Handling Correlation in Stacked Difference-in-Differences Estimates with Application to Medical Cannabis Policy

Nicholas J. Seewald

Department of Health Policy and Management
Johns Hopkins Bloomberg School of Public Health

Joint with K. Tormohlen, E.E. McGinty, and E.A. Stuart

Stuart Lab Meeting
August 30, 2022

Slides are online!

slides.nickseewald.com/stuartAug2022.pdf

Motivating Example: Medical Cannabis Laws and Opioid Prescribing

- **4x** increase in opioid prescribing in U.S. from 1999-2012
  - Opioid prescribing for chronic non-cancer pain has played a meaningful role
- Getting better: prescribing down since 2012, but still ~3x higher than 1999

Dart, R. C. et al. (2015). *New England Journal of Medicine*.

Motivating Example: Medical Cannabis Laws and Opioid Prescribing

- **4x** increase in opioid prescribing in U.S. from 1999-2012
  - Opioid prescribing for chronic non-cancer pain has played a meaningful role
- Getting better: prescribing down since 2012, but still ~3x higher than 1999

- Cannabis is a potentially effective treatment for chronic non-cancer pain, but evidence is limited.
- Patients with chronic non-cancer pain are eligible to use cannabis under all existing state medical cannabis laws

Dart, R. C. et al. (2015). *New England Journal of Medicine*.

Motivating Example: Medical Cannabis Laws and Opioid Prescribing

- **4x** increase in opioid prescribing in U.S. from 1999-2012
  - Opioid prescribing for chronic non-cancer pain has played a meaningful role
- Getting better: prescribing down since 2012, but still ~3x higher than 1999

- Cannabis is a potentially effective treatment for chronic non-cancer pain, but evidence is limited.
- Patients with chronic non-cancer pain are eligible to use cannabis under all existing state medical cannabis laws

Aim: Examine the effects of state medical cannabis laws on receipt of opioid and non-opioid treatment among patients with chronic non-cancer pain

Dart, R. C. et al. (2015). *New England Journal of Medicine*.

Motivating Example: Medical Cannabis Laws and Opioid Prescribing

Previous studies have found mixed results, but have key methodological limitations:

1. General population samples, and no individual-level data to identify individuals with chronic non-cancer pain
2. Policy endogeneity not addressed

Motivating Example: Medical Cannabis Laws and Opioid Prescribing

Previous studies have found mixed results, but have key methodological limitations:

1. **General population samples, and no individual-level data to identify individuals with chronic non-cancer pain**
2. Policy endogeneity not addressed

Individual-level data lets us identify the population, but adds methodological complexity in stacked difference-in-differences: existing methods assume comparison groups don't change across analyses.

Motivating Example: Medical Cannabis Laws and Opioid Prescribing

Our sample:

- 12 *treated* states that implemented a medical cannabis law between 2012 and 2018 and *do not also have recreational cannabis laws*
- 17 *comparison* states without medical or recreational cannabis laws

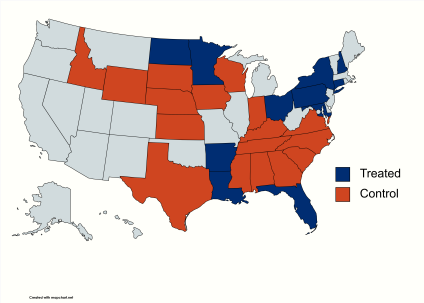

Motivating Example: Medical Cannabis Laws and Opioid Prescribing

Our sample:

- 12 *treated* states that implemented a medical cannabis law between 2012 and 2018 and *do not also have recreational cannabis laws*
- 17 *comparison* states without medical or recreational cannabis laws

Goal: Estimate the effect of implementing a medical cannabis law on opioid prescribing outcomes in each treatment state, relative to what would have happened in the absence of treatment.

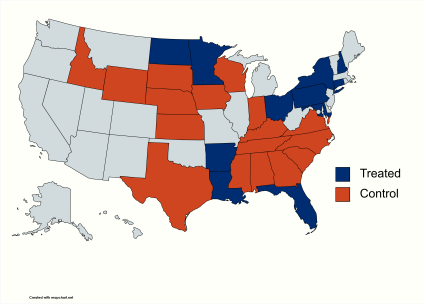

Difference-in-Differences: A Conceptual Introduction

- Compare change in outcome over time between treated and comparison groups
- Under assumption that treated group would look like comparison group in absence of treatment, can estimate causal treatment effect
  - This is called the *(counterfactual) parallel trends assumption*

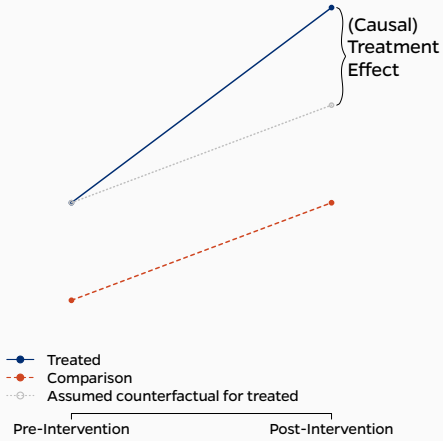

Difference-in-Differences: A Conceptual Introduction

- Parallel trends with two time periods:

$$\mathrm{E}[Y_2(0) - Y_1(0) \mid A = 1] = \mathrm{E}[Y_2(0) - Y_1(0) \mid A = 0]$$

- Strictest possible version with multiple time periods:

$$\mathrm{E}[Y_t(0) - Y_t'(0) \mid A = 1] = \mathrm{E}[Y_t(0) - Y_t'(0) \mid A = 0]$$

for all t in the post-tx period and t' in the pre-tx period

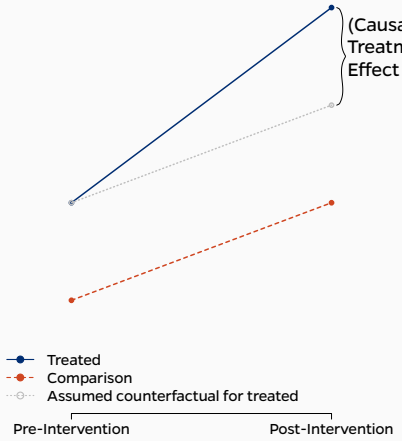

Difference-in-Differences: A Conceptual Introduction

Goal is to estimate the **average treated effect among the treated**:

$$\text{ATT}(t) = \text{E}\left[Y_t(1) - Y_t(0) \mid A = 1\right].$$

Under counterfactual parallel trends:

$$\begin{aligned}\text{ATT}(t) = &\left(\text{E}\left[Y_t \mid A = 1\right] - \text{E}\left[Y_{t'} \mid A = 1\right]\right) \\ &- \left(\text{E}\left[Y_t \mid A = 0\right] - \text{E}\left[Y_{t'} \mid A = 0\right]\right)\end{aligned}$$

for t' in the pre period, t in the post.

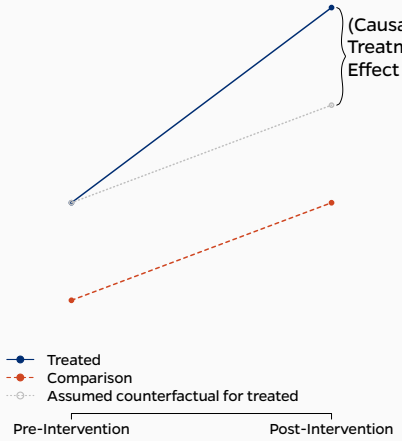

Difference-in-Differences: A Conceptual Introduction

- Using standard diff-in-diff to estimate an overall treatment effect under “staggered adoption” is problematic.
- Most common approach yields a biased treatment effect estimate partially based on inappropriate comparisons.
- Estimation procedures which get around this are available, but use aggregate data.

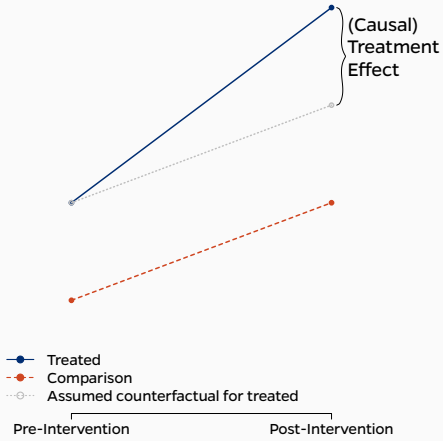

Goodman-Bacon, A. (2021). *Journal of Econometrics*.
Callaway, B. and Sant’Anna, P. H. C. (2021). *Journal of Econometrics*.

Difference-in-Differences: A Conceptual Introduction

How to get around this?

- We'll use standard diff-in-diff machinery to estimate a separate ATT *for each treated state*, then pool to get an average ATT.

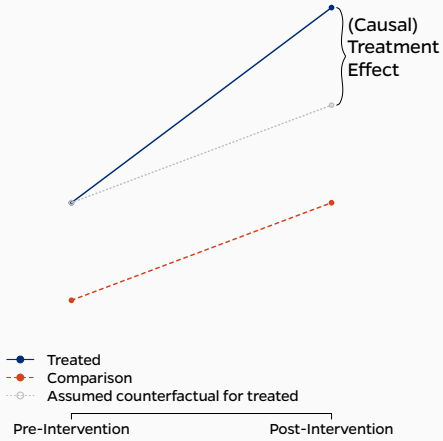

Medical Cannabis Study: Study Periods

- States implemented medical cannabis laws at different times
- Each state has its own 7-year study period anchored at implementation date
  - 4 years pre-law, 3 years post-law

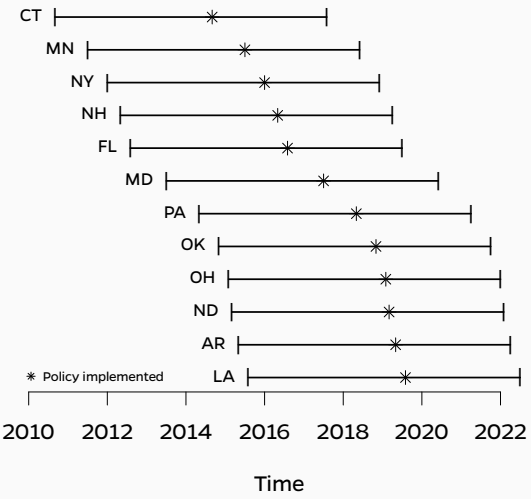

Medical Cannabis Study: State Cohorts

Data are individual-level commercial health insurance claims.

For each treatment state, we build a *cohort* of individuals in that state and the control states over the study period.

- Individuals 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: State Cohorts

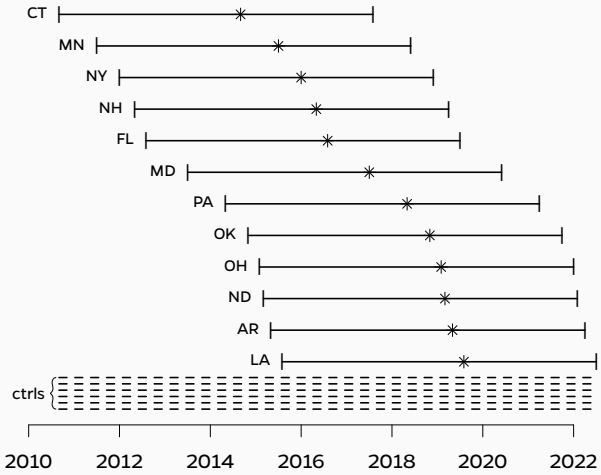

Medical Cannabis Study: State Cohorts

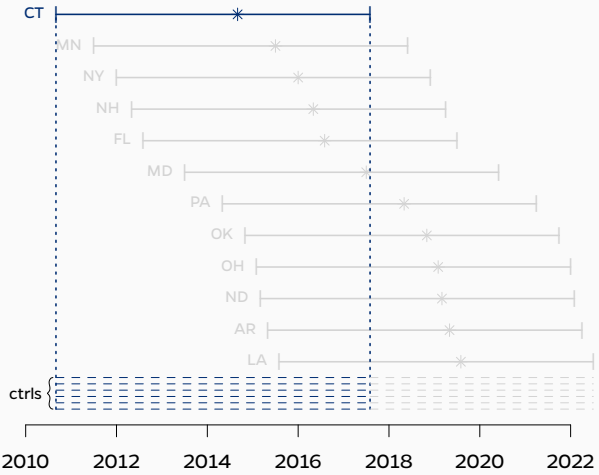

Medical Cannabis Study: State Cohorts

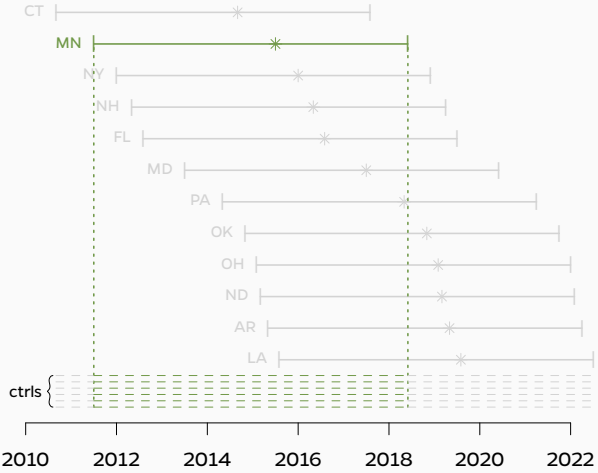

Medical Cannabis Study: State Cohorts

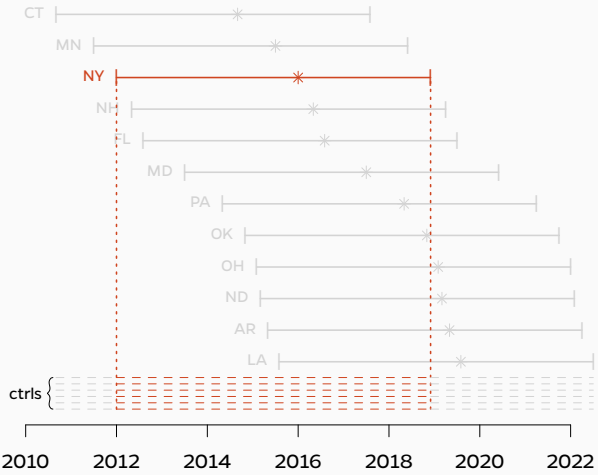

Shared Control Individuals

- 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!

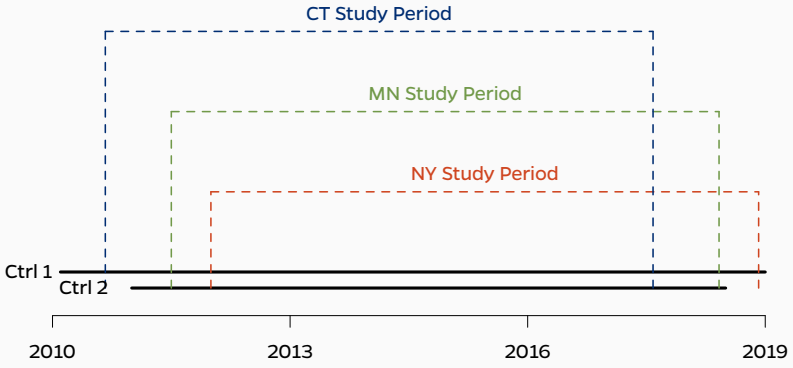

Shared Control Individuals

One cohort:

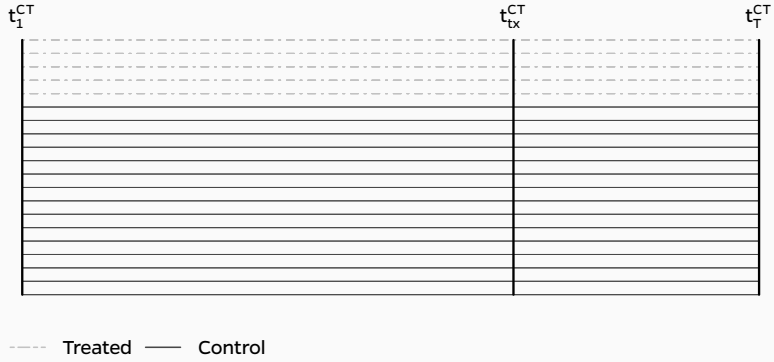

Shared Control Individuals

Two cohorts:

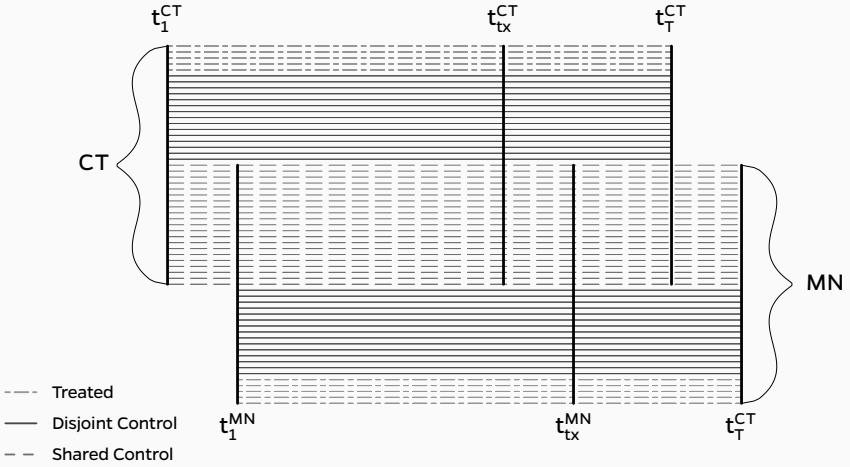

Handling Correlation Induced by Shared Control Indviduals

Goal: Estimate overall ATT, averaged across treated states.

- Correlation only an issue when pooling effect estimates
- Approach is for individual-level data
- **Big Idea:** Estimate pairwise correlation between estimates, then take inverse-variance weighted average.

A Common Approach to Diff-in-Diff

If only state s is treated, we could estimate its ATT as

$$\widehat{\text{ATT}}(s) = \left(\bar{Y}^{\text{tx}}_{s,\text{post}} - \bar{Y}^{\text{tx}}_{s,\text{pre}}\right) - \left(\bar{Y}^{\text{ctrl}}_{s,\text{post}} - \bar{Y}^{\text{ctrl}}_{s,\text{pre}}\right).$$

A Common Approach to Diff-in-Diff

If only state s is treated, we could estimate its ATT as

$$\widehat{\text{ATT}}(s) = \left(\bar{Y}^{\text{tx}}_{s,\text{post}} - \bar{Y}^{\text{tx}}_{s,\text{pre}}\right) - \left(\bar{Y}^{\text{ctrl}}_{s,\text{post}} - \bar{Y}^{\text{ctrl}}_{s,\text{pre}}\right).$$

Since there is only one treated unit, this is equivalent to fitting the linear model

$$\mathrm{E}\left[Y_{sit}\right] = \beta_{0,s} + \beta_{1,t} + \beta_2 A_{st};$$

then $\hat{\beta}_2 \equiv \widehat{\text{ATT}}(s)$.

A Common Approach to Diff-in-Diff

If only state s is treated, we could estimate its ATT as

$$\widehat{\text{ATT}}(s) = \left(\bar{Y}^{\text{tx}}_{s,\text{post}} - \bar{Y}^{\text{tx}}_{s,\text{pre}}\right) - \left(\bar{Y}^{\text{ctrl}}_{s,\text{post}} - \bar{Y}^{\text{ctrl}}_{s,\text{pre}}\right).$$

Since there is only one treated unit, this is equivalent to fitting the linear model

$$\mathrm{E}\left[Y_{sit}\right] = \beta_{0,s} + \beta_{1,t} + \beta_2 A_{st};$$

then $\hat{\beta}_2 \equiv \widehat{\text{ATT}}(s)$. Assuming states are mutually independent,

$$\begin{aligned}\text{Cov}\left(\widehat{\text{ATT}}(s), \widehat{\text{ATT}}(s')\right) = \text{Cov}\left(\bar{Y}^{\text{ctrl}}_{s,\text{post}}, \bar{Y}^{\text{ctrl}}_{s',\text{post}}\right) + \text{Cov}\left(\bar{Y}^{\text{ctrl}}_{s,\text{pre}}, \bar{Y}^{\text{ctrl}}_{s',\text{pre}}\right) \\ - \text{Cov}\left(\bar{Y}^{\text{ctrl}}_{s,\text{post}}, \bar{Y}^{\text{ctrl}}_{s',\text{pre}}\right) - \text{Cov}\left(\bar{Y}^{\text{ctrl}}_{s,\text{pre}}, \bar{Y}^{\text{ctrl}}_{s\prime,\text{post}}\right).\end{aligned}$$

A Common Approach to Diff-in-Diff

If only state s is treated, we could estimate its ATT as

$$\widehat{\text{ATT}}(s) = \left(\bar{Y}^{\text{tx}}_{s,\text{post}} - \bar{Y}^{\text{tx}}_{s,\text{pre}}\right) - \left(\bar{Y}^{\text{ctrl}}_{s,\text{post}} - \bar{Y}^{\text{ctrl}}_{s,\text{pre}}\right).$$

Since there is only one treated unit, this is equivalent to fitting the linear model

$$\mathrm{E}\left[Y_{sit}\right] = \beta_{0,s} + \beta_{1,t} + \beta_2 A_{st};$$

then $\hat{\beta}_2 \equiv \widehat{\text{ATT}}(s)$. Assuming states are mutually independent,

$$\begin{aligned}\text{Cov}\left(\widehat{\text{ATT}}(s), \widehat{\text{ATT}}(s')\right) = \boxed{\text{Cov}\left(\bar{Y}^{\text{ctrl}}_{s,\text{post}}, \bar{Y}^{\text{ctrl}}_{s',\text{post}}\right)} + \text{Cov}\left(\bar{Y}^{\text{ctrl}}_{s,\text{pre}}, \bar{Y}^{\text{ctrl}}_{s',\text{pre}}\right) \\ - \text{Cov}\left(\bar{Y}^{\text{ctrl}}_{s,\text{post}}, \bar{Y}^{\text{ctrl}}_{s',\text{pre}}\right) - \text{Cov}\left(\bar{Y}^{\text{ctrl}}_{s,\text{pre}}, \bar{Y}^{\text{ctrl}}_{s\prime,\text{post}}\right).\end{aligned}$$

Covariances with Shared Control Individuals

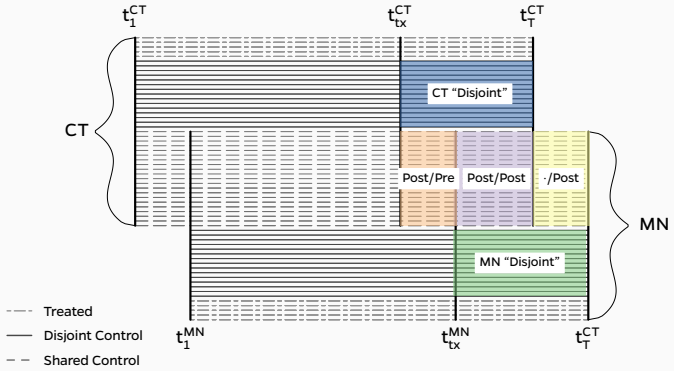

$$\text{Cov}\left(\bar{Y}^{\text{ctrl}}_{\text{CT,post}}, \bar{Y}^{\text{ctrl}}_{\text{MN}',\text{post}}\right) \approx \text{Cov}\left(\bar{Y}_{\text{CT Disjoint}} + \bar{Y}_{\text{Post/Pre}} + \bar{Y}_{\text{Post/Post}}, \bar{Y}_{\text{MN Disjoint}} + \bar{Y}_{\text{Post/Post}} + \bar{Y}_{\cdot/\text{Post}}\right)$$

Covariances with Shared Control Individuals

$$\operatorname{Cov}\left(\bar{Y}^{\text{ctrl}}_{\text{CT,post}}, \bar{Y}^{\text{ctrl}}_{\text{MN}',\text{post}}\right)$$

$$= \frac{1}{N^{\text{CT}}_{\text{ctrl}} N^{\text{MN}}_{\text{ctrl}} T^2_{\text{post}}} \operatorname{Cov}\left(\sum_{\text{CT Disjoint,post}} Y_{sit} + \sum_{\text{Post/Pre}} Y_{sit} + \sum_{\text{Post/Post}} Y_{sit}, \right.$$

$$\left. \sum_{\text{MN Disjoint}} Y_{sit} + \sum_{\text{Post/Post}} Y_{sit} + \sum_{\cdot/\text{Post}} Y_{sit} \right)$$

- This is a sum of *nine* covariances
- Assuming exchangeable within-person (i.e., longitudinal) correlation
- Computing these becomes a counting problem: how many people and timepoints contribute to each mean, and how?

When Does This Matter?

- A closed-form formula for the correlation between two effect estimates is available, but it's *messy*.
- Correlation between effect estimates depends on:
  - duration of pre- and post-treatment periods
  - delay between study period start times (Δ)
  - proportion of shared control individuals
  - within- and between-person correlations

When Does This Matter?

- In limited simulations, we see small but noticeable correlation between effect estimates (~10-15%)
  - Simple pre/post setting with 1-period unit gap in start times, all individuals are independent, exchangeable within-person correlation
- 10%+ correlations only with large proportion of shared control individuals ($\geq 75\%$)
- With two cohorts and when variance of estimates is constant, correlation increases variance of overall estimate by factor of $(\mathbf{1} + \rho)$ relative to if estimates were independent.

Ignoring this correlation leads to artificially small standard errors!

ρ is the correlation between estimates

Variance for Overall Treatment Effect

The overall treatment effect estimate is a precision-weighted average:

$$\widehat{\text{ATT}}_{\text{overall}} = \frac{1}{\text{tr}(\boldsymbol{\Sigma})} \sum_i \frac{1}{\Sigma_{ii}} \hat{\beta}_i$$

where $\hat{\boldsymbol{\beta}}$ is the vector of state-specific effect estimates and $\boldsymbol{\Sigma}$ is the (estimated, but taken as fixed) covariance of $\hat{\boldsymbol{\beta}}$.

Then

$$\text{Var}\left(\widehat{\text{ATT}}\right) = \frac{1}{\text{tr}(\boldsymbol{\Sigma})^2} \cdot \sum_{i,j} \frac{\Sigma_{ij}}{\Sigma_{ii}\Sigma_{jj}}$$

Conclusions

- Individual-level data is useful for identifying populations of interest in policy evaluation, but introduces methodological complexity.
- When using individual-level data that might be shared across cohorts in stacked diff-in-diff, it may be important to account for correlation between estimates
- A closed-form formula for induced correlation is available for select analyses

- Paper availiable on ArXiv soon!
  - Follow me on Twitter for updates: **@nickseewald**