

Using Individual-Level Data in Difference-in-Differences for Health Policy and Service Evaluation

Nicholas J. Seewald, PhD

Department of Health Policy and Management

Johns Hopkins Bloomberg School of Public Health

ALACRITY Center RIP Meeting
9 November 2022







https://slides.nickseewald.com/alacrityRIP.pdf

Amazing Collaborators!





Beth McGinty, PhD



Elizabeth Stuart, PhD



Kayla Tormohlen, PhD



Ian Schmid, ScM





- ▶ Postdoc in Health Policy and Management with Liz Stuart and Beth McGinty
- ► ALACRITY trainee
- Statistician by training
- My goal is to build statistical methods that help people do cool science

Motivation: Medical Cannabis Laws and Opioid Prescribing

- ▶ 4x increase in opioid prescribing in U.S. from 1999-2012
 - Meaningfully driven by opioid prescribing for chronic non-cancer pain
- Getting better: prescribing down since 2012 peak
 - ▶ 2020 had lowest opioid dispensing rate in previous 15 years (nationally, est. by IQVIA)
- Cannabis is a potentially effective treatment for chronic non-cancer pain, but evidence is limited

Motivation: Medical Cannabis Laws and Opioid Prescribing

Patients with chronic non-cancer pain are eligible to use cannabis under all existing state medical cannabis laws

• Question: What are the effects of state medical cannabis laws on receipt of opioid and non-opioid treatment among patients with chronic non-cancer pain?

Motivation: Medical Cannabis Laws and Opioid Prescribing

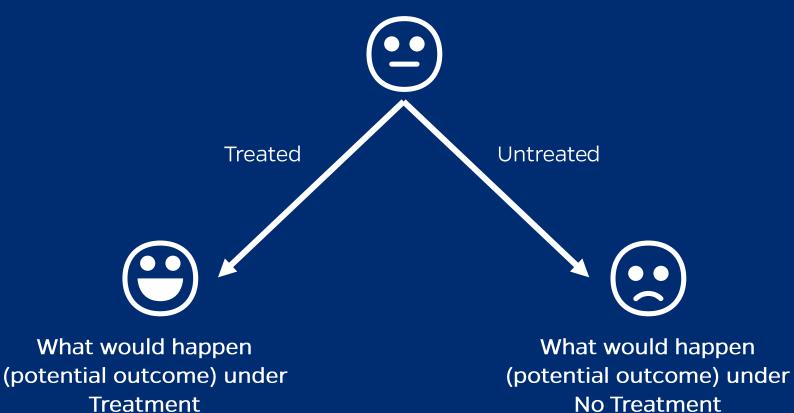
- Previous studies have found mixed results, but have key methodological limitations:
- 1. General population samples with no individual-level data to identify sample with chronic non-cancer pain
- 2. Policy endogeneity not addressed

Individual-level data lets us identify the population of interest but adds methodological complexity. My work addresses that complexity.

Potential Outcomes and Causal Inference

 $Y_{i,t}(1)$

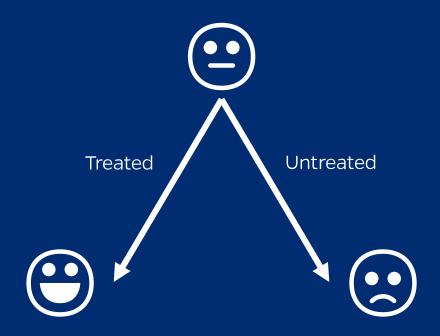




 $Y_{i,t}(0)$

Potential Outcomes and Causal Inference





What would happen (potential outcome) under Treatment $Y_{i,t}(1)$

What would happen (potential outcome) under No Treatment $Y_{i,t}(\mathbf{0})$

- $Y_{i,t}(1)$ is the outcome we would observe for unit i at time t had they received treatment
- $ightharpoonup Y_{i,t}(0)$ is the outcome we would observe for unit i at time t had they *not* received treatment
- An individual causal effect is $Y_{i,t}(1) Y_{i,t}(0) = \bigcirc \bigcirc$
- For each person, we can only observe $Y_{i,t}(1)$ OR $Y_{i,t}(0)$, not both!
- We can estimate average causal effects (over a population of interest)

Average treatment effect among the treated (ATT)



$$ATT = E[Y(1) - Y(0) | A = 1]$$

= $E[Y(1) | A = 1] - E[Y(0) | A = 1]$

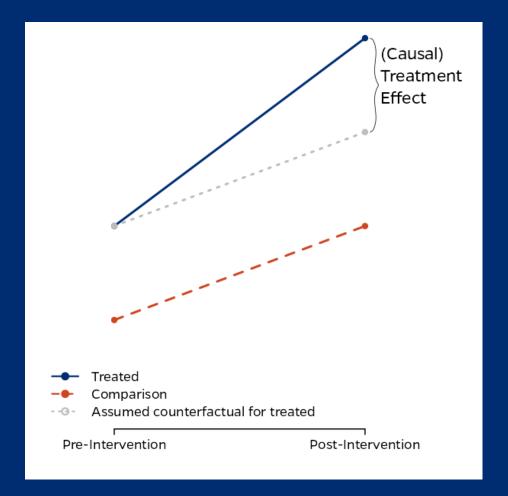
where A is treatment status.

The problem is that we can't directly estimate this from data: we don't observe Y(0) in the A=1 group.

Difference-in-Differences



- A commonly-used method for estimating the ATT in policy evaluation
- Idea: compare change in outcome over time between treated and comparison groups
- ▶ **Key assumption:** in the absence of treatment, the outcome evolution in the treated group would have looked like the outcome evolution in the comparison group
 - ▶ This is called the *counterfactual parallel trends assumption*.







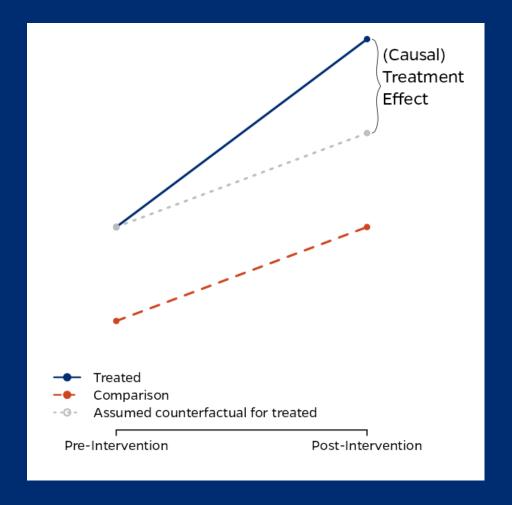
Under the counterfactual parallel trends assumption, we can write the ATT as

$$ATT$$

$$= (E[Y_{post} | A = 1] - E[Y_{pre} | A = 1])$$

$$- (E[Y_{post} | A = 0] - E[Y_{pre} | A = 0])$$

- ► The ATT is literally a difference in differences!
- We can estimate all these components from our data just using sample means or via regression.



Want more on Difference-in-Differences?

Check out Elizabeth Stone's lecture on the ALACRITY Center's YouTube channel!

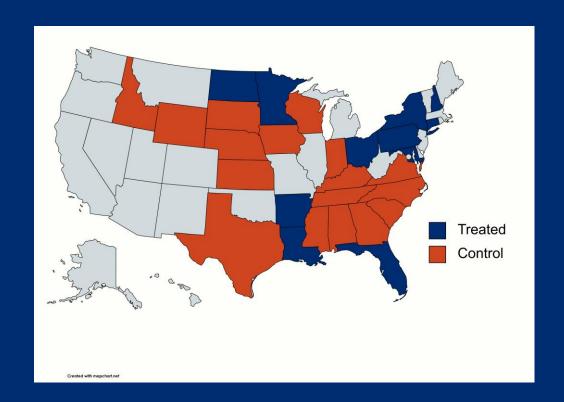


https://youtu.be/EYVYTIYGdq4

Medical Cannabis Study



- Our study:
 - ▶ 12 treated states that implemented medical cannabis law between 2012 and 2019, and do not also have recreational cannabis laws.
 - ▶ 17 *comparison* states without medical or recreational cannabis laws, 2012-2019
- ▶ **Goal:** Estimate the effect of implementing a medical cannabis law on opioid prescribing outcomes relative to what would have happened in the absence of treatment, on average across states.



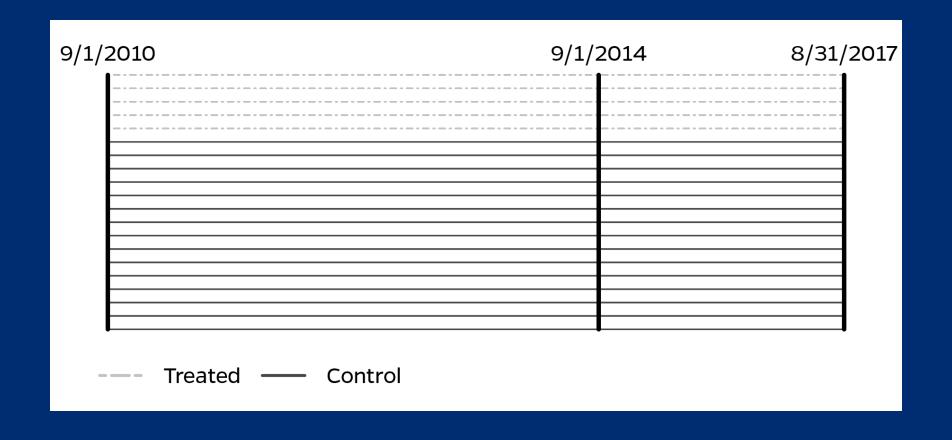
Medical Cannabis Study



- ► Data are individual-level longitudinal commercial health insurance claims from OptumLabs Data Warehouse, collected monthly.
- Each treated state has its own 7-year study period anchored at its law's implementation date
 - 4 years pre-law, 3 years post-law
- For each treated state, we build a cohort of individuals in that state and the control states over that state's study period.
 - Individuals are included if they have a chronic non-cancer pain diagnosis in the pre-law period *and* are continuously enrolled in commercial health insurance for the full study period.

Medical Cannabis Study: CT Cohort





Difference-in-Differences with Multiple Time Periods

- Previous difference-in-differences setup was for two units (one treated, one control) over two time periods (pre- / post-treatment).
- ▶ When there are more than two time periods, we need to adapt our setup!

$$ATT(t) = E[Y_t(1) - Y_t(0) \mid A = 1]$$

 Extend counterfactual parallel trends assumption to apply to all pairs of pre/post observations (can be relaxed)





Traditional way of estimating the ATT in multi-period difference-in-differences

$$Y_{uit} = \beta_{0u} + \beta_{1t} + \beta A_{ut} + \epsilon_{uit}$$

where u is unit, i is individual, t is time.

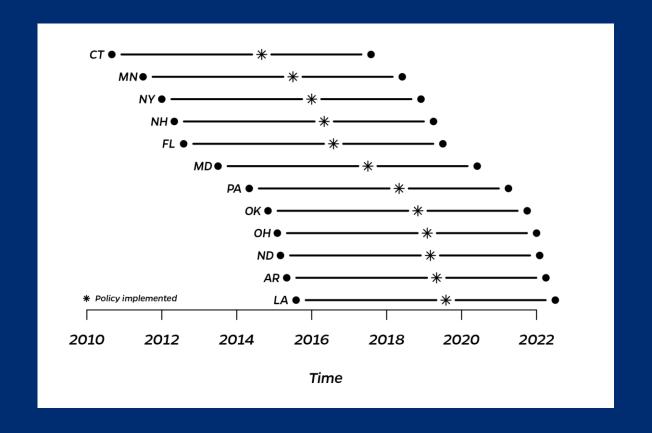
 $\triangleright \hat{\beta}$ is the estimate of the ATT (averaged over time):

$$\hat{\beta} \equiv (\bar{Y}_{\text{post,tx}} - \bar{Y}_{\text{pre,tx}}) - (\bar{Y}_{\text{post,ctrl}} - \bar{Y}_{\text{pre,ctrl}})$$





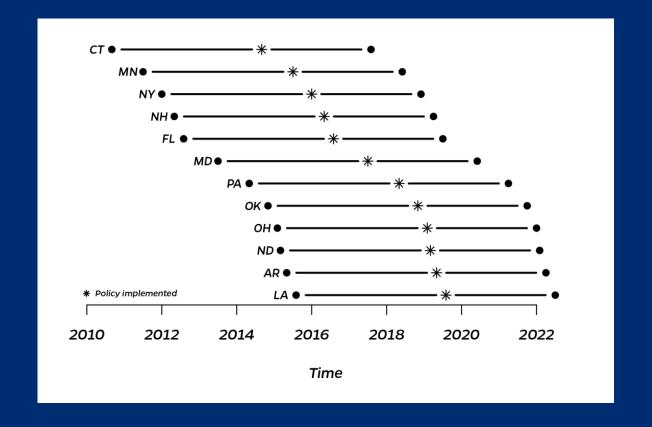
- Two-way fixed effects model can yield extremely biased effect estimates
 - when used with multiple treated units under "staggered adoption" and
 - when trying to estimate an overall effect across treated units
- Estimator implicitly (and inappropriately) adjusts for post-treatment information from some units



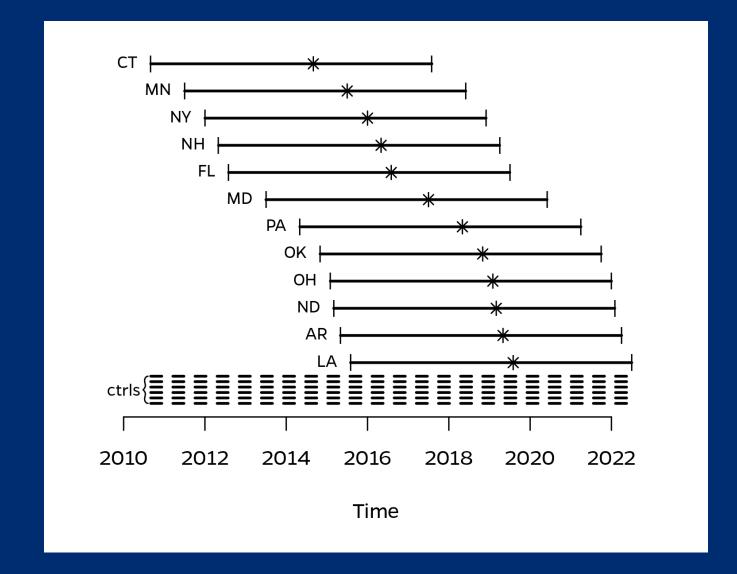




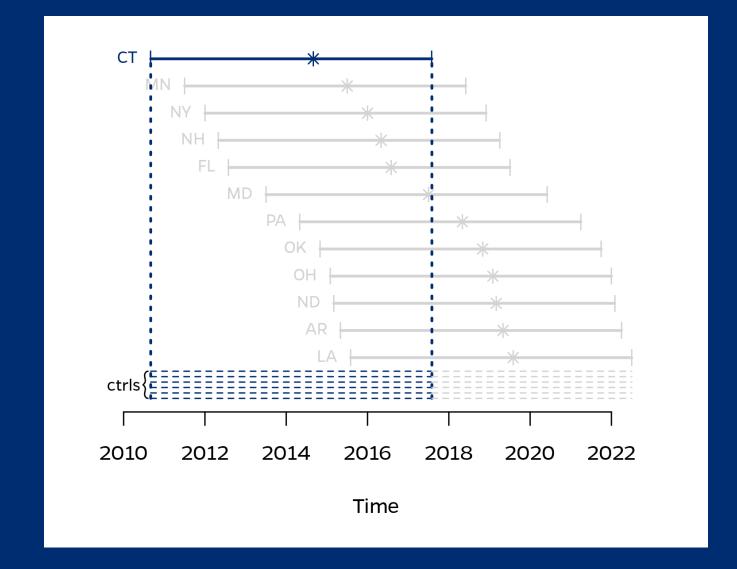
- A solution is to use **stacked** difference-indifferences
 - Also called "trial emulation"
- We'll estimate an ATT for each treated state, then pool the estimates to get an overall ATT



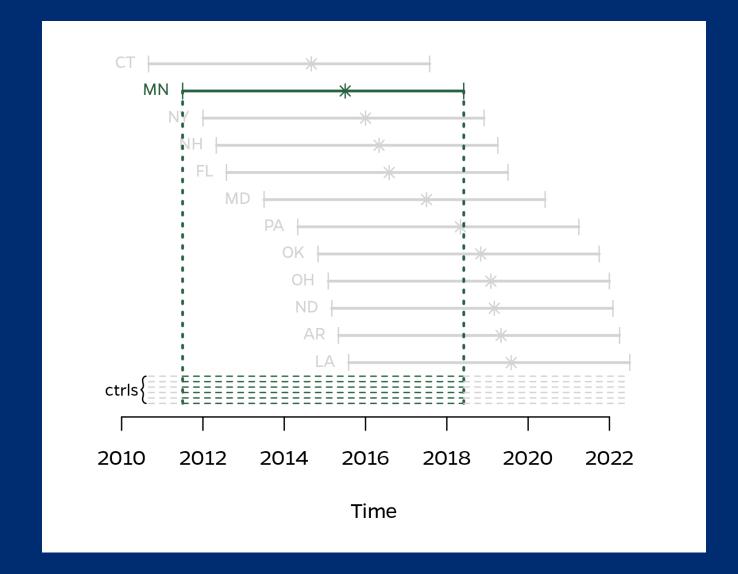




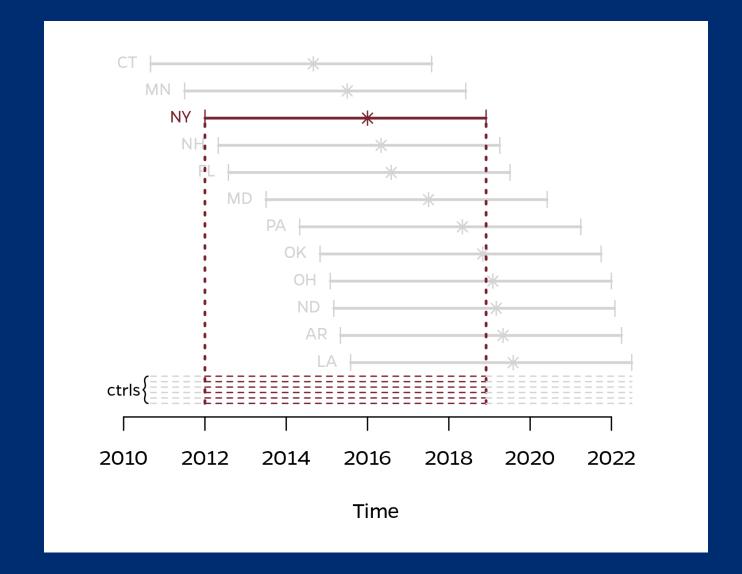








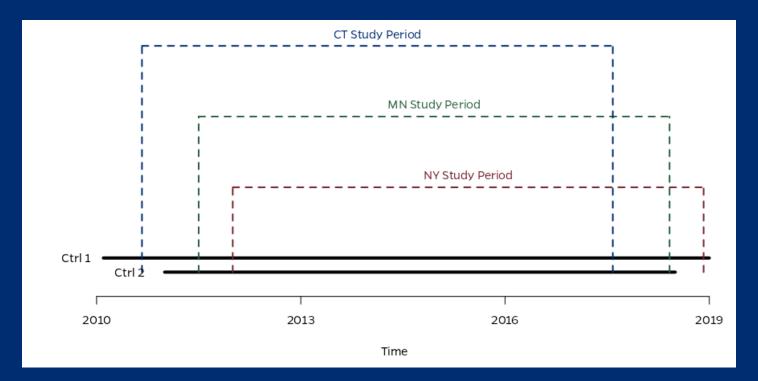






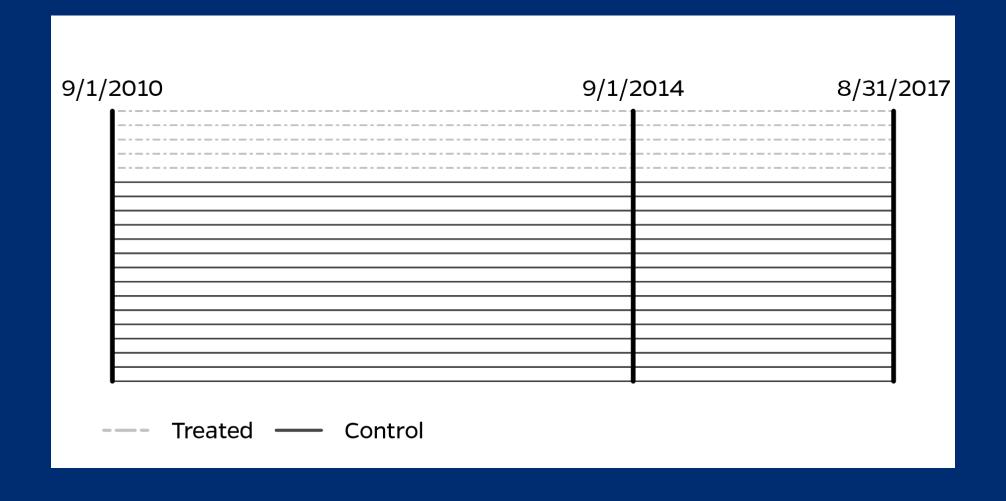


- ▶ Individuals in control states might appear in multiple cohorts.
 - "Ctrl 1" is in CT, MN, NY cohorts, but "Ctrl 2" is in MN cohort only.
- This induces correlation between treatment effect estimates for different cohorts!



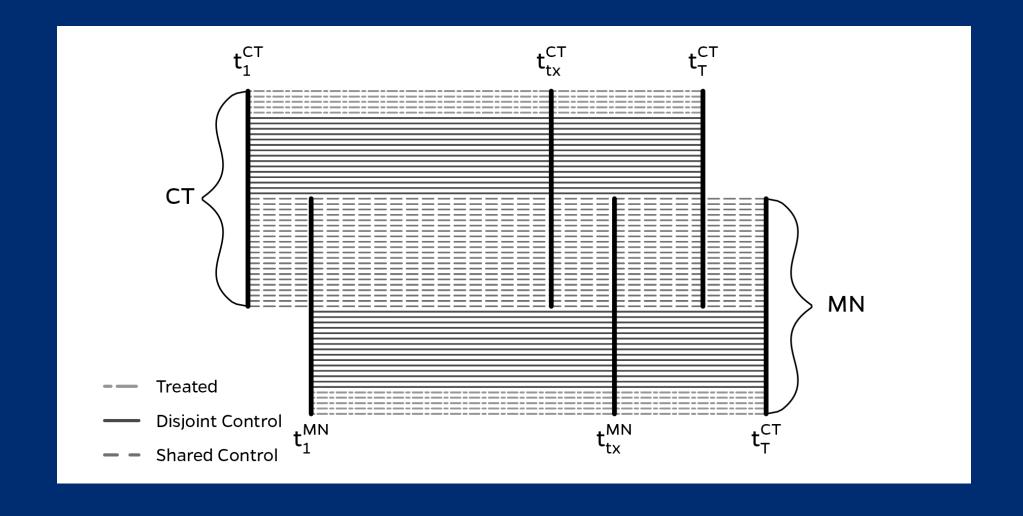
Medical Cannabis Study: CT Cohort





Shared Control Individuals





Correlation induced by shared control individuals



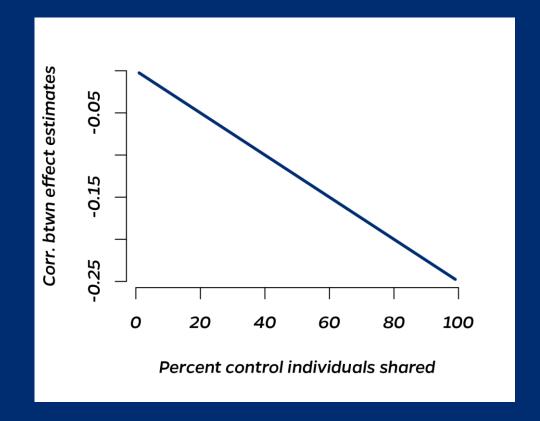
- ▶ Overall goal: Estimate an overall ATT across treated states
- Averaging correlated things changes the standard error of the average
- Only an issue when pooling effect estimates (not for single-state ATTs)
- This approach is for when individual-level data is available
- Things that affect correlation:
 - 1. Proportion of control individuals who are shared
 - 2. Duration of pre- and post-treatment periods
 - 3. Length of time between state implementation dates
 - 4. Within-person correlation in the outcome
 - 5. Between-person correlation in the outcome





- ► Very simple case:
 - One observation in each of pre- and posttreatment periods
 - One time unit between cohorts
- Key is percent of controls shared
- Correlation is negative!





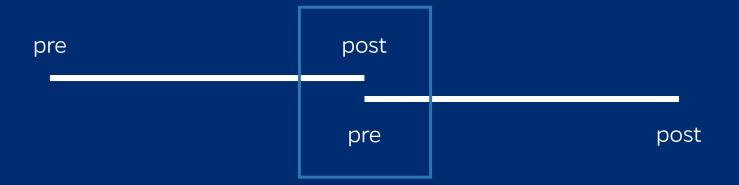




Remember the diff-in-diff ATT estimate for two timepoints:

$$\widehat{ATT} = (\overline{Y}_{post,tx} - \overline{Y}_{pre,tx}) - (\overline{Y}_{post,ctrl} - \overline{Y}_{pre,ctrl})$$

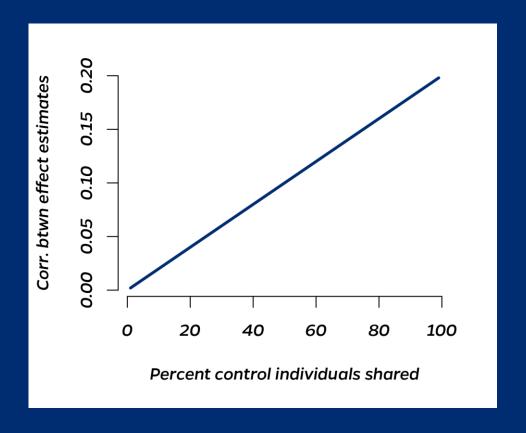
When there's a 1-unit gap in a 2-timepoint setting, one cohort's post period is the other's pre.







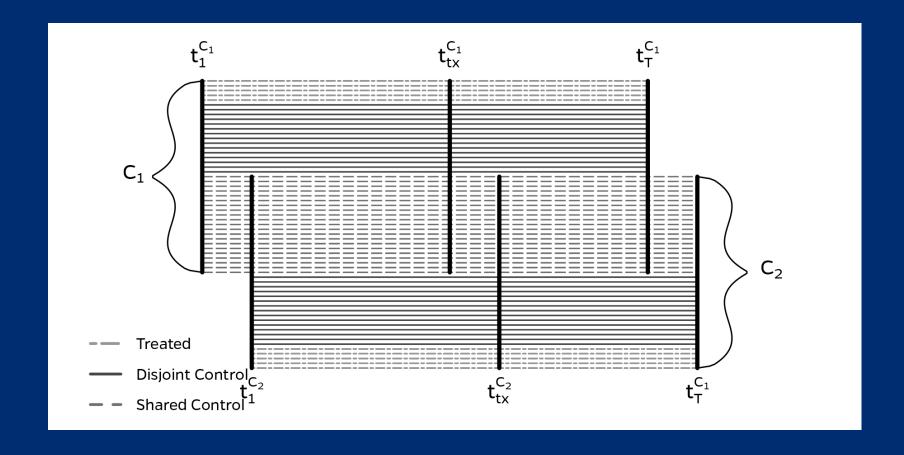
- ▶ 5 observations pre, 5 post
- Now correlation *increases* with the percent of shared controls and is positive.





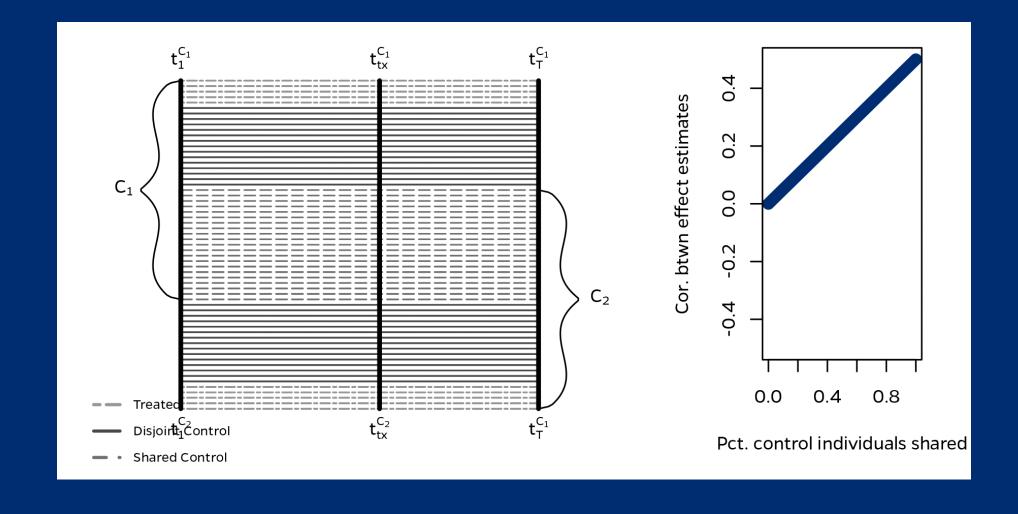


More opportunities for observations to be in both pre periods and both post periods!



Correlation depends on time gap





What do we do with this?



- ▶ Inverse-variance weighting to average estimates
 - Modifiable to incorporate correlation between estimates
- Overall average is not affected by correlation: only standard errors
 - ▶ Variance of a 2-state average changes by a factor of $(1 + \rho)$
- Example in simplified setting:
 - Correlation of -0.25 → 25% deflation in SE
 - Correlation of +0.25 > 25% inflation in SE

When do individual data matter for policy evaluation?



- Individual data let us better identify the population of interest
 - E.g., restrict to only people with chronic non-cancer pain
 - Can ensure constant sample composition over time (longitudinal data)
- Policies are implemented at "higher" levels (e.g., states)
 - Unclear that individual-level data is helpful in these analyses





- Individual level data is very big in these contexts (millions of people)
 - Rich data on individual trajectories seems useful!
 - Probably requires big data techniques (challenging)
- Aggregate level data is much smaller (e.g., 84 observations per state)
 - Significantly easier to work with!
 - Feels like there could be loss of statistical efficiency





Generate "fake" data from a model that we specify – we know the truth

$$Y_{uit} = \beta_0 + \beta_1 t + \beta_2 A_{st} + \beta_3(t) X_{uit} + b_{0si} + b_{0s} + \epsilon_{uit}$$

- Key here is that the relationship between Y_{uit} and covariate X_{uit} can vary with time.
- Analyze simulated data using a variety of methods
- Repeat 2000 times





$$Y_{uit} = \beta_0 + \beta_1 t + \beta_2 A_{st} + \beta_3(t) X_{uit} + b_{0si} + b_{0s} + \epsilon_{uit}$$

Initial scenarios:

- 1. No effect of X_{uit} (i.e., $\beta_3(t) = 0$)
- 2. Constant covariate X_{uit} with constant effect (i.e., $\beta_3(t) = \beta_3$)
- 3. Constant covariate X_{uit} with time-varying effect
- 4. Time-varying X_{uit} with constant effect

Do individual-level data improve statistical efficiency?



- So far, doesn't look like it.
- Models using individual-level data have slightly smaller SE's but at the cost of poor ergonomics
- Bigger takeaway might be about when to "cluster adjust" standard errors.
- Scenario at right is time-varying covariate with constant effect

Model	SE	RMSE	95% CI Covg.
Indiv. w/ OLS SE	0.017	0.016	96.8%
Indiv. w/ Indiv. SE	0.016	0.016	95.1%
Indiv. w/ State SE	0.015	0.019	91.4%
Agg. w/ OLS SE	0.018	0.018	95.0%
Agg. w/ State SE	0.018	0.13	91.7%

Do individual-level data improve statistical efficiency?



- Because policies are implemented at state level, individual-level data does not appear to improve statistical efficiency vs. aggregate-level data in evaluating state policy effects
 - Based on limited simulations!
- This is ongoing work still feels sort of counterintuitive.
 - We're actively trying to "break it" and find a situation where individual-level data does very well, but it's hard





- Individual-level data is useful for policy evaluation!
 - Helps identify population of interest
- When individual-level data is available, estimates from stacked difference-indifferences should be adjusted for correlation due to shared control individuals.
 - Sometimes sharing helps efficiency, sometimes it hurts
- Aggregate-level analyses (so far) don't seem to hurt anything
 - ▶ These are much easier to do fewer data





- Research reported in this presentation was supported by the National Institute on Drug Abuse of the National Institutes of Health under award number R01DA049789 (PI: McGinty)
- ▶ The content of this talk is solely the responsibility of the speaker and does not necessarily represent the official views of the National Institutes of Health, nor those of Johns Hopkins University.