

Panel Experiments and Dynamic Causal Effects: A Finite Population Perspective*

Iavor Bojinov[†] Ashesh Rambachan[‡] Neil Shephard[§]

December 14, 2020

Abstract

In panel experiments, we randomly expose multiple units to different interventions and measure their subsequent outcomes, sequentially repeating the procedure numerous times. Using the potential outcomes framework, we define finite population dynamic causal effects that capture the relative effectiveness of alternative treatment paths. For the leading example, known as the lag- p dynamic causal effects, we provide a nonparametric estimator that is unbiased over the randomization distribution. We then derive the finite population limiting distribution of our estimators as either the sample size or the duration of the experiment increases. Our approach provides a new technique for deriving finite population central limit theorems that exploits the underlying Martingale property of unbiased estimators. We further describe two methods for conducting inference on dynamic causal effects: a conservative test for weak null hypotheses of zero average causal effects using the limiting distribution and an exact randomization-based test for sharp null hypotheses. We also derive the finite population probability limit of commonly-used linear fixed effects estimators, showing that these estimators perform poorly in the presence of dynamic causal effects. We conclude with a simulation study and an empirical application in which we reanalyze a lab experiment on cooperation.

Keywords: Panel data, dynamic causal effects, potential outcomes, finite population, nonparametric.

*We thank Robert Minton, Karthik Rajkumar and Jonathan Roth for valuable comments and feedback. We especially thank James Andreoni and Larry Samuelson for kindly sharing their data. Finally, we are grateful to Gary Chamberlain for early conversations about this project. Any remaining errors are our own. Rambachan gratefully acknowledge financial support from the NSF Graduate Research Fellowship under Grant DGE1745303.

[†]Technology and Operations Management Unit, Harvard Business School: ibojinov@hbs.edu

[‡]Department of Economics, Harvard University: asheshr@g.harvard.edu

[§]Department of Economics and Department of Statistics, Harvard University: shephard@fas.harvard.edu

1 Introduction

Panel experiments, where we sequentially assign units to a random intervention, measure their response and repeat the procedure for a fixed period of time, form the basis of causal inference in many areas of biostatistics (e.g., [Murphy et al. \(2001\)](#)), epidemiology (e.g., [Robins \(1986\)](#)), and psychology (e.g., [Lillie et al. \(2011\)](#)). In experimental economics, many authors recognize the benefits of panel-based experiments, for instance [Bellemare, Bissonnette and Kroger \(2014, 2016\)](#) highlighted the potentially large gains in power and [Czibor, Jimenez-Gomez and List \(2019\)](#) emphasized that panel-based experiments may help uncover heterogeneity across units. Despite these benefits, panel experiments are used infrequently in part due to the lack of a formal statistical framework and concerns about how the impact of past treatments on subsequent outcomes may induce biases in conventional estimators ([Charness, Gneezy and Kuhn, 2012](#)). In practice, most authors typically assume away this complication by requiring that the outcomes only depend on contemporaneous treatment, what is often called the “no carryover assumption” (e.g., [Abadie et al. \(2017\)](#), [Athey and Imbens \(2018\)](#), [Athey et al. \(2018\)](#), [Imai and Kim \(2019\)](#), [Arkhangelsky and Imbens \(2019\)](#), [Imai and Kim \(2020\)](#), [de Chaisemartin and D’Haultfoeuille \(2020\)](#)). Even when researchers allow for carryover effects, they almost solely focus on incorporating the uncertainty due to sampling units from some large (potentially infinite) super-population as opposed to the design-based uncertainty, which arises due to the random assignment.¹

In this paper, we tackle these challenges by defining a variety of new panel-based dynamic causal estimands without evoking super-population assumptions nor restrictions on the extent to which treatments can impact subsequent outcomes. Our approach builds on the potential outcomes formulation of causal inference and takes a purely design-based perspective on uncertainty, allowing us to be agnostic to the outcomes model and avoid assumptions about hypothetical super-populations ([Neyman, 1923](#); [Kempthorne, 1955](#); [Cox, 1958](#); [Rubin, 1974](#)).

¹See [Abadie et al. \(2020\)](#) for a discussion of the difference between sampling-based and design-based uncertainty in the cross-sectional setting.

Our main estimands are various averages of lag- p dynamic causal effects, which capture how changes in the assignments affect outcomes after p periods.² For these estimands, we provide nonparametric estimators that are unbiased over the randomization distribution induced by the random design. By exploiting the underlying Martingale property of our unbiased estimators, we derive their finite population asymptotic distribution as either the number of sample periods, experimental units, or both increases. This is a new technique for proving finite population central limit theorems, which may be broadly useful and of independent interest to researchers.

Next, we describe two methods for conducting inference on dynamic causal effects. The first uses the limiting distribution to perform conservative, nonparametric inference on the weak null hypothesis of no average dynamic causal effect. The second provides an exact randomization test for the sharp null hypothesis of no dynamic causal effect at any point in time for all units. We then highlight the broader usefulness of our framework by deriving the finite population probability limit of a variety of standard linear estimation strategies commonly employed on panel data, such as the unit fixed effects estimator and the two-way fixed effects estimator. Our results show that such linear estimators are biased for the dynamic causal effects whenever there is a carryover effect and serial correlation in the assignment mechanism, underscoring the value of our proposed nonparametric estimator.

Finally, we illustrate our theoretical results in an extensive simulation study and apply our framework to reanalyze a panel-based experiment. The simulation study shows the validity of our finite population central limit theorems under a variety of assumptions about the underlying potential outcomes and assignment mechanism. We confirm that conservative tests based on the asymptotic approximation to the randomization distribution of our nonparametric estimator control size well and have good rejection rates against a variety of alternatives. We finish by reanalyzing an experiment conducted in [Andreoni and Samuelson \(2006\)](#), which stud-

²Our dynamic causal effects also provide a useful perspective to the vast literature on observational panel data in econometrics (e.g., the reviews of [Arellano \(2003\)](#), [Arellano and Bonhomme \(2012\)](#)) that tend to avoid formal discussions of causal effects in the modern sense that explicitly appeals to potential outcome functions to define counterfactuals (e.g., [Imbens and Rubin \(2015\)](#), [Hernan and Robins \(2019\)](#)) or directed acyclic graphs to define causal structures (e.g., [Pearl and Mackenzie \(2018\)](#)).

ies cooperative behavior in game theory. The experiment has a natural panel structure. Each participant played a twice-repeated prisoners’ dilemma many times, and the payoff structure of the game was randomly varied across each play. We confirm the authors’ original hypothesis that the payoff structure of the twice repeated prisoners’ dilemma has significant contemporaneous effects on cooperative behavior. Moreover, we provide suggestive evidence of dynamic causal effects in this experiment — the payoff structure of previously played games may affect cooperative behavior in the current game, which may be indicative of learning.

Our design-based framework provides a unified generalization of the finite population literature in cross-sectional causal inference (as reviewed in [Imbens and Rubin \(2015\)](#)) and time series experiments ([Bojinov and Shephard, 2019](#)) to panel experiments. Three crucial contributions differentiate our work from the existing literature. First, we focus on a much richer class of dynamic causal estimands, which answer a broader set of causal questions. Second, we derive two new finite population central limit theorems as the size of the population grows, and as both the duration and population size increase. Third, we compute the bias present in standard linear estimators in the presence of dynamic causal effects and serial correlation in the treatment assignment probabilities.

Our framework is also importantly distinct from earlier work by [Robins \(1986\)](#) and co-authors, that uses treatment paths for causal panel data analysis and solely focuses on providing super-population (or sampling-based) inference methods. In contrast, we avoid super-population arguments entirely and make our inference completely conditional on the potential outcomes. Avoiding super-populations arguments is often attractive in panel data applications. For example, a company only operates in a finite number of markets (e.g., states or cities within the United States) and can only conduct advertising or promotional experiments across markets. Such panel experiments are increasingly common in industry (e.g. [Bojinov, Sait-Jacques and Tingley, 2020](#); [Bojinov, Simchi-Levi and Zhao, 2020](#)).³

³Of course, in other cases, super-population arguments may be entirely natural, given the scientific question at hand. For example, in the mental healthcare digital experiments of [Boruvka et al. \(2018\)](#), it is compelling to use sampling-based arguments as the experimental units are drawn from a larger group of patients for whom we wish to make inference on as, if successful, the technology will be broadly rolled out.

Overview of the paper: In Section 2, we define potential outcome panels, for which we formally define a series of dynamic causal estimands of interest. In Section 3, we provide a nonparametric estimator for our dynamic causal estimands, derive their finite sample properties, and provide finite population central limit theorems that we use for inference. In Section 4, we obtain the finite population probability limits of standard linear estimation methods for potential outcome panels, such as the unit fixed effects estimator and the two-way fixed effects estimator. In Section 5, we present the results from a simulation study, and in Section 6, we use our framework to reanalyze a panel experiment conducted by Andreoni and Samuelson (2006). The appendix collects all non-trivial technical proofs as well as additional simulations and empirical results.

Notation: For an integer $t \geq 1$ and a variable A_t , we write $A_{1:t} := (A_1, \dots, A_t)$. We compactly write index sets as $[N] := \{1, \dots, N\}$ and $[T] := \{1, \dots, T\}$. Finally, for a random variable $A_{i,t}$ observed over $i \in [N]$ and $t \in [T]$, define its average over t as $\bar{A}_i := \frac{1}{T} \sum_{t=1}^T A_{i,t}$, its average over i as $\bar{A}_{\cdot t} := \frac{1}{N} \sum_{i=1}^N A_{i,t}$ and its average over both i and t as $\bar{A} := \frac{1}{NT} \sum_{t=1}^T \sum_{i=1}^N A_{i,t}$.

2 Potential outcome panel and dynamic causal effects

2.1 Assignment panels and potential outcomes

Consider a panel in which N units (e.g., individuals or firms), indexed by $i \in [N]$, are observed over T time periods, indexed by $t \in [T]$. For each unit i and period t , we allocate an assignment $W_{i,t} \in \mathcal{W}$. Throughout this paper, we assume that the assignment is a random variable and that the cardinality of \mathcal{W} is finite, $|\mathcal{W}| < \infty$. Whenever the assignment is binary, that is $\mathcal{W} = \{0, 1\}$, we follow convention and refer to “1” as treatment and “0” as control.

The **assignment path** for unit i is the sequence of assignments that are allocated to unit i over the entire sample period, denoted $W_{i,1:T} = (W_{i,1}, \dots, W_{i,T})' \in \mathcal{W}^T$. The **time- t cross-sectional assignment** describes the assignments allocated to all units at period t , denoted

$W_{1:N,t} = (W_{1,t}, \dots, W_{N,t})' \in \mathcal{W}^N$. The **assignment panel** is the $N \times T$ matrix $W_{1:N,1:T} \in \mathcal{W}^{N \times T}$ that summarizes the assignments given to all units over the entire sample period, where

$$W_{1:N,1:T} = \begin{pmatrix} W_{1:N,1} & \dots & W_{1:N,T} \end{pmatrix} = \begin{pmatrix} W'_{1,1:T} \\ \vdots \\ W'_{N,1:T} \end{pmatrix}.$$

Each column of $W_{1:N,1:T}$ is the cross-sectional assignment for a particular period and each row is the assignment path for a particular unit.

We next define a *potential outcome*, which describes what would be observed for a particular unit at a fixed point in time for every possible assignment panel.

Definition 1. *The **potential outcome** for unit- i at time- t along assignment panel $w_{1:N,1:T} \in \mathcal{W}^{N \times T}$ is written as $Y_{i,t}(w_{1:N,1:T})$.*

In principle, the potential outcome can depend upon the entire assignment panel allowing for arbitrary spillovers across units and time periods. The idea of defining potential outcomes as a function of assignment paths first appears in [Robins \(1986\)](#) and has been further developed in subsequent work such as [Robins \(1994\)](#), [Robins, Greenland and Hu \(1999\)](#), [Murphy et al. \(2001\)](#), [Boruvka et al. \(2018\)](#) and [Blackwell and Glynn \(2018\)](#).⁴ Our work differs from this approach by avoiding super-population arguments. Our estimands and inference procedures are conditioned on the potential outcomes and all uncertainty arises solely from the randomness in the assignment paths. Our work therefore extends the cross-sectional potential outcomes framework to the panel setting.

In econometrics, [Abadie et al. \(2017\)](#) highlight the appeal of this finite-sample, design-based perspective in panel data applications. However, the panel-based potential outcome model developed in that work contains no dynamics as the authors primarily focus on cross-sectional data with an underlying cluster structure.⁵ Our emphasis on the causal effects of

⁴[Hernan and Robins \(2019\)](#) provide a modern, textbook treatment of this literature.

⁵Similarly, [Athey and Imbens \(2018\)](#), [Athey et al. \(2018\)](#) and [Arkhangelsky and Imbens \(2019\)](#) also introduce a potential outcome model for panel data, but assume away carryover effects.

assignment paths are important contrasts to much of the existing literature on panel data analysis in econometrics. For example, [Arellano and Bonhomme \(2012\)](#) reviews non-linear panel data methods, which typically focus on estimating non-linear *contemporaneous* causal effects.⁶ Recently, [Heckman, Humphries and Veramendi \(2016\)](#), [Hull \(2018\)](#), and [Han \(2019\)](#) consider a potential outcome model similar to ours but again rely on super-population arguments, similar to [Robins \(1986\)](#), to perform inference.

2.2 The potential outcome panel model

We now define the potential outcomes panel model by developing assumptions on the potential outcomes that we maintain through the remainder of the paper. Our first assumption restricts the potential outcomes for a unit in a given period not to be affected by future assignments while allowing them to depend on the entire assignment panel up to that period.

Assumption 1. *The potential outcomes are **non-anticipating** if, for all $i \in [N]$, $t \in [T]$, and $w_{1:N,1:T}, \tilde{w}_{1:N,1:T} \in \mathcal{W}^{N \times T}$, $Y_{i,t}(w_{1:N,1:T}) = Y_{i,t}(\tilde{w}_{1:N,1:T})$ whenever $w_{1:N,1:t} = \tilde{w}_{1:N,1:t}$.*

Non-anticipation still allows an arbitrary dependence on past and contemporaneous assignments as well as the assignments of other units. Throughout the paper, we assume that the units are non-interfering, meaning there are no treatment spillovers across units ([Cox, 1958](#)). More formally, we assume that for all $i \in [N]$, $t \in [T]$, and $w_{1:N,1:T}, \tilde{w}_{1:N,1:T} \in \mathcal{W}^{N \times T}$, $Y_{i,t}(w_{1:N,1:T}) = Y_{i,t}(\tilde{w}_{1:N,1:T})$ whenever $w_{i,1:T} = \tilde{w}_{i,1:T}$.

Under Assumption 1 and the non-interference assumption, the potential outcome for unit i at time t only depends on the assignment path for unit i up to time t , allowing us to simplify the notation for the potential outcomes to be $Y_{i,t}(w_{i,1:t})$. Denote the collection of potential outcomes for unit i at time t for all possible assignment paths as $\mathbf{Y}_{i,t} = \{Y_{i,t}(w_{i,1:t}) : w_{i,1:t} \in \mathcal{W}^t\}$. Similarly, $\mathbf{Y}_{1:N,1:T} = \{\mathbf{Y}_{i,t} : i \in [N], t \in [T]\}$ denotes the collection of potential outcomes for all units across all time periods.

⁶For example, [Arellano, Blundell and Bonhomme \(2017\)](#) apply state-of-the-art non-linear panel data methods to estimate the contemporaneous causal effect of earnings fluctuations on household consumption.

Finally, to connect the observed outcomes with the potential outcomes, we assume every unit complies with the assignment.⁷ That is, for an observed assignment panel $w_{1:N,1:T}^{obs}$, the observed outcome panel is given by $y_{1:N,1:T}^{obs} = Y_{1:N,1:T}(w_{1:N,1:T}^{obs})$.

We say that a panel of units, treatments and outcomes in which the outcomes obey Assumption 1, and the units are both non-interfering and perfectly compliant with the treatment is a *potential outcome panel*. For the case where $N = 1$, the potential outcome panel reduces to the definition of a potential outcome time series in Bojinov and Shephard (2019). For $T = 1$, the potential outcome panel reduces to the canonical, cross-sectional potential outcome model (e.g., Holland (1986) and Chapter 1 of Imbens and Rubin (2015)).

2.2.1 Special case: linear potential outcome panel

The dynamic panel data literature mostly focuses on linear models, which are a special case of our general setup.

Definition 2. *A linear potential outcome panel is a potential outcome panel where*

$$Y_{i,t}(w_{i,1:t}) = \beta_{i,t,0}w_{i,t} + \dots + \beta_{i,t,t-1}w_{i,1} + \epsilon_{i,t} \quad \forall t \in [T] \text{ and } i \in [N],$$

and the non-stochastic coefficients $\beta_{i,t,0:t-1}$ and non-stochastic error $\epsilon_{i,t}$ do not depend upon treatments.

This model assumes that the potential outcome for unit i at time t is a linear function of unit i 's assignment path plus a non-stochastic error $\epsilon_{i,t}$ term that does not vary with the assignments. The error term is, therefore, much broader than a typical “error” that is used in many statistics and econometrics papers on panel data.

In some cases, we may wish to place further restrictions on the coefficients in a linear potential outcome panel. We formalize these restrictions with the following definition.

⁷In some applications, this assumption may be unrealistic. For example, in a panel-based clinical trial, we may worry that patients do not properly adhere to the assignment. In such cases, our analysis can be re-interpreted as focusing on dynamic intention-to-treat (ITT) effects.

Definition 3. For a linear potential outcome panel, the coefficients $\beta_{i,t,s}$ are **dynamic causal coefficients**. The dynamic causal coefficients are

- **time-invariant** if $\beta_{i,t,s} = \beta_{i,s}$, for all $t \in [T]$, and $s = 0, \dots, t-1$.
- **homogenous** if $\beta_{i,t,s} = \beta_{t,s}$, for all $i \in [N]$ and $s = 0, \dots, t-1$.

If $\beta_{i,t,s} = \beta_s$ for all $i \in [N], t \in [T], s = 0, \dots, t-1$ then the dynamic causal coefficients are homogeneous and time-invariant.

An example of a linear potential outcome panel is the autoregressive potential outcome panel.

Example 1. An autoregressive potential outcome panel is a potential outcome panel, in which the potential outcomes for any unit $i \in [N]$ obey, for all $t > 1$,

$$Y_{i,t}(w_{i,1:t}) = \phi_{i,t,1}Y_{i,t-1}(w_{i,1:t-1}) + \dots + \phi_{i,t,t-1}Y_{i,1}(w_{i,1}) + \beta_{i,t,0}w_{i,t} + \dots + \beta_{i,t,t-1}w_{i,1} + \epsilon_{i,t}$$

and $Y_{i,1}(w_{i,1}) = \beta_{i,1,0}w_{i,1} + \epsilon_{i,1}$ for $t = 1$, where the non-stochastic coefficients $\phi_{i,t,1:t-1}$, $\beta_{i,t,0:t-1}$ and non-stochastic errors $\epsilon_{i,1:t}$ do not depend on assignments. It is easy to rewrite the potential outcomes to produce the linear potential outcome given in Definition 2.

Example 1 allows for heterogeneity in the parameters across units as well as arbitrary dependence across units and time through $\epsilon_{i,t}$. It is therefore a vast generalization of the non-causal autoregressive econometric panel model associated with, for example, Nerlove (1971), Nickell (1981), Anderson and Hsiao (1982), Arellano and Bond (1991) and the review of Arellano (2003).⁸

⁸Allowing for heterogeneity in panel data models is useful in many empirical applications. For example, in many economic settings, there is extensive heterogeneity across units such as in modeling income processes (Browning, Ejrnaes and Alvarez, 2010) and estimating the dynamic response of consumption to earnings (Arellano, Blundell and Bonhomme, 2017). Time-varying heterogeneity is also an important feature. For example, it is a classic point of emphasis in studying human capital formation and education investments – see Ben-Porath (1967), Griliches (1977) and more recently, Cunha et al. (2006) and Cunha, Heckman and Schennach (2010).

2.3 Assignment mechanism assumptions

We now focus on a special class of assignment mechanism that allow us to define dynamic causal effects. The weakest of these requires that the assignment mechanism at time t only depends on past assignments and observed outcomes, so the assignment in a given period may not depend on future nor unobserved past potential outcomes.

Assumption 2. *The assignments are **sequentially randomized** if, for all $t \in [T]$ and any $w_{1:N,1:t-1} \in \mathcal{W}^{N \times (t-1)}$*

$$\Pr(W_{1:N,t} | W_{1:N,1:t-1} = w_{1:N,1:t-1}, \mathbf{Y}_{1:N,1:T}) = \Pr(W_{1:N,t} | W_{1:N,1:t-1} = w_{1:N,1:t-1}, Y_{1:N,1:t-1}(w_{1:N,1:t-1})).$$

Under Assumption 2, the assignment at any point in time may depend on the entire assignment panel up to the previous period as well as all prior, observed outcomes. This is the panel data analogue of an “unconfounded” or “ignorable” assignment mechanism in the literature on cross-sectional causal inference (as reviewed in Chapter 3 of [Imbens and Rubin \(2015\)](#)). This assumption is common in analyses of dynamic experiments in biostatistics and epidemiology (e.g., see [Robins \(1986\)](#); [Murphy \(2003\)](#)).

Depending on the design of the panel-based experiment, there are two important special cases of sequentially randomized assignments, which impose additional forms of conditional independence across assignments. Let $W_{-i,t} := (W_{1,t}, \dots, W_{i-1,t}, W_{i+1,t}, \dots, W_{N,t})$ and $\mathcal{F}_{1:N,t,T}$ be the filtration generated by $W_{1:N,1:t}$ and $\mathbf{Y}_{1:N,1:T}$.

Assumption 3. *The assignments are **contemporaneously independent** for unit i if, for all $t \in [T]$ and any $w_{1:N,1:t-1} \in \mathcal{W}^{N \times (t-1)}$*

$$\Pr(W_{i,t} | W_{-i,t}, \mathcal{F}_{1:N,t-1,T}) = \Pr(W_{i,t} | W_{1:N,1:t-1} = w_{1:N,1:t-1}, Y_{1:N,1:t-1}(w_{1:N,1:t-1})).$$

Assumption 4. *The assignments are **individualistic** for unit i if, for all $t \in [T]$ and any*

$$w_{1:N,1:t-1} \in \mathcal{W}^{N \times (t-1)}$$

$$\Pr(W_{i,t}|W_{-i,t}, \mathcal{F}_{1:N,t-1,T}) = \Pr(W_{i,t}|W_{i,1:t-1} = w_{i,1:t-1}, Y_{i,1:t-1}(w_{i,1:t-1})).$$

The assignments are contemporaneously independent if, conditional on all past assignments and outcomes, the time- t assignments are selected independently across units. Information from past observed outcomes and assignments across all units may still be used to determine the assignment probabilities at time t . For example, if we notice that, on average, assignment $w \in \mathcal{W}$ is outperforming all other assignments, we may increase the probability of administering assignment w to all subjects at time t . Assuming the assignments are individualistic further imposes that conditional on its own past assignments and outcomes, the assignment for unit i at time t is independent of the past assignments and outcomes of all other units. For example, the Bernoulli randomization mechanism, where $\Pr(W_{i,t}|W_{-i,t}, \mathcal{F}_{1:N,t-1,T}) = \Pr(W_{i,t})$ for all $i \in [N]$ and $t \in [T]$, is individualistic.

Remark 2.1. *Most adaptive experimental strategies (including multi-armed bandits) satisfy our individualistic, sequentially randomized treatment assignment assumptions (Robbins, 1952; Lai and Robbins, 1985). Although the standard approach for developing such strategies focuses on minimizing the utility shortfall across iterations of the algorithm, recently, there has been a growing interest in drawing inference based on the collected data (Hadad et al., 2019). In such adaptive experimental designs, the treatment probabilities are known to the researcher, and therefore our results in subsequent sections can be viewed as providing finite population techniques for drawing causal conclusions from adaptive experiments.*

2.4 Dynamic causal effects

For a potential outcome panel, a *dynamic causal effect* compares the potential outcomes for unit i at time t along different assignment paths. The dynamic causal effect of assignment path

$w_{i,1:t} \in \mathcal{W}^t$ relative to $\tilde{w}_{i,1:t} \in \mathcal{W}^t$ is denoted by

$$\tau_{i,t}(w_{i,1:t}, \tilde{w}_{i,1:t}) := Y_{i,t}(w_{i,1:t}) - Y_{i,t}(\tilde{w}_{i,1:t}), \quad (1)$$

Since we take a design-based perspective, we regard all the potential outcomes $\mathbf{Y}_{i,t}$ as fixed but unknown, or equivalently we condition on the set of all potential outcomes throughout our exposition.

Similar to the cross-sectional potential outcome model, we are more interested in averages of these dynamic causal effects. For example, we could average over units at a fixed time period t to get the average dynamic causal effect at time t , or we could average over time periods for a fixed unit i to get the average dynamic causal effect for unit i . We could also average over both units and time periods. In the rest of this section, we use these dynamic causal effects to build up causal estimands of interest.

2.4.1 Lag- p dynamic causal effects and average dynamic causal effects

Since the number of potential outcomes grows exponentially with the time period t , there is a considerable number of possible causal estimands. To make progress, we restrict our attention to a core class, referred to as the *lag- p dynamic causal effects*.

Definition 4. For $0 \leq p < t$ and $\mathbf{w}, \tilde{\mathbf{w}} \in \mathcal{W}^{p+1}$, the *i, t -th lag- p dynamic causal effect* is

$$\tau_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p) := \tau_{i,t}(\{w_{i,1:t-p-1}^{obs}, \mathbf{w}\}, \{w_{i,1:t-p-1}^{obs}, \tilde{\mathbf{w}}\}).$$

The i, t -th lag- p dynamic causal effect measures the difference between the outcomes from following assignment path \mathbf{w} from period $t - p$ to t compared to the alternative path $\tilde{\mathbf{w}}$, fixing the assignments for unit i to follow the observed path up to time $t - p - 1$.⁹ We use the bold notation for the assignments that define the dynamic causal effects to help differentiate them

⁹Defining dynamic causal effects conditional on the observed past assignments $w_{i,1:t-p-1}^{obs}$ is in line with the common focus on the average causal effect on the treated (e.g. see [Lechner \(2011\)](#)).

from other possible assignment paths. We intentionally drop the subscripts as we will average the lag- p dynamic causal effects across both time and units for a fixed \mathbf{w} and $\tilde{\mathbf{w}}$.

Example 2. *In the case of a linear potential outcome panel (Definition 2), i, t -th lag- p dynamic causal effects are linear functions of the difference between the assignment paths with $\tau_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p) = \sum_{s=0}^p \beta_{i,t,s}(w_{p+1-s} - \tilde{w}_{p+1-s})$, where both $\mathbf{w} = (w_1, \dots, w_{p+1})$ and $\tilde{\mathbf{w}} = (\tilde{w}_1, \dots, \tilde{w}_{p+1})$ are in \mathcal{W}^{p+1} .*

The i, t -th lag- p dynamic causal effects are the building blocks of many other interesting causal estimands. In particular, by restricting the paths \mathbf{w} and $\tilde{\mathbf{w}}$ to share some common features, we obtain the weighted average i, t -th lag- p dynamic causal effect.

Definition 5. *The **weighted average i, t -th lag- p, q dynamic causal effect** is defined as*

$$\tau_{i,t}^\dagger(\mathbf{w}, \tilde{\mathbf{w}}; p, q) := \sum_{\mathbf{v} \in \mathcal{W}^{p-q+1}} a_{\mathbf{v}} \{Y_{i,t}(w_{i,1:t-p-1}^{obs}, \mathbf{w}, \mathbf{v}) - Y_{i,t}(w_{i,1:t-p-1}^{obs}, \tilde{\mathbf{w}}, \mathbf{v})\},$$

where $\mathbf{w}, \tilde{\mathbf{w}} \in \mathcal{W}^q$, for integers p, q satisfying $0 \leq p < t$, $0 < q \leq p + 1$, while $\{a_{\mathbf{v}}\}$ are non-stochastic weights that satisfy $\sum_{\mathbf{v} \in \mathcal{W}^{p-q+1}} a_{\mathbf{v}} = 1$ and $a_{\mathbf{v}} \geq 0$ for all $\mathbf{v} \in \mathcal{W}^{p-q+1}$.

The weighted average i, t -th lag- p, q dynamic causal effect summarizes the effect of switching the assignment path between period $t-p$ and period $t-p+q$ from \mathbf{w} to $\tilde{\mathbf{w}}$ on outcomes at time t by averaging across all possible assignment paths from period $t-p+q+1$ to period t . In this sense, weighted average lag- p causal effects can be interpreted as finite-population causal analogues of the impulse response function, which is a well-known object of interest in econometrics and commonly used to summarize dynamic causal effects in panel analyses.¹⁰ The weights $a_{\mathbf{v}}$ are context specific and may be freely selected by the researcher. For a binary assignment, setting $N = q = 1$ gives us the special case of the weighted average i, t -th lag- p, q dynamic causal effect studied in [Bojinov and Shephard \(2019\)](#), and therefore, we study a much richer class of dynamic

¹⁰[Rambachan and Shephard \(2020\)](#) formally establish that a particular version of the weighted average lag- p causal effect is equivalent to the generalized impulse response function ([Koop, Pesaran and Potter, 1996](#)) in time series settings.

causal effects that summarizes both heterogeneity across units and time periods. As notation, whenever $q = 1$, we drop the q from the notation, simply writing $\tau_{i,t}^\dagger(w, \tilde{w}; p) := \tau_{i,t}^\dagger(w, \tilde{w}; p, 1)$.

Example 3. *Continuing with the linear potential outcome panel, the weighted average i, t -th lag- p, q dynamic causal effects are also linear functions of the difference between the assignment paths. That is, for $\mathbf{w}, \tilde{\mathbf{w}} \in \mathcal{W}^q$, $\tau_{i,t}^\dagger(\mathbf{w}, \tilde{\mathbf{w}}; p, q) = \sum_{s=0}^q \beta_{i,t,p+s}(w_{p+1-s} - \tilde{w}_{p+1-s})$.*

Example 4. *Assume the assignment is binary, $\mathcal{W} = \{0, 1\}$.*

Setting $p = 0$ gives us $\tau_{i,t}(1, 0; 0) = \tau_{i,t}^\dagger(1, 0; 0) = Y_{i,t}(w_{i,1:t-1}^{obs}, 1) + Y_{i,t}(w_{i,1:t-1}^{obs}, 0)$, the unit- i time- t contemporaneous causal effect that measures the instant impact of administering treatment as opposed to control on our outcome of interest.

Now set $p = 1$. Then, $\tau_{i,t}((1, 0), (0, 0); 1) = Y_{i,t}(w_{i,1:t-2}^{obs}, 1, 0) - Y_{i,t}(w_{i,1:t-2}^{obs}, 0, 0)$ measures the impact of giving a treatment as opposed to control at time $t - 1$ on the outcome, while the assignment at time t is, in both cases, zero. If we instead set the assignment at time t to be 1 we get, $\tau_{i,t}((1, 1), (0, 1); 1) = Y_{i,t}(w_{i,1:t-2}^{obs}, 1, 1) - Y_{i,t}(w_{i,1:t-2}^{obs}, 0, 1)$. The uniform weighted average i, t lag-1 dynamic causal effect is then,

$$\begin{aligned} \tau_{i,t}^\dagger(1, 0; 1) &= \frac{1}{2} [\{Y_{i,t}(w_{i,1:t-2}^{obs}, 1, 0) - Y_{i,t}(w_{i,1:t-2}^{obs}, 0, 0)\} + \{Y_{i,t}(w_{i,1:t-2}^{obs}, 1, 1) - Y_{i,t}(w_{i,1:t-2}^{obs}, 0, 1)\}] \\ &= \frac{1}{2} [\tau_{i,t}((1, 0), (0, 0); 1) + \tau_{i,t}((1, 1), (0, 1); 1)]. \end{aligned}$$

In other words, $\tau_{i,t}^\dagger(1, 0; 1)$ measures the average impact of changing the assignment at time $t - 1$ on the outcomes in period t .

Finally, $\tau_{i,t}((1, 1), (0, 0); 1) = Y_{i,t}(w_{i,1:t-2}^{obs}, 1, 1) - Y_{i,t}(w_{i,1:t-2}^{obs}, 0, 0)$ measures the impact of giving two consecutive treatments as opposed to controls on the outcome. The extreme version of this, $\tau_{i,t}((1, \dots, 1), (0, \dots, 0); t)$ is the commonly studied “total” causal effect estimand.

The main estimands of interest in this paper are averages of the dynamic causal effects that summarize how different assignments impact the experimental units.

Definition 6. For $p < T$ and $\mathbf{w}, \tilde{\mathbf{w}} \in \mathcal{W}^{p+1}$, the *time- t lag- p average dynamic causal effect* is

$$\bar{\tau}_t(\mathbf{w}, \tilde{\mathbf{w}}; p) := \frac{1}{N} \sum_{i=1}^N \tau_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p).$$

The *unit- i lag- p average dynamic causal effect* is

$$\bar{\tau}_i(\mathbf{w}, \tilde{\mathbf{w}}; p) := \frac{1}{T-p} \sum_{t=p+1}^T \tau_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p).$$

The *total lag- p average dynamic causal effect* is

$$\bar{\tau}(\mathbf{w}, \tilde{\mathbf{w}}; p) := \frac{1}{N(T-p)} \sum_{t=p+1}^T \sum_{i=1}^N \tau_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p).$$

These estimands extend to the weighted average i, t -th lag- p dynamic causal effect by analogously defining $\bar{\tau}^\dagger(\mathbf{w}, \tilde{\mathbf{w}}; p, q)$, $\bar{\tau}_i^\dagger(\mathbf{w}, \tilde{\mathbf{w}}; p, q)$, and $\bar{\tau}_t^\dagger(\mathbf{w}, \tilde{\mathbf{w}}; p, q)$.

Generally, dynamic causal effects allow us to ask and answer a much richer class of questions than typical cross-sectional experiments or panel experiments that make the no carryover assumption. For example, by computing the time- t lag- p average dynamic causal effects at different points in time, we can understand how the effect varies over time. By computing the unit- i lag- p average dynamic causal effects for different units, we may investigate heterogeneity across units.

Example 5. Assume a linear potential outcome panel, then, for $w, \tilde{w} \in \mathcal{W}$,

$$\begin{aligned} \tau_{i,t}(w, \tilde{w}; p) &= \beta_{i,t,p}(w - \tilde{w}), & \bar{\tau}_i(w, \tilde{w}; p) &= \bar{\beta}_{i,\cdot,p}(w - \tilde{w}), \\ \bar{\tau}_t(w, \tilde{w}; p) &= \bar{\beta}_{\cdot,t,p}(w - \tilde{w}), & \bar{\tau}(w, \tilde{w}; p) &= \bar{\beta}_p(w - \tilde{w}). \end{aligned}$$

3 Nonparametric estimation and inference

We now propose a nonparametric [Horvitz and Thompson \(1952\)](#) type estimator of the i, t -th lag- p dynamic causal effects and derive its properties. We then develop two techniques for performing inference. First, if the assignment mechanism is individualistic ([Assumption 4](#)) and further satisfies an additional assumption of probabilistic assignment that is defined below, we show that our proposed estimator is unbiased for the i, t -th lag- p dynamic causal effects and its related averages over the assignment mechanism. Additionally, as the population grows large, an appropriately scaled and centered version of our estimator for the average lag- p dynamic causal effects becomes approximately normally distributed. These limiting results are conditional on the potential outcomes, and so they are finite population central limit theorems in the spirit of [Freedman \(2008\)](#), and [Li and Ding \(2017\)](#). Using these finite population central limit theorems, we develop conservative tests for weak null hypotheses of no average dynamic causal effects. Second, if the assignment mechanism is probabilistic and sequentially randomized ([Assumption 1](#)), we show how to test the sharp null hypotheses of no treatment at any point in time using an exact randomization test.

3.1 Setup: adapted propensity score and probabilistic assignment

To make our notation more compact, we define the *adapted propensity score*, which captures the conditional probability of a given assignment path. For each i, t , and any $\mathbf{w} = (w_1, \dots, w_{p+1}) \in \mathcal{W}^{(p+1)}$, the adapted propensity score is

$$p_{i,t-p}(\mathbf{w}) := \Pr(W_{i,t-p:t} = \mathbf{w} | W_{i,1:t-p-1}, Y_{i,1:t}(W_{i,1:t-p-1}, \mathbf{w})), \quad (2)$$

and can be decomposed using individualistic assignment ([Assumption 4](#)) and the prediction decomposition.

Lemma 3.1. *For a potential outcome panel satisfying individualistic assignment ([Assumption](#)*

4) and any $\mathbf{w} \in \mathcal{W}^{(p+1)}$, the adapted propensity score can be factorized as

$$p_{i,t-p}(\mathbf{w}) = \Pr(W_{i,t-p} = w_1 | W_{i,1:t-p-1}, Y_{i,1:t-p-1}(W_{i,1:t-p-1})) \\ \times \prod_{s=1}^p \Pr(W_{i,t-p+s} = w_{s+1} | W_{i,1:t-p-1}, W_{i,t-p:t-p+s-1} = \mathbf{w}_{1:s}, Y_{i,1:t-p+s-1}(W_{i,1:t-p-1}, \mathbf{w}_{1:s})).$$

Proof. Use the prediction decomposition for assignments, given all outcomes,

$$p_{i,t-p}(\mathbf{w}) = \Pr(W_{i,t-p} = w_1 | W_{i,1:t-p-1}, Y_{i,1:t}(W_{i,1:t-p-1}, \mathbf{w})) \\ \times \prod_{s=1}^p \Pr(W_{i,t-p+s} = w_{s+1} | W_{i,1:t-p-1}, W_{i,t-p:t-p+s-1} = \mathbf{w}_{1:s}, Y_{i,1:t}(W_{i,1:t-p-1}, \mathbf{w})).$$

and then simplify using the individualistic assignment of treatments. \square

In panel experiments, even though we define the assignment mechanism, we only observe the outcomes along the realized assignment path $Y_{i,1:t}(w_{i,1:t}^{obs})$, and so it is generally not possible to use Lemma 3.1 to compute $p_{i,t-p}(\mathbf{w})$. We can, however, compute the adapted propensity score along the observed assignment path, $p_{i,t-p}(w_{i,t-p:t}^{obs})$.

We next assume that the assignment mechanism is *probabilistic*. This is a crucial assumption for design based inference, as it provides the only source of randomness as we treat the potential outcomes as unknown but fixed.

Assumption 5 (Probabilistic Assignment). *Consider a potential outcome panel. Assume that, for each $i \in [N]$, $t \in [T]$, there exists $c_{i,t}^L, c_{i,t}^U \in (0, 1)$ such that $c_{i,t}^L < p_{i,t-p}(\mathbf{w}) < c_{i,t}^U$ for all $\mathbf{w} \in \mathcal{W}^{(p+1)}$.*

This is also commonly known as the “overlap” or “common support” assumption in the literature on cross-sectional causal inference. From now on, all expectations, denoted by \mathbb{E} , are computed with respect to the probabilistic assignment mechanism. We write $\mathcal{F}_{i,t-p-1}$ as the filtration generated by $W_{i,1:t-p-1}$ and $\mathcal{F}_{1:N,t-p-1}$ as the filtration generated by $W_{1:N,1:t-p-1}$. For example, $\mathbb{E}[W_{i,t} | \mathcal{F}_{i,t-1}] = \sum_{w \in \mathcal{W}} w p_{i,t}(w)$. Since we treat the potential outcomes as fixed (or,

equivalently, we always condition on all of the potential outcomes), conditioning on $W_{i,1:t-p-1}$ is the same as conditioning on both $W_{i,1:t-p-1}$ and $Y_{i,1:t-p-1}(W_{i,1:t-p-1})$.

Remark 3.1. *When we develop our asymptotic arguments, we assume that there exists some $C^L > 0$ and $C^U < 1$ such that the bounds $c_{i,t}^L, c_{i,t}^U$ in Assumption 5 are bounded from above by C^U and below by C^L , for all $i \in [I]$ and $t \in [T]$. Crucially, C^L and C^U do not vary with N or T .*

3.2 Estimation of the i, t -th lag- p dynamic causal effect

For any $\mathbf{w}, \tilde{\mathbf{w}} \in \mathcal{W}^{(p+1)}$, recall the i, t -th lag- p dynamic causal effect is $\tau_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p) = Y_{i,t}(w_{i,1:t-p-1}^{obs}, \mathbf{w}) - Y_{i,t}(w_{i,1:t-p-1}^{obs}, \tilde{\mathbf{w}})$. Define the nonparametric estimator of $\tau_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p)$:

$$\hat{\tau}_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p) := \left\{ \frac{Y_{i,t}(w_{i,1:t-p-1}^{obs}, \mathbf{w}) \mathbb{1}(w_{i,t-p:t}^{obs} = \mathbf{w})}{p_{i,t-p}(\mathbf{w})} - \frac{Y_{i,t}(w_{i,1:t-p-1}^{obs}, \tilde{\mathbf{w}}) \mathbb{1}(w_{i,t-p:t}^{obs} = \tilde{\mathbf{w}})}{p_{i,t-p}(\tilde{\mathbf{w}})} \right\}, \quad (3)$$

where $\mathbb{1}\{A\}$ is an indicator function taking the value 1 if an event A is true and 0 otherwise. We show below that this estimator is conditionally unbiased for the i, t -th lag- p dynamic causal effect over the assignment mechanism, with a simple conditional covariance.

Crucially, under individualistic assignments (Assumption 4), the estimator simplifies to

$$\hat{\tau}_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p) = \frac{y_{i,t}^{obs} \{ \mathbb{1}(w_{i,t-p:t}^{obs} = \mathbf{w}) - \mathbb{1}(w_{i,t-p:t}^{obs} = \tilde{\mathbf{w}}) \}}{p_{i,t-p}(w_{i,t-p:t}^{obs})}, \quad (4)$$

which is computable as $p_{i,t-p}(w_{i,t-p:t}^{obs})$ is known.

Theorem 3.1. *Consider a potential outcome panel with an individualistic and probabilistic assignment mechanism (Assumptions 4-5). For any $\mathbf{w}, \tilde{\mathbf{w}} \in \mathcal{W}^{(p+1)}$,*

$$\mathbb{E}[\hat{\tau}_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p) | \mathcal{F}_{i,t-p-1}] = \tau_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p), \quad (5)$$

$$\text{Var}(\hat{\tau}_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p) | \mathcal{F}_{i,t-p-1}) = \gamma_{i,t}^2(\mathbf{w}, \tilde{\mathbf{w}}) - \tau_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p)^2, \quad (6)$$

where

$$\gamma_{i,t}^2(\mathbf{w}, \tilde{\mathbf{w}}; p) = \frac{Y_{i,t}(w_{i,1:t-p-1}^{obs}, \mathbf{w})^2}{p_{i,t-p}(\mathbf{w})} + \frac{Y_{i,t}(w_{i,1:t-p-1}^{obs}, \tilde{\mathbf{w}})^2}{p_{i,t-p}(\tilde{\mathbf{w}})}. \quad (7)$$

Further, for distinct $\mathbf{w}, \tilde{\mathbf{w}}, \bar{\mathbf{w}}, \hat{\mathbf{w}} \in \mathcal{W}^{(p+1)}$

$$\text{Cov}(\hat{\tau}_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p), \hat{\tau}_{i,t}(\bar{\mathbf{w}}, \hat{\mathbf{w}}; p) | \mathcal{F}_{i,t-p-1}) = -\tau_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p) \tau_{i,t}(\bar{\mathbf{w}}, \hat{\mathbf{w}}; p).$$

Finally, $\hat{\tau}_{i,t}(\mathbf{w}, \tilde{\mathbf{w}})$ and $\hat{\tau}_{j,t}(\mathbf{w}, \tilde{\mathbf{w}})$ are independent for $i \neq j$ conditional on $\mathcal{F}_{1:N,t-p-1}$,

Proof. Given in the Appendix. □

Theorem 3.1 states that for every i, t , the error in estimating $\tau_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p)$ is a martingale difference sequence (e.g., Hall and Heyde (1980)) through time and conditionally independent over the cross-section. As is common in potential outcome frameworks, the variance of $\hat{\tau}_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p)$ depends upon the potential outcomes under both the treatment and counterfactual (e.g. see Chapter 6 of Imbens and Rubin (2015) and Ding (2017)) and is generally not estimable. However, as shown in Theorem 3.1, the variance is bounded from above by $\gamma_{i,t}^2(\mathbf{w}, \tilde{\mathbf{w}}; p)$, which we can estimate by

$$\hat{\gamma}_{i,t}^2(\mathbf{w}, \tilde{\mathbf{w}}; p) = \frac{(y_{i,t}^{obs})^2 \{ \mathbb{1}(w_{i,t-p:t}^{obs} = \mathbf{w}) + \mathbb{1}(w_{i,t-p:t}^{obs} = \tilde{\mathbf{w}}) \}}{p_{i,t-p}(w_{i,t-p:t}^{obs})^2}. \quad (8)$$

Remark 3.2. Our bound on the variance is different from the typical Neyman variance bound, derived under the assumption of a completely randomized experiment (Imbens and Rubin, 2015, Chapter 5). In completely randomized experiments, there is a negative correlation between any two units' assignments since the total number of units assigned to each treatment is fixed. In our setup, all units' assignments are independent, precluding us from exploiting the negative correlation in deriving a bound. Other researchers have noted this difference; for example, both Aronow and Samii (2017) and Bojinov and Shephard (2019) use a similar bound exploiting the basic inequality $ab \leq \frac{1}{2}(a^2 + b^2)$ for any $a, b \in \mathcal{R}$.

The following lemma establishes that $\hat{\gamma}_{i,t}^2(\mathbf{w}, \tilde{\mathbf{w}}; p)$ is an unbiased estimator of $\gamma_{i,t}^2(\mathbf{w}, \tilde{\mathbf{w}}; p)$.

Lemma 3.2. Under the set up of Theorem 3.1, $\mathbb{E}[\hat{\gamma}_{i,t}^2(\mathbf{w}, \tilde{\mathbf{w}}; p) | \mathcal{F}_{i,t-p-1}] = \gamma_{i,t}^2(\mathbf{w}, \tilde{\mathbf{w}}; p)$.

Remark 3.3. Since the weighted average i, t -th lag- p, q dynamic causal effects (Definition 5) are linear combinations of the i, t -th lag- p dynamic causal effects, we can directly apply Theorem 3.1 and Lemma 3.2. We provide the details for the case when $q = 1$.

For $w, \tilde{w} \in \mathcal{W}$ and $\mathbf{v} \in \mathcal{W}^p$, a feasible nonparametric estimator of $\tau_{i,t}^\dagger(w, \tilde{w}; p)$ is

$$\hat{\tau}_{i,t}^\dagger(w, \tilde{w}; p) = \sum_{\mathbf{v} \in \mathcal{W}^p} a_{\mathbf{v}} \left\{ \frac{Y_{i,t}(w_{i,1:t-p-1}^{obs}, w, \mathbf{v}) \mathbb{1}(w_{i,t-p:t}^{obs} = (w, \mathbf{v}))}{p_{i,t-p}(w, \mathbf{v})} - \frac{Y_{i,t}(w_{i,1:t-p-1}^{obs}, \tilde{w}, \mathbf{v}) \mathbb{1}(w_{i,t-p:t}^{obs} = (\tilde{w}, \mathbf{v}))}{p_{i,t-p}(\tilde{w}, \mathbf{v})} \right\}.$$

Under an individualistic assignment mechanism, we can again simplify this to,

$$\hat{\tau}_{i,t}^\dagger(w, \tilde{w}; p) = \frac{a_{w_{i,t-p+1:t}^{obs}} y_{i,t}^{obs} \{ \mathbb{1}(w_{i,t-p}^{obs} = w) - \mathbb{1}(w_{i,t-p}^{obs} = \tilde{w}) \}}{p_{i,t-p}(w_{i,t-p:t}^{obs})}.$$

Again, this estimator is unbiased, over the randomization distribution, with variance that can be bounded from above. For uniform weights, the rest of the generalizations follow immediately by noticing that we can replace all instances of \mathbf{w} and $\tilde{\mathbf{w}}$ with (\mathbf{w}, \mathbf{v}) and $(\tilde{\mathbf{w}}, \mathbf{v})$.

3.3 Estimation of lag- p average causal effects

The martingale difference properties of the nonparametric estimator means that the cross-sectional and temporally averaged estimators

$$\hat{\tau}_t(\mathbf{w}, \tilde{\mathbf{w}}; p) := \frac{1}{N} \sum_{i=1}^N \hat{\tau}_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p) \quad (9)$$

$$\hat{\tau}_i(\mathbf{w}, \tilde{\mathbf{w}}; p) := \frac{1}{(T-p)} \sum_{t=p+1}^T \hat{\tau}_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p) \quad (10)$$

$$\hat{\tau}(\mathbf{w}, \tilde{\mathbf{w}}; p) := \frac{1}{N(T-p)} \sum_{i=1}^N \sum_{t=p+1}^T \hat{\tau}_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p) \quad (11)$$

are also unbiased for the average causal estimands $\bar{\tau}_t(\mathbf{w}, \tilde{\mathbf{w}}; p)$, $\bar{\tau}_i(\mathbf{w}, \tilde{\mathbf{w}}; p)$ and $\bar{\tau}(\mathbf{w}, \tilde{\mathbf{w}}; p)$, respectively.

Moreover, the martingale difference properties greatly ease the calculation of variances of cross-sectional and temporal averages, and allow us to apply a central limit theorem to appropriately scaled and centered versions of these estimators. In particular, write:

$$\sigma_{\cdot t}^2 := \frac{1}{N} \sum_{i=1}^N \{\gamma_{i,t}^2(\mathbf{w}, \tilde{\mathbf{w}}; p) - \tau_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p)^2\} \quad (12)$$

$$\sigma_i^2 := \frac{1}{(T-p)} \sum_{t=p+1}^T \{\gamma_{i,t}^2(\mathbf{w}, \tilde{\mathbf{w}}; p) - \tau_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p)^2\}, \quad (13)$$

$$\sigma^2 := \frac{1}{N(T-p)} \sum_{i=1}^N \sum_{t=p+1}^T \{\gamma_{i,t}^2(\mathbf{w}, \tilde{\mathbf{w}}; p) - \tau_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p)^2\}. \quad (14)$$

Theorem 3.2. *Consider a potential outcome panel with an individualistic and probabilistic assignment mechanism (Assumptions 4-5). Further assume that the potential outcomes are bounded. Then, for any $\mathbf{w}, \tilde{\mathbf{w}} \in \mathcal{W}^{(p+1)}$,*

$$\begin{aligned} \frac{\sqrt{N}\{\hat{\tau}_{\cdot t}(\mathbf{w}, \tilde{\mathbf{w}}; p) - \bar{\tau}_{\cdot t}(\mathbf{w}, \tilde{\mathbf{w}}; p)\}}{\sigma_{\cdot t}} &\xrightarrow{d} N(0, 1) \quad \text{as } N \rightarrow \infty, \\ \frac{\sqrt{T-p}\{\hat{\tau}_i(\mathbf{w}, \tilde{\mathbf{w}}; p) - \bar{\tau}_i(\mathbf{w}, \tilde{\mathbf{w}}; p)\}}{\sigma_i} &\xrightarrow{d} N(0, 1) \quad \text{as } T \rightarrow \infty, \\ \frac{\sqrt{N(T-p)}\{\hat{\tau}(\mathbf{w}, \tilde{\mathbf{w}}; p) - \bar{\tau}(\mathbf{w}, \tilde{\mathbf{w}}; p)\}}{\sigma} &\xrightarrow{d} N(0, 1) \quad \text{as } NT \rightarrow \infty. \end{aligned}$$

Proof. Given in the Appendix. □

Likewise, for bounded potential outcomes with an individualistic and probabilistic assignment mechanism, the scaled variances are

$$N \times \text{Var}(\hat{\tau}_{\cdot t}(w, \tilde{w}; p) | \mathcal{F}_{1:N, t-p-1}) = \mathbb{E}[\sigma_{\cdot t}^2 | \mathcal{F}_{1:N, t-p-1}], \quad (15)$$

$$(T-p) \times \text{Var}(\hat{\tau}_i(w, \tilde{w}; p) | \mathcal{F}_{i,0}) = \mathbb{E}[\sigma_i^2 | \mathcal{F}_{i,0}], \quad (16)$$

$$N(T-p) \times \text{Var}(\hat{\tau}(w, \tilde{w}; p) | \mathcal{F}_{1:N,0}) = \mathbb{E}[\sigma^2 | \mathcal{F}_{1:N,0}]. \quad (17)$$

Following the same logic as earlier, we can establish unbiased estimators of an upper-bound for the variance.

Lemma 3.3. *Under the set up of Theorem 3.2,*

$$\begin{aligned} \mathbb{E} \left[\left(\frac{1}{N} \sum_{i=1}^N \hat{\gamma}_{i,t}^2(\mathbf{w}, \tilde{\mathbf{w}}; p) \right) \mid \mathcal{F}_{1:N,t-p-1} \right] &= \frac{1}{N} \sum_{i=1}^N \gamma_{i,t}^2(\mathbf{w}, \tilde{\mathbf{w}}; p), \\ \mathbb{E} \left[\left(\frac{1}{(T-p)} \sum_{t=p+1}^T \hat{\gamma}_{i,t}^2(\mathbf{w}, \tilde{\mathbf{w}}; p) - \frac{1}{(T-p)} \sum_{t=p+1}^T \gamma_{i,t}^2(\mathbf{w}, \tilde{\mathbf{w}}; p) \right) \mid \mathcal{F}_{i,0} \right] &= 0, \\ \mathbb{E} \left[\left(\frac{1}{N(T-p)} \sum_{i=1}^N \sum_{t=p+1}^T \hat{\gamma}_{i,t}^2(\mathbf{w}, \tilde{\mathbf{w}}; p) - \frac{1}{N(T-p)} \sum_{i=1}^N \sum_{t=p+1}^T \gamma_{i,t}^2(\mathbf{w}, \tilde{\mathbf{w}}; p) \right) \mid \mathcal{F}_{1:N,0} \right] &= 0. \end{aligned}$$

This establishes feasible unbiased estimators for upper-bounds on σ_t^2 , σ_i^2 and σ^2 .

Notice that whenever we average over time, we can not take $\gamma_{i,t}^2(\mathbf{w}, \tilde{\mathbf{w}}; p)$ outside of the expectation as it is a function of past treatments. Lemma 3.3 shows that we can use these feasible estimators as conservative estimators of the scaled variances of the various lag- p average dynamic causal effects. The results in Theorem 3.2 and Lemma 3.3 naturally extend to the weighted average i, t -th lag- p, q dynamic causal effect from Definition 5 by using the estimator developed in Remark 3.3.

3.4 Confidence intervals and testing for lap- p average causal effects

Combining the estimators in Lemma (3.3) with the central limit theorems in Theorem 3.2, we can carry out conservative inference for $\bar{\tau}_t(\mathbf{w}, \tilde{\mathbf{w}}; p)$, $\bar{\tau}_i(\mathbf{w}, \tilde{\mathbf{w}}; p)$ and $\bar{\tau}(\mathbf{w}, \tilde{\mathbf{w}}; p)$. Such inference techniques can be used to provide conservative confidence intervals or carry out Neyman-type hypothesis testing of weak nulls that the average dynamic causal effects are zero. For example, these may be $H_0 : \bar{\tau}_i(\mathbf{w}, \tilde{\mathbf{w}}; p) = 0$ for $i = 3$ or $H_0 : \bar{\tau}_t(\mathbf{w}, \tilde{\mathbf{w}}; p) = 0$ for $t = 4$. Of course, each of these tests carry different causal interpretations and finding the appropriate null hypothesis is up to the practitioner.

An alternative is the more stringent Fisher-type, sharp nulls. An example of this would be $H_0 : \bar{\tau}_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p) = 0$, for all, $\mathbf{w}, \tilde{\mathbf{w}}, i \in [N]$ and specific $t = 4$. The key feature of the Fisher-type null is that it reveals all the potential outcomes $Y_{i,t}(w_{1:t-p-1}^{obs}, w) = y_{i,t}^{obs}$ for all w and i . Hence we can simulate for each i , the assignment path $W_{i,t-p:t} | W_{i,1:t-p-1}^{obs}, y_{i,1:t-p-1}^{obs}$ and

then compute the corresponding $\hat{\tau}_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p)$. This allows us to simulate the exact distribution of any test statistics under the sharp null and, by comparing it to the observed one, allows us to compute an exact p -value. Tests of these Fisher nulls can be inverted to construct confidence intervals. The validity of this exact randomization test only requires us to be able to simulate from the randomization distribution of the assignments paths. Therefore, such randomization tests may also be conducted if the treatment assignment mechanism is sequentially randomized (Assumption 1).

4 Estimation in a linear potential outcome panel

Much of the existing literature on causal inference from panel data in econometrics focuses on using models that assume the outcome is a linear function of the assignment path. In this section, we explore the properties of such standard methods for estimating causal coefficients in a linear potential outcome panel when there exists dynamic causal effects. We begin by analyzing a panel experiment as a repeated cross-section, estimating a separate linear model with the data in each period. We then consider the canonical unit fixed-effects estimator and two-way fixed effects estimator, highlighting the bias induced by the presence of dynamic causal effects. We consider limits as the number of units N grows large, holding the number of periods T fixed (i.e., focusing on small T , large N panel experiments). The available data at time t are the entire assignment panel $W_{1:N,1:t}$ and the corresponding outcomes $Y_{1:N,1:t}(W_{1:N,1:t})$.

Throughout this section, we adapt notation used in Wooldridge (2005) for analyzing panel fixed effects models. For a generic random variable $A_{i,t}$, we compactly write the within-period transformed variable as $\dot{A}_{i,t} = A_{i,t} - \bar{A}_{.t}$ and the within-unit transformed variable as $\check{A}_{i,t} = A_{i,t} - \bar{A}_i$. The within-unit and within-period transformed variable is $\dot{\check{A}}_{i,t} = (A_{i,t} - \bar{A}) - (\bar{A}_{.t} - \bar{A})$.

4.1 Estimation as a repeated cross-section

Denote $\dot{Y}_{1:N,t} = (\dot{Y}_{1,t}, \dots, \dot{Y}_{N,t})'$, $\dot{W}_{i,1:t} = (W_{i,t} - \bar{W}_{\cdot,t}, W_{i,t-1} - \bar{W}_{\cdot,t-1}, \dots, W_{i,1} - \bar{W}_{\cdot,1})'$ and $\dot{W}_{1:N,t} = (\dot{W}_{1,1:t}, \dots, \dot{W}_{N,1:t})'$. The least squares coefficient in the regression of $\dot{Y}_{1:N,t}$ on $\dot{W}_{1:N,t}$ is $\hat{\beta}_{1:N,t} = (\dot{W}_{1:N,t}' \dot{W}_{1:N,t})^{-1} \dot{W}_{1:N,t}' \dot{Y}_{1:N,t}$. Proposition 4.1 derives the finite population limiting distribution of $\hat{\beta}_{1:N,t}$ as the number of units grows large.

Proposition 4.1. *Assume a potential outcome panel and consider the “control” only path, for $0 \in \mathcal{W}$ let $\tilde{w}_{i,1:t} = \mathbf{0}$. Let $\dot{\mu}_{i,t}$ be the $t \times 1$ vector whose u -th element is $E[\dot{W}_{i,t-(u-1)} | \mathcal{F}_{1:N,0,T}]$ and $\Omega_{i,t}$ be the $t \times t$ matrix whose u, v -th element is $\text{Cov}(\dot{W}_{i,t-(u-1)}, \dot{W}_{i,t-(v-1)} | \mathcal{F}_{1:N,0,T})$. Additionally assume that:*

1. *The potential outcome panel is linear and homogeneous (Definitions 2-3), letting $\beta_t = (\beta_{t,0}, \dots, \beta_{t,t-1})$ be the vector of causal coefficients at time t .*
2. *$W_{i,1:t}$ is an individualistic stochastic assignment path and, over the randomization distribution, $\text{Var}(W_{i,t} | \mathcal{F}_{1:N,0,T}) = \sigma_{W,i,t}^2 < \infty$ for each $i \in [N]$, $t \in [T]$.*
3. *As $N \rightarrow \infty$,*
 - (a) *Non-stochastically, $N^{-1} \sum_{i=1}^N \Omega_{i,t} \rightarrow \Gamma_{2,t}$, where $\Gamma_{2,t}$ is positive definite.*
 - (b) *$N^{-1/2} \sum_{i=1}^N (\dot{W}_{i,1:t} - \dot{\mu}_{i,t}) \dot{Y}_{i,t}(\mathbf{0}) | \mathcal{F}_{1:N,0,T} \xrightarrow{d} N(0, \Gamma_{1,t})$.*
 - (c) *Non-stochastically, $N^{-1} \sum_{i=1}^N \dot{Y}_{i,t}(\mathbf{0}) \dot{\mu}_{i,t} \rightarrow \dot{\delta}_t$.*

Then, over the randomization distribution, as $N \rightarrow \infty$,

$$\sqrt{N}(\hat{\beta}_{1:N,t} - \beta_t - \Gamma_{2,t}^{-1} \dot{\delta}_t) | \mathcal{F}_{1:N,0,T} \xrightarrow{d} N(0, \Gamma_{2,t}^{-1} \Gamma_{1,t} \Gamma_{2,t}^{-1}).$$

Proof. Given in the Appendix. □

Typically we might expect that the asymptotic bias induced by $\Gamma_{2,t}^{-1} \dot{\delta}_t$ to be zero, as the deviations of the counterfactual for unit- i at time- t is unlikely to covary with the path of the

expected assignment, given we are conditioning on the potential outcomes. However, this condition depends on the particular assignment mechanism and needs to be checked.

4.2 Interpreting the unit fixed effects estimator

Our next result characterizes the finite population probability limit of the unit fixed effects estimator, $\hat{\beta}_{UFE} = \sum_{i=1}^N \sum_{t=1}^T \check{Y}_{i,t} \check{W}_{i,t} / \sum_{i=1}^N \sum_{t=1}^T \check{W}_{i,t}^2$, allowing for arbitrary heterogeneity in the causal coefficients across units and time periods.

Proposition 4.2. *Assume a potential outcome panel and consider the “control” only path, for $0 \in \mathcal{W}$ let $\check{w}_{i,1:t} = \mathbf{0}$. Denote $\text{Cov}(\check{W}_{i,t}, \check{W}_{i,s}) = \check{\sigma}_{W,i,t,s}$ and $\check{\mu}_{i,t} = \mathbb{E}[\check{W}_{i,t} | \mathcal{F}_{1:N,0,T}]$. Additionally, assume the potential outcome panel is linear (Definition 2), the assignment mechanism is individualistic and $\text{Var}(\check{W}_{i,t} | \mathcal{F}_{1:N,0,T}) = \check{\sigma}_{W,i,t}^2 < \infty$ for each $i \in [N]$, $t \in [T]$. Further assume that as $N \rightarrow \infty$, the following sequences converge non-stochastically:*

$$\begin{aligned} N^{-1} \sum_{i=1}^N \beta_{i,t,s} \check{\sigma}_{W,i,t,s} &\rightarrow \check{\kappa}_{W,\beta,t,s} \quad \forall t \in [T] \text{ \& } s \leq t, \\ N^{-1} \sum_{i=1}^N \check{\sigma}_{W,i,t}^2 &\rightarrow \check{\sigma}_{W,t}^2 \quad \forall t \in [T], \\ N^{-1} \sum_{i=1}^N \check{Y}_{i,t}(\mathbf{0}) \check{\mu}_{i,t} &\rightarrow \check{\delta}_t \quad \forall t \in [T]. \end{aligned}$$

Then, as $N \rightarrow \infty$,

$$\hat{\beta}_{UFE} \xrightarrow{p} \frac{\sum_{t=1}^T \check{\kappa}_{W,\beta,t,t}}{\sum_{t=1}^T \check{\sigma}_{W,t}^2} + \frac{\sum_{t=1}^T \sum_{s=1}^{t-1} \check{\kappa}_{W,\beta,t,s}}{\sum_{t=1}^T \check{\sigma}_{W,t}^2} + \frac{\sum_{t=1}^T \check{\delta}_t}{\sum_{t=1}^T \check{\sigma}_{W,t}^2}. \quad (18)$$

Proof. Given in the Appendix. □

Proposition 4.2 decomposes the finite population probability limit of the unit fixed effects estimator into three terms. The first term is an average of contemporaneous dynamic causal coefficients, describing how the contemporaneous causal coefficients covary with the within-unit transformed assignments over the assignment mechanism. The second term captures how past

causal coefficients covary with the within-unit transformed treatments and arises due to the presence of dynamic causal effects. The last term is an additional error that arises due to the possible relationship between the demeaned counterfactual $\nu_{i,t}(\mathbf{0})$ and the average, demeaned treatment assignment.

Proposition 4.2 is related to yet crucially different from results in Imai and Kim (2019), which show that the unit fixed effects estimator recover a weighted average of unit-specific contemporaneous causal effects if there are no carryover effects. In contrast, we establish that the unit fixed effects estimator does *not* recover a weighted average of unit-specific contemporaneous causal effects in the presence of carryover effects and persistence in the treatment path assignment mechanism. Additionally, Proposition 4.2 provides an exact decomposition of the finite population limiting bias in the unit fixed effects estimator in terms of the underlying dynamic causal effects.

Example 6. Consider an autoregressive potential outcome panel model (Example 1) with, for all $t > 1$,

$$Y_{i,t}(w_{i,1:t}) = \beta_0 w_{i,t} + \beta_1 w_{i,t-1} + \epsilon_{i,t}$$

and $Y_{i,1}(w_{i,1}) = \beta_0 w_{i,1} + \epsilon_{i,1}$ for $t = 1$. That is, there is no heterogeneity across units or time periods in the causal effects and no persistence. In this simple case, Proposition 4.2 implies

$$\hat{\beta}_{UFE} = \beta_0 + \beta_1 \frac{\sum_{t=2}^T \sigma_{\tilde{W},t,t-1}}{\sum_{t=1}^T \sigma_{\tilde{W},t}^2} + \frac{\sum_{t=1}^T \check{\delta}_t}{\sum_{t=1}^T \sigma_{\tilde{W},t}^2}.$$

In other words, the unit fixed effects estimator converges in probability to the contemporaneous dynamic causal coefficient β_0 plus a bias that depends on two terms. The first component of the bias depends on the lag-1 dynamic causal coefficient and the average covariance between assignments across periods. If there is serial correlation in the assignment mechanism across periods, this term will be non-zero.

4.3 Interpreting the two-way fixed effects estimator

It is increasingly common for researchers to estimate linear models with both unit and time fixed effects in panel data.¹¹ The two-way fixed-effect estimator is $\hat{\beta}_{TWFE} = \sum_{i=1}^N \sum_{t=1}^T \dot{Y}_{i,t} \dot{\tilde{W}}_{i,t} / \sum_{i=1}^N \sum_{t=1}^T \dot{\tilde{W}}_{i,t}^2$. We next derive the finite population probability limit of the two-way fixed effects estimator under the assumption of additive causal effects, allowing for arbitrary heterogeneity in the causal coefficients across units and time periods.

Proposition 4.3. *Assume a potential outcome panel and consider the “control” only path, for $0 \in \mathcal{W}$ let $\tilde{w}_{i,1:t} = \mathbf{0}$. Denote $E(\dot{\tilde{W}}_{i,t} | \mathcal{F}_{1:N,0,T}) = \dot{\mu}_{i,t}$ and $Cov(\dot{\tilde{W}}_{i,t}, \dot{\tilde{W}}_{i,s}) = \dot{\sigma}_{W,i,t,s}$. Additionally, assume that the potential outcome panel is linear (Definition 2), the assignment mechanism is individualistic and $Var(\dot{\tilde{W}}_{i,t} | \mathcal{F}_{1:N,0,T}) = \dot{\sigma}_{W,i,t}^2 < \infty$ for each $i \in [N]$, $t \in [T]$. Further assume that as $N \rightarrow \infty$, the following sequences converge non-stochastically*

$$\begin{aligned} N^{-1} \sum_{i=1}^N \beta_{i,t,s} \dot{\sigma}_{W,i,t,s} &\rightarrow \dot{\kappa}_{W,\beta,t,s} \quad \forall t \in [T] \text{ \& } s \leq t, \\ N^{-1} \sum_{i=1}^N \dot{\sigma}_{W,i,t}^2 &\rightarrow \dot{\sigma}_{W,t}^2 \quad \forall t \in [T], \\ N^{-1} \sum_{i=1}^N \dot{Y}_{i,t}(\mathbf{0}) \dot{\mu}_{i,t} &\rightarrow \dot{\delta}_t \quad \forall t \in [T]. \end{aligned}$$

Then, as $N \rightarrow \infty$,

$$\hat{\beta}_{TWFE} \xrightarrow{p} \frac{\sum_{t=1}^T \dot{\kappa}_{W,\beta,t,t}}{\sum_{t=1}^T \dot{\sigma}_{W,t}^2} + \frac{\sum_{t=1}^T \sum_{s=1}^{t-1} \dot{\kappa}_{W,\beta,t,s}}{\sum_{t=1}^T \dot{\sigma}_{W,t}^2} + \frac{\sum_{t=1}^T \dot{\delta}_t}{\sum_{t=1}^T \dot{\sigma}_{W,t}^2}$$

Proof. Given in the Appendix. □

¹¹This is often referred to as the “static” or “canonical” two-way fixed effects specification (e.g. Boryusak and Jaravel (2017), Allegretto et al. (2017), Goodman-Bacon (2018), Athey and Imbens (2018), Imai and Kim (2020)). A recent active area of research focuses on interpreting the “dynamic” two-way fixed effects, which additionally includes leads and lags of the assignment. See, for example, de Chaisemartin and D’Haultfoeuille (2020), Abraham and Sun (2020). Typically, both literatures focus on the special case where assignment is absorbing (meaning units receive the treatments at some period and forever after). In contrast, we place no restrictions on the stochastic assignment paths in our analysis, highlighting the roles of dynamic causal effects and serial correlation in the assignment mechanism in generating bias in the two-way fixed effects estimator.

Similar to our result for the unit fixed effects estimator, Proposition 4.3 shows that the two-way fixed effects estimand decomposes into three components under additive causal effects, where the interpretation of each component is similar to the unit fixed effects estimator.

5 Simulation Study

We now conduct a simulation study to investigate the finite sample properties of the asymptotic results presented in Section 3. We show that the finite population central limit theorems (Theorem 3.2) hold for a moderate number of treatment periods and experimental units. The proposed conservative tests also have correct size under the weak null of no average dynamic causal effects and reasonable rejection rates against a range of alternatives.

5.1 Simulation design

Throughout the simulation we generate the panel experiment using the autoregressive potential outcome panel model from Example 1,

$$\begin{aligned} Y_{i,t} &= \phi_{i,t,1}Y_{i,t-1}(w_{i,1:t-1}) + \dots, \phi_{i,t,t-1}Y_{i,1}(w_{i,1}) + \beta_{i,t,0}w_{i,t} + \dots + \beta_{i,t,t-1}w_{i,1} + \epsilon_{i,t} \quad \forall t > 1, \\ Y_{i,1}(w_{i,1}) &= \beta_{i,1,0}w_{i,1} + \epsilon_{i,1}, \end{aligned}$$

with $\phi_{i,t,1} \equiv \phi$, $\phi_{i,t,s} \equiv 0$ for $s > 1$, $\beta_{i,t,0} \equiv \beta$ and $\beta_{i,t,s} \equiv 0$ for $s > 0$. We vary the choice ϕ , which governs the persistence of the process, and β , which governs the size of the contemporaneous causal effects. We also vary the probability of treatment $p_{i,t-p}(w) = p(w)$ as well as the distribution of the errors $\epsilon_{i,t}$, which we will either sample from a standard normal or a Cauchy distribution.

In all simulations, we document the performance of our nonparametric estimators over the randomization distribution, meaning that we first generate the potential outcomes $\mathbf{Y}_{1:N,1:T}$ and then, holding these fixed, simulate over different assignment panels $W_{1:N,1:T}$.

5.2 Simulation results

5.2.1 Normal approximations and size control

Figure 1 plots the randomization distribution for $\hat{\tau}_t(1, 0; 0)$ under the null hypothesis of $\beta = 0$ for different combinations of the parameter $\phi \in \{0.25, 0.5, 0.75\}$ and treatment probability $p(w) \in \{0.25, 0.5, 0.75\}$. When the errors $\epsilon_{i,t}$ are normally distributed, the randomization distribution quickly converges to a normal distribution. As expected, when the errors are Cauchy distributed, the number of units must be quite large for the randomization distribution to become approximately normal. There is little difference in the results across the values of ϕ and $p(w)$. Testing based on the normal asymptotic approximation controls size effectively, staying close to the nominal 5% level (the exact rejection rates for the null hypothesis, $H_0 : \bar{\tau}_t(1, 0; 0) = 0$ are reported in Table A1).




Figure 1: Simulated randomization distribution for $\hat{\tau}_t(1, 0; 0)$ under different choices of the parameter ϕ (defined in Example 1) and treatment probability $p(w)$. The rows index the parameter ϕ , which ranges over values $\{0.25, 0.5, 0.75\}$. The columns index the treatment probability $p(w)$, which ranges over values $\{0.25, 0.5, 0.75\}$. Panel (a) plots the simulated randomization distribution with normally distributed errors $\epsilon_{i,t} \sim N(0, 1)$ and $N = 1000$. Panel (b) plots the simulated randomization distribution with Cauchy distribution errors $\epsilon_{i,t} \sim Cauchy$ and $N = 50,000$. Results are computed over 5,000 iterations.

Figure 2 plots the randomization distribution for $\hat{\tau}_i^\dagger(1, 0; 0)$. We see a similar pattern as before—when the errors are normally distributed, the randomization distribution converges quickly to a normal distribution, but it takes longer to do so when the errors are heavy-tailed. The null rejection rates for the hypothesis, $H_0 : \bar{\tau}_i(1, 0; 0) = 0$ are reported in Table A2 and, again, the test controls size well across a wide range of parameters.

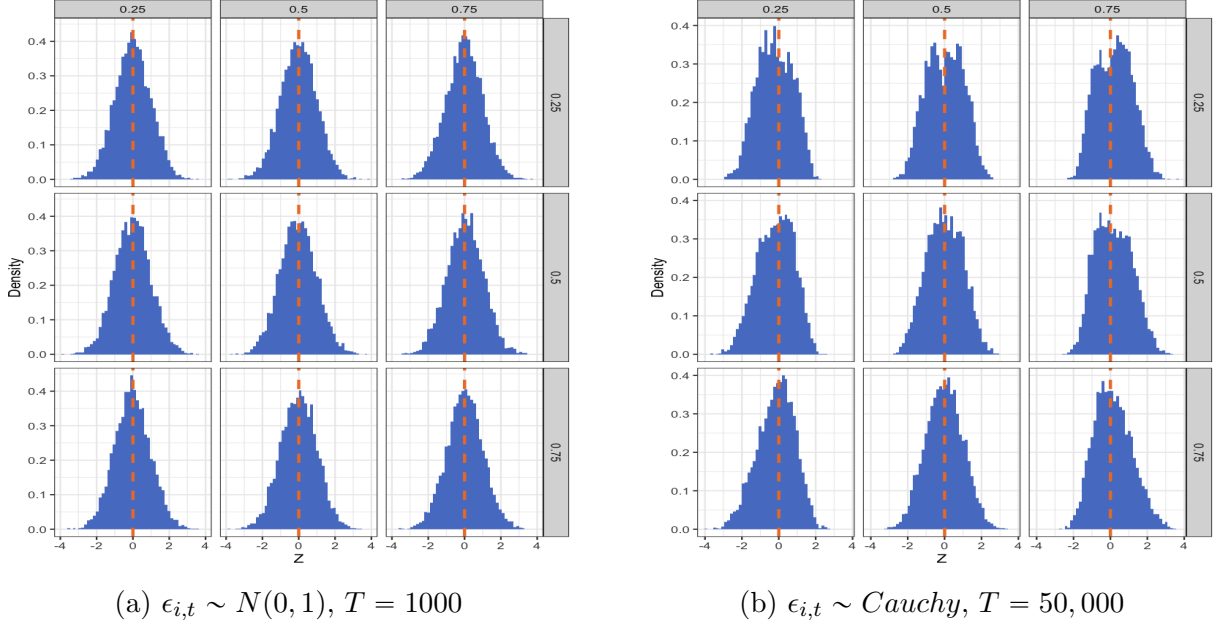




Figure 2: Simulated randomization distribution for $\hat{\tau}_i^\dagger(1, 0; 0)$ under different choices of the parameter ϕ (defined in Example 1) and treatment probability $p(w)$. The rows index the parameter ϕ , which ranges over values $\{0.25, 0.5, 0.75\}$. The columns index the treatment probability $p(w)$, which ranges over values $\{0.25, 0.5, 0.75\}$. Panel (a) plots the simulated randomization distribution with normally distributed errors $\epsilon_{i,t} \sim N(0, 1)$ and $T = 1000$. Panel (b) plots the simulated randomization distribution with Cauchy distribution errors $\epsilon_{i,t} \sim Cauchy$ and $T = 50,000$. Results are computed over 5,000 simulations.

Figure 3 plots the randomization distribution for $\hat{\tau}_i^\dagger(1, 0; 1)$. We present results for the case with $N = 100$, $T = 10$ and $N = 500$, $T = 100$ but note that the results are similar when the roles of N, T are reversed. The null rejection rates for the hypothesis, $H_0 : \bar{\tau}_i^\dagger(1, 0; 1) = 0$ are reported in Table A3.

Figures A1-A3 in the Appendix provides quantile-quantile plots of the simulated randomization distributions to further illustrate the quality of the normal approximations. We also plot the randomization distributions for $\hat{\tau}_t^\dagger(1, 0; 1)$, $\hat{\tau}_i^\dagger(1, 0; 1)$ and $\hat{\tau}^\dagger(1, 0; 1)$ in Figures A4-A6



(a) $\epsilon_{i,t} \sim N(0, 1)$, $N = 100$, $T = 10$



(b) $\epsilon_{i,t} \sim Cauchy$, $N = 500$, $T = 10$

Figure 3: Simulated randomization distribution for $\hat{\tau}^\dagger(1, 0; 1)$ under different choices of the parameter ϕ (defined in Example 1) and treatment probability $p(w)$. The rows index the parameter ϕ , which ranges over values $\{0.25, 0.5, 0.75\}$. The columns index the treatment probability $p(w)$, which ranges over values $\{0.25, 0.5, 0.75\}$. Panel (a) plots the simulated randomization distribution with normally distributed errors $\epsilon_{i,t} \sim N(0, 1)$ and $N = 100, T = 10$. Panel (b) plots the simulated randomization distribution with Cauchy distribution errors $\epsilon_{i,t} \sim Cauchy$ and $N = 500, T = 10$. Results are computed over 5,000 simulations.

respectively.

5.2.2 Rejection rate

Focusing on simulations with normally distributed errors, we next investigate the rejection rate of statistical tests based on the normal asymptotic approximations. To do so, we generate potential outcomes $\mathbf{Y}_{1:N,1:T}$ under different values of β , which governs the magnitude of the contemporaneous causal effect. As we vary $\beta = \{-1, -0.9, \dots, 0.9, 1\}$, we also vary the parameter $\phi \in \{0.25, 0.5, 0.75\}$ and probability of treatment $p(w) \in \{0.25, 0.5, 0.75\}$ to investigate how rejection varies across a range of parameter values. We report the fraction of tests that reject the null hypothesis of zero average dynamic causal effects.

First, we investigate the rejection rate of the statistical test based on the normal asymptotic approximation for $H_0 : \bar{\tau}_t(1, 0; 0) = 0$ and $H_0 : \bar{\tau}_t^\dagger(1, 0; 1) = 0$. Figure 4 plots rejection rate curves against the null hypotheses as the parameter β varies for different choices of the parameter ϕ and treatment probability $p(w)$. The rejection rate against $H_0 : \bar{\tau}_t(1, 0; 0) = 0$ quickly converges to one as β moves away from zero across a range of simulations. This is encouraging as it indicates that the conservative variance bound still leads to informative tests. Unsurprisingly, when $\phi = 0.25$, the rejection rate against $H_0 : \bar{\tau}_t^\dagger(1, 0; 1) = 0$ is relatively low – lower values of ϕ imply less persistence in the causal effects across periods. When $\phi = 0.75$, there is substantial persistence in the causal effects across periods and we observe that the rejection rate curves look similar. Visually, it appears that much of the variation in rejection rates is driven by variation in the causal effects β . To confirm this, we project the rejection rates onto fixed effects for each possible value of the parameter ϕ , $p(w)$ and β , finding that roughly 70% of the variance in rejection rates across simulations at $p = 0$ and roughly 36% of the variance in rejection rates across simulations at $p = 1$ is driven by variation in the causal effects β .

Next, we investigate the rejection rate of the statistical test based on the normal asymptotic approximation for $H_0 : \bar{\tau}_i^\dagger(1, 0; 0) = 0$ and $H_0 : \bar{\tau}_i^\dagger(1, 0; 1) = 0$, plotting the rejection rates in




Figure 4: Rejection probabilities for a test of the null hypothesis $H_0 : \bar{\tau}_t(1,0;0) = 0$ and $H_0 : \bar{\tau}_t(1,0;1) = 0$ as the parameter β varies under different choices of the parameter ϕ and treatment probability $p(w)$. The rejection rate curve against $H_0 : \bar{\tau}_t(1,0;0) = 0$ is plotted in blue and the rejection rate curve against $H_0 : \bar{\tau}_t(1,0;1) = 0$ is plotted in orange. The rows index the parameter ϕ , which ranges over values $\{0.25, 0.5, 0.75\}$. The columns index the treatment probability $p(w)$, which ranges over values $\{0.25, 0.5, 0.75\}$. The simulations are conducted with normally distributed errors $\epsilon_{i,t} \sim N(0, 1)$ and $N = 1000$. Results are averaged over 5000 simulations.

Figure 5. Once again, we observe that the rejection rate against $H_0 : \bar{\tau}_i(1,0;0) = 0$ quickly converges to one as β moves away from zero across a range of simulations. Moreover, when the persistence of the causal effects is low ($\phi = 0.25$), the rejection rate against $H_0 : \bar{\tau}_i(1,0;1) = 0$ is low. Similarly, we find that much of the variation in rejection rates is driven by variation in the causal effects β – estimating the same projection as before, we find 65% of the variance in rejection rates across simulations at $p = 0$ and roughly 32% of the variance in rejection rates across simulations at $p = 1$ are driven by variation in the causal effects β .

Finally, we investigate the rejection rate of the statistical test based on the normal asymptotic approximation for $H_0 : \bar{\tau}^\dagger(1,0;0) = 0$ and $H_0 : \bar{\tau}^\dagger(1,0;1) = 0$. Figure 6 plots rejection rate curves against the null hypotheses as the parameter β varies for different choices of the parameter ϕ and treatment probability $p(w)$. The qualitative patterns are similar as before.¹²

¹²Estimating the same projection of rejection rates onto fixed effects for each possible value of $\phi, p(w)$ and β , we find 68% of the variance in rejection rates across simulations at $p = 0$ and roughly 32% of the variance in rejection rates across simulations at $p = 1$ are driven by variation in the causal effects β .

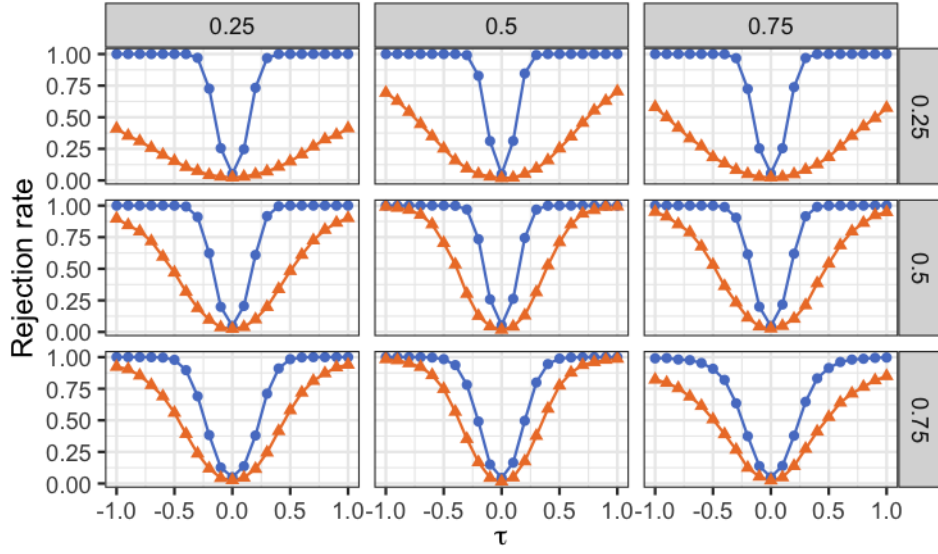


Figure 5: Rejection probabilities for a test of the null hypothesis $H_0 : \bar{\tau}_i^\dagger(1, 0; 0) = 0$ and $H_0 : \bar{\tau}_i^\dagger(1, 0; 1) = 0$ as the parameter β varies under different choices of the parameter ϕ and treatment probability $p(w)$. The rejection rate curve against $H_0 : \bar{\tau}_i^\dagger(1, 0; 0) = 0$ is plotted in blue and the rejection rate curve against $H_0 : \bar{\tau}_i^\dagger(1, 0; 1) = 0$ is plotted in orange. The rows index the parameter ϕ , which ranges over values $\{0.25, 0.5, 0.75\}$. The columns index the treatment probability $p(w)$, which ranges over values $\{0.25, 0.5, 0.75\}$. The simulations are conducted with normally distributed errors $\epsilon_{i,t} \sim N(0, 1)$ and $T = 1000$. Results are averaged over 5000 simulations.




Figure 6: Rejection probabilities for a test of the null hypothesis $H_0 : \bar{\tau}^\dagger(1,0;0) = 0$ and $H_0 : \bar{\tau}^\dagger(1,0;1) = 0$ as the parameter β varies under different choices of the parameter ϕ and treatment probability $p(w)$. The rejection rate curve against $H_0 : \bar{\tau}^\dagger(1,0;0) = 0$ is plotted in blue and the rejection rate curve against $H_0 : \bar{\tau}^\dagger(1,0;1) = 0$ is plotted in orange. The rows index the parameter ϕ , which ranges over values $\{0.25, 0.5, 0.75\}$. The columns index the treatment probability $p(w)$, which ranges over values $\{0.25, 0.5, 0.75\}$. The simulations are conducted with normally distributed errors $\epsilon_{i,t} \sim N(0, 1)$ and $N = 100$, $T = 10$. Results are averaged over 5000 simulations.

6 Empirical application in experimental economics

We now apply our methods to reanalyze an experiment from [Andreoni and Samuelson \(2006\)](#) that tests a game-theoretic model of “rational cooperation” in a lab environment. Specifically, [Andreoni and Samuelson \(2006\)](#) studied how variations in the payoff structure of a two-player, twice-played prisoners’ dilemma affect the choices of players. The payoffs of the games were determined by two parameters $x_1, x_2 \geq 0$ such that $x_1 + x_2 = 10$. In each period, both players simultaneously select either C (cooperate) or D (defect) and subsequently received the payoffs associated with these choices. Table 1 summarizes the exact payoff structure; for example, if the players select (C, C) in period one, they receive $(3x_1, 3x_2)$, respectively. The game had two stages to allow the authors to estimate how changing the relative payoffs between period one and period two impacts the players’ behavior. Let $\lambda = \frac{x_2}{x_1 + x_2} \in [0, 1]$ govern the relative payoffs between the two periods of the prisoners’ dilemma; when $\lambda = 0$, all payoffs occurred in period

one and when $\lambda = 1$, all payoffs occurred in period two. The authors predicted that when λ is large, players will cooperate more often in period one compared to when λ is small.

	C	D		C	D
C	$(3x_1, 3x_1)$	$(0, 4x_1)$	C	$(3x_2, 3x_2)$	$(0, 4x_2)$
D	$(4x_1, 0)$	(x_1, x_1)	D	$(4x_2, 0)$	(x_2, x_2)
Period one			Period two		

Table 1: Stage games from twice-played prisoners’ dilemma in the experiment conducted by [Andreoni and Samuelson \(2006\)](#), where the parameters satisfy $x_1, x_2 \geq 0$, $x_1 + x_2 = 10$ and $\lambda = \frac{x_1}{x_1 + x_2}$. The choice C denotes “cooperate” and the choice D “defect.”

To investigate this hypothesis, [Andreoni and Samuelson \(2006\)](#) conducted a panel-based experiment. In each session of the experiment, 22 subjects were recruited to play 20 rounds of the twice-played prisoners’ dilemma in Table 1. In each round, participants were randomly matched into pairs, and each pair was then randomly assigned a value λ from the set $\{0, 0.1, \dots, 0.9, 1\}$ with equal probability. The authors conducted the experiment over five sessions for a total sample of 110 participants and 110 plays of the stage game. Since each play of the stage game involves two participants, we observe the 220 choices total.

The [Andreoni and Samuelson \(2006\)](#) experiment is a natural application for the methods we developed in this paper for three reasons. First, since each subject plays the twice-played prisoners’ dilemma many times under several randomly assigned payoff structures, the experiment has a panel structure. Second, the sequential nature of the games leaves open the possibility that past assignments impact future actions as participants learn about the structure of the game over time; in other words, there may exist dynamic causal effects that could bias standard methods for estimating causal effects from panel experiments. Finally, the authors analyzed the experimental data using regression models with unit-level fixed effects, which we showed may be biased in the presence of dynamic causal effects in Section 4.

In our notation, the outcome of interest Y is an indicator that equals one whenever the participant cooperated in period one of the stage game, $N = 110$, and $T = 20$. The assignment W is binary and equals one whenever the assigned value λ is greater than 0.6, meaning that the

	Counts		Mean
	0	1	
Observed treatment, $W_{i,t}$	1136	1064	0.484
Observed outcome, $Y_{i,t}$	521	1679	0.763

Table 2: Summary statistics for the experiment in [Andreoni and Samuelson \(2006\)](#). The treatment $W_{i,t}$ equals one when the assigned value of λ is larger than 0.6. The outcome $Y_{i,t}$ equals one whenever the participant cooperates in period one of the twice-repeated prisoners’ dilemma. There are 110 participants and 1110 plays of the stage game in the experiment. Since each play of the stage game involves two participants, we observe 2220 choices total.

payoffs are more concentrated in period two than period one of the stage game. For a given pair of subjects, the assignment mechanism is Bernoulli with probability $p = 5/11$ for treatment and $p = 6/11$ for control. Table 2 summarizes the observed assignments and observed outcomes in the experiment.

One potential complication that may arise from the subjects playing against each other in the stage game is possible spillovers across units. The impact of such spillovers is, however, unlikely to be substantial as the matches are anonymous, and no players play each other more than once. Following [Andreoni and Samuelson \(2006\)](#), we, therefore, ignore this concern in our analysis. In Appendix C, we report additional results in which the outcome of interest Y is a player’s total payoff in the stage game.

6.1 Analysis of unit and time-specific average dynamic causal effects

We begin by estimating the unit-specific average dynamic causal effects. To illustrate the individual estimates, we focus on two randomly selected units in the experiment and construct estimates of their average i, t -th lag-0 dynamic causal effect, $\tau_{i,t}(1, 0; 0)$ (Definition 5). Figure 7 shows the nonparametric estimates $\hat{\tau}_{i,t}(1, 0; 0)$ for $t \in [T]$, for the two units. The figure also contains the nonparametric estimate of the average unit- i lag-0 dynamic causal effect $\bar{\tau}_i(1, 0; 0) = \frac{1}{T} \sum_{t=1}^T \hat{\tau}_{i,t}(1, 0; 0)$. The result shows that the point estimate of the average unit- i lag-0 dynamic causal effect is positive for both units, suggesting that a larger value of λ in the current game increases the likelihood of cooperation for both units. Since each unit only plays




Figure 7: Estimates of the weighted average i, t -th lag-0 dynamic causal effect (Definition 5) of $W = \mathbb{1}\{\lambda \geq 0.6\}$ on cooperation in period one for two units in the experiment of Andreoni and Samuelson (2006). The solid black line plots the nonparametric estimator $\hat{\tau}_{i,t}(1, 0; 0)$ given in Remark 3.3. The dashed black line plots the running average of the period-specific estimator for each unit; that is, for each $t \in [T]$, $\frac{1}{t} \sum_{s=1}^t \hat{\tau}_{i,s}(1, 0; 0)$. The dashed red line plots the estimated weighted average unit- i lag-0 dynamic causal effect, $\hat{\bar{\tau}}_i(1, 0; 0) = \frac{1}{T} \sum_{t=1}^T \hat{\tau}_{i,t}(1, 0; 0)$.

a total of twenty rounds, the estimated variance of these unit-specific estimators is quite large.

We next estimate period-specific, weighted average dynamic causal effects that pool information across units in order to gain precision. For each time period $t \in [T]$, we construct estimates based on the nonparametric estimator of the weighted average time- t , lag- p dynamic causal effect $\bar{\tau}_t^\dagger(1, 0; p) = \frac{1}{N} \sum_{i=1}^N \tau_{i,t}^\dagger(1, 0; p)$ for $p = 0, 1, 2, 3$. For each value of p , the dashed black line in Figure 8 plots the estimates $\hat{\bar{\tau}}_t^\dagger(1, 0; p)$ and the grey region plots a 95% pointwise conservative confidence band for the period-specific weighted average dynamic causal effects. For each value of p , there appears to be some heterogeneity in the period-specific weighted causal dynamic causal effects across time periods.

To further investigate these dynamic causal effects, the solid blue line in Figure 8 plots the nonparametric estimator the total lag- p weighted average causal effect $\bar{\tau}^\dagger(1, 0; p)$ for $p = 0, 1, 2, 3$, which further pools information across all units and time periods. The dashed blue lines plot the conservative confidence interval for the total lag- p weighted average causal effect.

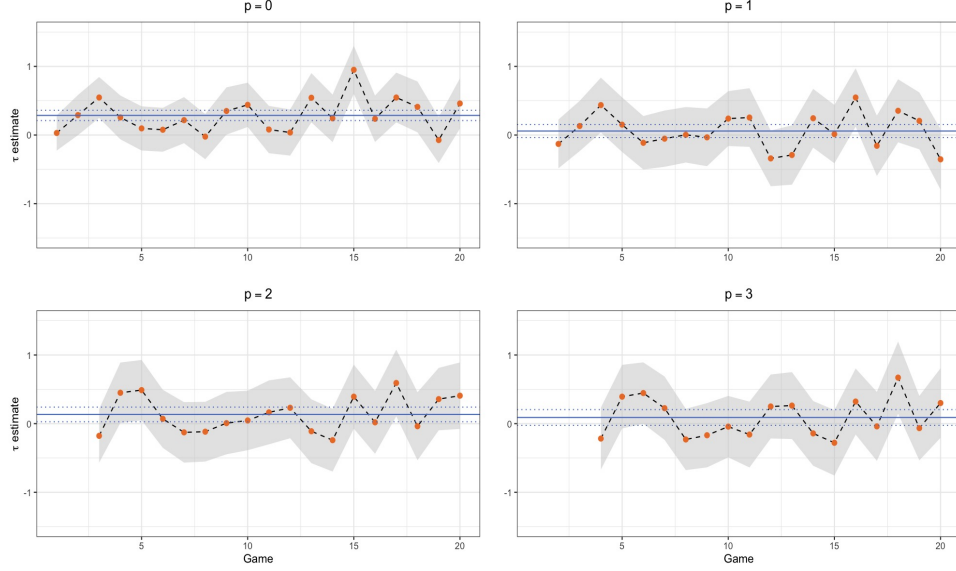


Figure 8: Estimates of the time- t lag- p weighted average dynamic causal effect, $\bar{\tau}_t^\dagger(1, 0; p)$ of $W = \mathbb{1}\{\lambda \geq 0.6\}$ on cooperation in period one based on the experiment of Andreoni and Samuelson (2006) for each time period $t \in [T]$ and $p = 0, 1, 2, 3$. The black dashed line plots the nonparametric estimator of the time- t lag- p weighted average dynamic causal effect, $\hat{\tau}_t^\dagger(1, 0; p)$, for each period $t \in [T]$. The grey region plots the 95% point-wise confidence band for $\bar{\tau}_t^\dagger(1, 0; p)$ based on the conservative estimator of the asymptotic variance of the nonparametric estimator (Theorem 3.2). The solid blue line plots the nonparametric estimator of the total lag- p weighted average dynamic causal effect, $\hat{\tau}^\dagger(1, 0; p)$ and the dashed blue lines plot the 95% confidence interval for $\bar{\tau}^\dagger(1, 0; p)$ based on the conservative estimator of the asymptotic variance of the nonparametric estimator.

As can be seen, the weak null hypothesis that $\bar{\tau}^\dagger(1, 0; 0) = 0$ can be soundly rejected, indicating that the treatment has a positive contemporaneous effect on cooperation in period one of the stage game and confirming the hypothesis of Andreoni and Samuelson (2006). Table 3 summarizes these estimates of the total lag- p weighted average causal effects. Interestingly, the point estimates are positive at for $p = 1, 2, 3$, suggesting there may be dynamic causal effects.

6.2 Exact randomization inference on total average dynamic causal effects

We further unpack these results using randomization tests based on the sharp null of no dynamic causal effects. We construct the randomization distribution for the nonparametric estimator of the total lag- p weighted average dynamic causal effect $\hat{\tau}^\dagger(1, 0; p)$ for $p = 0, 1, 2, 3$ under

	lag- p			
	0	1	2	3
Point estimate, $\hat{\tau}^\dagger(1, 0; p)$	0.285	0.058	0.134	0.089
Conservative p-value	0.000	0.226	0.013	0.126
Randomization p-value	0.000	0.263	0.012	0.114

Table 3: Estimates of the total lag- p weighted average dynamic causal effect for $p = 0, 1, 2, 3$. The conservative p-value reports the p-value associated with testing the weak null hypothesis of no average dynamic causal effects, $H_0 : \bar{\tau}^\dagger(1, 0; p) = 0$, using the conservative estimator of the asymptotic variance of the nonparametric estimator (Theorem 3.2). The randomization p-value reports the p-value associated with randomization test of the sharp null of dynamic causal effects, $H_0 : \tau_{i,t}(w, \tilde{w}; p) = 0$ for all $i \in [N], t \in [T]$. The randomization p-values are constructed based on 10,000 draws.

the sharp null hypothesis of no lag- p dynamic dynamical causal effects for all units and time periods; $H_0 : \tau_{i,t}(w, \tilde{w}; p) = 0$ for all $i \in [N], t \in [T]$. Under this sharp null hypothesis, all relevant potential outcomes are known and we can construct the randomization distribution by redrawing the entire assignment panel according to the known assignment mechanism. When redrawing assignment paths, we do so in a manner that respects the realized pairs of subjects in the experiment, meaning that subjects that are paired in the same round receive the same assignment.

Figure 9 plots the randomization distributions for $p = 0, 1, 2, 3$ along with the point estimate $\hat{\tau}^\dagger(1, 0; p)$ at the realized assignment panel. The randomization distributions appear to be smooth and symmetric around zero. The p-value for the randomization test at $p = 0$ is approximately zero, strongly rejected the sharp null of no contemporaneous dynamic causal effects for all units. This further confirms the hypothesis of Andreoni and Samuelson (2006) that higher values of λ induce more cooperation in the twice-repeated prisoners' dilemma. There is some suggestive evidence of dynamic causal effects. While we are unable to reject the sharp null of no dynamic causal effects at lags $p = 1, 3$, we are able to reject at the 5% level for $p = 2$ (p-value equals 0.012). This suggests that there may have been dynamic causal effects of the treatment on cooperative behavior across rounds of the twice-repeated prisoners' dilemma. For example, the treatment may induce participants to learn about the value of cooperation, thereby producing persistent effects. Table 3 summarizes randomization p-values for the total

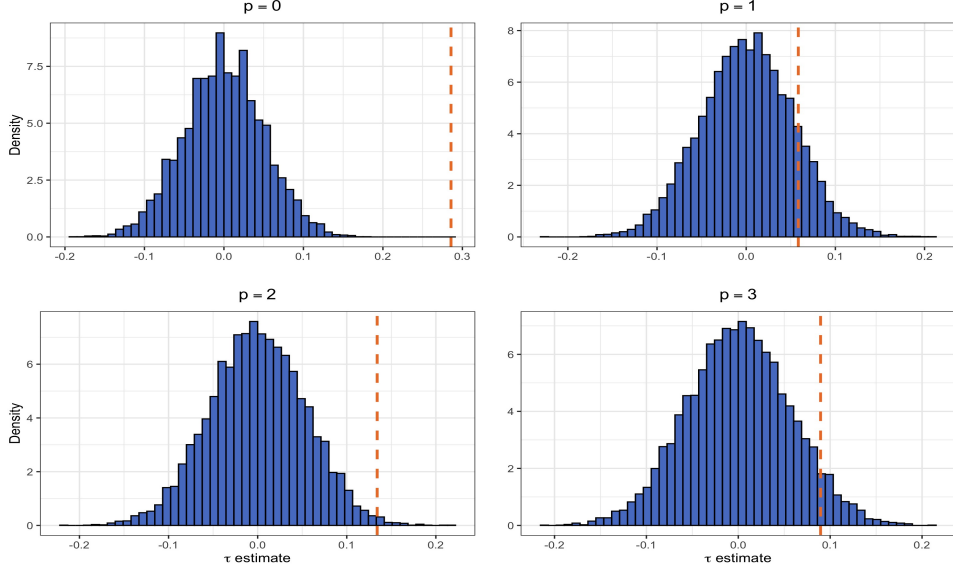


Figure 9: Estimated randomization distribution of the nonparametric estimator of the total lag- p weighted average dynamic causal effect, $\hat{\tau}^\dagger(1, 0; p)$, under the sharp null of no dynamic causal effect, $\tau_{i,t}(w, w; p) = 0$ for all $i \in [N], t \in [T]$. The dashed orange line plots the estimate, $\hat{\tau}^\dagger(1, 0; p)$ at the realized assignments in the experiment of [Andreoni and Samuelson \(2006\)](#). The estimated randomization distributions are constructed based on 10,000 draws.

lag- p weighted average causal effects.

7 Conclusion

In this paper, we developed a potential outcome model for studying dynamic causal effects in a panel experiment. Our work provided the first formal framework for incorporating design-based uncertainty in panel experiments—meaning that we treated the potential outcome as fixed, and the only source of randomization came from the random assignments. We defined new panel-based dynamic causal estimands such as the lag- p dynamic causal effect and introduced an associated nonparametric estimator. We showed that this estimator is unbiased for lag- p dynamic causal effects over the randomization distribution, and we derived its finite population asymptotic distribution. We then developed tools to conduct inference on these dynamic causal effects—introducing both an asymptotically conservative test for Neyman-type weak nulls and a randomization-based test for Fisher-type sharp nulls. We also derived the finite population

probability limit of the linear unit fixed effects estimator and two-way fixed effects estimator, showing that these estimators are asymptotically biased for the contemporaneous causal effects in the presence of dynamic causal effects and persistence in the treatment assignment mechanism. Finally, we illustrated our results in an extensive simulation study, and we reanalyzed a panel experiment conducted by [Andreoni and Samuelson \(2006\)](#).

References

- Abadie, Alberto, Susan C Athey, Guido W Imbens, and Jeffrey Wooldridge. 2017. “When Should You Adjust Standard Errors for Clustering?” NBER Working Paper No. 24003.
- Abadie, Alberto, Susan C Athey, Guido W Imbens, and Jeffrey Wooldridge. 2020. “Sampling-based vs. Design-based Uncertainty in Regression Analysis.” *Econometrica*, 88(1): 265–296.
- Abraham, Sarah, and Liyang Sun. 2020. “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects.”
- Allegretto, Sylvia, Arindajit Dube, Michael Reich, and Ben Zipperer. 2017. “Credible Research Designs for Minimum Wage Studies: A Response to Neumark, Salas, and Wascher.” *ILR Review*, 70(3): 559–592.
- Anderson, T.W., and C. Hsiao. 1982. “Formulation and Estimation of Dynamic Models Using Panel Data.” *Journal of Econometrics*, 18: 47–82.
- Andreoni, James, and Larry Samuelson. 2006. “Building Rational Cooperation.” *Journal of Economic Theory*, 127: 117–154.
- Arellano, M. 2003. *Panel Data Econometrics*. Oxford:Oxford University Press.
- Arellano, Manuel, and Stephanne Bonhomme. 2012. “Nonlinear Panel Data Analysis.” *Annual Review of Economics*, 3: 395–424.
- Arellano, Manuel, and Steve R Bond. 1991. “Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations.” *Review of Economic Studies*, 58: 277–297.
- Arellano, Manuel, Richard Blundell, and Stephanne Bonhomme. 2017. “Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework.” *Econometrica*, 85: 693–734.
- Arkhangelsky, Dmitry, and Guido Imbens. 2019. “Double-Robust Identification for Causal Panel Data Models.” arXiv preprint arXiv:1909.09412.
- Aronow, Peter M, and Cyrus Samii. 2017. “Estimating average causal effects under general interference, with application to a social network experiment.” *The Annals of Applied Statistics*, 11(4): 1912–1947.
- Athey, Susan, and Guido Imbens. 2018. “Design-based Analysis in Difference-In-Differences Settings with Staggered Adoption.” arXiv preprint arXiv:1808.05293.
- Athey, Susan, Mohsen Bayati, Nikolay Doudchenko, Guido Imbens, and Khashayar Koshravi. 2018. “Matrix Completion Methods for Causal Panel Data Models.” arXiv preprint arXiv 1710.10251.

- Bellemare, Charles, Luc Bissonnette, and Sabine Kroger.** 2014. “Statistical Power of Within and Between-Subjects Designs in Economic Experiments.” IZA Working Paper No. 8583.
- Bellemare, Charles, Luc Bissonnette, and Sabine Kroger.** 2016. “Simulating power of economic experiments: the powerBBK package.” *Journal of the Economic Science Association*, 2: 157–168.
- Ben-Porath, Yoram.** 1967. “The Production of Human Capital and the Life Cycle of Earnings.” *Journal of Political Economy*, 75: 352–365.
- Blackwell, Matthew, and Adam Glynn.** 2018. “How to make causal inferences with time-series and cross-sectional data.” *American Political Science Review*, 112: 1067–1082.
- Bojinov, Iavor, and Neil Shephard.** 2019. “Time series experiments and causal estimands: exact randomization tests and trading.” *Journal of the American Statistical Association*, 114(528): 1665–1682.
- Bojinov, Iavor, David Simchi-Levi, and Jinglong Zhao.** 2020. “Design and Analysis of Switchback Experiments.” arXiv preprint arXiv:2009.00148.
- Bojinov, Iavor, Guillaume Sait-Jacques, and Martin Tingley.** 2020. “Avoid the Pitfalls of A/B Testing.” *Harvard Business Review*, 98(2): 48–53.
- Boruvka, A, D Almirall, K Witkiwitz, and Susan A. Murphy.** 2018. “Assessing time-varying causal effect moderation in mobile health.” *Journal of the American Statistical Association*, 113: 1112–1121.
- Boryusak, Kirill, and Xavier Jaravel.** 2017. “Revisiting Event Study Designs, with an Application to the Estimation of the Marginal Propensity to Consume.”
- Browning, Martin, Mette Ejrnaes, and Javier Alvarez.** 2010. “Modelling income processes with lots of heterogeneity.” *Review of Economic Studies*, 77: 1353–1381.
- Charness, Gary, Uri Gneezy, and Michael A. Kuhn.** 2012. “Experimental methods: Between-subject and within-subject design.” *Journal of Economic and Business Organization*, 81(1): 1–8.
- Cox, David R.** 1958. *Planning of Experiments*. Oxford, United Kingdom:Wiley.
- Cunha, Flavio, James J. Heckman, and Susanne M. Schennach.** 2010. “Estimating the Technology of Cognitive and Noncognitive Skill Formation.” *Econometrica*, 78(3): 883–931.
- Cunha, Flavio, James J. Heckman, Lance Lochner, and Dimitriy V. Masterov.** 2006. “Chapter 12 Interpreting the Evidence on Life Cycle Skill Formation.” In *Handbook of the Economics of Education*. Vol. 1, , ed. E. Hanushek and F. Welch, 697–812. Elsevier.
- Czibor, Eszter, David Jimenez-Gomez, and John A. List.** 2019. “The Dozen Things Experimental Economists Should Do (More of).” *Southern Economic Journal*, 86(2): 371–432.

- de Chaisemartin, Clment, and Xavier D’Haultfoeuille.** 2020. “Two-way fixed effects estimators with heterogeneous treatment effects.” *American Economic Review*, 110(9): 2964–96.
- Ding, Peng.** 2017. “A Paradox from Randomization-Based Causal Inference.” *Statistical Science*, 32: 331–345.
- Freedman, David A.** 2008. “On regression adjustments to experimental data.” *Advances in Applied Mathematics*, 40(2): 180–193.
- Goodman-Bacon, Andrew.** 2018. “Difference-in-Differences with Variation in Treatment Timing.” NBER Working Paper No. 25018.
- Griliches, Zvi.** 1977. “Estimating the Returns to Schooling: Some Econometric Problems.” *Econometrica*, 45: 1–22.
- Hadad, Vitor, David A Hirshberg, Ruohan Zhan, Stefan Wager, and Susan Athey.** 2019. “Confidence intervals for policy evaluation in adaptive experiments.” *arXiv preprint arXiv:1911.02768*.
- Hall, Peter, and Chris C Heyde.** 1980. *Martingale Limit Theory and its Applications*. San Diego, California, USA:Academic Press.
- Han, Sukhjin.** 2019. “Identification in Nonparametric Models for Dynamic Treatment Effects.” *Journal of Econometrics*. Forthcoming.
- Heckman, James J., John Eric Humphries, and Gregory Veramendi.** 2016. “Dynamic treatment effects.” *Journal of Econometrics*, 191: 276–292.
- Hernan, M A, and J M Robins.** 2019. *Causal Inference*. Boca Raton, Florida, USA:Chapman & Hall. Forthcoming.
- Holland, Paul W.** 1986. “Statistics and Causal Inference.” *Journal of the American Statistical Association*, 81: 945–960.
- Horvitz, D G, and D J Thompson.** 1952. “A generalization of sampling without replacement from a finite universe.” *Journal of the American Statistical Association*, 47: 663–685.
- Hull, Peter.** 2018. “Estimating Treatment Effects in Mover Designs.” Unpublished paper: University of Chicago.
- Imai, Kosuke, and In Song Kim.** 2019. “When Should We Use Unit Fixed Effects Regression Models for Causal Inference with Longitudinal Data?” *American Journal of Political Science*, 63: 467–490.
- Imai, Kosuke, and In Song Kim.** 2020. “On the Use of Two-way Fixed Effects Regression Models for Causal Inference with Panel Data.” Unpublished paper: Harvard University.

- Imbens, Guido W, and Donald B Rubin.** 2015. *Causal Inference for Statistics, Social and Biomedical Sciences: An Introduction*. Cambridge, United Kingdom:Cambridge University Press.
- Kempthorne, O.** 1955. “The randomization theory of experimental inference.” *Journal of the American Statistical Association*, 50: 946–967.
- Koop, Gary, M. Hashem Pesaran, and Simon M. Potter.** 1996. “Impulse response analysis in nonlinear multivariate models.” *Journal of Econometrics*, 74: 119–147.
- Lai, Tze Leung, and Herbert Robbins.** 1985. “Asymptotically efficient adaptive allocation rules.” *Advances in applied mathematics*, 6(1): 4–22.
- Lechner, Michael.** 2011. “The Estimation of Causal Effects by Difference-in-Difference Methods.” *Foundations and Trends in Econometrics*, 4: 165–224.
- Lillie, Elizabeth O, Bradley Patay, Joel Diamant, Brian Issell, Eric J Topol, and Nicholas J Schork.** 2011. “The n-of-1 clinical trial: the ultimate strategy for individualizing medicine?” *Personalized medicine*, 8(2): 161–173.
- Li, Xinran, and Peng Ding.** 2017. “General Forms of Finite Population Central Limit Theorems with Applications to Causal Inference.” *Journal of the American Statistical Association*, 112(520): 1759–1769.
- Murphy, S. A., M. J. van der Laan, J. M. Robins, and Conduct Problems Prevention Research Group.** 2001. “Marginal Mean Models for Dynamic Regimes.” *Journal of the American Statistical Association*, 96: 1410–1423.
- Murphy, Susan A.** 2003. “Optimal Dynamic Treatment Regimes.” *Journal of the Royal Statistical Society B*, 65: 331–366.
- Nerlove, Marc.** 1971. “Further evidence on the estimation of dynamic economic relations from a time series of cross-sections.” *Econometrica*, 39: 359–387.
- Neyman, Jerzy.** 1923. “On the Application of Probability Theory to Agricultural Experiments. Essay on Principles. Section 9.” *Statistical Science*, 5: 465–472. Originally published 1923, republished in 1990, translated by Dorota M. Dabrowska and Terence P. Speed.
- Nickell, Steve J.** 1981. “Biases in dynamic models with fixed effects.” *Econometrica*, 49: 1417–1426.
- Pearl, Judea, and Dana Mackenzie.** 2018. *The Book of Why: The New Science of Cause and Effect*. Basic Books.
- Rambachan, Ashesh, and Neil Shephard.** 2020. “Econometric analysis of potential outcomes time series: instruments, shocks, linearity and the causal response function.” arXiv preprint arXiv:1903.01637.
- Robbins, Herbert.** 1952. “Some aspects of the sequential design of experiments.” *Bulletin of the American Mathematical Society*, 58(5): 527–535.

- Robins, James M.** 1986. “A New approach to causal inference in mortality studies With sustained exposure periods: application to control of the healthy Worker Survivor Effect.” *Mathematical Modelling*, 7: 1393–1512.
- Robins, James M.** 1994. “Correcting for non-compliance in randomization trials using structural nested mean models.” *Communications in Statistics — Theory and Methods*, 23: 2379–2412.
- Robins, James M, Sander Greenland, and Fu-Chang Hu.** 1999. “Estimation of the causal effect of a time-varying exposure on the marginal mean of a repeated binary outcome.” *Journal of the American Statistical Association*, 94: 687–700.
- Rubin, Donald B.** 1974. “Estimating causal effects of treatments in randomized and nonrandomized studies.” *Journal of Educational Psychology*, 66: 688–701.
- Wooldridge, Jeffrey M.** 2005. “Fixed-Effects and Related Estimators for Correlated Random-Coefficient and Treatment Effect Panel Data Models.” *Review of Economics and Statistics*, 87: 385390.

Panel-Based Experiments and Dynamic Causal Effects: A Finite Population Perspective

Online Appendix

Iavor Bojinov Ashesh Rambachan Neil Shephard

A Proofs of Main Results

Proof of Theorem 3.1

We begin the proof with a Lemma that will be used later on.

Lemma A.1. *Assume a potential outcome panel obeys Assumptions 4-5. Define, for any $\mathbf{w} \in \mathcal{W}^{(p+1)}$, the random function $Z_{i,t-p:t}(\mathbf{w}) := p_{i,t-p}(\mathbf{w})^{-1} \mathbb{1}\{W_{i,t-p:t} = \mathbf{w}\}$. Then, over the randomization mechanism, $\mathbb{E}(Z_{i,t-p:t}(\mathbf{w})|\mathcal{F}_{i,t-p-1}) = 1$ and $\text{Var}(Z_{i,t-p:t}(\mathbf{w})|\mathcal{F}_{i,t-p-1}) = p_{i,t-p}(\mathbf{w})^{-1}(1 - p_{i,t-p}(\mathbf{w}))$, and $\text{Cov}(Z_{i,t-p:t}(\mathbf{w}), Z_{i,t-p:t}(\tilde{\mathbf{w}})|\mathcal{F}_{i,t-p-1}) = -1$ for all $\mathbf{w} \neq \tilde{\mathbf{w}}$. Moreover, $Z_{i,t-p:t}(\mathbf{w})$ and $Z_{j,t-p:t}(\mathbf{w})$ are, conditioning on $\mathcal{F}_{1:N,t-p-1}$, independent for $i \neq j$.*

Proof. The expectation is by construction, the variance comes from the variance of a Bernoulli trial. The conditional independence is by the individualistic assignment assumption. \square

For any $\mathbf{w}, \tilde{\mathbf{w}} \in \mathcal{W}^{(p+1)}$, let $u_{i,t-p}(\mathbf{w}, \tilde{\mathbf{w}}; p) = \hat{\tau}_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p) - \tau_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p)$ be the estimation error. Now

$$u_{i,t-p}(\mathbf{w}, \tilde{\mathbf{w}}; p) = Y_{i,t}(w_{i,1:t-p-1}^{obs}, \mathbf{w})(Z_{i,t-p:t}(\mathbf{w}) - 1) - Y_{i,t}(w_{i,1:t-p-1}^{obs}, \tilde{\mathbf{w}})(Z_{i,t-p:t}(\tilde{\mathbf{w}}) - 1).$$

Hence the zero condition expectation follows using Lemma A.1. Then,

$$\begin{aligned} \text{Var}(u_{i,t-p}(\mathbf{w}, \tilde{\mathbf{w}}; p)|\mathcal{F}_{i,t-p-1}) &= Y_{i,t}(w_{i,1:t-p-1}^{obs}, \mathbf{w})^2 \text{Var}(Z_{i,t-p:t}(\mathbf{w})|\mathcal{F}_{i,t-p-1}) \\ &\quad + Y_{i,t}(w_{i,1:t-p-1}^{obs}, \tilde{\mathbf{w}})^2 \text{Var}(Z_{i,t-p:t}(\tilde{\mathbf{w}})|\mathcal{F}_{i,t-p-1}) \\ &\quad - 2Y_{i,t}(w_{i,1:t-p-1}^{obs}, \mathbf{w})Y_{i,t}(w_{i,1:t-p-1}^{obs}, \tilde{\mathbf{w}}) \text{Cov}(Z_{i,t-p:t}(\mathbf{w}), Z_{i,t-p:t}(\tilde{\mathbf{w}})|\mathcal{F}_{i,t-p-1}) \\ &= Y_{i,t}(w_{i,1:t-p-1}^{obs}, \mathbf{w})^2 p_{i,t-p}(\mathbf{w})^{-1}(1 - p_{i,t-p}(\mathbf{w})) \\ &\quad + Y_{i,t}(w_{i,1:t-p-1}^{obs}, \tilde{\mathbf{w}})^2 p_{i,t-p}(\tilde{\mathbf{w}})^{-1}(1 - p_{i,t-p}(\tilde{\mathbf{w}})) \\ &\quad - 2Y_{i,t}(w_{i,1:t-p-1}^{obs}, \mathbf{w})Y_{i,t}(w_{i,1:t-p-1}^{obs}, \tilde{\mathbf{w}}). \end{aligned}$$

Simplifying gives the result on the variance of the estimation error. Then,

$$\begin{aligned} \text{Cov}(u_{i,t-p}(\mathbf{w}, \tilde{\mathbf{w}}; p), u_{i,t-p}(\bar{\mathbf{w}}, \hat{\mathbf{w}}; p)|\mathcal{F}_{i,t-p-1}) &= Y_{i,t}(w_{i,1:t-p-1}^{obs}, \mathbf{w})Y_{i,t}(w_{i,1:t-p-1}^{obs}, \bar{\mathbf{w}}) \text{Cov}(Z_{i,t-p:t}(\mathbf{w}), Z_{i,t-p:t}(\bar{\mathbf{w}})|\mathcal{F}_{i,t-p-1}) \\ &\quad - Y_{i,t}(w_{i,1:t-p-1}^{obs}, \mathbf{w})Y_{i,t}(w_{i,1:t-p-1}^{obs}, \hat{\mathbf{w}}) \text{Cov}(Z_{i,t-p:t}(\mathbf{w}), Z_{i,t-p:t}(\hat{\mathbf{w}})|\mathcal{F}_{i,t-p-1}) \\ &\quad - Y_{i,t}(w_{i,1:t-p-1}^{obs}, \tilde{\mathbf{w}})Y_{i,t}(w_{i,1:t-p-1}^{obs}, \bar{\mathbf{w}}) \text{Cov}(Z_{i,t-p:t}(\tilde{\mathbf{w}}), Z_{i,t-p:t}(\bar{\mathbf{w}})|\mathcal{F}_{i,t-p-1}) \\ &\quad - Y_{i,t}(w_{i,1:t-p-1}^{obs}, \tilde{\mathbf{w}})Y_{i,t}(w_{i,1:t-p-1}^{obs}, \hat{\mathbf{w}}) \text{Cov}(Z_{i,t-p:t}(\tilde{\mathbf{w}}), Z_{i,t-p:t}(\hat{\mathbf{w}})|\mathcal{F}_{i,t-p-1}) \end{aligned}$$

$$\begin{aligned}
&= -Y_{i,t}(w_{i,1:t-p-1}^{obs}, \mathbf{w})Y_{i,t}(w_{i,1:t-p-1}^{obs}, \bar{\mathbf{w}}) + Y_{i,t}(w_{i,1:t-p-1}^{obs}, \mathbf{w})Y_{i,t}(w_{i,1:t-p-1}^{obs}, \hat{\mathbf{w}}) \\
&+ Y_{i,t}(w_{i,1:t-p-1}^{obs}, \tilde{\mathbf{w}})Y_{i,t}(w_{i,1:t-p-1}^{obs}, \bar{\mathbf{w}}) - Y_{i,t}(w_{i,1:t-p-1}^{obs}, \tilde{\mathbf{w}})Y_{i,t}(w_{i,1:t-p-1}^{obs}, \hat{\mathbf{w}})
\end{aligned}$$

Finally, conditional independence of the errors follows due to the individualistic assignment of treatments. \square

Proof of Theorem 3.2

Only the third results requires a new proof. In particular, the first result is a reinterpretation of the classic cross-sectional result using a triangular array central limit theorem, for the usual Lindeberg condition must hold due to the bounded potential outcomes and the treatments being probabilistic. The second result follows from results in [Bojinov and Shephard \(2019\)](#), who use a martingale difference array central limit theorem.

The third result, which holds for NT going to infinity, can be split into three parts. For NT to go to infinity we must have either: (i) T goes to infinity with N finite, (ii) N goes to infinity with T finite, or (iii) both N and T go to infinity. In the case (i), we apply the martingale difference CLT but now where each time period we have preaveraged the cross-sectional errors over the N terms. The preaverage is still a martingale difference, so the technology is the same. In the case (ii) we preaverage the time aspect. Then we are back to a standard triangular array CLT. As we have both (i) and (ii), then (iii) must hold. \square

Proof of Proposition 4.1

Under linear potential outcomes,

$$Y_{i,t}(W_{i,1:t}) - Y_{i,t}(\tilde{W}_{i,1:t}) = \sum_{s=0}^{t-1} \beta_{i,t,s}(W_{i,t-s} - \tilde{W}_{i,t-s}).$$

Focus on the counterfactual $\tilde{W}_{i,1:t} = \mathbf{0}$, then

$$Y_{i,t} = Y_{i,t}(W_{i,1:t}) = \bar{Y}_t(\mathbf{0}) + \sum_{s=0}^{t-1} \beta_{i,t,s}W_{i,t-s} + \dot{Y}_{i,t}(\mathbf{0}).$$

Therefore, the within-period transformed outcome equals

$$\dot{Y}_{i,t} = Y_{i,t} - \bar{Y}_t = \sum_{s=0}^{t-1} \{\beta_{i,t,s}W_{i,t-s} - \frac{1}{N} \sum_{j=1}^N \beta_{j,t,s}W_{j,t-s}\} + \dot{Y}_{i,t}(\mathbf{0}).$$

Further imposing homogeneity, it simplifies to

$$\dot{Y}_{i,t} = \sum_{s=0}^{t-1} \{\beta_{t,s}(W_{i,t-s} - \frac{1}{N} \sum_{j=1}^N W_{j,t-s})\} + \dot{Y}_{i,t}(\mathbf{0}).$$

Stacking everything across units, this becomes $\dot{Y}_{1:N,t} = \dot{W}_{1:N,t}\beta_t + \dot{Y}_{1:N,t}(\mathbf{0})$, and so the linear projection coefficient is given by

$$\hat{\beta}_t = (\dot{W}'_{1:N,t}\dot{W}_{1:N,t})^{-1}\dot{W}'_{1:N,t}\dot{Y}_{1:N,t} = \beta_t + (\dot{W}'_{1:N,t}\dot{W}_{1:N,t})^{-1}\dot{W}'_{1:N,t}\dot{Y}_{1:N,t}(\mathbf{0}).$$

The important unusual point here is that $\dot{Y}_{1:N,t}(\mathbf{0})$ is non-stochastic and that $\dot{W}_{1:N,t}$ is random, exactly the opposite of the case often discussed in the statistical analysis of linear regression. Now

$$\frac{1}{N}\dot{W}'_{1:N,t}\dot{W}_{1:N,t} = \frac{1}{N}\sum_{i=1}^N\dot{W}_{i,1:t}\dot{W}'_{i,1:t},$$

and

$$\frac{1}{N}\sum_{i=1}^N\dot{W}_{i,1:t}\dot{Y}_{i,t}(\mathbf{0}) = \frac{1}{N}\sum_{i=1}^N(\dot{W}_{i,1:t} - \mu_{i,t})\dot{Y}_{i,t}(\mathbf{0}) + \frac{1}{N}\sum_{i=1}^N\mu_{i,t}\dot{Y}_{i,t}(\mathbf{0}).$$

Then, under the assumption of individualistic assignment (Assumption 4),

$$\frac{1}{N}\sum_{i=1}^N\dot{W}_{i,1:t}\dot{W}'_{i,1:t}|\mathcal{F}_{1:N,0,T} \xrightarrow{p} \Gamma_{2,t},$$

recalling $\dot{Y}_{i,t}(\mathbf{0})$ is non-stochastic and applying Assumptions 3(b) and 3(c), then Slutsky's theorem delivers the result stated in the paper. \square

Proof of Proposition 4.2

Begin by writing the observed outcomes as

$$Y_{i,t} = Y_{i,t}(\mathbf{0}) + \sum_{s=1}^t \beta_{i,t,t-s} W_{i,s}.$$

Similarly, write $\bar{Y}_{i\cdot} = \bar{Y}_{i\cdot}(\mathbf{0}) + \bar{\beta}\bar{W}_{i\cdot}$, where $\bar{\beta}\bar{W}_{i\cdot} = \frac{1}{T}\sum_{t=1}^T\sum_{s=1}^t\beta_{i,t,t-s}W_{i,s}$. The transformed outcome can be then written as

$$\check{Y}_{i,t} = \sum_{s=1}^t \beta_{i,t,t-s} W_{i,s} - \bar{\beta}\bar{W}_{i\cdot} + \check{Y}_{i,t}(\mathbf{0}).$$

Consider the numerator of the unit fixed effects estimator. Substituting in, we arrive at

$$\begin{aligned} \frac{1}{NT}\sum_{i=1}^N\sum_{t=1}^T\check{Y}_{i,t}\widetilde{W}_{i,t} &= \frac{1}{NT}\sum_{i=1}^N\sum_{t=1}^T\beta_{i,t,0}W_{i,t}\widetilde{W}_{i,t} + \frac{1}{NT}\sum_{i=1}^N\sum_{t=1}^T\left(\sum_{s=1}^{t-1}\beta_{i,t,t-s}W_{i,s}\widetilde{W}_{i,t}\right) + \frac{1}{NT}\sum_{i=1}^N\sum_{t=1}^T\check{Y}_{i,t}(\mathbf{0})\widetilde{W}_{i,t} \\ &= \frac{1}{T}\sum_{t=1}^T\left(\frac{1}{N}\sum_{i=1}^N\beta_{i,t,0}W_{i,t}\widetilde{W}_{i,t}\right) + \frac{1}{T}\sum_{t=1}^T\sum_{s=1}^{t-1}\left(\frac{1}{N}\sum_{i=1}^N\beta_{i,t,t-s}W_{i,s}\widetilde{W}_{i,t}\right) + \frac{1}{T}\sum_{t=1}^T\left(\frac{1}{N}\sum_{i=1}^N\check{Y}_{i,t}(\mathbf{0})\widetilde{W}_{i,t}\right). \end{aligned}$$

Therefore, for fixed T as $N \rightarrow \infty$,

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \left(\frac{1}{N} \sum_{i=1}^N \beta_{i,t,0} W_{i,t} \widetilde{W}_{i,t} \right) &\xrightarrow{p} \frac{1}{T} \sum_{t=1}^T \check{\kappa}_{W,\beta,t,t}, \\ \frac{1}{T} \sum_{t=1}^T \sum_{s=1}^{t-1} \left(\frac{1}{N} \sum_{i=1}^N \beta_{i,t,t-s} W_{i,s} \widetilde{W}_{i,t} \right) &\xrightarrow{p} \frac{1}{T} \sum_{t=1}^T \sum_{s=1}^{t-1} \check{\kappa}_{W,\beta,t,s}, \\ \frac{1}{T} \sum_{t=1}^T \left(\frac{1}{N} \sum_{i=1}^N \check{Y}_{i,t}(\mathbf{0}) \widetilde{W}_i \right) &\xrightarrow{p} \frac{1}{T} \sum_{t=1}^T \check{\delta}_t. \end{aligned}$$

Similarly, the denominator converges to $\frac{1}{NT} \sum_{t=1}^T \sum_{i=1}^N \widetilde{W}_{i,t}^2 \xrightarrow{p} \frac{1}{T} \sum_{t=1}^T \check{\sigma}_{W,t}^2$. The result then follows by Slutsky. \square

Proof of Proposition 4.3

Begin by writing

$$Y_{i,t} = Y_{i,t}(\mathbf{0}) + \sum_{s=1}^t \beta_{i,t,t-s} W_{i,s}.$$

Then, $\bar{Y}_t = \bar{Y}_t(\mathbf{0}) + \overline{\beta W}_t$, $\bar{Y}_i = \bar{Y}_i(\mathbf{0}) + \overline{\beta W}_i$, and $\bar{Y} = \bar{Y}(\mathbf{0}) + \overline{\beta W}$. Therefore,

$$\dot{\check{Y}}_{i,t} = \dot{\check{Y}}_{i,t}(\mathbf{0}) + \left(\sum_{s=1}^t \beta_{i,t,t-s} W_{i,s} - \overline{\beta W} \right) - (\overline{\beta W}_t - \overline{\beta W}) - (\overline{\beta W}_i - \overline{\beta W}).$$

Consider the numerator of the unit fixed effects estimator. Substituting in,

$$\begin{aligned} \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \dot{\check{Y}}_{i,t} \dot{\check{W}}_{i,t} &= \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \beta_{i,t,0} W_{i,t} \dot{\check{W}}_{i,t} + \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \sum_{s=1}^{t-1} \beta_{i,t,t-s} W_{i,s} \dot{\check{W}}_{i,t} + \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \dot{\check{Y}}_{i,t}(\mathbf{0}) \dot{\check{W}}_{i,t} \\ &= \frac{1}{T} \sum_{t=1}^T \left(\frac{1}{N} \sum_{i=1}^N \beta_{i,t,0} W_{i,t} \dot{\check{W}}_{i,t} \right) + \frac{1}{T} \sum_{t=1}^T \left(\frac{1}{N} \sum_{i=1}^N \sum_{s=1}^{t-1} \beta_{i,t,t-s} W_{i,s} \dot{\check{W}}_{i,t} \right) + \frac{1}{T} \sum_{t=1}^T \left(\frac{1}{N} \sum_{i=1}^N \dot{\check{Y}}_{i,t}(\mathbf{0}) \dot{\check{W}}_{i,t} \right). \end{aligned}$$

Therefore,

$$\begin{aligned} \frac{1}{N} \sum_{i=1}^N \beta_{i,t,0} W_{i,t} \dot{\check{W}}_{i,t} &\xrightarrow{p} \dot{\check{\kappa}}_{W,\beta,t,t}, \\ \frac{1}{N} \sum_{i=1}^N \sum_{s=1}^{t-1} \beta_{i,t,t-s} W_{i,s} \dot{\check{W}}_{i,t} &\xrightarrow{p} \sum_{s=1}^{t-1} \dot{\check{\kappa}}_{W,\beta,t,s}, \\ \frac{1}{N} \sum_{i=1}^N \dot{\check{Y}}_{i,t}(\mathbf{0}) \dot{\check{W}}_{i,t} &\xrightarrow{p} \dot{\check{\delta}}_t. \end{aligned}$$

A similar argument applies to the denominator and the result follows. \square

B Additional simulation results

		$p(w)$		
		0.25	0.5	0.75
	0.25	0.044	0.049	0.050
ϕ	0.5	0.048	0.050	0.049
	0.75	0.050	0.051	0.057

(a) $\epsilon_{i,t} \sim N(0, 1), N = 1000$

		$p(w)$		
		0.25	0.5	0.75
	0.25	0.031	0.031	0.034
ϕ	0.5	0.048	0.039	0.043
	0.75	0.052	0.047	0.057

(b) $\epsilon_{i,t} \sim Cauchy, N = 50,000$

Table A1: Null rejection rate for the test of the null hypothesis $H_0 : \bar{\tau}_t(1, 0; 0) = 0$ based upon the normal asymptotic approximation to the randomization distribution of $\hat{\tau}_t(1, 0; 0)$. Panel (a) reports the null rejection probabilities in simulations with $\epsilon_{i,t} \sim N(0, 1)$ and $N = 1000$. Panel (b) reports the null rejection probabilities in simulations with $\epsilon_{i,t} \sim Cauchy$ and $N = 50,000$. Results are computed over 5,000 simulations. See Section 5 of the main text for further details.

		$p(w)$		
		0.25	0.5	0.75
	0.25	0.044	0.046	0.052
ϕ	0.5	0.050	0.054	0.050
	0.75	0.046	0.049	0.054

(a) $\epsilon_{i,t} \sim N(0, 1), T = 1000$

		$p(w)$		
		0.25	0.5	0.75
	0.25	0.031	0.031	0.034
ϕ	0.5	0.048	0.039	0.043
	0.75	0.052	0.047	0.057

(b) $\epsilon_{i,t} \sim Cauchy, T = 50,000$

Table A2: Null rejection rate for the test of the null hypothesis $H_0 : \bar{\tau}_i(1, 0; 0) = 0$ based upon the normal asymptotic approximation to the randomization distribution of $\hat{\tau}_i(1, 0; 0)$. Panel (a) reports the null rejection probabilities in simulations with $\epsilon_{i,t} \sim N(0, 1)$ and $T = 1000$. Panel (b) reports the null rejection probabilities in simulations with $\epsilon_{i,t} \sim Cauchy$ and $T = 50,000$. Results are computed over 5,000 simulations. See Section 5 of the main text for further details.

		$p(w)$		
		0.25	0.5	0.75
	0.25	0.050	0.047	0.048
ϕ	0.5	0.052	0.052	0.050
	0.75	0.050	0.049	0.048

(a) $\epsilon_{i,t} \sim N(0, 1), N = 100, T = 10$

		$p(w)$		
		0.25	0.5	0.75
	0.25	0.028	0.029	0.032
ϕ	0.5	0.046	0.039	0.044
	0.75	0.055	0.044	0.054

(b) $\epsilon_{i,t} \sim Cauchy, N = 500, T = 100$

Table A3: Null rejection rate for the test of the null hypothesis $H_0 : \bar{\tau}(1, 0; 0) = 0$ based upon the normal asymptotic approximation to the randomization distribution of $\hat{\tau}(1, 0; 0)$. Panel (a) reports the null rejection probabilities in simulations with $\epsilon_{i,t} \sim N(0, 1)$ and $N = 100, T = 10$. Panel (b) reports the null rejection probabilities in simulations with $\epsilon_{i,t} \sim Cauchy$ and $N = 500, T = 100$. Results are computed over 5,000 simulations. See Section 5 of the main text for further details.

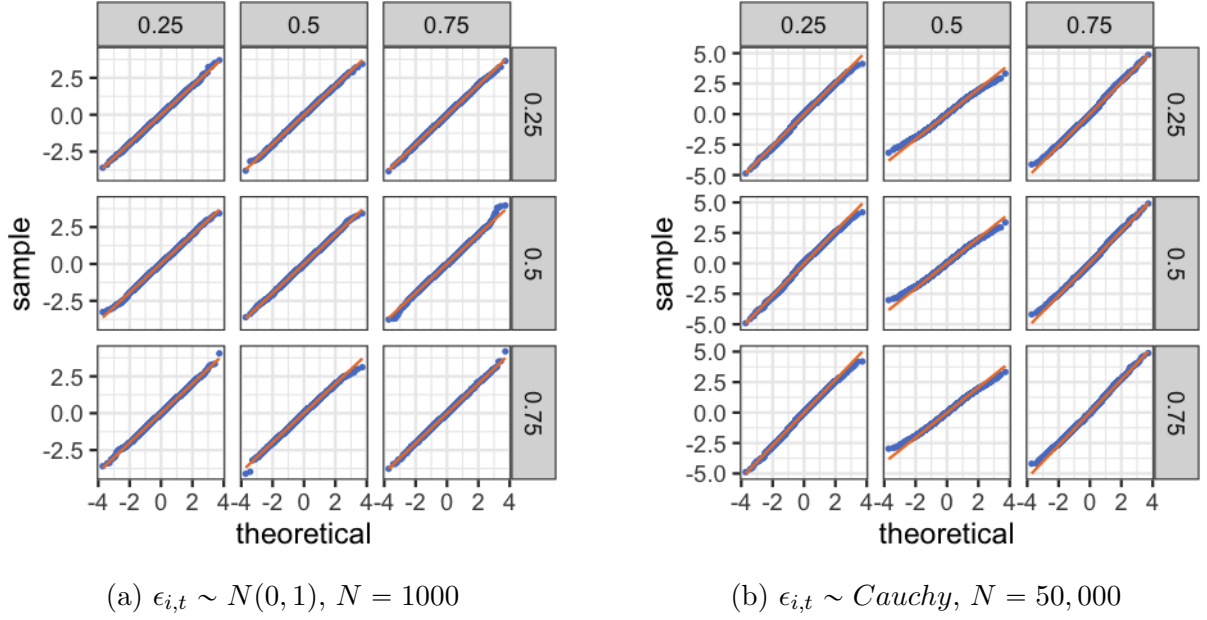




Figure A1: Quantile-quantile plots for the simulated randomization distribution for $\hat{\tau}_t(1; 0; 0)$ under different choices of the parameter ϕ (defined in Example 1) and treatment probability $p(w)$. The quantile-quantile plots compare the quantiles of the simulated randomization distribution (y-axis) against the quantiles of a standard normal random variable (x-axis). The 45 degree line is plotted in solid orange. The rows index the parameter ϕ , which ranges over values $\{0.25, 0.5, 0.75\}$. The columns index the treatment probability $p(w)$, which ranges over values $\{0.25, 0.5, 0.75\}$. Panel (a) plots the quantile-quantile plots for simulated randomization distribution with normally distributed errors $\epsilon_{i,t} \sim N(0, 1)$ and $N = 1000$. Panel (b) plots the quantile-quantile plots simulated randomization distribution with Cauchy distribution errors $\epsilon_{i,t} \sim Cauchy$ and $N = 50,000$. Results are computed over 5,000 simulations. See Section 5 of the main text for further details.

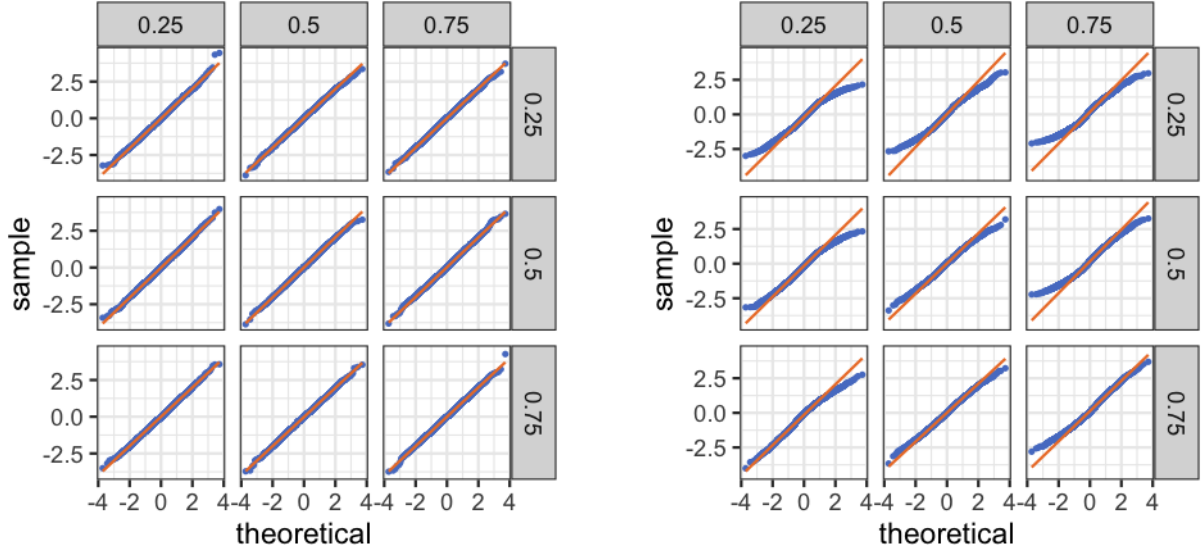


(a) $\epsilon_{i,t} \sim N(0, 1)$, $T = 1000$



(b) $\epsilon_{i,t} \sim Cauchy$, $T = 50,000$


Figure A2: Quantile-quantile plots for the simulated randomization distribution for $\hat{\tau}_i(1, 0; 0)$ under different choices of the parameter ϕ (defined in Example 1) and treatment probability $p(w)$. The quantile-quantile plots compare the quantiles of the simulated randomization distribution (y-axis) against the quantiles of a standard normal random variable (x-axis). The 45 degree line is plotted in solid orange. The rows index the parameter ϕ , which ranges over values $\{0.25, 0.5, 0.75\}$. The columns index the treatment probability $p(w)$, which ranges over values $\{0.25, 0.5, 0.75\}$. Panel (a) plots the quantile-quantile plots for simulated randomization distribution with normally distributed errors $\epsilon_{i,t} \sim N(0, 1)$ and $T = 1000$. Panel (b) plots the quantile-quantile plots simulated randomization distribution with Cauchy distribution errors $\epsilon_{i,t} \sim Cauchy$ and $T = 50,000$. Results are computed over 5,000 simulations. See Section 5 of the main text for further details.




(a) $\epsilon_{i,t} \sim N(0, 1)$, $N = 100$, $T = 10$

(b) $\epsilon_{i,t} \sim Cauchy$, $N = 500$, $T = 100$

Figure A3: Quantile-quantile plots for the simulated randomization distribution for $\hat{\tau}(1, 0; 0)$ under different choices of the parameter ϕ (defined in Example 1) and treatment probability $p(w)$. The quantile-quantile plots compare the quantiles of the simulated randomization distribution (y-axis) against the quantiles of a standard normal random variable (x-axis). The 45 degree line is plotted in solid orange. The rows index the parameter ϕ , which ranges over values $\{0.25, 0.5, 0.75\}$. The columns index the treatment probability $p(w)$, which ranges over values $\{0.25, 0.5, 0.75\}$. Panel (a) plots the quantile-quantile plots for simulated randomization distribution with normally distributed errors $\epsilon_{i,t} \sim N(0, 1)$ and $N = 100, T = 10$. Panel (b) plots the quantile-quantile plots simulated randomization distribution with Cauchy distribution errors $\epsilon_{i,t} \sim Cauchy$ and $N = 500, T = 100$. Results are computed over 5,000 simulations. See Section 5 of the main text for further details.




(a) $\epsilon_{i,t} \sim N(0, 1)$, $N = 1000$




(b) $\epsilon_{i,t} \sim Cauchy$, $N = 50,000$

Figure A4: Simulated randomization distribution for $\hat{\tau}_{t,t}^{\dagger}(1, 0; 1)$ under different choices of the parameter ϕ (defined in Example 1) and treatment probability $p(w)$. The rows index the parameter ϕ , which ranges over values $\{0.25, 0.5, 0.75\}$. The columns index the treatment probability $p(w)$, which ranges over values $\{0.25, 0.5, 0.75\}$. Panel (a) plots the simulated randomization distribution with normally distributed errors $\epsilon_{i,t} \sim N(0, 1)$ and $N = 1000$. Panel (b) plots the simulated randomization distribution with Cauchy distribution errors $\epsilon_{i,t} \sim Cauchy$ and $N = 50,000$. Results are computed over 5,000 simulations. See Section 5 of the main text for further details.




(a) $\epsilon_{i,t} \sim N(0, 1)$, $T = 1000$




(b) $\epsilon_{i,t} \sim Cauchy$, $T = 50,000$

Figure A5: Simulated randomization distribution for $\hat{\tau}_i^\dagger(1, 0; 1)$ under different choices of the parameter ϕ (defined in Example 1) and treatment probability $p(w)$. The rows index the parameter ϕ , which ranges over values $\{0.25, 0.5, 0.75\}$. The columns index the treatment probability $p(w)$, which ranges over values $\{0.25, 0.5, 0.75\}$. Panel (a) plots the simulated randomization distribution with normally distributed errors $\epsilon_{i,t} \sim N(0, 1)$ and $T = 1000$. Panel (b) plots the simulated randomization distribution with Cauchy distribution errors $\epsilon_{i,t} \sim Cauchy$ and $T = 50,000$. Results are computed over 5,000 simulations. See Section 5 of the main text for further details.




(a) $\epsilon_{i,t} \sim N(0, 1)$, $N = 100$, $T = 10$




(b) $\epsilon_{i,t} \sim Cauchy$, $N = 500$, $T = 100$

Figure A6: Simulated randomization distribution for $\hat{\tau}^\dagger(1, 0; 1)$ under different choices of the parameter ϕ (defined in Example 1) and treatment probability $p(w)$. The rows index the parameter ϕ , which ranges over values $\{0.25, 0.5, 0.75\}$. The columns index the treatment probability $p(w)$, which ranges over values $\{0.25, 0.5, 0.75\}$. Panel (a) plots the simulated randomization distribution with normally distributed errors $\epsilon_{i,t} \sim N(0, 1)$ and $N = 100, T = 10$. Panel (b) plots the simulated randomization distribution with Cauchy distribution errors $\epsilon_{i,t} \sim Cauchy$ and $N = 500, T = 10$. Results are computed over 5,000 simulations. See Section 5 of the main text for further details.



(a) $\epsilon_{i,t} \sim N(0, 1)$, $N = 1000$



(b) $\epsilon_{i,t} \sim Cauchy$, $N = 50,000$

Figure A7: Quantile-quantile plots for the simulated randomization distribution for $\hat{\tau}_{\cdot,t}^{\dagger}(1; 0; 1)$ under different choices of the parameter ϕ (defined in Example 1) and treatment probability $p(w)$. The quantile-quantile plots compare the quantiles of the simulated randomization distribution (y-axis) against the quantiles of a standard normal random variable (x-axis). The 45 degree line is plotted in solid orange. The rows index the parameter ϕ , which ranges over values $\{0.25, 0.5, 0.75\}$. The columns index the treatment probability $p(w)$, which ranges over values $\{0.25, 0.5, 0.75\}$. Panel (a) plots the quantile-quantile plots for simulated randomization distribution with normally distributed errors $\epsilon_{i,t} \sim N(0, 1)$ and $T = 1000$. Panel (b) plots the quantile-quantile plots simulated randomization distribution with Cauchy distribution errors $\epsilon_{i,t} \sim Cauchy$ and $T = 50,000$. Results are computed over 5,000 simulations. See Section 5 of the main text for further details.

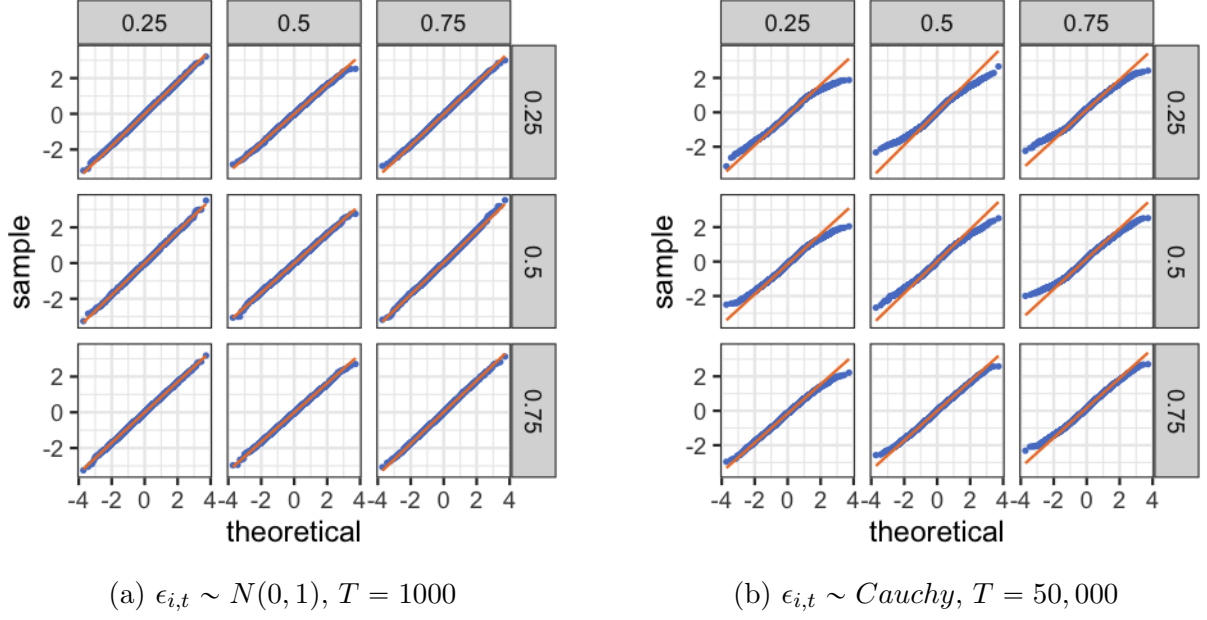


Figure A8: Quantile-quantile plots for the simulated randomization distribution for $\hat{\tau}_i^\dagger(1; 0; 1)$ under different choices of the parameter ϕ (defined in Example 1) and treatment probability $p(w)$. The quantile-quantile plots compare the quantiles of the simulated randomization distribution (y-axis) against the quantiles of a standard normal random variable (x-axis). The 45 degree line is plotted in solid orange. The rows index the parameter ϕ , which ranges over values $\{0.25, 0.5, 0.75\}$. The columns index the treatment probability $p(w)$, which ranges over values $\{0.25, 0.5, 0.75\}$. Panel (a) plots the quantile-quantile plots for simulated randomization distribution with normally distributed errors $\epsilon_{i,t} \sim N(0, 1)$ and $T = 1000$. Panel (b) plots the quantile-quantile plots simulated randomization distribution with Cauchy distribution errors $\epsilon_{i,t} \sim Cauchy$ and $T = 50,000$. Results are computed over 5,000 simulations. See Section 5 of the main text for further details.

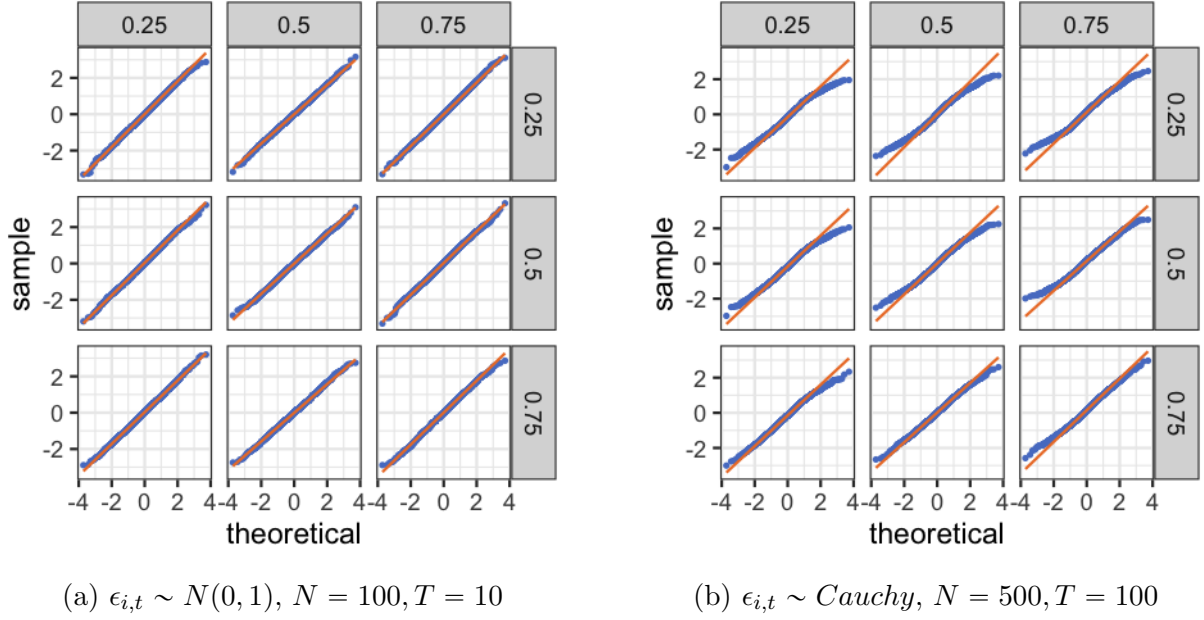


Figure A9: Quantile-quantile plots for the simulated randomization distribution for $\hat{\tau}^\dagger(1, 0; 1)$ under different choices of the parameter ϕ (defined in Example 1) and treatment probability $p(w)$. The quantile-quantile plots compare the quantiles of the simulated randomization distribution (y-axis) against the quantiles of a standard normal random variable (x-axis). The 45 degree line is plotted in solid orange. The rows index the parameter ϕ , which ranges over values $\{0.25, 0.5, 0.75\}$. The columns index the treatment probability $p(w)$, which ranges over values $\{0.25, 0.5, 0.75\}$. Panel (a) plots the quantile-quantile plots for simulated randomization distribution with normally distributed errors $\epsilon_{i,t} \sim N(0, 1)$ and $T = 1000$. Panel (b) plots the quantile-quantile plots simulated randomization distribution with Cauchy distribution errors $\epsilon_{i,t} \sim Cauchy$ and $T = 50,000$. Results are computed over 5,000 simulations. See Section 5 of the main text for further details.

C Additional empirical results

As in Section 6, we begin by estimating unit-specific, weighted average dynamic causal effects to investigate the causal effect of $W = 1 \{ \lambda \geq 0.6 \}$ on the total payoffs earned. We focus on the same two units as in Figure 7, Figure A10 plots the nonparametric estimates $\hat{\tau}_{i,t}(1, 0; 0)$ for $t \in [T]$ for the total payoffs outcome.




Figure A10: Estimates of the weighted average i, t -th lag-0 dynamic causal effect (Definition 5) of $W = 1 \{ \lambda \geq 0.6 \}$ on total stage game payoffs for two units in the experiment of Andreoni and Samuelson (2006). The solid black line plots the nonparametric estimator $\hat{\tau}_{i,t}(1, 0; 0)$. The dashed black line plots the running average of the period-specific estimator for each unit; that is, for each $t \in [T]$, $\frac{1}{t} \sum_{s=1}^t \hat{\tau}_{i,s}(1, 0; 0)$. The dashed red line plots the estimated unit- i lag-0 average weighted dynamic causal effect, $\hat{\tau}_i(1, 0; 0) = \frac{1}{T} \sum_{t=1}^T \hat{\tau}_{i,t}(1, 0; 0)$.

We next estimate period-specific, weighted average dynamic causal effects for each time period $t \in [T]$ and $p = 0, 1, 2, 3$. The results are plotted in Figure A11. It also plots the nonparametric estimator the total lag- p weighted average causal effect $\tau^\dagger(1, 0; p)$ for $p = 0, 1, 2, 3$. While of course the units are different, the qualitative results are unchanged from Section 6. We find strong evidence of a contemporaneous causal effect on the total payoffs but mixed evidence on dynamic causal effects.

Finally, Figure A12 plots the randomization distributions under the sharp null of no dynamic causal effects for $p = 0, 1, 2, 3$ along with the point estimate $\hat{\tau}^\dagger(1, 0; p)$ at the realized treatment panel. As before, the randomization distributions appear to be smooth and symmetric around zero. We reject the sharp null of no dynamic causal effects at the 5% level for $p = 0$ (p-value is 0.0071) but are unable to do so for $p = 1, 2, 3$.

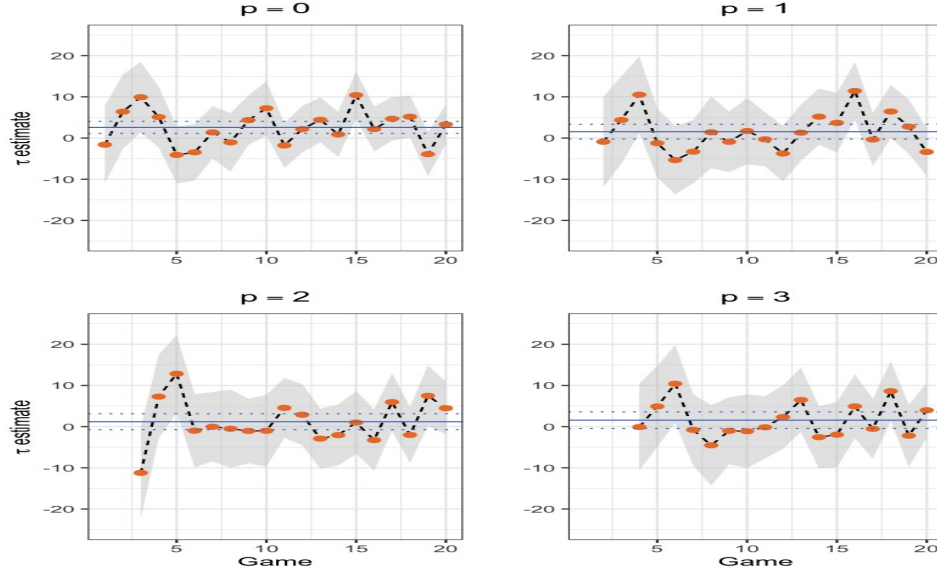


Figure A11: Estimates of the time- t lag- p weighted average dynamic causal effect, $\bar{\tau}_t^\dagger(1, 0; p)$ of $W = 1\{\lambda \geq 0.6\}$ on total payoffs based on the experiment of [Andreoni and Samuelson \(2006\)](#) for each time period $t \in [T]$ and $p = 0, 1, 2, 3$. The black dashed line plots the nonparametric estimator of the time- t lag- p weighted average dynamic causal effect, $\hat{\tau}_t^\dagger(1, 0; p)$, for each period $t \in [T]$. The grey region plots the 95% point-wise confidence interval for $\bar{\tau}_t^\dagger(1, 0; p)$ based on the conservative estimator of the asymptotic variance of the nonparametric estimator (Theorem 3.2). The solid blue line plots the nonparametric estimator of the total lag- p weighted average dynamic causal effect, $\hat{\tau}^\dagger(1, 0; p)$ and the dashed blue lines plot the 95% confidence interval for $\bar{\tau}^\dagger(1, 0; p)$ based on the conservative estimator of the asymptotic variance of the nonparametric estimator.

	lag- p			
	0	1	2	3
Point estimate, $\hat{\tau}^\dagger(1, 0; p)$	2.580	1.538	1.192	1.568
Conservative p-value	0.000	0.092	0.226	0.127
Randomization p-value	0.007	0.116	0.213	0.107

Table A4: Estimates of the total lag- p weighted average dynamic causal effect for $p = 0, 1, 2, 3$. The conservative p-value reports the p-value associated with testing the weak null hypothesis of no average dynamic causal effects, $H_0 : \bar{\tau}^\dagger(1, 0; 0) = 0$, using the conservative estimator of the asymptotic variance of the nonparametric estimator (Theorem 3.2). The randomization p-value reports the p-value associated with randomization test of the sharp null of dynamic causal effects, $H_0 : \tau_{i,t}(w, \tilde{w}; 0) = 0$ for all $i \in [N], t \in [T]$. The randomization p-values are constructed based on 10,000 draws.

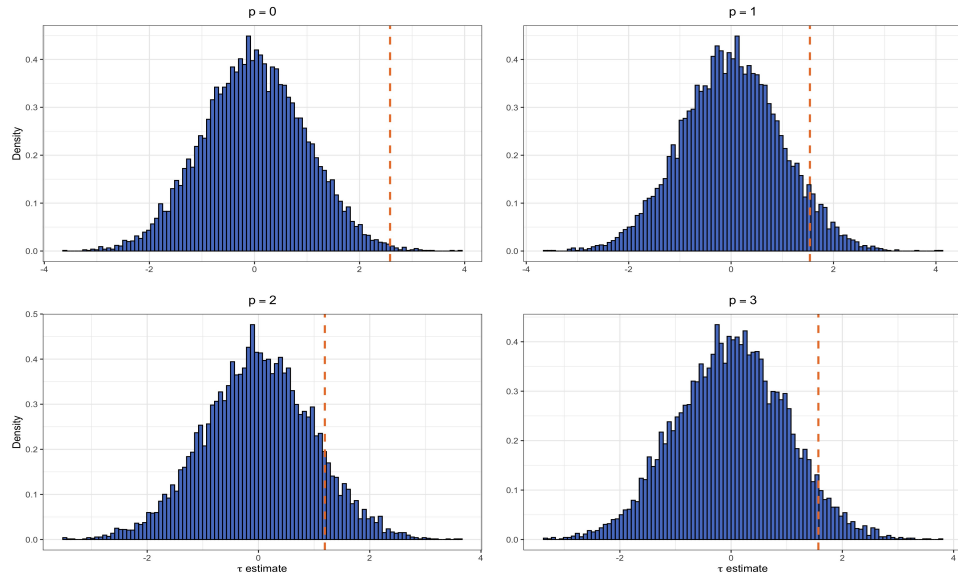


Figure A12: Estimated randomization distribution of the nonparametric estimator of the total lag- p weighted average dynamic causal effect, $\hat{\tau}^\dagger(1, 0; p)$, under the sharp null of no dynamic causal effect, $\tau_{i,t}(w, w; p) = 0$ for all $i \in [N], t \in [T]$. The dashed orange line plots the estimate, $\hat{\tau}^\dagger(1, 0; p)$ at the realized treatments in the experiment of [Andreoni and Samuelson \(2006\)](#). The estimated randomization distributions are constructed based on 10,000 draws.