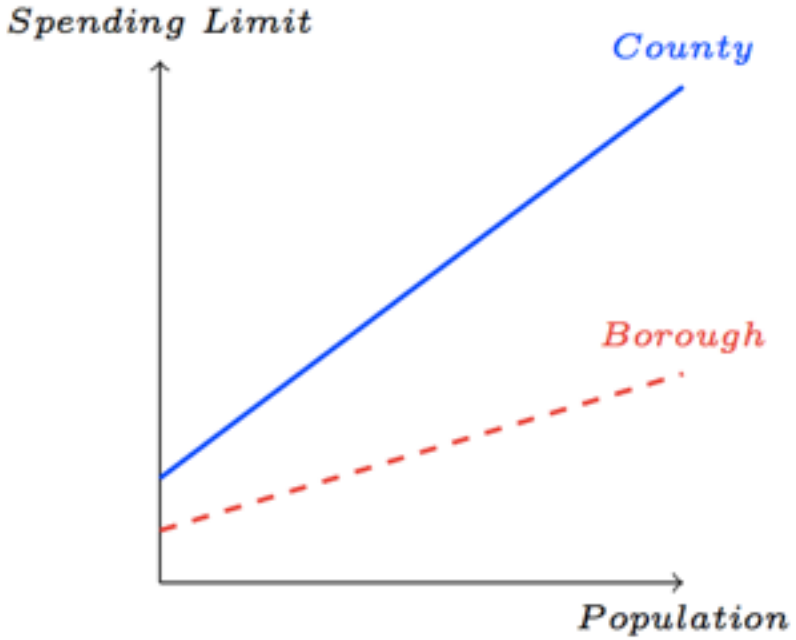

Example study: Fouirnaies, "How Do Campaign Spending Limits Affect Electoral Competition? Evidence from Great Britain 1885-2010"

Question: how does spending limit affect avg margin between winner and loser?

Basic approach:

$$Win\ Margin_{it} = \alpha_i + \delta_t + \beta_1 Limit_{it} + f(population_{it}) + \varepsilon_{it}$$

So, can think of this as:

- LSDV with dummies for constituency and year
- Regression of "demeaned" (by constituency and year) win margin on "demeaned" (by constituency and year) spending limits
- Regression of residuals from [regression of win margin on constituency and year dummies] on residuals from [regression of spending limits on constituency and year dummies]

Example study: Fouirnaies, "How Do Campaign Spending Limits Affect Electoral Competition? Evidence from Great Britain 1885-2010"

Relationship between spending limits and WinMargin:

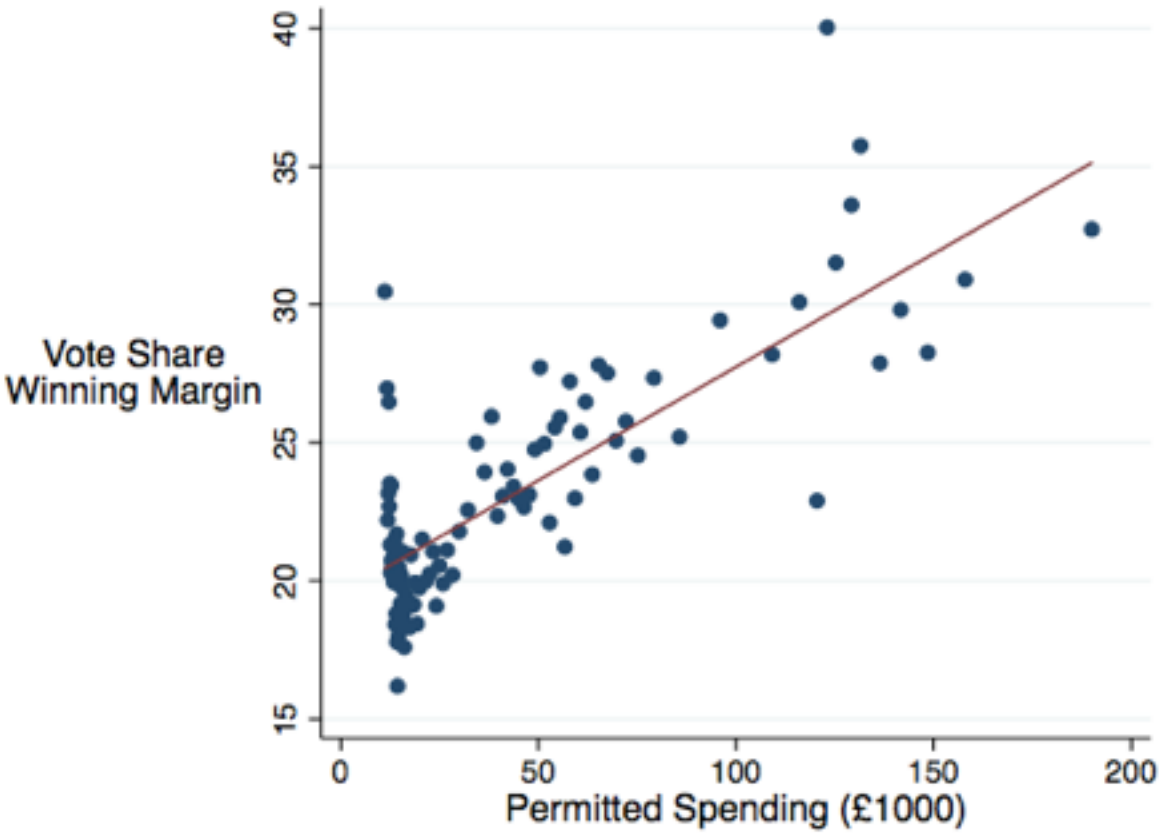

Example study: Fouirnaies, "How Do Campaign Spending Limits Affect Electoral Competition? Evidence from Great Britain 1885-2010"

y-axis: residuals from regressing WinMargin on constituency and year dummies

x-axis: residuals from regressing spending limits on constituency and year dummies

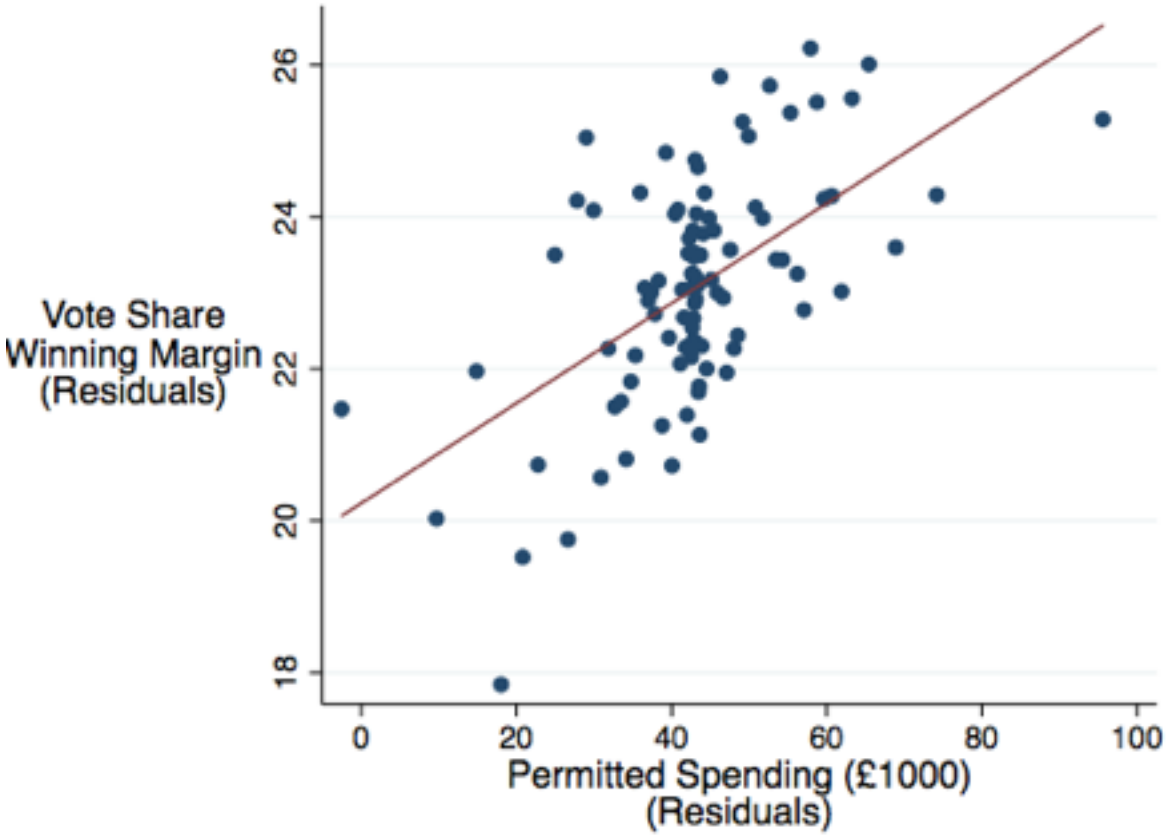

Example study: Fouirnaies, "How Do Campaign Spending Limits Affect Electoral Competition? Evidence from Great Britain 1885-2010"

What is the CIA here?

"In absence of treatment, affected constituencies would have evolved as unaffected [constituencies]."

i.e. no time-varying confounders.

i.e. changes in spending limits not systematically related to changes in competition, controlling for changes in other places

Example: Levitt (1994) 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:

$$DemCongVoteShare_i = \beta_0 + \beta_1 (DemSpend_i - RepSpend_i) + \beta_2 DemPresVoteShare_i + \varepsilon_i$$

where

$DemCongVoteShare_i$: Vote share for Democratic congressional candidate in district i

$DemSpend_i$, $RepSpend_i$: Spending by Democratic and Republican congressional candidates in district i

$DemPresVoteShare_i$: Vote share for Democratic presidential candidate in district i

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

Example: Levitt (1994) on effects of campaign spending (2)

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

Unmeasured confounder: Candidate quality (i.e. attractiveness)

Research design: Two-way fixed effects, where groups are (i) *pairs* of candidates and (ii) years. (Also includes scandal and incumbency dummies.)

Ways of stating the CIA:

- Levitt (782-783): "An individual candidate's quality must be constant over time."
- Election-relevant features of pairs of candidates are fixed over time.
- The same pair of candidates in a different year is comparable, after controlling for nationwide year-to-year swings in electoral outcomes.
- Variation in spending over time within a given pair of candidates is unrelated to potential outcomes conditional on the year (and other controls).

To discuss:

- How might this CIA be violated?
- Do pairs of candidates who only appear once in the dataset contribute anything to the estimation of the effect?

Fixed effects and first-differences

Consider situation where (as in simulation) $x_{j(i)}$ is a time-invariant confounder:

$$y_{it} = \beta_0 + \beta_1 d_{it} + \beta_2 x_{j(i)} + \varepsilon_{it}$$

Above we addressed this with LSDV/group fixed effects, which we showed was equivalent to a deviation-from-means regression:

$$\tilde{y}_{it} = \beta_0 + \beta_1 \tilde{d}_{it} + \varepsilon_{it}$$

Another way to address time-invariant confounders in panel data: **first-differences**. Calculate $\Delta y_{it} = y_{it} - y_{i,t-1}$, etc and estimate

$$\Delta y_{it} = \beta_0 + \beta_1 \Delta d_{it} + \Delta \varepsilon_{it}$$

As explained in MHE (224):

- Algebraically equivalent when just two periods
- Not otherwise, but consistent (i.e. as sample size increases, both converge to truth)

Example: Ansell (2014) on effect of house prices on political preferences

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:

$$SupportForRedistribution_{\mathbf{i}} = \beta_0 + \beta_{\mathbf{1}}\, PriceOfHouse_{\mathbf{i}} + \beta_{\mathbf{2}}\, Income_{\mathbf{i}} + \beta_{\mathbf{3}}\, Age_{\mathbf{i}}\, \varepsilon_{\mathbf{i}}$$

- Would you expect $\beta_{\mathbf{1}}$ to be positive or negative?
- What CIA is necessary to interpret that coefficient causally?
- Why might this CIA be violated?

Example: Ansell (2014) on effect of house prices on political preferences

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?

Research design: First-difference regression in seven-wave British Household Panel Survey

The CIA: Confounding variables are constant within individuals between two waves of panel.

How might this be violated?

Not all panel FE analysis is causal

Example: Fowler (2015), "Do elections select better representatives?"

Question: Why do incumbents win more electoral support than non-incumbents?

Research design: Descriptive decomposition.

- Estimate

  $$DemVoteShare_{\mathbf{i}} = \beta_0 + \beta_{\mathbf{1}} \, (DemIncumbent_{\mathbf{i}} - RepIncumbent_{\mathbf{i}}) + \varepsilon_{\mathbf{i}}$$

- Estimate same thing with state-decade and year fixed effects:

  $$DemVoteShare_{\mathbf{i}} = \alpha_{\mathbf{j}} + \gamma_t + \beta_{\mathbf{1}} \, (DemIncumbent_{\mathbf{i}} - RepIncumbent_{\mathbf{i}}) + \varepsilon_{\mathbf{i}}$$

- Interpret difference in $\beta_{\mathbf{1}}$ as a measure of "party match" component of incumbent success (tendency of incumbents to belong to locally-popular party)
- (Estimate "officeholder benefits" via RDD)

The CIA: Not relevant, because no causal claims.

How do you test the CIA?

Because of the fundamental problem of causal inference, CIA is **always** an assumption; it cannot be directly tested.

Diff-in-diff is a special case of panel fixed-effects regression that allows for a very nice indirect test.

Why is it possible?

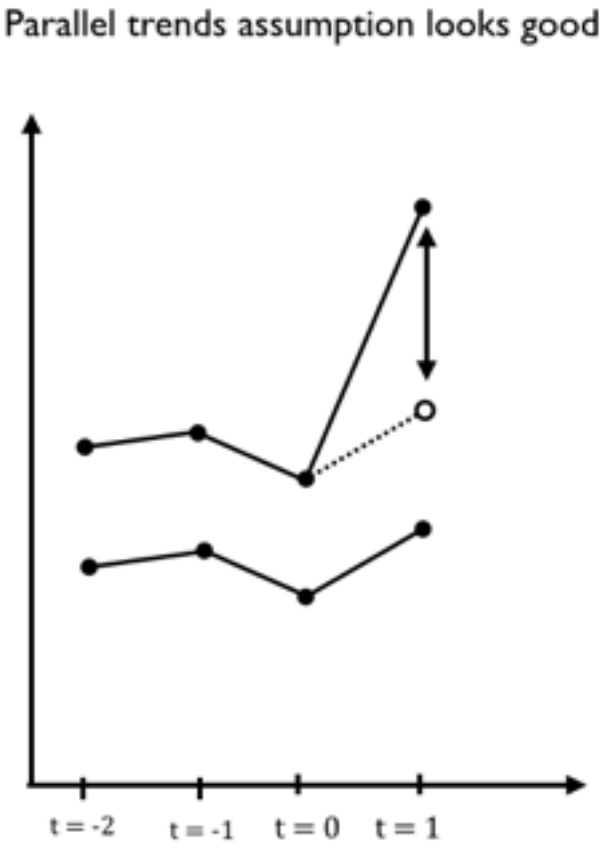

How do you test the CIA? (2)

Transparent indirect tests of the CIA: a characteristic of my favorite research designs.

e.g. regression discontinuity designs (RDD) (at right: Lee (2008), "Randomized experiments from non-random selection in U.S. House elections")

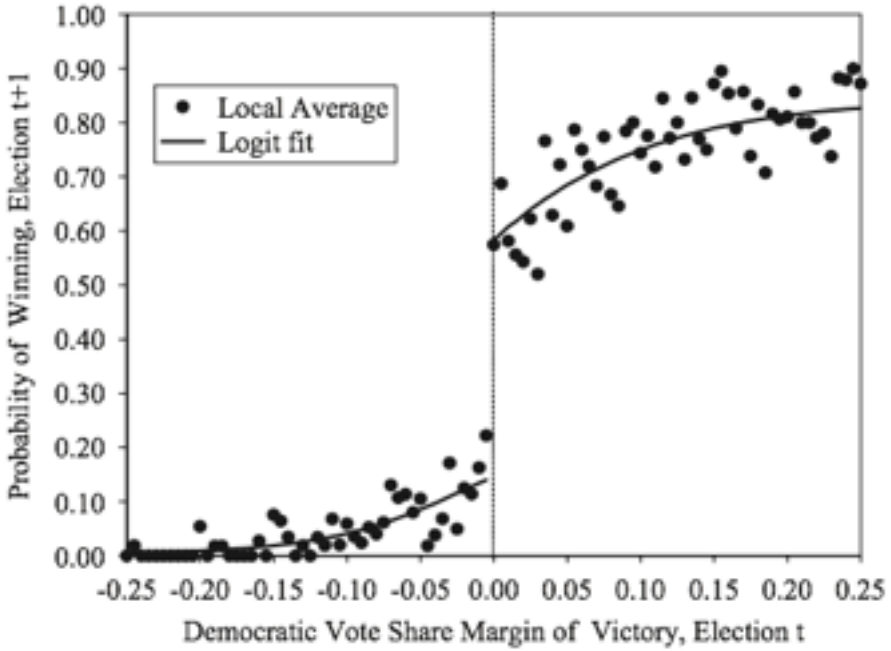

b

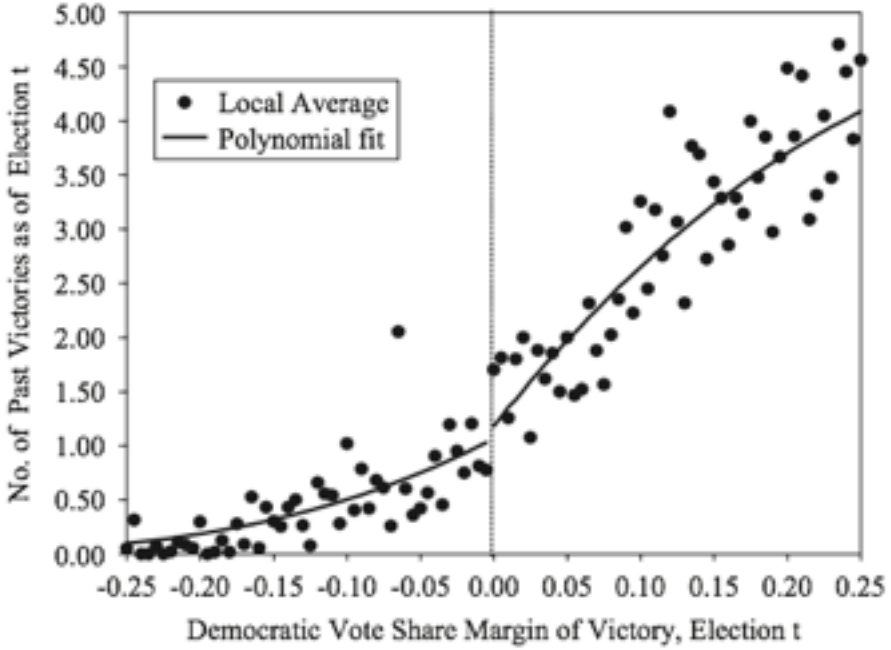

How do you test the CIA? (3)

What do these cases have in common? The fact that we slightly **extend** the CIA and test the additional implications.

- **Diff-in-diff:**
  - **CIA**: Parallel trends assumption applies between pre- and post-treatment (i.e. potential outcomes are same for treatment group and control group post-treatment, after adjusting for time-invariant difference)
  - **Extension to CIA**: Parallel trends assumption should apply between pre-pre-treatment and pre-treatment period too
- RDD:
  - **CIA:** At 0% threshold of vote share margin, potential outcomes are independent of treatment status (i.e. difference in outcomes between winners and losers of dead heat is due to treatment)
  - **Extension to CIA**: At 0% threshold of vote share margin, pre-treatment covariates should also be independent of treatment status

How do you test the CIA? (4)

- **Generalization of diff-in-diff (panel fixed effects design):**
  - **CIA**: No time-variant unit specific confounders
  - **Extension to CIA**: When including dummies for "2 years before treatment", "1 year before treatment", etc in regression, lagged treatment has no effect (see Kuziemko & Werker 2006 for example; also MHE 237-238)

General pattern:

- does something that should have no effect have an effect ("placebo")
- does treatment affect something it shouldn't? ("placebo outcome", falsification test)

As we add complexity from simple diff-in-diff (esp when we go beyond binary treatment), testing the CIA becomes harder and harder.

A problem with transparent designs: You get asked for placebo tests!

Wrapping up

If there are unobserved confounders, but you think they may be fixed within groups (or within units over time), fixed effects (or first differences) can help.

All causal inference requires a conditional independence assumption.

If panel methods are especially credible, it is because

- CIA may be particularly credible within units over time
- some cases (e.g. simple diff-in-diff) allow for transparent indirect tests of CIA

Next time: lagged dependent variables, synthetic control method, random effects, standard errors (or a subset)