

NBER WORKING PAPER SERIES

DESIGN-BASED ANALYSIS IN DIFFERENCE-IN-DIFFERENCES SETTINGS
WITH STAGGERED ADOPTION

Susan Athey
Guido W. Imbens

Working Paper 24963
<http://www.nber.org/papers/w24963>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2018

We are grateful for comments by participations in the conference in honor of Gary Chamberlain at Harvard in May 2018, and in particular by Gary Chamberlain. Gary's insights over the years have greatly affected our thinking on these problems. We also wish to thank Sylvia Kloskin and Michael Pollmann for superb research assistance. This research was generously supported by ONR grant N00014-17-1-2131. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w24963.ack>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2018 by Susan Athey and Guido W. Imbens. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Design-based Analysis in Difference-In-Differences Settings with Staggered Adoption
Susan Athey and Guido W. Imbens
NBER Working Paper No. 24963
August 2018
JEL No. C01,C23,C31

ABSTRACT

In this paper we study estimation of and inference for average treatment effects in a setting with panel data. We focus on the setting where units, e.g., individuals, firms, or states, adopt the policy or treatment of interest at a particular point in time, and then remain exposed to this treatment at all times afterwards. We take a design perspective where we investigate the properties of estimators and procedures given assumptions on the assignment process. We show that under random assignment of the adoption date the standard Difference-In-Differences estimator is an unbiased estimator of a particular weighted average causal effect. We characterize the properties of this estimand, and show that the standard variance estimator is conservative.

Susan Athey
Graduate School of Business
Stanford University
655 Knight Way
Stanford, CA 94305
and NBER
athey@stanford.edu

Guido W. Imbens
Graduate School of Business
Stanford University
655 Knight Way
Stanford, CA 94305
and NBER
Imbens@stanford.edu

1 Introduction

In this paper we study estimation of and inference for average treatment effects in a setting with panel data. We focus on the setting where units, e.g., individuals, firms, or states, adopt the policy or treatment of interest at a particular point in time, and then remain exposed to this treatment at all times afterwards. The adoption date at which units are first exposed to the policy may, but need not, vary by unit. We refer to this as a *staggered adoption design* (SAD), such designs are sometimes also referred to as event study designs. An early example is Athey and Stern [1998] where adoption of an enhanced 911 technology by counties occurs over time, with the adoption date varying by county. This setting is a special case of the general Difference-In-Differences (DID) set up (e.g., Card [1990], Meyer et al. [1995], Angrist and Pischke [2008], Angrist and Krueger [2000], Abadie et al. [2010], Borusyak and Jaravel [2016], Athey and Imbens [2006], Card and Krueger [1994], Freyaldenhoven et al. [2018], de Chaisemartin and D’Haultfoeuille [2018], Abadie [2005]) where, at least in principle, units can switch back and forth between being exposed or not to the treatment. In this SAD setting we are concerned with identification issues as well as estimation and inference. In contrast to most of the DID literature, e.g., Bertrand et al. [2004], Shah et al. [1977], Conley and Taber [2011], Donald and Lang [2007], Stock and Watson [2008], Arellano [1987, 2003], Abraham and Sun [2018], Wooldridge [2010], de Chaisemartin and D’Haultfoeuille [2017, 2018], we take a *design-based* perspective where the stochastic nature and properties of the estimators arises from the stochastic nature of the assignment of the treatments, rather than a *sampling-based* perspective where the uncertainty arises from the random sampling of units from a large population. Such a design perspective is common in the analysis of randomized experiments, e.g., Neyman [1923/1990], Rosenbaum [2002, 2017]. See also Aronow and Samii [2016], Abadie et al. [2016, 2017] for this approach in cross-section regression settings. This perspective is particularly attractive in the current setting when the sample comprises the entire population, e.g., all states of the US, or all countries of the world. Our critical assumptions involve restrictions on the assignment process as well as exclusion restrictions, but in general do not involve functional form assumptions. Commonly made common trend assumptions (de Chaisemartin and D’Haultfoeuille [2018], Abraham and Sun [2018]) follow from some of our assumptions, but are not the starting point.

As in Abraham and Sun [2018] we set up the problem with the adoption date, rather than the actual exposure to the intervention, as the basic treatment defining the potential outcomes. We

consider assumptions under which this discrete multivalued treatment (the adoption date) can be reduced to a binary one, defined as the indicator whether or not the treatment has already been adopted. We then investigate the interpretation of the standard DID estimator under assumptions about the assignment of the adoption date and under various exclusion restrictions. We show that under a random adoption date assumption, the standard DID estimator can be interpreted as the weighted average of several types of causal effects; within our framework, these concern the impact of different types of changes in the adoption date of the units. We also consider design-based inference for this estimand. We derive the exact variance of the DID estimator in this setting. We show that under a random adoption date assumption the standard Liang-Zeger (LZ) variance estimator (Liang and Zeger [1986], Bertrand et al. [2004]), or the clustered bootstrap, are conservative. For this case we propose an improved (but still conservative) variance estimator.

Our paper is most closely related to a very interesting set of recent papers on DID methods that explicitly focus on issues with heterogeneous treatment effects (Abraham and Sun [2018], de Chaisemartin and D’Haultfoeuille [2018], Han [2018], Goodman-Bacon [2017], Callaway and Sant’Anna [2018], Hull [2018], and Borusyak and Jaravel [2016]). Among other things these papers derive interpretations of the DID estimator as weighted averages of causal effects and bias terms under various assumptions. In many cases they find that these interpretations involve weighted averages of basic average causal effects with potentially negative weights and propose alternative estimators that do not involve negative weights.

2 Set Up

Using the potential outcome framework for causal inference, we consider a setting with a population of N units. Each of these N units are characterized by a set of potential outcomes in T periods for $T + 1$ treatment levels, $Y_{it}(a)$. Here $i \in \{1, \dots, N\}$ indexes the units, $t \in \mathbb{T} = \{1, \dots, T\}$ indexes the time periods, and the argument of the potential outcome function $Y_{it}(\cdot)$, $a \in \mathbb{A} = \mathbb{T} \cup \{\infty\} = \{1, \dots, T, \infty\}$ indexes the discrete treatment, the date that the binary policy was first adopted by a unit. Units can adopt the policy at any of the time periods $1, \dots, T$, or not adopt the policy at all during the period of observation, in which case we code the adoption date as ∞ . Once a unit adopts the treatment, it remains exposed to the

treatment for all periods afterwards. This set up is like that in Abraham and Sun [2018], and in contrast to most of the DID literature where the binary indicator whether a unit is exposed to the treatment in the current period indexes the potential outcomes. We observe for each unit in the population the adoption date $A_i \in \mathbb{A}$ and the sequence of T realized outcomes, Y_{it} , for $t \in \mathbb{T}$, where

$$Y_{it} \equiv Y_{it}(A_i),$$

is the realized outcome for unit i at time t . We may also observe pre-treatment characteristics, denoted by the K -component vector X_i , although for most of the discussion we abstract from their presence. Let \mathbf{Y} , \mathbf{A} , and \mathbf{X} denote the $N \times T$, $N \times 1$, and $N \times K$ matrices with typical elements Y_{it} , A_i , and X_{ik} respectively. Implicitly we have already made a *sutva*-type assumption (Rubin [1978], Imbens and Rubin [2015]) that units are not affected by the treatments (adoption dates) for other units. Our design-based analysis views the potential outcomes $Y_{it}(a)$ as deterministic, and only the adoption dates A_i , as well as functions thereof such as the realized outcomes as stochastic. Distributions of estimators will be fully determined by the adoption date distribution, with the number of units N and the number of time periods T fixed, unless explicitly stated otherwise. Following the literature we refer to this as a randomization, or designed-based, distribution (Rosenbaum [2017], Imbens and Rubin [2015], Abadie et al. [2017]), as opposed to a sampling-based distribution.

In many cases the units themselves are clusters of units of a lower level of aggregation. For example, the units may be states, and the outcomes could be averages of outcomes for individuals in that state, possibly of samples drawn from subpopulations from these states. In such cases N and T may be as small as 2, although in many of the cases we consider N will be at least moderately large. This distinction between cases where Y_{it} is itself an average over basic units or not, affects some, but not all, of the formal statistical analyses. It may make some of the assumptions more plausible, and it may affect the inference, especially if individual level outcomes and covariates are available.

Define $W(a, t) = \mathbf{1}_{a \leq t}$ to be the binary indicator for the adoption date a preceding t , and define W_{it} to be the indicator for the the policy having been adopted by unit i prior to, or at,

time t :

$$W_{it} \equiv W(A_i, t) = \mathbf{1}_{A_i \leq t},$$

so that the $N \times T$ matrix \mathbf{W} with typical element W_{it} has the form:

$$\mathbf{W}_{N \times T} = \begin{pmatrix} 0 & 0 & 0 & 0 & \dots & 0 & \text{(never adopter)} \\ 0 & 0 & 0 & 0 & \dots & 1 & \text{(late adopter)} \\ 0 & 0 & 0 & 0 & \dots & 1 & \\ 0 & 0 & 1 & 1 & \dots & 1 & \\ 0 & 0 & 1 & 1 & \dots & 1 & \text{(medium adopter)} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \\ 0 & 1 & 1 & 1 & \dots & 1 & \text{(early adopter)} \end{pmatrix}$$

Let $N_a \equiv \sum_{i=1}^N \mathbf{1}_{A_i=a}$ be the number of units in the sample with adoption date a , and define $\pi_a \equiv N_a/N$, for $a \in \mathbb{A}$, as the fraction of units with adoption date equal to a , and $\Pi_t \equiv \sum_{s=1}^t \pi_s$, for $t \in \mathbb{T}$, as the fraction of units with an adoption date on or prior to t .

Also define $\bar{Y}_t(a)$ to be the population average of the potential outcome in period t for adoption date a :

$$\bar{Y}_t(a) \equiv \frac{1}{N} \sum_{i=1}^N Y_{it}(a), \quad \text{for } t \in \mathbb{T}, a \in \mathbb{A}.$$

Define the average causal effect of adoption date a' relative to a , on the outcome in period t , as

$$\tau_{t,aa'} \equiv \bar{Y}_t(a') - \bar{Y}_t(a) = \frac{1}{N} \sum_{i=1}^N \left\{ Y_{it}(a') - Y_{it}(a) \right\}.$$

Abraham and Sun [2018] focus on slightly different building blocks, what they call $CATT_{a,t}$, which, for $0 \leq t \leq T - a$, are the super-population equivalent of $(1/N_a) \sum_{i|A_i=a} \{Y_{ia+t}(a) - Y_{ia+t}(\infty)\}$. The average causal effects $\tau_{t,aa'}$ are the building blocks of many of the estimands we

consider later. A particularly interesting average effect is

$$\tau_{t,\infty 1} = \frac{1}{N} \sum_{i=1}^N \left(Y_{it}(1) - Y_{it}(\infty) \right),$$

the average effect of switching the entire population from never adopting the policy ($a = \infty$), to adopting the policy in the first period ($a = 1$). Formally there is nothing special about the particular average effect $\tau_{t,\infty 1}$ relative to any other $\tau_{t,aa'}$, but $\tau_{t,\infty 1}$ will be useful as a benchmark. Part of the reason is that for all t and i the comparison $Y_{it}(1) - Y_{it}(\infty)$ is between potential outcomes for adoption prior to or at time t (namely adoption date $a = 1$) and potential outcomes for adoption later than t (namely, never adopting, $a = \infty$). In contrast, any other average effect $\tau_{t,aa'}$ will for some t involve comparing potential outcomes neither of which correspond to having adopted the treatment yet, or comparing potential outcomes both of which correspond to having adopted the treatment already. Therefore, $\tau_{t,\infty 1}$ reflects more on the effect of having adopted the policy than any other $\tau_{t,aa'}$.

3 Assumptions

We consider three sets of assumptions. The first set, containing only a single assumption, is about the design, that is, the assignment of the treatment, here the adoption date, conditional on the potential outcomes and possibly pretreatment variables. We refer to this as a design assumption because it can be guaranteed by design. The second set of assumptions is about the potential outcomes, and rules out the presence of certain treatment effects. These exclusion restrictions are substantive assumptions, and they cannot be guaranteed by design. The third set of assumptions consists of four auxiliary assumptions, two about homogeneity of certain causal effects, one about sampling from a large population, and one about an outcome model in a large population. The nature of these three sets of assumptions, and their plausibility, is very different, and it is in our view useful to carefully distinguish between them. The current literature often combines various parts of these assumptions implicitly in the notation used and in assumptions about the statistical models for the realized outcomes.

3.1 The Design Assumption

The first assumption is about the assignment process for the adoption date A_i . Our starting point is to assume that the adoption date is completely random:

Assumption 1. (RANDOM ADOPTION DATE) *For some set of positive integers N_a , for $a \in \mathbb{A}$,*

$$\text{pr}(\mathbf{A} = \mathbf{a}) = \left(\frac{N!}{\prod_{a \in \mathbb{A}} N_a!} \right)^{-1},$$

for all N -vectors \mathbf{a} such that for all $a \in \mathbb{A}$, $\sum_{i=1}^N \mathbf{1}_{a_i=a} = N_a$.

This assumption is obviously very strong. However, without additional assumptions that restrict either the potential outcomes, or expand what we observe, for example by including pre-treatment variables or covariates, this assumption has no testable implications in a setting with exchangeable units.

Lemma 1. (NO TESTABLE RESTRICTIONS) *Suppose all units are exchangeable. Then Assumption 1 has no testable implications for the joint distribution of (\mathbf{Y}, \mathbf{A}) .*

All proofs are given in the Appendix.

Hence, if we wish to relax the assumptions, we need to bring in additional information. Such additional information can come in the form of pretreatment variables, that is, variables that are known not to be affected by the treatment. In that case we can relax the assumption by requiring only that the adoption date is completely random within subpopulations with the same values for the pre-treatment variables. Additional information can also come in the form of limits on the treatment effects. The implications of such restrictions on the ability to relax the random adoption assumption is more complex, as discussed in more detail in Section 3.2.

Under Assumption 1 the marginal distribution of the adoption dates is fixed, and so also the fraction π_a is fixed in the repeated sampling thought experiment. This part of the set up is similar in spirit to fixing the number of treated units in the sample in a completely randomized experiment. It is convenient for obtaining finite sample results. Note that it implies that the adoption dates for units i and j are not independent. Note also that in the standard framework where the uncertainty arises solely from random sampling, this fraction does not remain constant in the repeated sampling thought experiment.

An important role is played by what we label the *adjusted treatment*, adjusted for unit and time period averages:

$$\dot{W}_{it} \equiv W_{it} - \bar{W}_{\cdot t} - \bar{W}_{i \cdot} + \bar{W},$$

where $\bar{W}_{\cdot t}$, $\bar{W}_{i \cdot}$, and \bar{W} are averages over units, time periods, and both, respectively:

$$\bar{W}_{\cdot t} \equiv \frac{1}{N} \sum_{i=1}^N W_{it} = \frac{1}{N} \sum_{i=1}^N \mathbf{1}_{A_i \leq t} = \frac{1}{N} \sum_{i=1}^N \sum_{s \leq t} \mathbf{1}_{A_i = s} = \sum_{s \leq t} \frac{1}{N} \sum_{i=1}^N \mathbf{1}_{A_i = s} = \sum_{s \leq t} \pi_s,$$

$$\bar{W}_{i \cdot} \equiv \frac{1}{T} \sum_{t=1}^T W_{it} = \mathbf{1}_{A_i \leq T} \frac{T+1-A_i}{T},$$

and

$$\bar{W} \equiv \frac{1}{T} \sum_{t=1}^T \bar{W}_{\cdot t} = \frac{1}{T} \sum_{t=1}^T \sum_{s \leq t} \pi_s = \frac{1}{T} \sum_{t=1}^T (T+1-t)\pi_t,$$

where, with some minor abuse of notation, we adopt the convention that $a\mathbf{1}_{a \leq T}$ is zero if $a = \infty$. Note that under Assumption 1, $ow_{i,t}$ and \bar{W} are non-stochastic. Using these representations we can write the adjusted treatment indicator as

$$\dot{W}_{it} = g(t, A_i),$$

where

$$g(t, a) \equiv \left(\mathbf{1}_{a \leq t} - \sum_{s \leq t} \pi_s \right) + \frac{1}{T} \left(a\mathbf{1}_{a \leq T} - \sum_{s=1}^T s\pi_s \right) + \frac{T+1}{T} (\mathbf{1}_{a=\infty} - \pi_\infty). \quad (3.1)$$

Because the marginal distribution of A_i is fixed under Assumption 1, the sum $\sum_{i,t} \dot{W}_{it}^2$ is non-stochastic under this assumption, even though \dot{W}_{it} and thus \dot{W}_{it}^2 are stochastic. This fact enables us to derive exact finite sample results for the standard DID estimator as discussed in Section 4. This is similar in spirit to the derivation of the exact variance for the estimator for the average treatment effect in completely randomized experiments when we fix the number of treated and controls.

3.2 Exclusion Restrictions

The next two assumptions concern the potential outcomes. Their formulation does not involve the assignment mechanism, that is, the distribution of the adoption date. In essence these are exclusion restrictions, assuming that particular causal effects are absent. Collectively these two assumptions imply that we can think of the treatment as a binary one, the only relevant component of the adoption date being whether a unit is exposed to the treatment at the time we measure the outcome. Versions of such assumptions are also considered in Borusyak and Jaravel [2016], de Chaisemartin and D’Haultfœuille [2018], Abraham and Sun [2018] and Imai and Kim [2016], where in the latter a graphical approach is taken in the spirit of the work by Pearl [2000].

The first of the two assumptions, and likely the more plausible of the two in practice, rules out effects of future adoption dates on current outcomes. More precisely, it assumes that if the policy has not been adopted yet, the exact future date of the adoption has no causal effect on potential outcomes for the current period.

Assumption 2. (NO ANTICIPATION) *For all units i , all time periods t , and for all adoption dates a , such that $a > t$,*

$$Y_{it}(a) = Y_{it}(\infty).$$

We can also write this assumption as requiring that for all (i, t, a) ,

$$Y_{it}(a) = \mathbf{1}_{a \leq t} Y_{it}(a) + \mathbf{1}_{a > t} Y_{it}(\infty), \quad \text{or} \quad \mathbf{1}_{a > t} (Y_{it}(a) - Y_{it}(\infty)) = 0,$$

with the last representation showing most clearly how the assumption rules out certain causal effects. Note that this assumption does not involve the adoption date, and so does not restrict the distribution of the adoption dates. Violations of this assumption may arise if the policy is anticipated prior to its implementation.

The next assumption is arguably much stronger. It asserts that for potential outcomes in period t it does not matter how long the unit has been exposed to the treatment, only whether the unit is exposed at time t .

Assumption 3. (INVARIANCE TO HISTORY) *For all units i , all time periods t , and for all*

adoption dates a , such that $a \leq t$,

$$Y_{it}(a) = Y_{it}(1).$$

This assumption can also be written as

$$Y_{it}(a) = \mathbf{1}_{a \leq t} Y_{it}(1) + \mathbf{1}_{a > t} Y_{it}(a), \quad \text{or} \quad \mathbf{1}_{a \leq t} (Y_{it}(a) - Y_{it}(1)) = 0,$$

with again the last version of the assumption illustrating the exclusion restriction in this assumption. Again, the assumption does not rule out any correlation between the potential outcomes and the adoption date, only that there is no causal effect of an early adoption versus a later adoption on the outcome in period t , as long as adoption occurred before or on period t .

In general, this assumption is very strong. However, there are important cases where it may be more plausible. Suppose the units are clusters of individuals, where in each period we observe different sets of individuals. To be specific, suppose the the units are states, the time periods are years, and outcome is the employment rate for twenty-five year olds, and the treatment is the presence or absence of some regulation, say a subsidy for college tuition. In that case it may well be reasonable to assume that the educational choices for students graduating high school in a particular state depends on what the prevailing subsidy is, but much less on the presence of subsidies in previous years.

If both the exclusion restrictions, that is, both Assumptions 2 and 3, hold, then the potential outcome $Y_{it}(a)$ can be indexed by the binary indicator $W(a, t) = \mathbf{1}_{a \leq t}$:

Lemma 2. (BINARY TREATMENT) *Suppose Assumptions 2 and 3 hold. Then for all units i , all time periods t and adoption dates $a > a'$, (i)*

$$Y_{it}(a') - Y_{it}(a) = \mathbf{1}_{a' \leq t < a} (Y_{it}(1) - Y_{it}(\infty)),$$

so that,

$$Y_{it}(a) = Y_{it}(\infty) + \mathbf{1}_{a \leq t} (Y_{it}(1) - Y_{it}(\infty)) = \begin{cases} Y_{it}(\infty) & \text{if } a \leq t \\ Y_{it}(1) & \text{otherwise,} \end{cases}$$

and, for all time periods t , and adoption dates $a > a'$, (ii)

$$\tau_{t,aa'} = \tau_{t,\infty} \mathbf{1}_{a' \leq t < a} = \begin{cases} \tau_{t,\infty} & \text{if } a' \leq t < a, \\ 0 & \text{otherwise.} \end{cases}$$

If these two assumptions hold, we can therefore simplify the notation for the potential outcomes and focus on $Y_{it}(1)$ and $Y_{it}(\infty)$.

Note that these two assumptions are substantive, and cannot be guaranteed by design. This in contrast to the Assumption 1, which can be guaranteed by randomization of the adoption date. It is also important to note that in many empirical studies Assumptions 2 and 3 are made, often implicitly by writing a model for realized outcome Y_{it} that depends solely on the contemporaneous treatment exposure W_{it} , and not on the actual adoption date A_i or treatment exposure $W_{it'}$ in other periods t' . In the current discussion we want to be explicit about the fact that this restriction is an assumption, and that it does not automatically hold. Note that the assumption does not restrict the time series dependence between the potential outcomes.

It is trivial to see that without additional information, the exclusion restrictions in Assumptions 2 and 3 have no testable implications because they impose restrictions on pairs of potential outcomes that cannot be observed together. However, in combination with random assignment, 2 and 3, there are testable implications as long as $T \geq 2$ and there is some variation in the adoption date.

Lemma 3. (TESTABLE RESTRICTIONS FROM THE EXCLUSION RESTRICTIONS) (i) *Assumptions 2 and 3 jointly have no testable implications for the joint distribution of (\mathbf{Y}, \mathbf{W}) .*

(ii) *Suppose $T \geq 2$, and $\pi_2, \pi_\infty > 0$. Then the combination of Assumptions 1–3 impose testable restrictions on the joint distribution of (\mathbf{Y}, \mathbf{W}) .*

3.3 Auxiliary Assumptions

In this section we consider four auxiliary assumptions that are convenient for some analyses, and in particular can have implications for the variance of specific estimators, but that are not essential in many cases. These assumptions are often made in empirical analyses without researchers explicitly discussing them.

The first of these assumptions assumes that the effect of adoption date a' , relative to adoption

date a , on the outcome in period t , is the same for all units.

Assumption 4. (CONSTANT TREATMENT EFFECT OVER UNITS) *For all units i, j and for all time periods t and all adoption dates a and a'*

$$Y_{it}(a) - Y_{it}(a') = Y_{jt}(a) - Y_{jt}(a').$$

The second assumption restricts the heterogeneity of the treatment effects over time.

Assumption 5. (CONSTANT TREATMENT EFFECT OVER TIME) *For all units i and all time periods t and t'*

$$Y_{it}(1) - Y_{it}(\infty) = Y_{it'}(1) - Y_{it'}(\infty).$$

We only restrict the time variation for comparisons of the adoption dates 1 and ∞ because we typically use this assumption in combination with Assumptions 2 and 3. In that case we obtain a constant binary treatment effect set up, as summarized in the following Lemma.

Lemma 4. (BINARY TREATMENT AND CONSTANT TREATMENT EFFECTS) *Suppose Assumptions 2-5 hold. Then for all t and $a' < a$*

$$Y_{it}(a') - Y_{it}(a) = \mathbf{1}_{a' \leq t < a} \tau_{1\infty}.$$

The final assumption allows us to view the potential outcomes as random by postulating a large population from which the sample is drawn.

Assumption 6. (RANDOM SAMPLING) *The sample can be viewed as a random sampling from an infinitely large population, with joint distribution for $(A_i, Y_{it}(a), a \in \mathbb{A}, t \in \mathbb{T})$ denoted by $f(a, y_1(1), \dots, y_T(\infty))$.*

Under this assumption we can put additional structure on average potential outcomes.

Assumption 7. (ADDITIVITY)

$$\mathbb{E}[Y_{it}(\infty)] = \alpha_i + \beta_t.$$

4 Difference-In-Differences Estimators: Interpretation and Inference

In this section we consider the standard DID set up (e.g., Meyer et al. [1995], Bertrand et al. [2004], Angrist and Pischke [2008], Donald and Lang [2007], de Chaisemartin and D’Haultfœuille [2018]). In the simplest setting with N units and T time periods, without additional covariates, the realized outcome in period t for unit i is modeled as

$$Y_{it} = \alpha_i + \beta_t + \tau W_{it} + \varepsilon_{it}. \quad (4.1)$$

In this model there are unit effects α_i and time effects β_t , but both are additive with interactions between them ruled out. The effect of the treatment is implicitly assumed to be additive and constant across units and time periods.

We interpret the DID estimand under the randomized adoption date assumption, leading to a different setting from that considered in de Chaisemartin and D’Haultfœuille [2018], Abraham and Sun [2018], Goodman-Bacon [2017]. We also derive its variance and show that in general it is lower than the standard random-sampling based variance. Finally we propose a variance estimator that is smaller than the regular variance estimators such as the Liang-Zeger and clustered bootstrap variance estimators.

4.1 Difference-In-Differences Estimators

Consider the least squares estimator for τ based on the specification in (4.1):

$$\left(\hat{\tau}_{\text{did}}, \{\hat{\alpha}_i\}_{i=2}^N, \{\hat{\beta}_t\}_{t=1}^T \right) = \arg \min_{\tau, \{\alpha_i\}_{i=2}^N, \{\beta_t\}_{t=1}^T} \sum_{i=1}^N \sum_{t=1}^T (Y_{it} - \alpha_i - \beta_t - \tau W_{it})^2.$$

It is convenient to write $\hat{\tau}_{\text{did}}$ in terms of the adjusted treatment indicator \dot{W}_{it} as

$$\hat{\tau}_{\text{did}} = \frac{\sum_{i,t} \dot{W}_{it} Y_{it}}{\sum_{i,t} \dot{W}_{it}^2}.$$

The primary question of interest in this section concerns the properties of the estimator $\hat{\tau}_{\text{did}}$. This includes the interpretation of its expectation under various sets of assumptions, and its variance. Mostly we focus on exact properties in finite samples.

In order to interpret the expected value of $\hat{\tau}_{\text{did}}$ we consider some intermediate objects. Define, for all adoption dates $a \in \mathbb{A}$, and all time periods $t \in \mathbb{T}$ the average of the outcome in period t for units with adoption date a :

$$\bar{Y}_{t,a} = \begin{cases} \frac{1}{N_a} \sum_{i:A_i=a} Y_{it} & \text{if } N_a > 0, \\ 0 & \text{otherwise.} \end{cases}$$

Under Assumption 1 the stochastic properties of these averages are well-defined because the N_a are fixed over the randomization distribution. The averages are stochastic because the realized outcomes depend on the adoption date. Define also the following two difference between outcome averages:

$$\hat{\tau}_{t,aa'} = \bar{Y}_{t,a'} - \bar{Y}_{t,a}.$$

In general these differences do not have a causal interpretation. Such an interpretation requires some assumptions, for example, on random assignment of the adoption date.

EXAMPLE: To facilitate the interpretation of some of the results it is useful to consider a special case where the results from completely randomized experiments directly apply. Suppose $\mathbb{T} = \{1, 2\}$, and $\mathbb{A} = \{2, \infty\}$, with a fraction $\pi = \pi_2 = 1 - \pi_\infty$ adopting the policy in the second period. Suppose also that $Y_{i1}(a) = 0$ for all i and a . Then the DID estimator is

$$\hat{\tau}_{\text{did}} = \hat{\tau}_{2,2\infty} = \bar{Y}_{2,2} - \bar{Y}_{2,\infty} = \frac{1}{N_2} \sum_{i:A_i=2} Y_{i2} - \frac{1}{N_\infty} \sum_{i:A_i=\infty} Y_{i2},$$

the simple difference in means for the second period outcomes for adopters and non-adopters. Under Assumption 1, the standard results for the variance of the difference in means for a randomized experiments apply (e.g., Neyman [1923/1990], Imbens and Rubin [2015]), and the

exact variance of $\hat{\tau}_{\text{did}}$ is,

$$\begin{aligned} \mathbb{V}(\hat{\tau}_{\text{did}}) &= \frac{1}{N_2(N-1)} \sum_{i=1}^N \{Y_{i2}(2) - \bar{Y}_2(2)\}^2 + \frac{1}{N_\infty(N-1)} \sum_{i=1}^N \{Y_{i2}(\infty) - \bar{Y}_2(\infty)\}^2 \\ &\quad - \frac{1}{N(N-1)} \sum_{i=1}^N \{(Y_{i2}(2) - \bar{Y}_2(2)) - (Y_{i2}(\infty) - \bar{Y}_2(\infty))\}^2. \end{aligned}$$

The standard Neyman estimator for this variance ignores the third term, and uses unbiased estimators for the first two terms, leading to:

$$\hat{\mathbb{V}}(\hat{\tau}_{\text{did}}) = \frac{1}{N_2(N_2-1)} \sum_{i:A_i=2} \{Y_{i2} - \bar{Y}_{2,2}\}^2 + \frac{1}{N_\infty(N_\infty-1)} \sum_{i:A_i=\infty} \{Y_{i2} - \bar{Y}_{2,\infty}\}^2.$$

□

4.2 The Interpretation of Difference-In-Differences Estimators

The following weights play an important role in the interpretation of the DID estimand:

$$\gamma_{t,a} \equiv \frac{\pi_a g(t,a)}{\sum_{t' \in \mathbb{T}} \sum_{a' \in \mathbb{A}} \pi_{a'} g(t', a')^2}, \quad \gamma_{t,+} \equiv \sum_{a \leq t} \gamma_{t,a}, \quad \text{and} \quad \gamma_{t,-} \equiv \sum_{a > t} \gamma_{t,a}, \quad (4.2)$$

with $g(a, t)$ as defined in (3.1). Note that these weights are non-stochastic, that is, fixed over the randomization distribution.

EXAMPLE (CTD): Continuing the example with two periods and adoption in the second period or never, we have in that case

$$\gamma_{t,a} = \begin{cases} 0 & \text{if } (t, a) = (1, 1), \\ 0 & \text{if } (t, a) = (2, 1), \\ -1 & \text{if } (t, a) = (1, 2), \\ 1 & \text{if } (t, a) = (2, 2), \\ 1 & \text{if } (t, a) = (1, \infty), \\ -1 & \text{if } (t, a) = (2, \infty), \end{cases} \quad \gamma_{t,+} = \begin{cases} 0 & \text{if } t = 1, \\ 1 & \text{if } t = 2, \end{cases} \quad \text{and} \quad \gamma_{t,-} = \begin{cases} 0 & \text{if } t = 1, \\ -1 & \text{if } t = 2. \end{cases}$$

□

The weights $\gamma_{t,a}$ have some important properties,

$$\sum_{t \in \mathbb{T}} \gamma_{t,+} = 1 \quad \sum_{t \in \mathbb{T}} \gamma_{t,-} = -1, \quad \text{and} \quad \sum_{t=1}^T \sum_{a \in \mathbb{A}} \gamma_{t,a} = \sum_{t \in \mathbb{T}} \gamma_{t,+} + \sum_{t \in \mathbb{T}} \gamma_{t,-} = 0.$$

Now we can state the first main result of the paper.

Lemma 5. *We can write $\hat{\tau}_{\text{did}}$ as*

$$\hat{\tau}_{\text{did}} = \sum_{t \in \mathbb{T}} \sum_{a \in \mathbb{A}} \gamma_{t,a} \bar{Y}_{t,a} = \sum_{t \in \mathbb{T}} \gamma_{t,+} \hat{\tau}_{t,\infty 1} + \sum_{t \in \mathbb{T}} \sum_{a > t} \gamma_{t,a} \hat{\tau}_{t,\infty a} - \sum_{t \in \mathbb{T}} \sum_{a \leq t} \gamma_{t,a} \hat{\tau}_{t,a1}. \quad (4.3)$$

Comment 1. Alternative characterizations of the DID estimator or estimand as a weighted average of potentially causal comparisons are presented in Abraham and Sun [2018], de Chaisemartin and D’Haultfœuille [2018], Han [2018], Goodman-Bacon [2017], and Borusyak and Jaravel [2016]). The characterizations differ in terms of the building blocks that are used in the representation and the assumptions made. Like our representation, the representation in Abraham and Sun [2018] is in terms of average causal effects of different adoption dates, but it imposes no-anticipation. Goodman-Bacon [2017] presents the DID estimator in terms of basic two-group DID estimators. Like our representation, the Goodman-Bacon [2017] is mechanical and does not rely on any assumptions. To endow the building blocks and the representation itself with a causal interpretation requires some assumption on, for example, the assignment mechanism. □

Comment 2. The lemma implies that the DID estimator has an interpretation as a weighted average of simple estimators for the causal effect of changes in adoption dates, the $\hat{\tau}_{t,aa'}$. Moreover, the estimator can be written as the sum of three averages of these $\hat{\tau}_{t,aa'}$. The first is a weighted average of the $\hat{\tau}_{t,\infty 1}$, which are all averages of switching from never adopting to adopting in the first period, meaning that these are averages of changes in adoption dates that involve switching from not being treated at time t to being treated at time t . The sum of the weights for these averages is one, although not all the weights are necessarily non-negative. The second sum is a weighted sum of $\hat{\tau}_{t,\infty a}$, for $a > t$, so that the causal effect always involves changing the adoption date from never adopting to adopting some time after t , meaning that the comparison is between potential outcomes neither of which involves being treated at the time. The sum of the weights for these averages is one again. The third sum is a weighted sum of $\hat{\tau}_{t,a1}$, for $a \leq t$,

so that the causal effect always involves changing the adoption date from adopting prior to, or at time, t relative to adopting at the initial time, meaning that the comparison is between potential outcomes both of which involves being treated at the time. These weights sum to minus one. \square

If we are willing to make the random adoption date assumption we can give this representation a causal interpretation:

Theorem 1. *Suppose Assumption 1 holds. Then (i):*

$$\mathbb{E}[\hat{\tau}_{t,aa'}] = \tau_{t,aa'},$$

and (ii)

$$\mathbb{E}[\hat{\tau}_{\text{did}}] = \sum_{t \in \mathbb{T}} \gamma_{t,+} \tau_{t,\infty 1} + \sum_{t \in \mathbb{T}} \sum_{a > t} \gamma_{t,a} \tau_{t,\infty a} - \sum_{t \in \mathbb{T}} \sum_{a \leq t} \gamma_{t,a} \tau_{t,a1}.$$

Suppose also Assumption 2 holds. Then (iii):

$$\mathbb{E}[\hat{\tau}_{\text{did}}] = \sum_{t \in \mathbb{T}} \sum_{a \leq t} \gamma_{t,a} \tau_{t,\infty a}.$$

Suppose also Assumption 3 holds. Then (iv):

$$\mathbb{E}[\hat{\tau}_{\text{did}}] = \sum_{t=1}^T \gamma_{t,+} \tau_{t,\infty 1}.$$

Suppose also Assumption 5 holds. Then (v):

$$\mathbb{E}[\hat{\tau}_{\text{did}}] = \tau_{\infty 1}.$$

Part (iii) of the theorem where we make the no-anticipation assumption is closely related to one of the results in Abraham and Sun [2018], who make a super-population common trend assumption that, in the super-population context, weakens our random adoption date assumption. Part (iv) of the theorem, where we assume both the exclusion restrictions so that the treatment is effectively a binary one, is related to the results in de Chaisemartin and D'Haultfœuille [2018],

although unlike those authors we do not restrict the trends in the potential outcomes.

Without either Assumptions 2 or 3, the estimand τ_{did} has a causal interpretation, but it is not clear it is a very interesting one concerning the receipt of the treatment. With the no-anticipation assumption (Assumption 2), the interpretation, as given in part (iii) of the theorem, is substantially more interesting. Now the estimand is a weighted average of $\tau_{t,\infty a}$ for $a \leq t$, with weights summing to one. These $\tau_{t,\infty a}$ are the average causal effect of changing the adoption date from never adopting to some adoption date prior to, or equal to, time t , so that the average always involves switching from not being exposed to the treatment to being exposed to the treatment.

4.3 The Randomization Variance of the Difference-In-Differences Estimators

In this section we derive the randomization variance for $\hat{\tau}_{\text{did}}$ under the randomized adoption date assumption. We do not rely on other assumptions here, although they may be required for making the estimand a substantively interesting one. The starting point is the representation $\hat{\tau}_{\text{did}} = \sum_{t,a} \gamma_{t,a} \bar{Y}_{t,a}$. Because under Assumption 1 the weights $\gamma_{t,a}$ are fixed, the variance is

$$\mathbb{V}(\hat{\tau}_{\text{did}}) = \sum_{t,a} \gamma_{t,a}^2 \mathbb{V}(\bar{Y}_{t,a}) + \sum_{(t,a) \neq (t',a')} \gamma_{t,a} \gamma_{t',a'} \mathbb{C}(\bar{Y}_{t,a}, \bar{Y}_{t',a'}).$$

Note that the $\gamma_{t,a}$ are known. Working out the variance $\mathbb{V}(\bar{Y}_{t,a})$, and finding an unbiased estimator for it, is straightforward. It is more challenging to infer the covariance terms $\mathbb{C}(\bar{Y}_{t,a}, \bar{Y}_{t',a'})$, and even more difficult to estimate them. In general that is not possible. Note that for a sampling-based variance the $\gamma_{t,a}$ are not fixed, because in different samples the fractions with a particular adoption date will be stochastic. This in general leads to a larger variance, as we verify in the simulations.

Define

$$Y_i(a) = \sum_{t=1}^T \gamma_{t,a} Y_{it}(a), \quad \bar{Y}(a) = \sum_{t=1}^T \gamma_{t,a} \bar{Y}_t(a) \quad \text{and} \quad \bar{Y}_a = \sum_{t=1}^T \gamma_{t,a} \bar{Y}_{t,a}.$$

Now we can write $\hat{\tau}^{\text{did}}$ as

$$\hat{\tau}^{\text{did}} = \sum_{a \in \mathbb{A}} \sum_{t \in \mathbb{T}} \gamma_{t,a} \bar{Y}_{t,a} = \sum_{a \in \mathbb{A}} \bar{Y}_a.$$

Define also

$$S_a^2 = \frac{1}{N-1} \sum_{i=1}^N (Y_i(a) - \bar{Y}(a))^2,$$

and

$$V_{a,a'}^2 = \frac{1}{N-1} \sum_{i=1}^N \left\{ (Y_i(a) - \bar{Y}(a)) + (Y_i(a') - \bar{Y}(a')) \right\}^2.$$

Theorem 2. *Suppose Assumptions 1 holds. Then the exact variance of $\hat{\tau}_{\text{did}}$ over the randomization distribution is*

$$\mathbb{V}(\hat{\tau}_{\text{did}}) = \sum_{a \in \mathbb{A}} S_a^2 \left(\frac{1}{N_a} + \frac{T-1}{N} \right) - \sum_{a \in \mathbb{A}} \sum_{a' \in \mathbb{A}, a' > a} \frac{V_{a,a'}^2}{N},$$

with

$$\mathbb{V}(\hat{\tau}_{\text{did}}) \leq \sum_{a \in \mathbb{A}} S_a^2 / N_a.$$

Comment (ctd): In our two period example with some units adopting in the second period and the others not at all, and $Y_{i1}(a) = 0$, we have

$$\gamma_1 = \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \quad \gamma_2 = \begin{pmatrix} -1 \\ 1 \end{pmatrix}, \quad \text{and} \quad \gamma_\infty = \begin{pmatrix} 1 \\ -1 \end{pmatrix}.$$

$$S_1^2 = 0,$$

$$S_{,2}^2 = \frac{1}{N-1} \sum_{i=1}^N \left(Y_{i2}(2) - \frac{1}{N} \sum_{j=1}^N Y_{j2}(2) \right)^2,$$

$$S_{\infty}^2 = \frac{1}{N-1} \sum_{i=1}^N \left(Y_{i2}(\infty) - \frac{1}{N} \sum_{j=1}^N Y_{j2}(\infty) \right)^2,$$

$$V_{1,2}^2 = 0, \quad S_{1,\infty}^2 = 0,$$

$$V_{2,\infty}^2 = \frac{1}{N-1} \sum_{i=1}^N \left(\left(Y_{i2}(2) - \frac{1}{N} \sum_{j=1}^N Y_{j2}(2) \right) - \left(Y_{i2}(\infty) - \frac{1}{N} \sum_{j=1}^N Y_{j2}(\infty) \right) \right)^2,$$

so that in this special

$$\begin{aligned} \mathbb{V}(\hat{\tau}_{\text{did}}) &= \frac{1}{N_2(N-1)} \sum_{i=1}^N \left(Y_{i2}(2) - \frac{1}{N} \sum_{j=1}^N Y_{j2}(2) \right)^2 \\ &\quad + \frac{1}{N_{\infty}(N-1)} \sum_{i=1}^N \left(Y_{i2}(\infty) - \frac{1}{N} \sum_{j=1}^N Y_{j2}(\infty) \right)^2 \\ &\quad - \frac{1}{N(N-1)} \sum_{i=1}^N \left(\left(Y_{i2}(2) - \frac{1}{N} \sum_{j=1}^N Y_{j2}(2) \right) - \left(Y_{i2}(\infty) - \frac{1}{N} \sum_{j=1}^N Y_{j2}(\infty) \right) \right)^2, \end{aligned}$$

which agrees with the Neyman variance for a completely randomized experiment. \square

4.4 Estimating the Randomization Variance of the Difference-In-Differences Estimators

In this section we discuss estimating the variance of the DID estimator. In general there is no unbiased estimator for $\mathbb{V}(\hat{\tau}_{\text{did}})$. This is not surprising, because there is no such estimator for the simple difference in means estimator in a completely randomized experiment, and this corresponds to the special case with $T = 1$. However, it turns out that just like in the simple randomized experiment case, there is a conservative variance estimator. In the current case it is based on using unbiased estimators for the terms involving $S_{\gamma_a, a}^2$, and ignoring the terms involving $V_{\gamma_a, a, \gamma_{a'}, a'}^2$. Because the latter are non-negative, and enter with a minus sign, ignoring them leads to an upwardly biased variance estimator. One difference with the simple randomized experiment case is that there is no simple case with constant treatment effects such that the variance estimator is unbiased.

Next, define the estimated variance of this by adoption date:

$$s_a^2 \equiv \frac{1}{N_a - 1} \sum_{i:A_i=a} (Y_i - \bar{Y}_a)^2.$$

Now we can characterize the proposed variance estimator as

$$\hat{\mathbb{V}}_{\text{did}} \equiv \sum_{a \in \mathbb{A}} \frac{s_a^2}{N_a}.$$

Theorem 3. *Suppose Assumption 1 holds. Then*

$$\mathbb{E} \left[\hat{\mathbb{V}}_{\text{did}} \right] \geq \mathbb{V}(\hat{\tau}_{\text{did}}),$$

so that $\hat{\mathbb{V}}_{\text{did}}$ is a conservative variance estimator for $\hat{\tau}_{\text{did}}$.

There are two important issues regarding this variance estimator. The first is its relation to the standard variance estimator for DID estimators. The second is whether one can improve on this variance estimator given that in general it is conservative.

The relevant variance estimators are the Liang-Zeger clustered variance estimator and the clustered bootstrap (Bertrand et al. [2004], Liang and Zeger [1986]). Both have large sample justifications under random sampling from a large population, so they are in general not equal to the variance estimator here. In large samples both the Liang-Zeger and bootstrap variance will be more conservative than $\hat{\mathbb{V}}_{\text{did}}$ because they also take into account variation in the weights $\gamma_{t,a}$. These weights are kept fixed under the randomization scheme, because that keeps fixed the marginal distribution of the adoption dates. In contrast, under the Liang-Zeger calculations and the clustered bootstrap, the fraction of units with a particular adoption date varies, and that introduces additional uncertainty.

The second issue is whether we can improve on the conservative variance estimator $\hat{\mathbb{V}}_{\text{did}}$. In general there is only a limited ability to do so. Note, for example, that in the two period example this variance reduces to the Neyman variance in randomized experiments. In that case we know we can improve on this variance a little bit exploiting heteroskedasticity, e.g., Aronow et al. [2014], but in general those gains are modest.

5 Some Simulations

The goal is to compare the exact variance, and the corresponding estimator in the paper to the two leading alternatives, the Liang-Zeger (stata) clustered standard errors and the clustered bootstrap. We want to confirm settings where the proposed variance estimator differs from the Liang-Zeger clustered variance, and settings where it is the same. We have N units, observed for T time periods. We focus primarily on the case with $T = 3$. The adoption date is randomly assigned, with $\pi_I = (\pi_1, \pi_2, \pi_3, \pi_\infty) = (0, 0.67, 0, 0.33)$, and $\pi_{II} = (\pi_1, \pi_2, \pi_3, \pi_\infty) = (0, 0.5, 0.4, 0.1)$.

We consider two designs for the potential outcome distributions in the population, the $Y_i(a)$ for $a \in \{1, 2, 3, \infty\}$. In design A the potential outcomes, are generated as

$$\begin{pmatrix} Y_{i1}(2) \\ Y_{i1}(3) \\ Y_{i1}(\infty) \\ Y_{i2}(2) \\ Y_{i2}(3) \\ Y_{i2}(\infty) \\ Y_{i3}(2) \\ Y_{i3}(3) \\ Y_{i3}(\infty) \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} 0 \\ 0 \\ 0 \\ 4 \\ 3 \\ 3 \\ 2 \\ 2 \\ 1 \end{pmatrix}, \sigma^2 \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \right).$$

In this design the treatment effect is constant, and depends only on whether the adoption date precedes the potential outcome date, or

$$Y_{it}(a) = 1_{a \leq t} + \varepsilon_{it},$$

where the ε_{it} are correlated over time.

In design B the potential outcomes are generated as

$$\begin{pmatrix} Y_{i1}(2) \\ Y_{i1}(3) \\ Y_{i1}(\infty) \\ Y_{i2}(2) \\ Y_{i2}(3) \\ Y_{i2}(\infty) \\ Y_{i3}(2) \\ Y_{i3}(3) \\ Y_{i3}(\infty) \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} 0 \\ 0 \\ 0 \\ 2 \\ 1 \\ 1 \\ 2 \\ 11 \\ 1 \end{pmatrix}, \sigma^2 \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 10 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \right).$$

Here the treatment effects depend on the treatment having been adopted, but the effect differs by the adoption date.

In design C the potential outcomes are generated with positive correlations between the potential outcomes as

$$\begin{pmatrix} Y_{i1}(2) \\ Y_{i1}(3) \\ Y_{i1}(\infty) \\ Y_{i2}(2) \\ Y_{i2}(3) \\ Y_{i2}(\infty) \\ Y_{i3}(2) \\ Y_{i3}(3) \\ Y_{i3}(\infty) \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} 0 \\ 0 \\ 0 \\ 2 \\ 1 \\ 1 \\ 2 \\ 11 \\ 1 \end{pmatrix}, \sigma^2 \begin{pmatrix} 1 & 0.9 & 0.9 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.9 & 1 & 0.9 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.9 & 0.9 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0.9 & 0.9 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.9 & 1 & 0.9 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.9 & 0.9 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0.9 & 0.9 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0.9 & 1 & 0.9 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0.9 & 0.9 & 1 \end{pmatrix} \right).$$

In design D the potential outcomes are generated with negative correlations between the potential outcomes as

$$\begin{pmatrix} Y_{i1}(2) \\ Y_{i1}(3) \\ Y_{i1}(\infty) \\ Y_{i2}(2) \\ Y_{i2}(3) \\ Y_{i2}(\infty) \\ Y_{i3}(2) \\ Y_{i3}(3) \\ Y_{i3}(\infty) \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} 0 \\ 0 \\ 0 \\ 2 \\ 1 \\ 1 \\ 2 \\ 11 \\ 1 \end{pmatrix}, \sigma^2 \begin{pmatrix} 1 & -0.4 & -0.4 & 0 & 0 & 0 & 0 & 0 & 0 \\ -0.4 & 1 & -0.4 & 0 & 0 & 0 & 0 & 0 & 0 \\ -0.4 & -0.4 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & -0.4 & -0.4 & 0 & 0 & 0 \\ 0 & 0 & 0 & -0.4 & 1 & -0.4 & 0 & 0 & 0 \\ 0 & 0 & 0 & -0.4 & -0.4 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & -0.4 & -0.4 \\ 0 & 0 & 0 & 0 & 0 & 0 & -0.4 & 1 & -0.4 \\ 0 & 0 & 0 & 0 & 0 & 0 & -0.4 & -0.4 & 1 \end{pmatrix} \right).$$

For a particular design, eg $(A, 2)$ draw the four sets of three-component vectors of potential outcomes for each unit (the three components corresponding to the three time periods), one set for each of the values of $a \in \{1, 2, 3, \infty\}$. We keep these sets of potential outcomes fixed across all simulations for a given design. Then for each simulation draw the adoption date according to the distribution for that design, keeping the fraction of units with a particular adoption date fixed.

We want to look at variances and the corresponding confidence intervals based on four methods for estimating the variance for the DID estimator. The confidence intervals are Normal-distribution based, simply equal to the point estimates plus and minus 1.96 times the square root of the variances. We can write $\hat{\tau}_{\text{did}}$ as a regression estimator with NT observations, and $N + T$ regressors. Let with $j = 1, \dots, NT$. For observation j , $T_j \in \{1, \dots, T\}$ denotes the time period the observation is from, and $N_j \in \{1, \dots, N\}$ denotes the unit it corresponds to. Now let $Y_j = Y_{N_j, T_j}$ and $W_j = W_{N_j, T_j}$, so that the regression function can be written as

$$Y_j = \mu + \sum_{n=1}^{N-1} \alpha_n \mathbf{1}_{N_j=n} + \sum_{t=1}^{T-1} \beta_t \mathbf{1}_{T_j=t} + \tau W_j + \varepsilon_j = Y_j = X_j^\top \theta + \varepsilon_j,$$

where $X_j = (1, \mathbf{1}_{N_j=1}, \dots, \mathbf{1}_{N_j=N-1}, \mathbf{1}_{T_j=1}, \dots, \mathbf{1}_{T_j=T-1}, W_j)$, and $\theta = (\mu, \alpha_1, \dots, \alpha_{N-1}, \beta_1, \dots, \beta_{T-1}, \tau)$.

We compare five variances. The first is exact randomization-based variance,

$$\mathbb{V}_{\text{did}} = \mathbb{V}(\hat{\tau}) = \sum_{a \in \mathbb{A}} \frac{S_{\gamma_{a,a}}^2}{N_a} - \sum_{a \in \mathbb{A}} \sum_{a' \in \mathbb{A}, a' > a} \frac{V_{\gamma_{a,a}, \gamma_{a',a'}}^2}{N}.$$

The other four are estimators of the variance.

First, the feasible conservative variance estimator $\widehat{\mathbb{V}}_{\text{did}}$.

Second, the standard Liang-Zeger clustered variance. Start with the representation $Y_j = X_j^\top \theta + \varepsilon_j$. Let $\hat{\varepsilon}_j = Y_j - X_j^\top \hat{\theta}$ be the residual from this regression. Calculate the variance as

$$\widehat{\mathbb{V}}_{\text{LZ}} = \left(\sum_{j=1}^J X_j X_j^\top \right)^{-1} \left(\sum_{n=1}^N \left(\sum_{j:N_j=n} X_j \hat{\varepsilon}_j \right) \left(\sum_{j:N_j=n} X_j \hat{\varepsilon}_j \right)^\top \right) \left(\sum_{j=1}^J X_j X_j^\top \right)^{-1},$$

and get the corresponding variance estimator for $\hat{\tau}_{\text{did}}$.

Third, the clustered bootstrap, $\widehat{\mathbb{V}}_{B1}$. Draw bootstrap samples based on drawing units, with all time periods for each unit drawn. Note that this explicitly changes from bootstrap sample to bootstrap sample the fraction of units with a particular adoption date.

Fourth, a modification of the clustered bootstrap, $\widehat{\mathbb{V}}_{B2}$, where we fix the fraction of units with each value for the adoption date.

In Table 1 we report the results. For each of the five variances we report the average of variance, and the coverage rate for the 95% confidence interval.

We see that the standard Liang-Zeger and the clustered bootstrap ($\widehat{\mathbb{V}}_{B1}$) substantially overestimate the variance in Design B. The fixed adoption date bootstrap ($\widehat{\mathbb{V}}_{B2}$) and the proposed variance estimator ($\widehat{\mathbb{V}}_{\text{did}}$) have the appropriate coverage.

6 Conclusion

We develop a design-based approach to Difference-In-Differences estimation in a setting with staggered adoption. We characterize what the standard DID estimator is estimating under a random adoption date assumption, and what the variance of the standard estimator is. We show that the standard DID estimator is a weighted average of different types of causal effects, for example, the effect of changing from never adopting to adopting in the first period, or

Table 1: : SIMULATIONS

Design	π	N	V_{did}	Cov	\hat{V}_{did}	Cov	\hat{V}_{LZ}	Cov	\hat{V}_{B1}	Cov	\hat{V}_{B2}	Cov
A	I	30	0.144	0.951	0.239	0.979	0.214	0.974	0.232	0.975	0.219	0.973
B	I	30	0.111	0.947	0.187	0.986	0.163	0.978	0.182	0.982	0.172	0.978
C	I	30	0.201	0.953	0.217	0.947	0.181	0.925	0.211	0.942	0.200	0.932
D	I	30	0.064	0.949	0.265	1.000	0.230	0.999	0.257	1.000	0.244	0.999
A	II	30	0.112	0.946	0.165	0.972	0.146	0.966	0.158	0.969	0.142	0.956
B	II	30	0.085	0.947	0.139	0.973	0.268	0.999	0.269	0.999	0.119	0.962
C	II	30	0.184	0.949	0.191	0.939	0.279	0.983	0.285	0.981	0.162	0.920
D	II	30	0.081	0.950	0.164	0.992	0.285	1.000	0.280	0.999	0.142	0.987
A	I	150	0.027	0.953	0.047	0.991	0.045	0.989	0.047	0.989	0.046	0.989
B	I	150	0.022	0.955	0.041	0.994	0.039	0.992	0.041	0.992	0.041	0.992
C	I	150	0.035	0.956	0.038	0.960	0.036	0.956	0.037	0.955	0.037	0.954
D	I	150	0.019	0.950	0.044	0.997	0.044	0.997	0.044	0.996	0.043	0.995
A	II	150	0.020	0.952	0.033	0.989	0.033	0.989	0.033	0.987	0.032	0.987
B	II	150	0.021	0.945	0.036	0.985	0.053	0.997	0.052	0.997	0.035	0.984
C	II	150	0.034	0.952	0.035	0.953	0.051	0.985	0.052	0.983	0.034	0.947
D	II	150	0.016	0.950	0.028	0.990	0.044	0.998	0.044	0.998	0.028	0.987

changing from never adopting to adopting later. In this approach the standard Liang-Zeger and clustered bootstrap variance estimators are unnecessarily conservative, and we propose an improved variance estimator.

APPENDIX

Proof of Lemma 1: Let \mathbf{Y}^p denote the $N \times (T \cdot (T + 1))$ dimensional matrix with all the potential outcomes. Because the units are exchangeable we can write the joint distribution of the potential outcomes and \mathbf{A} as

$$f(\mathbf{Y}^p, \mathbf{A}) = \prod_{i=1}^N f(\mathbf{Y}_i^p, A_i).$$

Now we shall construct a distribution $f(\mathbf{Y}_i^p, A_i)$ that satisfies two conditions. First, A_i is independent of all the potential outcomes and second, the implied distribution for the adoption date and the realized outcome is consistent with the actual distribution. To do so we assume independence of the sets potential outcomes $Y_{i1}(a), \dots, f_{iT}(a)$ for different a , and assume that

$$f(Y_{i1}(a), \dots, f_{iT}(a)) = f(Y_{i1}(a), \dots, f_{iT}(a) | A_i = a) = f(Y_{i1}, \dots, f_{iT} | A_i = a).$$

□

Proof of Lemma 2: By Assumption 2 we have

$$Y_{it}(a) = \mathbf{1}_{a \leq t} Y_{it}(a) + \mathbf{1}_{a > t} Y_{it}(\infty),$$

and by Assumption 3 we have

$$Y_{it}(a) = \mathbf{1}_{a \leq t} Y_{it}(1) + \mathbf{1}_{a > t} Y_{it}(a).$$

Combining the two assumptions implies

$$Y_{it}(a) = \mathbf{1}_{a \leq t} Y_{it}(1) + \mathbf{1}_{a > t} Y_{it}(\infty).$$

Hence

$$\begin{aligned} Y_{it}(a') - Y_{it}(a) &= \mathbf{1}_{a' \leq t} Y_{it}(1) + \mathbf{1}_{a' > t} Y_{it}(\infty) - (\mathbf{1}_{a \leq t} Y_{it}(1) + \mathbf{1}_{a > t} Y_{it}(\infty)) \\ &= \mathbf{1}_{a' \leq t < a} (Y_{it}(1) - Y_{it}(\infty)), \end{aligned}$$

which proves part (i).

For part (ii)

$$\begin{aligned}
\tau_{t,aa'} &= \frac{1}{N} \sum_{i=1}^N (Y_{it}(a') - Y_{it}(a)) \\
&= \frac{1}{N} \sum_{i=1}^N \mathbf{1}_{a' \leq t < t} (Y_{it}(1) - Y_{it}(\infty)) \\
&= \mathbf{1}_{a' \leq t < t} \frac{1}{N} \sum_{i=1}^N (Y_{it}(1) - Y_{it}(\infty)) = \mathbf{1}_{a \leq t < a'} \tau_{t,\infty 1}.
\end{aligned}$$

□

Proof of Lemma 3: Part (i) follows directly from the fact that the exclusion restrictions place restrictions only on potential outcomes that cannot be observed together.

Let us turn to part (ii). By assumption

$$Y_{it}(a) \perp\!\!\!\perp A_i,$$

which as a special case includes

$$Y_{i1}(\infty) \perp\!\!\!\perp A_i.$$

Hence

$$Y_{i1}(\infty) \perp\!\!\!\perp A_i \mid A_i \geq 2$$

which implies

$$Y_{i1} \perp\!\!\!\perp A_i \mid A_i \geq 2$$

and thus

$$Y_{i1} \perp\!\!\!\perp A_i \mid A_i \in \{2, \infty\},$$

which is a testable restriction. □

Proof of Lemma 4: By Assumptions 2 and 3 we have

$$Y_{it}(a) - Y_{it}(\infty) = \mathbf{1}_{a \leq t} \left(Y_{it}(1) - Y_{it}(\infty) \right).$$

By Assumptions 4 and 5, $Y_{it}(1) - Y_{it}(\infty) = \tau_{1\infty}$, so that

$$Y_{it}(a) - Y_{it}(\infty) = \mathbf{1}_{a \leq t} \tau_{1\infty}.$$

□

Proof of Lemma 5: Using the definition for $g(t, a)$, we can write $\hat{\tau}_{\text{did}}$ as

$$\begin{aligned} \hat{\tau}_{\text{did}} &= \frac{\sum_{i,t} \dot{W}_{it} Y_{it}}{\sum_{i,t} \dot{W}_{it}^2} = \frac{\sum_{t \in \mathbb{T}} \sum_{a \in \mathbb{A}} \sum_{i: A_i = a} \dot{W}_{it} Y_{it}}{N \sum_{t \in \mathbb{T}} \sum_{a \in \mathbb{A}} \pi_a g(a, t)^2} = \frac{\sum_{t \in \mathbb{T}} \sum_{a \in \mathbb{A}} \sum_{i: A_i = a} g(a, t) Y_{it}}{N \sum_{t \in \mathbb{T}} \sum_{a \in \mathbb{A}} \pi_a g(a, t)^2} \\ &= \frac{\sum_{t \in \mathbb{T}} \sum_{a \in \mathbb{A}} \sum_{i: A_i = a} g(a, t) N_a \bar{Y}_{t,a}}{N \sum_{t \in \mathbb{T}} \sum_{a \in \mathbb{A}} \pi_a g(a, t)^2} \\ &= \frac{\sum_{t \in \mathbb{T}} \sum_{a \in \mathbb{A}} g(a, t) \pi_a \bar{Y}_{t,a}}{\sum_{t \in \mathbb{T}} \sum_{a \in \mathbb{A}} \pi_a g(a, t)^2} = \sum_{t,a} \gamma_{t,a} \bar{Y}_{t,a}, \end{aligned}$$

where $\gamma_{t,a}$ is as given in (4.2). □

Proof of Theorem 1: First consider part (i). We will show that

$$\mathbb{E}[\bar{Y}_{ta}] = \bar{Y}_t(a),$$

which in turn implies the result in (i). We can write

$$\mathbb{E}[\bar{Y}_{ta}] = \mathbb{E} \left[\frac{1}{N_a} \sum_{i=1}^N \mathbf{1}_{A_i = a} Y_{it} \right] = \mathbb{E} \left[\frac{1}{N_a} \sum_{i=1}^N \mathbf{1}_{A_i = a} Y_{it}(a) \right].$$

By Assumption 1 this is equal to

$$\frac{1}{N_a} \sum_{i=1}^N \mathbb{E}[\mathbf{1}_{A_i = a} Y_{it}(a)] = \frac{1}{N_a} \sum_{i=1}^N \frac{N_a}{N} Y_{it}(a) = \frac{1}{N} \sum_{i=1}^N Y_{it}(a) = \bar{Y}_t(a),$$

which is the desired result.

Next consider part (ii). By Lemma 5,

$$\hat{\tau}_{\text{did}} = \sum_{t \in \mathbb{T}} \gamma_{t,+} \hat{\tau}_{t,\infty 1} + \sum_{t \in \mathbb{T}} \sum_{a > t} \gamma_{t,a} \hat{\tau}_{t,\infty a} - \sum_{t \in \mathbb{T}} \sum_{a \leq t} \gamma_{t,a} \hat{\tau}_{t,a1},$$

so that

$$\mathbb{E}[\hat{\tau}_{\text{did}}] = \mathbb{E} \left[\sum_{t \in \mathbb{T}} \gamma_{t,+} \hat{\tau}_{t,\infty 1} + \sum_{t \in \mathbb{T}} \sum_{a > t} \gamma_{t,a} \hat{\tau}_{t,\infty a} - \sum_{t \in \mathbb{T}} \sum_{a \leq t} \gamma_{t,a} \hat{\tau}_{t,a1} \right],$$

which by Assumption 1 is equal to

$$\sum_{t \in \mathbb{T}} \gamma_{t,+} \mathbb{E}[\hat{\tau}_{t,\infty 1}] + \sum_{t \in \mathbb{T}} \sum_{a > t} \gamma_{t,a} \mathbb{E}[\hat{\tau}_{t,\infty a}] - \sum_{t \in \mathbb{T}} \sum_{a \leq t} \gamma_{t,a} \mathbb{E}[\hat{\tau}_{t,a1}].$$

This in turn, by part (i), is equal to

$$\sum_{t \in \mathbb{T}} \gamma_{t,+} \tau_{t,\infty 1} + \sum_{t \in \mathbb{T}} \sum_{a > t} \gamma_{t,a} \tau_{t,\infty a} - \sum_{t \in \mathbb{T}} \sum_{a \leq t} \gamma_{t,a} \tau_{t,a1},$$

which finishes the proof of part (ii).

Next consider part (iii). If Assumption 2 holds, then for all $a > t$, $\tau_{t,\infty a} = 0$, so that

$$\begin{aligned} \mathbb{E}[\hat{\tau}_{\text{did}}] &= \sum_{t \in \mathbb{T}} \gamma_{t,+} \tau_{t,\infty 1} - \sum_{t \in \mathbb{T}} \sum_{a \leq t} \gamma_{t,a} \tau_{t,a1} \\ &= \sum_{t \in \mathbb{T}} \sum_{a \leq t} \gamma_{t,a} \tau_{t,\infty a}. \end{aligned}$$

Next consider part (iv). If also Assumption 3 holds, then also for all $a \leq t$, $\tau_{t,a1} = 0$, so that

$$\begin{aligned} \mathbb{E}[\hat{\tau}_{\text{did}}] &= \sum_{t \in \mathbb{T}} \gamma_{t,+} \tau_{t,\infty 1} + \sum_{t \in \mathbb{T}} \sum_{a > t} \gamma_{t,a} \tau_{t,\infty a} - \sum_{t \in \mathbb{T}} \sum_{a \leq t} \gamma_{t,a} \tau_{t,a1} \\ &= \sum_{t \in \mathbb{T}} \gamma_{t,+} \tau_{t,\infty 1}. \end{aligned}$$

Finally, consider part (v). This follows directly from part (iv) in combination with the constant treatment effect assumption (Assumption 5). \square

Next we give a preliminary result.

Lemma A.1. *Suppose that Assumption 1 holds. Then (i) the variance of \bar{Y}_a is*

$$\mathbb{V}(\bar{Y}_a) = \frac{S_a^2}{N_a} \left(1 - \frac{N_a}{N}\right),$$

(ii), the covariance of \bar{Y}_a and $\bar{Y}_{a'}$ is

$$\mathbb{C}(\bar{Y}_a, \bar{Y}_{a'}) = -\frac{1}{2N} (S_a^2 + S_{a'}^2 - S_{aa'}^2) = \frac{1}{2N} (S_a^2 + S_{a'}^2 - V_{aa'}^2),$$

(iii), the variance of the sum of the \bar{Y}_a is

$$\mathbb{V}\left(\sum_{a \in \mathbb{A}} \bar{Y}_a\right) = \sum_{a \in \mathbb{A}} S_a^2 \left(\frac{1}{N_a} + \frac{T-1}{N}\right) - \frac{1}{2N} \sum_{a, a': a \neq a'} V_{aa'}^2,$$

and *(iv),*

$$\mathbb{V}\left(\sum_{a \in \mathbb{A}} \bar{Y}_a\right) \leq \sum_{a \in \mathbb{A}} \frac{S_a^2}{N_a}.$$

Proof of Lemma A.1: Part (i) follows directly from the variance of a sample average with random sampling from a finite population.

Next consider part (ii). Define

$$S_{aa'}^2 = \frac{1}{N-1} \sum_{i=1}^N \{(Y_i(a') - \bar{Y}(a')) - (Y_i(a) - \bar{Y}(a))\}^2.$$

Recall that the variance of the difference between $\bar{Y}_{a'}$ and \bar{Y}_a is

$$\mathbb{V}(\bar{Y}_{a'} - \bar{Y}_a) = \frac{S_a^2}{N_a} + \frac{S_{a'}^2}{N_{a'}} - \frac{S_{aa'}^2}{N},$$

from the results in Neyman [1923/1990], Imbens and Rubin [2015] for completely randomized experiments with a binary treatment. In general it is also true that

$$\mathbb{V}(\bar{Y}_{a'} - \bar{Y}_a) = \mathbb{V}(\bar{Y}_a) + \mathbb{V}(\bar{Y}_{a'}) - 2\mathbb{C}(\bar{Y}_a, \bar{Y}_{a'}).$$

Combining these two characterizations of the variance of the standard estimator for the average treat-

ment effect, it follows that the covariance is equal to

$$\begin{aligned}
\mathbb{C}(\bar{Y}_a, \bar{Y}_{a'}) &= \frac{1}{2} \{ \mathbb{V}(\bar{Y}_a) + \mathbb{V}(\bar{Y}_{a'}) - \mathbb{V}(\bar{Y}_{a'} - \bar{Y}_a) \} \\
&= \frac{1}{2} \left\{ \frac{S_a^2}{N_a} \left(1 - \frac{N_a}{N} \right) + \frac{S_{a'}^2}{N_{a'}} \left(1 - \frac{N_{a'}}{N} \right) - \left\{ \frac{S_a^2}{N_a} + \frac{S_{a'}^2}{N_{a'}} - \frac{S_{aa'}^2}{N} \right\} \right\} \\
&= -\frac{1}{2N} \{ S_a^2 + S_{a'}^2 - S_{aa'}^2 \} \\
&= -\frac{1}{2N} \{ S_a^2 + S_{a'}^2 + V_{aa'}^2 - 2S_a^2 - 2S_{a'}^2 \} \\
&= \frac{1}{2N} \{ S_a^2 + S_{a'}^2 - V_{aa'}^2 \}.
\end{aligned}$$

Next, consider part (iii). Using the result in part (ii),

$$\begin{aligned}
\mathbb{V} \left(\sum_{a \in \mathbb{A}} \bar{Y}_a \right) &= \sum_{a \in \mathbb{A}} \mathbb{V}(\bar{Y}_a) + \sum_{a, a': a \neq a'} \mathbb{C}(\bar{Y}_a, \bar{Y}_{a'}) \\
&= \sum_{a \in \mathbb{A}} \frac{S_a^2}{N_a} \left(1 - \frac{N_a}{N} \right) + \frac{1}{2N} \sum_{a, a': a \neq a'} \{ S_a^2 + S_{a'}^2 - V_{aa'}^2 \} \\
&= \sum_{a \in \mathbb{A}} S_a^2 \left(\frac{1}{N_a} - \frac{1}{N} + \frac{T}{N} \right) - \frac{1}{2N} \sum_{a, a': a \neq a'} V_{aa'}^2 \\
&= \sum_{a \in \mathbb{A}} S_a^2 \left(\frac{1}{N_a} + \frac{T-1}{N} \right) - \frac{1}{2N} \sum_{a, a': a \neq a'} V_{aa'}^2.
\end{aligned}$$

Finally, consider part (iv). The third term, the sum of $V_{aa'}^2$ terms is not directly estimable. Because it has a negative sign, we need to find a lower bound on this sum. A trivial lower bound is zero, but we can do better. We will show that

$$\frac{1}{2N} \sum_{a, a': a \neq a'} V_{aa'}^2 \geq \sum_{a \in \mathbb{A}} S_a^2 \frac{T-1}{N}. \tag{A.1}$$

This in turn implies

$$-\frac{1}{2N} \sum_{a, a': a \neq a'} V_{aa'}^2 \leq -\sum_{a \in \mathbb{A}} S_a^2 \frac{T-1}{N},$$

and thus

$$\begin{aligned}
\mathbb{V}\left(\sum_{a \in \mathbb{A}} \bar{Y}_a\right) &= \sum_{a \in \mathbb{A}} S_a^2 \left(\frac{1}{N_a} + \frac{T-1}{N}\right) - \frac{1}{2N} \sum_{a \neq a'} V_{aa'}^2 \\
&\leq \sum_{a \in \mathbb{A}} S_a^2 \left(\frac{1}{N_a} + \frac{T-1}{N}\right) - \sum_{a \in \mathbb{A}} S_a^2 \frac{T-1}{N} \\
&= \sum_{a \in \mathbb{A}} \frac{S_a^2}{N_a}.
\end{aligned}$$

The last inequality to prove is (A.1). First,

$$\begin{aligned}
V_{aa'}^2 &= \frac{1}{N-1} \sum_{i=1}^N ((Y_i(a') - \bar{Y}(a')) + (Y_i(a) - \bar{Y}(a)))^2 \\
&= \frac{1}{N-1} \sum_{i=1}^N \left\{ (Y_i(a') - \bar{Y}(a'))^2 + (Y_i(a) - \bar{Y}(a))^2 + 2(Y_i(a') - \bar{Y}(a'))(Y_i(a) - \bar{Y}(a))^2 \right\} \\
&= \frac{1}{N} (S_a^2 + S_{a'}^2 + 2\mathbb{C}(Y_i(a), Y_i(a'))).
\end{aligned}$$

Hence

$$\begin{aligned}
\frac{1}{2N} \sum_{a \neq a'} V_{aa'}^2 &= \frac{1}{2N} \sum_{a, a': a \neq a'} \{S_a^2 + S_{a'}^2 + 2\mathbb{C}(Y_i(a), Y_i(a'))\} \\
&= \sum_{a \in \mathbb{A}} S_a^2 \frac{T}{N} + \frac{1}{N} \sum_{a \neq a'} \mathbb{C}(Y_i(a), Y_i(a')).
\end{aligned} \tag{A.2}$$

Next,

$$0 \leq \mathbb{V}\left(\sum_{a \in \mathbb{A}} Y_i(a)\right) = \sum_{a \in \mathbb{A}} \mathbb{V}(Y_i(a)) + \sum_{a, a': a \neq a'} \mathbb{C}(Y_i(a), Y_i(a')).$$

Therefore

$$\sum_{a, a': a \neq a'} \mathbb{C}(Y_i(a), Y_i(a')) \geq - \sum_{a \in \mathbb{A}} \mathbb{V}(Y_i(a)) = - \sum_{a \in \mathbb{A}} S_a^2. \tag{A.3}$$

Combining (A.2) and (A.3) we get the bound

$$\begin{aligned} \frac{1}{2N} \sum_{a,a':a \neq a'} V_{aa'}^2 &= \sum_{a \in \mathbb{A}} S_a^2 \frac{T}{N} + \frac{1}{N} \sum_{a,a':a \neq a'} \mathbb{C}(Y_i(a), Y_i(a')) \\ &\geq \sum_{a \in \mathbb{A}} S_a^2 \frac{T}{N} - \sum_{a \in \mathbb{A}} S_a^2 = \sum_{a \in \mathbb{A}} S_a^2 \frac{T-1}{N}, \end{aligned}$$

which proves (A.1). \square

Proof of Theorem 2: This follows directly from the results in Lemma A.1. \square

Proof of Theorem 3: By Assumption 1 it follows that

$$\mathbb{E} [s_{\gamma_a, a}^2] = S_{\gamma_a, a}^2.$$

This implies that

$$\mathbb{E} [\hat{\mathbb{V}}_{\text{did}}] = \mathbb{E} \left[\sum_{a \in \mathbb{A}} s_{\gamma_a, a}^2 / N_a \right] = \sum_{a \in \mathbb{A}} S_{\gamma_a, a}^2 / N_a \geq \mathbb{V}(\hat{\tau}_{\text{did}}),$$

where the inequality is by Theorem 2. \square

References

- Alberto Abadie. Semiparametric difference-in-differences estimators. The Review of Economic Studies, 72(1):1–19, 2005.
- Alberto Abadie, Alexis Diamond, and Jens Hainmueller. Synthetic control methods for comparative case studies: Estimating the effect of California’s tobacco control program. Journal of the American Statistical Association, 105(490):493–505, 2010.
- Alberto Abadie, Susan Athey, Guido Imbens, and Jeffrey Wooldridge. Clustering as a design problem. 2016.
- Alberto Abadie, Susan Athey, Guido W Imbens, and Jeffrey M Wooldridge. Sampling-based vs. design-based uncertainty in regression analysis. arXiv preprint arXiv:1706.01778, 2017.
- Sarah Abraham and Liyang Sun. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. 2018.
- Joshua Angrist and Alan Krueger. Empirical strategies in labor economics. Handbook of Labor Economics, 3, 2000.
- Joshua Angrist and Steve Pischke. Mostly Harmless Econometrics: An Empiricists’ Companion. Princeton University Press, 2008.
- Manuel Arellano. Computing robust standard errors for within group estimators. Oxford bulletin of Economics and Statistics, 49(4):431–434, 1987.
- Manuel Arellano. Panel data econometrics. Oxford university press, 2003.
- P. Aronow, D. Green, and D. Lee. Sharp bounds on the variance in randomized experiments. Annals of Statistics, 42(3):850–871, 2014.
- Peter M. Aronow and Cyrus Samii. Does regression produce representative estimates of causal effects? American Journal of Political Science, 60(1):250–267, 2016.
- Susan Athey and Guido Imbens. Identification and inference in nonlinear difference-in-differences models. Econometrica, 74(2):431–497, 2006.

- Susan Athey and Scott Stern. An empirical framework for testing theories about complementarity in organizational design. Technical report, National Bureau of Economic Research, 1998.
- Marianne Bertrand, Esther Duflo, and Sendhil Mullainathan. How much should we trust differences-in-differences estimates? The Quarterly Journal of Economics, 119(1):249–275, 2004.
- Kirill Borusyak and Xavier Jaravel. Revisiting event study designs. 2016.
- Brantly Callaway and Pedro HC Sant’Anna. Difference-in-differences with multiple time periods and an application on the minimum wage and employment. arXiv preprint arXiv:1803.09015, 2018.
- David Card. The impact of the mariel boatlift on the miami labor market. Industrial and Labor Relation, 43(2):245–257, 1990.
- David Card and Alan Krueger. Minimum wages and employment: Case study of the fast-food industry in new jersey and pennsylvania. American Economic Review, 84(4):772–793, 1994.
- Timothy G Conley and Christopher R Taber. Inference with difference in differences with a small number of policy changes. The Review of Economics and Statistics, 93(1):113–125, 2011.
- Clément de Chaisemartin and Xavier D’Haultfœuille. Fuzzy differences-in-differences. The Review of Economic Studies, 85(2):999–1028, 2017.
- Clément de Chaisemartin and Xavier D’Haultfœuille. Two-way fixed effects estimators with heterogeneous treatment effects. 2018.
- Stephen G Donald and Kevin Lang. Inference with difference-in-differences and other panel data. The review of Economics and Statistics, 89(2):221–233, 2007.
- Simon Freyaldenhoven, Christian Hansen, and Jesse Shapiro. Pre-event trends in the panel event-study design. Technical report, Brown University Working Paper, 2018.
- Andrew Goodman-Bacon. Difference-in-differences with variation in treatment timing. Technical report, Working Paper, 2017.
- Sukjin Han. Identification in nonparametric models for dynamic treatment effects. 2018.
- Peter Hull. Estimating treatment effects in mover designs. arXiv preprint arXiv:1804.06721, 2018.

- Kosuke Imai and In Song Kim. When Should We Use Linear Fixed Effects Regression Models for Causal Inference with Longitudinal Data? PhD thesis, Working paper, Princeton University, Princeton, NJ, 2016.
- Guido W Imbens and Donald B Rubin. Causal Inference in Statistics, Social, and Biomedical Sciences. Cambridge University Press, 2015.
- Kung-Yee Liang and Scott L Zeger. Longitudinal data analysis using generalized linear models. Biometrika, 73(1):13–22, 1986.
- Bruce D Meyer, W Kip Viscusi, and David L Durbin. Workers’ compensation and injury duration: evidence from a natural experiment. The American Economic Review, pages 322–340, 1995.
- Jerzey Neyman. On the application of probability theory to agricultural experiments. essay on principles. section 9. Statistical Science, 5(4):465–472, 1923/1990.
- Judea Pearl. Causality: Models, Reasoning, and Inference. Cambridge University Press, New York, NY, USA, 2000. ISBN 0-521-77362-8.
- Paul R Rosenbaum. Observational studies. In Observational Studies. Springer, 2002.
- Paul R Rosenbaum. Observation and Experiment: An Introduction to Causal Inference. Harvard University Press, 2017.
- Donald B Rubin. Bayesian inference for causal effects: The role of randomization. The Annals of statistics, pages 34–58, 1978.
- Bbabubhai V Shah, Mary Margaret Holt, and Ralph E Folsom. Inference about regression models from sample survey data. Bulletin of the International Statistical Institute, 47(3):43–57, 1977.
- James H Stock and Mark W Watson. Heteroskedasticity-robust standard errors for fixed effects panel data regression. Econometrica, 76(1):155–174, 2008.
- Jeffrey M Wooldridge. Econometric analysis of cross section and panel data. MIT press, 2010.