

A Service of

ZBW

Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre for Economics

Bojinov, Iavor; Rambachan, Ashesh; Shephard, Neil G.

Article

Panel experiments and dynamic causal effects: A finite population perspective

Quantitative Economics

Provided in Cooperation with: The Econometric Society

Suggested Citation: Bojinov, Iavor; Rambachan, Ashesh; Shephard, Neil G. (2021) : Panel experiments and dynamic causal effects: A finite population perspective, Quantitative Economics, ISSN 1759-7331, The Econometric Society, New Haven, CT, Vol. 12, Iss. 4, pp. 1171-1196, https://doi.org/10.3982/QE1744

This Version is available at: https://hdl.handle.net/10419/253613

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.



WWW.ECONSTOR.EU

https://creativecommons.org/licenses/by-nc/4.0/

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.





1759-7331/20211171

Panel experiments and dynamic causal effects: A finite population perspective

IAVOR BOJINOV Technology and Operations Management Unit, Harvard Business School

> ASHESH RAMBACHAN Department of Economics, Harvard University

NEIL SHEPHARD Department of Economics and Department of Statistics, Harvard University

In panel experiments, we randomly assign units to different interventions, measuring their outcomes, and repeating the procedure in several periods. Using the potential outcomes framework, we define finite population dynamic causal effects that capture the relative effectiveness of alternative treatment paths. For a rich class of dynamic causal effects, we provide a nonparametric estimator that is unbiased over the randomization distribution and derive its finite population limiting distribution as either the sample size or the duration of the experiment increases. We develop two methods for inference: a conservative test for weak null hypotheses and an exact randomization test for sharp null hypotheses. We further analyze the finite population probability limit of linear fixed effects estimators. These commonly-used estimators do not recover a causally interpretable estimand if there are dynamic causal effects and serial correlation in the assignments, highlighting the value of our proposed estimator.

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

JEL CLASSIFICATION. C14, C21, C23.

1. Introduction

Panel experiments, where we randomly assign units to different interventions, measuring their response, and repeating the procedure in several periods, 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 eco-

Iavor Bojinov: ibojinov@hbs.edu

Ashesh Rambachan: asheshr@g.harvard.edu

Neil Shephard: shephard@fas.harvard.edu

© 2021 The Authors. Licensed under the Creative Commons Attribution-NonCommercial License 4.0. Available at http://qeconomics.org. https://doi.org/10.3982/QE1744

We thank Isaiah Andrews, Robert Minton, Karthik Rajkumar, and Jonathan Roth for helpful discussions. We thank two anonymous referees for valuable and constructive comments. 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 acknowledges financial support from the NSF Graduate Research Fellowship under Grant DGE1745303.

nomics, 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, 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 commonly focus on incorporating the uncertainty due to sampling units from some superpopulation 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 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 (Neyman (1923), Kempthorne (1955), Cox (1958), Rubin (1974)). Our main estimands are various averages of lag-p dynamic causal effects, which capture how changes in the assignments affect outcomes after p periods. 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.

We develop two methods for conducting nonparametric inference on these dynamic causal effects. The first uses the limiting distribution to perform conservative tests on weak null hypotheses of no average dynamic causal effects. The second provides exact randomization tests for sharp null hypotheses of no dynamic causal effects. We then highlight the usefulness of our framework by deriving the finite population probability limit of commonly used linear estimation strategies, such as the unit fixed effects estimator and the two-way fixed effects estimator. Such estimators are biased for a contemporaneous causal effect whenever there exists carryover effects and serial correlation in the assignment mechanism, underscoring the value of our proposed nonparametric estimator.

Finally, we illustrate our theoretical results in a simulation study and apply our framework to reanalyze a panel-based experiment. The simulation study illustrates our finite population central limit theorems under a variety of assumptions about the underlying potential outcomes and assignment mechanism. We confirm that conservative

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

tests based on the limiting distribution of our nonparametric estimator control size well and have good rejection rates against a variety of alternatives. We finish by reanalyzing a panel experiment conducted in Andreoni and Samuelson (2006), which studies cooperative behavior in game theory and is a natural application of our methods. Participants in the experiment played a twice-repeated prisoners' dilemma many times, and payoff structure of the game was randomly varied across plays. The sequential nature of the experiment raises the possibility that past assignments may impact future actions as participants learn about the structure of the game over time. For example, the random variation in the payoff structure may induce participants to explore possible strategies. This motivates us to analyze the experiment using our methods that are robust to possible dynamic causal effects. 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 such 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 by summarizing heterogeneity across both units and time periods. 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 foundational work by Robins (1986) and co-authors that uses treatment paths for causal panel data analysis and focuses on providing superpopulation (or sampling-based) inference methods. In contrast, we avoid superpopulation arguments entirely. Our estimands and inference procedures are conditioned on the potential outcomes and all uncertainty arises solely from the randomness in assignments. Avoiding superpopulation 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)).² In econometrics, Abadie et al. (2017) highlighted the appeal of this 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. Similarly,

²Of course, in other applications, superpopulation arguments may be entirely natural. 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.

Athey and Imbens (2018), Athey et al. (2018), and Arkhangelsky and Imbens (2019) also introduced a potential outcome model for panel data, but assume away carryover effects. Heckman, Humphries, and Veramendi (2016), Hull (2018), and Han (2019) considered a potential outcome model similar to ours but again rely on superpopulation arguments to perform inference. Additionally, an influential literature in econometrics focuses on estimating dynamic causal effects in panel data under rich models that allow heterogeneity across units, but does not introduce potential outcomes to define counterfactuals and also relies on super-population arguments for inference (e.g., see Arellano and Bonhomme (2016), Arellano, Blundell, and Bonhomme (2017), and the review in Arellano and Bonhomme (2012)).

Notation

For an integer $t \ge 1$ and a variable A_t , we write $A_{1:t} := (A_1, \ldots, A_t)$. We compactly write index sets as $[N] := \{1, \ldots, N\}$ and $[T] := \{1, \ldots, T\}$. Finally, for a 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) are observed over T time periods. For each unit $i \in [N]$ and period $t \in [T]$, we allocate an assignment $W_{i,t} \in W$. The assignment is a random variable and we assume $|W| < \infty$. For a binary assignment $W = \{0, 1\}$, we refer to "1" as treatment and "0" as control.

The assignment path for unit *i* is the sequence of assignments allocated to unit *i*, denoted $W_{i,1:T} := (W_{i,1}, \ldots, W_{i,T})' \in W^T$. The cross-sectional assignment at time-*t* describes all assignments allocated at period *t*, denoted $W_{1:N,t} := (W_{1,t}, \ldots, W_{N,t})' \in W^N$. The assignment panel is the $N \times T$ matrix $W_{1:N,1:T} \in W^{N \times T}$ that summarizes the assignments given to all units over the sample period, where $W_{1:N,1:T} := (W_{1:N,1}, \ldots, W_{1:N,T}) = (W'_{1:1:T}, \ldots, W'_{N,1:T})'$.

A *potential outcome* describes what would be observed for a particular unit at a fixed point in time along any assignment path.

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

In principle, the potential outcome can depend upon the entire assignment path allowing for arbitrary spillovers across time periods. Definition 1 imposes that there are no treatment spillovers across units (Cox (1958)).³

³The 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).

Quantitative Economics 12 (2021)

2.2 The potential outcome panel model

We now define the potential outcomes panel model by restricting the potential outcomes for a unit in a given period not to be affected by future assignments.

ASSUMPTION 1. The potential outcomes are nonanticipating if, for all $i \in [N]$, $t \in [T]$, and $w_{i,1:T}, \tilde{w}_{i,1:T} \in W^T$, $Y_{i,t}(w_{i,1:T}) = Y_{i,t}(\tilde{w}_{i,1:T})$ whenever $w_{i,1:t} = \tilde{w}_{i,1:t}$.

Nonanticipation still allows an arbitrary dependence on past and contemporaneous assignments, and arbitrary heterogeneity across units and time periods.⁴ Under Assumption 1, 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 write the potential outcomes as $Y_{i,t}(w_{i,1:t})$. As notation, let $\mathbf{Y}_{i,t} = \{Y_{i,t}(w_{i,1:t}) : w_{i,1:t} \in \mathcal{W}^t\}$ denote the collection of potential outcomes for unit *i* at time *t* and $\mathbf{Y}_{1:N,1:T} = \{\mathbf{Y}_{i,t} : i \in [N], t \in [T]\}$ denote the collection of potential outcomes for all units across all time periods. Along an assignment panel $w_{1:N,1:t} \in \mathcal{W}^{N \times t}$ up to time *t*, let $Y_{1:N,1:t}(w_{1:N,1:t})$ denote the associated $N \times t$ matrix of outcomes for all units up to time *t*.

To connect the observed outcomes with the potential outcomes, we assume every unit complies with the assignment.⁵ For all $i \in [N]$, $t \in [T]$, the observed outcomes for unit *i* are $y_{i,1:T}^{obs} = Y_{i,1:T}(w_{i,1:T}^{obs})$, where $w_{i,1:T}^{obs}$ is the observed assignment path for unit *i*.

A panel of units, assignments and outcomes in which the units are noninterfering and compliant with the assignments and the outcomes obey Assumption 1 is a *potential outcome panel*. For N = 1, the potential outcome panel reduces to the potential outcome time series model in Bojinov and Shephard (2019). For T = 1, the potential outcome panel reduces to the cross-sectional potential outcome model (e.g., Holland (1986) and Imbens and Rubin (2015)).

2.3 Assignment mechanism assumptions

We focus on randomized experiments in which the assignment mechanisms for each period only depend on past assignments and observed outcomes, but not on future potential outcomes nor unobserved past potential outcomes.

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

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

= $Pr(W_{1:N,t}|W_{1:N,1:t-1} = w_{1:N,1:t-1}, Y_{1:N,1:t-1}(w_{1:N,1:t-1})).$

⁴Allowing for rich heterogeneity in panel data models is often useful in many economic applications. For example, there is extensive heterogeneity across units in income processes (Browning, Ejrnaes, and Alvarez (2010)) and 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; see Ben-Porath (1967), Griliches (1977), and more recently, Cunha et al. (2006) and Cunha, Heckman, and Schennach (2010).

⁵In 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 reinterpreted as focusing on dynamic intention-to-treat (ITT) effects.

It is common to focus on sequentially randomized assignments in biostatistics and epidemiology (Robins (1986), Murphy (2003)). 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)).⁶ Since future potential outcomes and counterfactual past potential outcomes are unobservable, any feasible assignment mechanism must be sequentially randomized.

An important special case imposes further conditional independence structure 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 $\mathbf{Y}_{1:N,1:T}$.

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

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

An individualistic assignment mechanism 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 as well as all other contemporaneous assignments. For example, the Bernoulli assignment 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.

EXAMPLE 1. Consider a food delivery firm that is testing the effectiveness of a new pricing policy across ten major U.S. cities (Kastelman and Ramesh (2018), Sneider and Tang (2018)). Each city is an experimental unit, and the intervention administers the appropriate pricing policy for a duration of one hour. The outcome is the total revenue generated during each hour of the experiment, $t \in [T]$ and from city $i \in [N]$. The firm wishes to learn the best policy for each city and the best overall policy across all cities. To do so, it may conduct a panel experiment with an individualistic treatment assignment in which the probability a particular pricing policy is administered in a given city over the next hour depends on prior observed revenue in that city in earlier hours of the experiment.

REMARK 2.1. Many adaptive experimental strategies (such as the one described in Example 1), in which a series of units are sequentially exposed to random treatments whose probability vary depending on the past observed data, satisfy our individualistic sequentially randomized assignment assumptions (e.g., Robbins (1952), Lai, Leung, and Robbins (1985)). Such experiments are widely used by technology companies to quickly discern user preferences in recommendation algorithms (Li et al. (2010), Li, Karatzoglou, and Gentile (2016)) and by academics interested in improving their power against a particular hypothesis (van der Laan Mark (2008)). There has been a growing interest in drawing causal inferences based on the collected data in such adaptive experimental designs (Hadad et al. (2021), Zhang, Janson, and Murphy (2020)). Since the assignment

⁶If the researcher further observes characteristics $X_{i,t}$ that are causally unaffected by the assignments, then the definition of a sequentially randomized assignment mechanism can be modified to additionally condition on past and contemporaneous values of the characteristics $X_{1:N,1:t}$.

probabilities are known to the researcher, our results can be viewed as providing finite population techniques for drawing causal conclusions from adaptive experiments. In the special case of our framework where N = 1, $t \in [T]$ indexes individuals arriving over time and there no carryover effects, our results in the subsequent section are the finite population analogue of the inference results in Hadad et al. (2021).⁷

Our finite population central limit theorems require that the assignment mechanism be individualistic. In a nonindividualistic assignment mechanism, the past outcomes of other units may affect the contemporaneous assignment of a given unit, which introduces complex dependence structure across units. A similar difficulty arises in the growing literature on relaxing the noninterference assumptions in cross-sectional experiments, where researchers allow one unit's potential outcomes to depend on another unit's assignments (e.g., see Sävje, Aronow, and Hudgens (2019)). To derive the asymptotic distribution of causal estimators in such settings, researchers typically require the assignment mechanism to be independent (Chin (2018)) or at least have only limited dependence structure across units (Aronow and Samii (2017)).

2.4 Dynamic causal effects

A *dynamic causal effect* compares the potential outcomes for unit *i* at time *t* along different assignment paths, which we denote by $\tau_{i,t}(w_{i,1:t}, \tilde{w}_{i,1:t}) := Y_{i,t}(w_{i,1:t}) - Y_{i,t}(\tilde{w}_{i,1:t})$ for assignment paths $w_{i,1:t}, \tilde{w}_{i,1:t} \in W^t$. 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 $0 \le p < t$ and $\mathbf{w}, \tilde{\mathbf{w}} \in \mathcal{W}^{p+1}$, the *i*, *t*th lag- *p* dynamic causal effect is

$$\tau_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p) := \begin{cases} \tau_{i,t} \left(\{ w_{i,1:t-p-1}^{\text{obs}}, \mathbf{w} \}, \{ w_{i,1:t-p-1}^{\text{obs}}, \tilde{\mathbf{w}} \} \right) & \text{if } p < t-1, \\ \tau_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}) & \text{otherwise.} \end{cases}$$

The *i*, *t*th lag-*p* dynamic causal effect measures the difference between the outcomes from following assignment path **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. Generally, when $N \gg T$ we recommend setting p = t - 1, removing the dependence on the observed path.⁸

⁷The setup with N = 1 was developed in Bojinov and Shephard (2019), but this connection to adaptive experiments has not been previously made.

⁸In a time series experiment with N = 1, Bojinov and Shephard (2019) introduced defining causal effects that depend on the observed assignment path because most potential outcomes are unobserved since there is only one experimental unit in their setting. In our more general panel experiments setting, an analogous problem arises when *T* is of a similar order as *N*.

By further restricting the paths **w** and $\tilde{\mathbf{w}}$ to share common features, we obtain the weighted average *i*, *t*th lag-*p* dynamic causal effect.

DEFINITION 5. For integers p, q satisfying $0 \le p < t$, $0 < q \le p + 1$, the weighted average i, tth lag-p, q dynamic causal effect is

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

where $\mathbf{w}, \tilde{\mathbf{w}} \in \mathcal{W}^q$ and $\{a_{\mathbf{v}}\}$ are nonstochastic weights chosen by the researcher that satisfy $\sum_{\mathbf{v} \in \mathcal{W}^{p-q+1}} a_{\mathbf{v}} = 1$ and $a_{\mathbf{v}} \ge 0$ for all $\mathbf{v} \in \mathcal{W}^{p-q+1}$.

The weighted average *i*, *t*th lag-*p*, *q* dynamic causal effect summarizes the *ceteris paribus*, average causal effect of switching the assignment path between period t - p and period t - p + q from **w** to $\tilde{\mathbf{w}}$ on outcomes at time *t*.⁹ In this sense, the weighted average lag-*p*, *q* causal effect is a finite-population causal generalization of an impulse response function, which is a common estimand of interest in existing econometric research.¹⁰ Whenever q = 1, we drop the *q* from the notation, simply writing $\tau_{i,t}^{\dagger}(w, \tilde{w}; p) := \tau_{i,t}^{\dagger}(w, \tilde{w}; p, 1)$.

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 $\mathbf{w}, \tilde{\mathbf{w}} \in \mathcal{W}^{p+1}$,

- 1. the time-t lag-p average dynamic causal effect is $\bar{\tau}_{\cdot t}(\mathbf{w}, \tilde{\mathbf{w}}; p) := \frac{1}{N} \sum_{i=1}^{N} \tau_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p)$.
- 2. the unit-i lag-p average dynamic causal effect is $\bar{\tau}_{i.}(\mathbf{w}, \tilde{\mathbf{w}}; p) := \frac{1}{T-p} \sum_{t=p+1}^{T} \tau_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p)$.
- 3. 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 $\bar{\tau}_{i}^{\dagger}(\mathbf{w}, \tilde{\mathbf{w}}; p, q), \bar{\tau}_{i}^{\dagger}(\mathbf{w}, \tilde{\mathbf{w}}; p, q)$, and $\bar{\tau}^{\dagger}(\mathbf{w}, \tilde{\mathbf{w}}; p, q)$.

We can augment any of the above averages to incorporate nonstochastic weights. For example, we could define $\{c_{i,t}\}_{i=1}^{N}$ the weights and consider the weighted time-*t* lag-*p* average dynamic causal effect $\frac{1}{N} \sum_{i=1}^{N} c_{i,t} \tau_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p)$. These weights, for instance, could be used to adjust for different assignment path probabilities up to time t - p - 1, which are nonstochastic since the assignment mechanism is known.

⁹For a binary assignment, setting N = q = 1 gives us a special case that was studied in Bojinov and Shephard (2019).

¹⁰For time series experiments, Rambachan and Shephard (2020) show that a particular version of the weighted average lag-p, 1 causal effect is equivalent to the generalized impulse response function (Koop, Hashem Pesaran, and Potter (1996)).

3. Nonparametric estimation and inference

In this section, we develop a nonparametric Horvitz and Thompson (1952) type estimator of the *i*, *t*th lag-*p* dynamic causal effects and derive its properties. If the assignment mechanism is individualistic (Definition 3) and probabilistic (defined below), our proposed estimator is unbiased for the *i*, *t*th lag-*p* dynamic causal effects and its related averages over the assignment mechanism. An appropriately scaled and centered version of our estimator for the average lag-*p* dynamic causal effects becomes approximately normally distributed as either the number of units or time periods grows large. These limiting results are finite population central limit theorems in the spirit of Freedman (2008), and Li and Ding (2017).

3.1 Setup: Adapted propensity score and probabilistic assignment

For each *i*, *t*, and any $\mathbf{w} = (w_1, \ldots, w_{p+1}) \in W^{(p+1)}$, the *adapted propensity score* summarizes the conditional probability of a given assignment path and is given by $p_{i,t-p}(\mathbf{w}) :=$ $\Pr(W_{i,t-p:t} = \mathbf{w}|W_{i,1:t-p-1}, Y_{i,1:t}(W_{i,1:t-p-1}, \mathbf{w}))$. Even though the assignment mechanism is known, we only observe the outcomes along the realized assignment path $Y_{i,1:t}(w_{i,1:t}^{obs})$, and so it is not possible to compute $p_{i,t-p}(\mathbf{w})$ for all assignment paths. However, we can compute the adapted propensity score along the observed assignment path, $p_{i,t-p}(w_{i,t-p:t}^{obs})$ (see Appendix B in the Online Supplementary Material (Bojinov, Rambachan, and Shephard (2021) for further discussion).

We next assume that the assignment mechanism is probabilistic.

ASSUMPTION 2 (Probabilistic assignment). Consider a potential outcome panel. There exists C^L , $C^U \in (0, 1)$ such that $C^L < p_{i,t-p}(\mathbf{w}) < C^U$ for all $i \in [N]$, $t \in [T]$ and $\mathbf{w} \in W^{(p+1)}$.

This is also commonly known as the "overlap" or "common support" assumption.

All expectations, denoted by $\mathbb{E}[\cdot]$, 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 $W_{1:N,1:t-p-1}$. Since we 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 $Y_{i,1:t-p-1}(W_{i,1:t-p-1})$.

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

For any $\mathbf{w}, \tilde{\mathbf{w}} \in \mathcal{W}^{(p+1)}$, the nonparametric estimator of $\tau_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p)$ is

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

where $\mathbb{1}\{A\}$ is an indicator function for an event *A*. Under individualistic assignments (Definition 3), the estimator simplifies to $\hat{\tau}_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p) = \frac{y_{i,t}^{\text{obs}}\{\mathbb{1}(w_{i,t-p:t}^{\text{obs}} = \mathbf{w}) - \mathbb{1}(w_{i,t-p:t}^{\text{obs}} = \tilde{\mathbf{w}})\}}{p_{i,t-p}(w_{i,t-p:t}^{\text{obs}})}$.

THEOREM 3.1. Consider a potential outcome panel with an assignment mechanism that is individualistic (Definition 3) and probabilistic (Assumption 2). For any $\mathbf{w}, \tilde{\mathbf{w}} \in \mathcal{W}^{(p+1)}$,

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

$$\operatorname{Var}\left(\hat{\tau}_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p) | \mathcal{F}_{i,t-p-1}\right) = \gamma_{i,t}^{2}(\mathbf{w}, \tilde{\mathbf{w}}) - \tau_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p)^{2} := \sigma_{i,t}^{2},$$

where

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

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

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

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 $\mathcal{F}_{1:N,t-p-1}$.

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 through time and conditionally independent across units. The variance of $\hat{\tau}_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p)$ depends upon the potential outcomes under both the treatment and counterfactual and is generally not estimable. However, its variance is bounded from above by $\gamma_{i,t}^2(\mathbf{w}, \tilde{\mathbf{w}}; p)$, which we can estimate by $\hat{\gamma}_{i,t}^2(\mathbf{w}, \tilde{\mathbf{w}}; p) = \frac{(y_{i,t}^{\text{obs}})^2 \{\mathbbm{1}(w_{i,t-p:t}^{\text{obs}} = \mathbf{w}) + \mathbbm{1}(w_{i,t-p:t}^{\text{obs}} = \mathbf{w})\}}{p_{i,t-p}(w_{i,t-p:t}^{\text{obs}})^2}$. The following proposition establishes that $\hat{\gamma}_{i,t}^2(\mathbf{w}, \tilde{\mathbf{w}}; p)$ is an unbiased estimator of $\gamma_{i,t}^2(\mathbf{w}, \tilde{\mathbf{w}}; p)$ and its error in estimating $\gamma_{i,t}^2(\mathbf{w}, \tilde{\mathbf{w}}; p)$ is also a martingale difference sequence through time and conditionally independent across units.

PROPOSITION 3.1. Under the setup of Theorem 3.1, $\mathbb{E}[\hat{\gamma}_{i,t}^2(\mathbf{w}, \tilde{\mathbf{w}}; p) | \mathcal{F}_{i,t-p-1}] = \gamma_{i,t}^2(\mathbf{w}, \tilde{\mathbf{w}}; p)$. Additionally, $\hat{\gamma}_{i,t}^2(\mathbf{w}, \tilde{\mathbf{w}}; p)$ and $\hat{\gamma}_{j,t}^2(\mathbf{w}, \tilde{\mathbf{w}}; p)$ are independent for $i \neq j$ conditional on $F_{1:N,t-p-1}$.

The variance bound $\gamma_{i,t}^2(\mathbf{w}, \tilde{\mathbf{w}}; p)$ is different from the typical Neyman variance bound, derived under the assumption of a completely randomized experiment (Imbens and Rubin (2015), Chapter 5). In a completely randomized experiment, there is a negative correlation between any two units' assignments since the total number of units assigned to each treatment is fixed. In our setting, all units' assignments are conditionally independent under individualistic assignments, precluding us from exploiting the negative correlation in deriving a bound.

REMARK 3.1. 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 Proposition 3.1. We provide the details for the case when q = 1.

Quantitative Economics 12 (2021)

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

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

Under an individualistic assignment mechanism, this estimator simplifies to

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

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

3.3 Estimation of lag-p average causal effects

The martingale difference properties of the nonparametric estimator means that the averaged plug-in estimators

$$\begin{split} \hat{\bar{\tau}}_{\cdot t}(\mathbf{w}, \tilde{\mathbf{w}}; p) &\coloneqq \frac{1}{N} \sum_{i=1}^{N} \hat{\tau}_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p), \\ \hat{\bar{\tau}}_{i}(\mathbf{w}, \tilde{\mathbf{w}}; p) &\coloneqq \frac{1}{(T-p)} \sum_{t=p+1}^{T} \hat{\tau}_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p), \\ \hat{\bar{\tau}}(\mathbf{w}, \tilde{\mathbf{w}}; p) &\coloneqq \frac{1}{N(T-p)} \sum_{i=1}^{N} \sum_{t=p+1}^{T} \hat{\tau}_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p) \end{split}$$

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. We next derive the limiting distribution of appropriately scaled and centered versions of these averaged estimators.

THEOREM 3.2. Consider a potential outcome panel with an individualistic (Definition 3) and probabilistic assignment mechanism (Assumption 2). Further assume that the potential outcomes are bounded.¹¹ Then, for any $\mathbf{w}, \tilde{\mathbf{w}} \in \mathcal{W}^{(p+1)}$,

$$\frac{\sqrt{N}\left\{\hat{\bar{\tau}}_{.t}(\mathbf{w},\tilde{\mathbf{w}};p)-\bar{\tau}_{.t}(\mathbf{w},\tilde{\mathbf{w}};p)\right\}}{\sigma_{.t}} \stackrel{d}{\to} N(0,1) \quad as N \to \infty,$$

¹¹Assuming the potential outcomes are bounded is a common simplifying assumption made in deriving finite population central limit theorems. As discussed in Li and Ding (2017), this assumption can often be replaced by a finite-population analogue of the Lindeberg condition in analyses of cross-sectional, randomized experiments.

$$\frac{\sqrt{T-p}\left\{\hat{\tau}_{i.}(\mathbf{w},\tilde{\mathbf{w}};p)-\bar{\tau}_{i.}(\mathbf{w},\tilde{\mathbf{w}};p)\right\}}{\sigma_{i.}} \xrightarrow{d} N(0,1) \quad as \ T \to \infty,$$
$$\frac{\sqrt{N(T-p)}\left\{\hat{\tau}(\mathbf{w},\tilde{\mathbf{w}};p)-\bar{\tau}(\mathbf{w},\tilde{\mathbf{w}};p)\right\}}{\sigma} \xrightarrow{d} N(0,1) \quad as \ NT \to \infty,$$

where $\sigma_{\cdot t}$, $\sigma_{i\cdot}$, and σ are the square root of the appropriate averages of $\sigma_{i,t}^2$, defined in *Theorem* 3.1.

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

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

Following the same logic as earlier, we can establish unbiased and consistent estimators of the variance bounds of the averaged estimators.

PROPOSITION 3.2. Under the setup of Theorem 3.2, for any $\mathbf{w}, \tilde{\mathbf{w}} \in \mathcal{W}^{(p+1)}$,

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

Moreover,

$$\begin{aligned} \frac{1}{N}\sum_{i=1}^{N}\hat{\gamma}_{i,t}^{2}(\mathbf{w},\tilde{\mathbf{w}};p) - \frac{1}{N}\sum_{i=1}^{N}\gamma_{i,t}^{2}(\mathbf{w},\tilde{\mathbf{w}};p) \xrightarrow{p} 0 \quad as \ N \to \infty, \\ \frac{1}{(T-p)}\sum_{t=p+1}^{T}\hat{\gamma}_{i,t}^{2}(\mathbf{w},\tilde{\mathbf{w}};p) - \frac{1}{(T-p)}\sum_{t=p+1}^{T}\gamma_{i,t}^{2}(\mathbf{w},\tilde{\mathbf{w}};p) \xrightarrow{p} 0 \quad as \ T \to \infty, \\ \frac{1}{N(T-p)}\sum_{i=1}^{N}\sum_{t=p+1}^{T}\hat{\gamma}_{i,t}^{2}(\mathbf{w},\tilde{\mathbf{w}};p) - \frac{1}{N(T-p)}\sum_{i=1}^{N}\sum_{t=p+1}^{T}\gamma_{i,t}^{2}(\mathbf{w},\tilde{\mathbf{w}};p) \xrightarrow{p} 0 \quad as \ NT \to \infty. \end{aligned}$$

Proposition 3.2 shows that increasing the lag p increases our estimator's variance, highlighting an important trade-off: increasing the lag p reduces the dependence on the observed treatment path at the cost of increased variance. Striking the correct balance depends on the context and the design of the experiment.

Theorem 3.2 and Proposition 3.2 naturally extend to the weighted average i, tth lagp, q dynamic causal effect from Definition 5 by using the estimator developed in Remark 3.1.

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

Combining the variance bound estimators in Proposition 3.2 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 techniques can be used to construct conservative confidence intervals or tests of weak null hypotheses 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.

Alternatively, we may construct exact tests for sharp null hypotheses. An example of such a sharp null hypothesis is $H_0: \bar{\tau}_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p) = 0$, for all, $\mathbf{w}, \tilde{\mathbf{w}}, i \in [N]$, and specific t = 4. Since all potential outcomes are known under such sharp null hypotheses, we can simulate the assignment path $W_{i,t-p:t}|W_{i,1:t-p-1}^{obs}$, $y_{i,1:t-p-1}^{obs}$ for each unit *i* and compute $\hat{\tau}_{i,t}(\mathbf{w}, \tilde{\mathbf{w}}; p)$ at each draw. Therefore, we may simulate the exact distribution of any test statistics under the sharp null hypothesis and compute an exact *p*-value for the observed test statistic. These randomization tests only require 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 (Definition 2).

4. Estimation in a linear potential outcome panel

This section explore the properties of commonly used linear estimators, such as the canonical unit fixed-effects estimator and two-way fixed effects estimator, under the potential outcomes panel model. We establish that if there are dynamic causal effects and serial correlation in the treatment assignment mechanism, both the unit fixed-effects estimator and the two-way fixed effects estimator are asymptotically biased for a weighted average of contemporaneous causal effects. In Appendix B, we consider analyzing the panel experiment as a repeated cross-section, estimating a separate linear model in each period t.

Throughout this section, we further assume that the potential outcomes themselves are a linear function of the assignment path.

DEFINITION 7. A linear potential outcome panel is a potential outcome panel where

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

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

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 $\dot{A}_{i,t} = (A_{i,t} - \bar{A}) - (\bar{A}_{.t} - \bar{A}) - (\bar{A}_{.t} - \bar{A}) - (\bar{A}_{.t} - \bar{A})$.

4.1 Interpreting the unit fixed effects estimator

Our next result characterizes the finite population probability limit of the unit fixed effects estimator, $\hat{\beta}_{UFE} = \sum_{i=1}^{N} \sum_{t=1}^{T} \check{Y}_{i,t} \check{W}_{i,t} / \sum_{i=1}^{N} \sum_{t=1}^{T} \check{W}_{i,t}^{2}$, under the linear potential outcome panel model. Define $\text{Cov}(\check{W}_{i,t}, \check{W}_{i,s} | \mathcal{F}_{1:N,0,T}) := \check{\sigma}_{W,i,t,s}$ and $\check{\mu}_{i,t} := \mathbb{E}[\check{W}_{i,t} | \mathcal{F}_{1:N,0,T}]$.

PROPOSITION 4.1. Assume a linear potential outcome panel and that the assignment mechanism is individualistic (Definition 3) with $\operatorname{Var}(\widetilde{W}_{i,t}|\mathcal{F}_{1:N,0,T}) := \check{\sigma}_{W,i,t}^2 < \infty$ for each $i \in [N], t \in [T]$. Further assume that as $N \to \infty$, the following sequences converge non-stochastically:

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

Then, as $N \rightarrow \infty$ *,*

$$\hat{\beta}_{\text{UFE}} \xrightarrow{p} \frac{\sum\limits_{t=1}^{T} \check{\kappa}_{W,\beta,t,t}}{\sum\limits_{t=1}^{T} \check{\sigma}_{W,t}^{2}} + \frac{\sum\limits_{t=1}^{T} \sum\limits_{s=1}^{t-1} \check{\kappa}_{W,\beta,t,s}}{\sum\limits_{t=1}^{T} \check{\sigma}_{W,t}^{2}} + \frac{\sum\limits_{t=1}^{T} \check{\delta}_{t}}{\sum\limits_{t=1}^{T} \check{\sigma}_{W,t}^{2}}.$$

Proposition 4.1 decomposes the finite population probability limit of the unit fixed effects estimator into three terms. The first term is a weighted 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 assignments causal effects. The last term is an additional error that 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 $\check{Y}_{i,t}(\mathbf{0})$ and the average, demeaned treatment assignment. A sufficient condition for the last term to be equal zero is for the counterfactual outcomes to be time invariant $Y_{i,t}(\mathbf{0}) = \alpha_i$, in which case $\check{Y}_{i,t}(\mathbf{0}) = 0$ for all $i \in [N]$, $t \in [T]$. Therefore, the last term is zero whenever unit fixed effects correctly summarize the variation in the "control-only" counterfactual outcomes across units and time.

Quantitative Economics 12 (2021)

Proposition 4.1 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.

EXAMPLE 2. Consider a linear outcome panel model 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. Assume $\operatorname{Var}(\widetilde{W}_{i,t} | \mathcal{F}_{1:N,0,T}) = \widetilde{\sigma}_{W,t}^2$ for all t and $\operatorname{Cov}(\widetilde{W}_{i,t}, \widetilde{W}_{i,t-1} | \mathcal{F}_{1:N,0,T}) = \widetilde{\sigma}_{W,t,t-1}$ for all t > 1 are constant across units. In this case, Proposition 4.1 implies

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

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 covariance between assignments across periods.

4.2 Interpreting the two-way fixed effects estimator

Consider the two-way fixed-effect estimator is

$$\hat{\beta}_{\text{TWFE}} = \sum_{i=1}^{N} \sum_{t=1}^{T} \dot{\check{Y}}_{i,t} \dot{\check{W}}_{i,t} / \sum_{i=1}^{N} \sum_{t=1}^{T} \dot{\check{W}}_{i,t}^{2}.$$

Define $E(\dot{\widetilde{W}}_{i,t}|\mathcal{F}_{1:N,0,T}) := \dot{\breve{\mu}}_{i,t}$ and $Cov(\dot{\widetilde{W}}_{i,t}, \dot{\widetilde{W}}_{i,s}) := \dot{\breve{\sigma}}_{W,i,t,s}$.

PROPOSITION 4.2. Assume a linear potential outcome panel and assume that the assignment mechanism is individualistic and $\operatorname{Var}(\overset{\cdot}{W}_{i,t}|\mathcal{F}_{1:N,0,T}) := \overset{\cdot}{\sigma}_{W,i,t}^2 < \infty$ for each $i \in [N]$, $t \in [T]$. Further assume that as $N \to \infty$, the following sequences converge nonstochastically:

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

1186 Bojinov, Rambachan, and Shephard

Then, as $N \to \infty$,

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

Similar to Proposition 4.1, the two-way fixed effects estimand can be decomposed into three components under the linear potential outcome panel model, where the interpretation of each component is similar to the unit fixed effects estimator. A simple sufficient condition for the last term to equal zero is for counterfactual outcome to be additively separable into a time-specific and unit-specific effect, $Y_{i,t}(\mathbf{0}) = \alpha_i + \lambda_t$ for all $i \in [N], t \in [T]$. Therefore, the last term is zero whenever unit and time fixed effects are correctly summarize the variation in the "control-only" counterfactual outcomes across units and time.

An active literature in econometrics analyzes the two-way fixed effects estimator under various identifying assumptions. For example, de Chaisemartin and D'Haultfoeuille (2020) ruled out carryover effects and decomposed the two-way fixed effects estimand under a "common-trends" assumption that restricts how the potential outcomes under control evolve over time across groups. Abraham and Sun (2020) decomposed the twoway fixed effects estimand in staggered designs (meaning units receive the treatments at some period and forever after) under a common-trends assumption. Boryusak and Jaravel (2017), Athey and Imbens (2018), and Goodman-Bacon (2018) also provided a decomposition of the two-way fixed effects estimand in staggered designs. Proposition 4.2 provides a decomposition in panel experiments without restrictions on the carryover effects, whereas these existing decompositions are useful in observational settings where other identifying assumptions may be plausible.

5. SIMULATION STUDY

We conduct a simulation study to investigate the finite sample properties of the asymptotic results presented in Section 3. These simulations 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 for the weak null of no average dynamic causal effects have correct size and reasonable rejection rates against a range of alternatives.

5.1 Simulation design

We generate the potential outcomes for the panel experiment using an autoregressive model,

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

 $Y_{i,1}(w_{i,1}) = \beta_{i,1,0}w_{i,1} + \epsilon_{i,1}$ with $\phi_{i,t,1} = \phi$, $\phi_{i,t,s} = 0$ for s > 1, $\beta_{i,t,0} = \beta$ and $\beta_{i,t,s} = 0$ 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 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 either sample from a standard normal or Cauchy distribution.

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 simulate over different assignment panels $W_{1:N,1:T}$, holding the potential outcomes fixed. In the main text, we focus on evaluating the properties of our estimator for the total average dynamic causal effect $\hat{\tau}(1,0;0)$. Appendix C explores the properties of our estimators for the time-*t* average $\hat{\tau}_{.t}(1,0;0)$ and the unit-*i* average $\hat{\tau}_{i.}(1,0;0)$, as well as our estimators of the lag-1 weighted average dynamic causal effects $\hat{\tau}_{.t}^{\dagger}(1,0;1), \hat{\tau}_{i.}^{\dagger}(1,0;1)$, and $\hat{\tau}^{\dagger}(1,0;1)$.

5.2 Normal approximations and size control

Figure 1 plots the randomization distribution for the estimator of the total average dynamic causal effect $\hat{\tau}(1, 0; 0)$. We present results for the case with N = 100, T = 10, and N = 500, T = 100 (the results are similar when the roles of N, T are reversed). When the errors $\epsilon_{i,t}$ are normally distributed, the randomization distribution quickly converges to



FIGURE 1. Simulated randomization distribution for $\hat{\tau}(1, 0; 0)$ under different choices of the parameter ϕ (defined in equation (1)) and treatment probability p(w). The rows index the parameter $\phi \in \{0.25, 0.5, 0.75\}$. The columns index the treatment probability $p(w) \in \{0.25, 0.5, 0.75\}$. 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.

TABLE 1. 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.

		p(w)			p(w)		
		0.25	0.5	0.75	0.25	0.5	0.75
		(a) $\epsilon_{i,t}$ ~	N(0, 1), N = 10	0, $T = 10$	(b) $\epsilon_{i,t} \sim 0$	Cauchy, $N = 50$	00, $T = 100$
ϕ	0.25	0.050	0.047	0.048	0.028	0.029	0.032
	0.5	0.052	0.052	0.050	0.046	0.039	0.044
	0.75	0.050	0.049	0.048	0.055	0.044	0.054

a normal distribution. When the errors are Cauchy distributed, the total number of units and time periods 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). Appendix C provides quantile-quantile plots of the simulated randomization distributions to further illustrate the quality of the normal approximations. Testing based on the normal asymptotic approximation controls size effectively, staying close to the nominal 5% level (see Table 1).

5.3 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.

Figure 2 plots rejection rate curves against the weak null hypotheses $H_0: \bar{\tau}(1, 0; 0) = 0$ and $H_0: \bar{\tau}^{\dagger}(1, 0; 1) = 0$ as the parameter β varies for different choices of the parameter ϕ and treatment probability p(w). The rejection rate against $H_0: \bar{\tau}(1, 0; 0) = 0$ quickly converges to one as β moves away from zero across a range of simulations, indicating that the conservative variance bound still leads to informative tests. When $\phi = 0.25$, the rejection rate against $H_0: \bar{\tau}^{\dagger}(1, 0; 1) = 0$ 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 looks similar.

Appendix C analyzes the rejection rate curves against the weak null hypothesis on the time-*t* average dynamic causal effects with N = 100 units and the unit-*i* average dynamic causal effect with T = 100 time periods. The conservative tests can have low power against these unit-specific or time period-specific weak null hypotheses in small



FIGURE 2. Rejection probabilities for a test of the null hypothesis $H_0: \bar{\tau}(1,0;0) = 0$ 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 $H_0: \bar{\tau}(1,0;0) = 0$ is plotted by the circles and the rejection rate curve against $H_0: \bar{\tau}^{\dagger}(1,0;1) = 0$ is plotted by the triangles. The rows index the parameter $\phi \in \{0.25, 0.5, 0.75\}$. The columns index the treatment probability $p(w) \in \{0.25, 0.5, 0.75\}$. 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.

experiments with few units or few time periods. Unless researchers are analyzing a panel experiment with a large cross-sectional or time dimension, we recommend that researchers focus on analyzing total lag-p dynamic causal effects, which enables them to improve power by pooling information across both units and time periods.

6. Empirical application in experimental economics

We apply our methods to reanalyze a panel experiment from Andreoni and Samuelson (2006) that tests a game-theoretic model of "rational cooperation" and studied how variation in the payoff structure of a two-player, twice-played prisoners' dilemma affects the choices of players.

The payoffs of the game were determined by two parameters $x_1, x_2 \ge 0$ 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 2 summarizes the payoff structure. 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 develop a model of rational cooperation that predicts when λ is large; players will cooperate more often in period one compared to when λ is small.

To investigate this hypothesis, Andreoni and Samuelson (2006) conducted a panelbased experiment. In each session of the experiment, 22 subjects were recruited to play 20 rounds of the twice-played prisoners' dilemma in Table 2. In each round, participants were randomly matched into pairs, and each pair was then randomly assigned

TABLE 2. Stage games from twice-played prisoners' dilemma in the experiment conducted by Andreoni and Samuelson (2006), where the parameters satisfy $x_1, x_2 \ge 0, x_1 + x_2 = 10$, and $\lambda = \frac{x_1}{x_1+x_2}$. The choice *C* denotes "cooperate" and the choice *D* "defect."

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

 $\lambda \in \{0, 0.1, \dots, 0.9, 1\}$ with equal probability. The authors conducted the experiment over five sessions for a total sample of 110 participants and we observe 2200 choices total.

This panel experiment is a natural application of our methods. The sequential nature of the experiment raises the possibility that past assignments may impact future actions as participants learn about the structure of the game over time. For example, random variation in the payoff structure may induce players to explore the strategy space. Additionally, the authors originally analyzed the experiment using regression models with unit-level fixed effects, which may be biased in the presence of dynamic causal effects even if the potential outcomes are linear as discussed in Section 4.

In our analysis, 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 $W \in W = \{0, 1\}$ is binary and equals one whenever the assigned value λ is greater than 0.6, meaning that the payoffs are more concentrated in period two than period one of the stage game. We binarize the assignment in this manner to keep its cardinality (and, therefore, the number of possible assignment paths) manageable, while continuing to test the authors' core prediction on cooperative behavior. For a given pair of subjects, the assignment mechanism is Bernoulli with probability p = 5/11 for treatment and p = 6/11 for control.¹² Table 3 summarizes the observed assignments and observed outcomes in the experiment.

TABLE 3. 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 we observe 2220 choices total.

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

¹²One potential complication that may arise from the subjects playing against each other in the stage game is possible spillovers or interference 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. We ignore this concern in our analysis.

TABLE 4. 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, $H_0 : \overline{\tau}^{\dagger}(1, 0; p) = 0$, 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.

		laş	8- <i>p</i>	
	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

6.1 Inference on total lag-p weighted average dynamic causal effects

We analyze the total lag-p weighted average causal effect $\bar{\tau}^{\dagger}(1,0;p)$ for p = 0, 1, 2, 3, which pools information across all units and time periods to investigate dynamic causal effects.¹³ Based on the conservative test in Section 3.4, the weak null hypothesis $\bar{\tau}^{\dagger}(1,0;0) = 0$ 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 4 summarizes these estimates of the total lag-p weighted average causal effects. Interestingly, the point estimates are positive at p = 1, 2, 3, suggesting there may be dynamic causal effects 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.

We further investigate 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 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 $i \in [N]$, $t \in [T]$.¹⁴ Table 4 summarizes randomization p-values for the total lag-p weighted average causal effects. The p-value for the randomization test at p = 0 is approximately zero, strongly rejecting the sharp null of no contemporaneous dynamic causal effects for all units and again confirming the hypothesis of Andreoni and Samuelson (2006).

¹³Appendix D investigates unit-specific and period-specific weighted average lag-p dynamic causal effects. Since there are only N = 110 units and T = 20 periods in the experiment, these estimates are noisier than our estimates of the total lag-p weighted average dynamic causal effects.

¹⁴When simulating the randomization distribution, we redraw assignment paths 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.

7. Conclusion

This paper developed a potential outcome model for studying dynamic causal effects in a panel experiment. We defined new panel-based dynamic causal estimands such as the lag-p dynamic causal effect and introduced an associated nonparametric estimator. Our proposed estimator is unbiased for lag-p dynamic causal effects over the randomization distribution, and we derived its finite population asymptotic distribution. We developed tools for inference on these dynamic causal effects—a conservative test for weak nulls and an exact randomization test for sharp nulls. We showed that the linear unit fixed effects estimator and two-way fixed effects estimator are asymptotically biased for the contemporaneous causal effects in the presence of dynamic causal effects and persistence in the assignment mechanism. Finally, we illustrated our results through a simulation study and analyzed a panel experiment on rational cooperation in games.

References

Abadie, A., S. C. Athey, G. W. Imbens, and J. Wooldridge (2017), "When should you adjust standard errors for clustering?" NBER Working Paper No. 24003. [1172, 1173]

Abadie, A., S. C. Athey, G. W. Imbens, and J. Wooldridge (2020), "Sampling-based vs. design-based uncertainty in regression analysis." *Econometrica*, 88 (1), 265–296. [1172]

Abraham, S. and L. Sun (2020), "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects." arXiv preprint. Available at arXiv:1804.05785. [1186]

Andreoni, J. and L. Samuelson (2006), "Building rational cooperation." *Journal of Economic Theory*, 127, 117–154. [1173, 1189, 1190, 1191]

Arellano, M., R. Blundell, and S. Bonhomme (2017), "Earnings and consumption dynamics: A nonlinear panel data framework." *Econometrica*, 85, 693–734. [1174, 1175]

Arellano, M. and S. Bonhomme (2012), "Nonlinear panel data analysis." *Annual Review* of *Economics*, 3, 395–424. [1174]

Arellano, M. and S. Bonhomme (2016), "Nonlinear panel data estimation via quantile regressions." *The Econometrics Journal*, 19 (3), C61–C94. [1174]

Arkhangelsky, D. and G. Imbens (2019), "Double-robust identification for causal panel data models." arXiv preprint. Available at arXiv:1909.09412. [1172, 1174]

Aronow, P. M. and C. 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. [1177]

Athey, S., M. Bayati, N. Doudchenko, G. Imbens, and K. Koshravi (2018), "Matrix completion methods for causal panel data models." arXiv preprint. Available at arXiv:1710.10251. [1172, 1174] Quantitative Economics 12 (2021)

Athey, S. and G. Imbens (2018), "Design-based analysis in difference-in-differences settings with staggered adoption." arXiv preprint. Available at arXiv:1808.05293. [1172, 1174, 1186]

Bellemare, C., L. Bissonnette, and S. Kroger (2014), "Statistical power of within and between-subjects designs in economic experiments." IZA Working Paper No. 8583. [1172]

Bellemare, C., L. Bissonnette, and S. Kroger (2016), "Simulating power of economic experiments: The powerBBK package." *Journal of the Economic Science Association*, 2, 157–168. [1172]

Ben-Porath, Y. (1967), "The production of human capital and the life cycle of earnings." *Journal of Political Economy*, 75, 352–365. [1175]

Blackwell, M. and A. Glynn (2018), "How to make causal inferences with time-series and cross-sectional data." *American Political Science Review*, 112, 1067–1082. [1174]

Bojinov, I., A. Rambachan, and N. Shephard (2021), "Supplement to 'Panel experiments and dynamic causal effects: A finite population perspective'." *Quantitative Economics Supplemental Material*, 12, https://doi.org/10.3982/QE1744. [1179]

Bojinov, I., G. Sait-Jacques, and M. Tingley (2020), "Avoid the Pitfalls of A/B Testing." *Harvard Business Review*, 98 (2), 48–53. [1173]

Bojinov, I. and N. Shephard (2019), "Time series experiments and causal estimands: Exact randomization tests and trading." *Journal of the American Statistical Association*, 114 (528), 1665–1682. [1173, 1175, 1177, 1178]

Bojinov, I., D. Simchi-Levi, and J. Zhao (2020), "Design and analysis of switchback experiments." arXiv preprint. Available at arXiv:2009.00148. [1173]

Boruvka, A., D. Almirall, K. Witkiwitz, and S. A. Murphy (2018), "Assessing time-varying causal effect moderation in mobile health." *Journal of the American Statistical Association*, 113, 1112–1121. [1173, 1174]

Boryusak, K. and X. Jaravel (2017), "Revisiting event study designs, with an application to the estimation of the marginal propensity to consume." Available at https://ssrn.com/abstract=2826228. [1186]

Browning, M., M. Ejrnaes, and J. Alvarez (2010), "Modelling income processes with lots of heterogeneity." *Review of Economic Studies*, 77, 1353–1381. [1175]

Charness, G., U. Gneezy, and M. A. Kuhn (2012), "Experimental methods: Betweensubject and within-subject design." *Journal of Economic and Business Organization*, 81 (1), 1–8. [1172]

Chin, A. (2018), "Central limit theorems via Stein's method for randomized experiments under interference." arXiv preprint. Available at arXiv:1804.03105. [1177]

Cox, D. R. (1958), Planning of Experiments. Wiley, Oxford, UK. [1172, 1174]

Cunha, F., J. J. Heckman, L. Lochner, and D. V. Masterov (2006), "Chapter 12 interpreting the evidence on life cycle skill formation." In *Handbook of the Economics of Education*, Vol. 1 (E. Hanushek and F. Welch, eds.), 697–812, Elsevier. [1175]

Cunha, F., J. J. Heckman, and S. M. Schennach (2010), "Estimating the technology of cognitive and noncognitive skill formation." *Econometrica*, 78 (3), 883–931. [1175]

Czibor, E., D. Jimenez-Gomez, and J. A. List (2019), "The dozen things experimental economists should do (more of)." *Southern Economic Journal*, 86 (2), 371–432. [1172]

de Chaisemartin, C. and X. D'Haultfoeuille (2020), "Two-way fixed effects estimators with heterogeneous treatment effects." *American Economic Review*, 110 (9), 2964–2996. [1172, 1186]

Freedman, D. A. (2008), "On regression adjustments to experimental data." *Advances in Applied Mathematics*, 40 (2), 180–193. [1179]

Goodman-Bacon, A. (2018), "Difference-in-differences with variation in treatment timing." NBER Working Paper No. 25018. [1186]

Griliches, Z. (1977), "Estimating the returns to schooling: Some econometric problems." *Econometrica*, 45, 1–22. [1175]

Hadad, V., D. A. Hirshberg, R. Zhan, S. Wager, and S. Athey (2021), "Confidence intervals for policy evaluation in adaptive experiments." *Proceedings of the National Academy of Sciences*, 118 (15), 2014602118. [1176, 1177]

Han, S. (2019), "Identification in nonparametric models for dynamic treatment effects." *Journal of Econometrics*. (Forthcoming). [1174]

Heckman, J. J., J. E. Humphries, and G. Veramendi (2016), "Dynamic treatment effects." *Journal of Econometrics*, 191, 276–292. [1174]

Holland, P. W. (1986), "Statistics and causal inference." *Journal of the American Statistical Association*, 81, 945–960. [1175]

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. [1179]

Hull, P. (2018), "Estimating treatment effects in mover designs." arXiv preprint. Available at arXiv:1804.06721. [1174]

Imai, K. and I. S. 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. [1172, 1185]

Imai, K. and I. S. Kim (2020), "On the use of two-way fixed effects regression models for causal inference with panel data." *Political Analysis*. (Forthcoming). [1172]

Imbens, G. W. and D. B. Rubin (2015), *Causal Inference for Statistics, Social and Biomedical Sciences: An Introduction.* Cambridge University Press, Cambridge, UK. [1173, 1175, 1176, 1180] Kastelman, D. and R. Ramesh (2018), "Switchback tests and randomized experimentation under network effects at DoorDash." https://medium.com/@DoorDash/ switchback-tests-and-randomized-experimentation-under-network-effects-atdoordash-f1d938ab7c2a. [1176]

Kempthorne, O. (1955), "The randomization theory of experimental inference." *Journal of the American Statistical Association*, 50, 946–967. [1172]

Koop, G., M. Hashem Pesaran, and S. M. Potter (1996), "Impulse response analysis in nonlinear multivariate models." *Journal of Econometrics*, 74, 119–147. [1178]

Lai, T. Leung, and H. Robbins (1985), "Asymptotically efficient adaptive allocation rules." *Advances in Applied Mathematics*, 6 (1), 4–22. [1176]

Li, L., W. Chu, J. Langford, and R. E. Schapire (2010), "A contextual-bandit approach to personalized news article recommendation." *Proceedings of the 19th International Conference on World Wide Web*, 661–670. [1176]

Li, S., A. Karatzoglou, and C. Gentile (2016), "Collaborative filtering bandits." *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 539–548. [1176]

Li, X. and P. 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. [1179, 1181]

Lillie, E. O., B. Patay, J. Diamant, B. Issell, E. J. Topol, and N. J. Schork (2011), "The n-of-1 clinical trial: The ultimate strategy for individualizing medicine?" *Personalized Medicine*, 8 (2), 161–173. [1171]

Murphy, S. A. (2003), "Optimal dynamic treatment regimes." *Journal of the Royal Statisticsl Society B*, 65, 331–366. [1176]

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. [1171, 1174]

Neyman, J. (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. [1172]

Rambachan, A. and N. Shephard (2020), "Econometric analysis of potential outcomes time series: Instruments, shocks, linearity and the causal response function." arXiv preprint. Available at arXiv:1903.01637. [1178]

Robbins, H. (1952), "Some aspects of the sequential design of experiments." *Bulletin of the American Mathematical Society*, 58 (5), 527–535. [1176]

Robins, J. 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. [1171, 1173, 1174, 1176] Robins, J. M. (1994), "Correcting for non-compliance in randomization trials using structural nested mean models." *Communications in Statistics—Theory and Methods*, 23, 2379–2412. [1174]

Robins, J. M., S. Greenland, and F.-C. 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. [1174]

Rubin, D. B. (1974), "Estimating causal effects of treatments in randomized and nonrandomized studies." *Journal of Educational Psychology*, 66, 688–701. [1172]

Sävje, F., P. M. Aronow, and M. G. Hudgens (2019), "Average treatment effects in the presence of unknown interference." arXiv preprint. Available at arXiv:1711.06399. [1177]

Sneider, C. and Y. Tang (2018), "Experiment rigor for switchback experiment analysis." https://doordash.engineering/2019/02/20/experiment-rigor-for-switchbackexperiment-analysis/. [1176]

van der Laan Mark, J. (2008), "The construction and analysis of adaptive group sequential designs." Available at https://biostats.bepress.com/ucbbiostat/paper232/. [1176]

Wooldridge, J. M. (2005), "Fixed-effects and related estimators for correlated randomcoefficient and treatment effect panel data models." *Review of Economics and Statistics*, 87, 385–390. [1183]

Zhang, K. W., L. Janson, and S. A. Murphy (2020), "Inference for batched bandits." arXiv preprint. Available at arXiv:2002.03217. [1176]

Co-editor Andres Santos handled this manuscript.

Manuscript received 12 October, 2020; final version accepted 18 May, 2021; available online 28 May, 2021.