

LONGITUDINAL CAUSAL INFERENCE WITH SELECTIVE ELIGIBILITY: APPLICATION TO PRETRIAL RISK ASSESSMENT

BY ZHICHAO JIANG^{1,a}, ELI BEN-MICHAEL^{2,b}, D. JAMES GREINER^{3,c}, RYAN HALEN^{3,d}
AND KOSUKE IMAI^{4,e}

¹*School of Mathematics, Sun Yat-sen University, jiangzhch7@mail.sysu.edu.cn*

²*Department of Statistics & Data Science and Heinz College of Information Systems & Public Policy, Carnegie Mellon University, ebenmichael@cmu.edu*

³*Access to Justice Lab, Harvard Law School, jgreiner@law.harvard.edu, rhale@law.harvard.edu*

⁴*Department of Government, Department of Statistics, Harvard University, imai@harvard.edu*

Dropout presents a major challenge for causal inference in longitudinal studies with time-varying treatments. We focus on selective eligibility—an important but overlooked source of nonignorable dropout that occurs when a unit’s prior treatment history affects its eligibility for future treatments. Our motivating application evaluates a pretrial risk assessment instrument in the criminal justice system, where cases were randomly assigned to receive a risk assessment report. Selective eligibility arises because individuals can be re-arrested over time, and rearrest depends on prior judicial decisions, which may be influenced by earlier treatment assignments. We develop a general methodological framework for longitudinal causal inference under selective eligibility. By conditioning on subgroups of units who would remain eligible for treatment under a given treatment history, we define time-specific eligible treatment effects and the expected total outcome under a given treatment sequence. Assuming a generalized form of sequential ignorability, we derive nonparametric identification formulas and the efficient influence functions for both estimands, yielding doubly robust estimators. Applying this methodology, we find that risk assessment reports increase alignment between judicial decisions and algorithmic recommendations, particularly when reports are provided repeatedly. However, the reports have a limited impact on downstream criminal justice outcomes.

1. Introduction. Dropout presents a major challenge for longitudinal causal inference with time-varying treatments. In many applications, dropout prevents the definition of subsequent treatment and outcome variables, effectively truncating the observed data. This contrasts with conventional censoring in survival analysis, where outcomes are assumed to exist but may be unobserved. While existing methods separately address time-varying treatments and dropout, they typically cannot handle both issues simultaneously. For example, the causal inference literature has established identification formulas and estimation strategies for time-varying treatment effects but without dropout (e.g., [Robins \(1986\)](#), [Robins and Hernán \(2009\)](#), [Richardson and Robins \(2013\)](#), [Hernán and Robins \(2020\)](#)). In addition, the survival analysis literature offers various model-based methods to account for dropout, often referred to as semicompeting events. However, the treatment is assumed to be time-invariant (e.g., [Xu et al. \(2022\)](#), [Baer, Strawderman and Ertefaie \(2023\)](#), [Janvin et al. \(2024\)](#), [Rytgaard and van der Laan \(2024\)](#), [Comment et al. \(2025\)](#)).

When treatments are time-varying *and* dropout is present, existing methods are not directly applicable because standard causal estimands are not well defined for those who have

Received June 2025; revised February 2026.

Key words and phrases. Dropout, efficient influence function, sequential ignorability, stochastic intervention, truncation by death.

dropped out and are no longer eligible for subsequent treatments. In this paper we examine selective eligibility, an important yet overlooked source of nonignorable dropout in a discrete time longitudinal setting. Specifically, we empirically evaluate the impact of a pretrial risk assessment instrument, called the Public Safety Assessment (PSA), in the criminal justice system. We analyze the data from our own randomized controlled trial (RCT) in which the provision of the PSA is randomized across cases (Greiner et al. (2020), Imai et al. (2023)). A primary goal of the PSA is to help judges make release decisions that address concerns about public safety while avoiding unnecessary incarceration during the pretrial period. We investigate how the use of this PSA affects both judicial decision-making and subsequent pretrial outcomes.

The problem of selective eligibility arises in this context because an arrestee would become eligible for the subsequent randomized PSA provision only if they are rearrested. Moreover, rearrest can depend on the judge's previous decision, which in turn may be influenced by the provision of the PSA. Previous analyses of this RCT focused upon first arrest cases to avoid selective eligibility (Imai et al. (2023), Ben-Michael et al. (2025a,b)). We tackle this challenge by developing a general methodological framework for longitudinal causal inference with selective eligibility.

A key feature of selective eligibility is that it causes dropout *before* the administration of treatment. It implies that both treatment and outcome are defined only for treatment-eligible units.¹ Such selective eligibility arises in a range of real-world settings beyond the PSA trial. For instance, in multidose vaccine trials, participants must remain uninfected and medically eligible after receiving an earlier dose to obtain subsequent doses (e.g., Baden et al. (2021)). In a job training program with multiple interventions over time, participants must remain actively engaged in the program to become eligible for later stages (e.g., Andersson et al. (2024)). In school-based behavioral interventions, eligibility for follow-up sessions may depend on continued enrollment or behavioral assessments (e.g., Lochman and Wells (2004)). In all of these examples, one's eligibility status is updated over time, based on prior treatment and outcomes, creating a longitudinal structure in which both treatment assignment and outcome measurement are selectively undefined for ineligible units.

Existing literature primarily addresses a different type of nonignorable treatment-induced dropout, where dropout occurs between treatment administration and outcome measurement. In this case the outcome is observed only for a subset of units that received treatment. This type of dropout is commonly handled using the principal stratification framework (e.g., Kurland et al. (2009), Lee, Daniels and Sargent (2010), Lee and Daniels (2013), Shardell, Hicks and Ferrucci (2015), Shardell and Ferrucci (2018), Josefsson et al. (2016), Josefsson and Daniels (2021), Comment et al. (2025), Grossi et al. (in press)), which extends the classic method for truncation by death problems (e.g., Zhang and Rubin (2003), Rubin (2006), Imai (2008), Ding et al. (2011), Wang, Richardson and Zhou (2017), Yang and Ding (2018)). In addition, we emphasize that among these works, only Josefsson and Daniels (2021) consider time-varying treatments while the others focus on time-invariant treatments.

We develop a general methodological framework for selective eligibility (Section 3). We define the average treatment effect at each time period by focusing on a group of units who become eligible for treatment given a specific treatment history. This causal estimand, which we call the average eligible treatment effect (ETE), characterizes treatment effect heterogeneity across time periods and treatment histories. We also introduce the expected total

¹Boruvka et al. (2018) address the related issue of treatment availability in the context of mobile health applications. They consider the setting where some units may not be able to receive the treatment, even when assigned to the treatment condition. Since these units remain in the study and take up the control condition, this is a problem of noncompliance rather than that of dropout.

outcomes (ETO) under a given treatment history. We consider stochastic interventions for which deterministic interventions are a special case. Stochastic interventions can deal with a relatively large number of time periods because they enable us to specify high-dimensional treatment strategies with low-dimensional parameters (e.g., [Kennedy \(2019\)](#), [Papadogeorgou et al. \(2022\)](#), [Schindl, Shen and Kennedy \(2024\)](#), [Wu et al. \(2024\)](#)).

We develop a statistical methodology for estimating the ETE and ETO by extending the sequential ignorability assumption (see [Hernán and Robins \(2020\)](#), for a review) to longitudinal studies with selective eligibility. Importantly, this assumption does not require the standard overlap condition for the eligibility indicator. Moreover, the ignorability assumption applies to treatment assignment among eligible units and does not extend to the eligibility indicator itself, thereby allowing eligibility status to depend on unobserved variables. Under this assumption we first derive two identification formulae based on the outcome regression and inverse probability weighting. The former leads to an estimator that requires the correct specification of the outcome regression and treatment eligibility models, while the latter implies an estimator based on the correctly specified model for the propensity score. The existence of multiple estimators leads to the question of how to derive a more efficient estimator by combining them.

We address this question by deriving the efficient influence function (EIF; [Bickel et al. \(1993\)](#)) for each causal estimand. The EIFs motivate new estimators that combine nonparametric models of outcome regression, eligibility probability, and propensity score. The proposed estimators are doubly robust; they are consistent so long as either the propensity score model or the outcome regression and eligibility models are correctly specified. This result extends the classic doubly robust estimator ([Bang and Robins \(2005\)](#), [Hernán and Robins \(2020\)](#)) to longitudinal studies with selective eligibility. Section 4 presents simulation studies to evaluate the finite sample performance of the proposed estimators.

Using data from our RCT, we apply the proposed methodology to estimate the effect of PSA provision on: (i) whether the judge's decision agrees with the algorithm's recommendation, (ii) the subsequent negative outcomes, and (iii) the quality of the judge's decision (Section 5). We consider estimating ETEs as well as ETOs under two strategies: one that always provides the PSA and one that never does. Our findings indicate that, first, providing the PSA causes the judge to agree with the algorithmic recommendation more often. Our analysis also yields suggestive evidence that the rate of agreement further increases when the PSA is provided for multiple rearrest incidents for the same arrestee. Second, we find minimal impact on all subsequent negative outcomes. This holds for ETEs across most time periods and treatment histories. Indeed, we find only a small difference in the ETO between the two strategies, further corroborating the null effect of PSA provision. Third, we find little evidence that PSA provision improves the judge's classification ability, as measured by misclassification rates based on the arrestee's underlying risk.

2. Pretrial risk assessment experiment. In this section we describe an RCT of a particular pretrial risk assessment instrument used in the U.S. criminal justice system called the Public Safety Assessment (PSA). We briefly describe the background of this original RCT and discuss methodological challenges that motivate the proposed methodology. [Greiner et al. \(2020\)](#) and [Imai et al. \(2023\)](#) provide additional details of this study.

2.1. Background. In many U.S. jurisdictions, individuals who have been arrested attend first appearance hearings. During this hearing a judge decides whether to release an arrestee, pending the disposition of any criminal charges, and the conditions to impose for their release (e.g., cash bail and supervision). Depending on the jurisdiction, a deciding judge balances a variety of factors, but always at least weighs the costs of incarceration against the risk of

the following negative outcomes if the arrestee is released: (1) failing to appear (FTA) at subsequent court hearings, (2) engaging in a new criminal activity (NCA), and (3) engaging in a new violent criminal activity (NVCA).

The PSA is an algorithmic recommendation system designed to help judges make these pretrial release decisions. It uses nine factors, including an arrestee's age and criminal history, while excluding other demographic factors such as race and gender. The PSA consists of three separate risk scores for FTA, NCA, and NVCA. The information about the formulae used to calculate these scores is publicly available (<https://perma.cc/RE5D-QJZN>). Finally, a decision making framework, which varies from jurisdiction to jurisdiction, combines the risk scores and other factors to produce overall release and monitoring recommendations for the judge.

2.2. The experiment. We conducted a randomized experiment in Dane County, Wisconsin, in which a single judge adjudicated over 80% of hearings, to examine how the PSA affects the judge's pretrial release decisions. We randomized the provision of the PSA across cases, observed the judge's decisions, and collected the outcome variables, that is, the occurrence of an FTA, NCA, or NVCA. Each arrest case entering the system was assigned a unique case number by a court employee. Multiple cases involving the same individual were treated as distinct. Treatment assignment was based on the last digit of the case number: if the digit was even, the case was assigned to the treatment group; otherwise, it was assigned to the control group. For each case assigned to the treatment condition, the judge received a PSA report that contained the three risk scores, the values of the nine factors supporting those scores, and the recommendation generated by the Dane County decision-making framework. For cases randomized to the control condition, the judge received no such report. The occurrence of an FTA, NCA, or NVCA was measured during the pretrial period, defined as the time between the judge's initial release decision and the disposition of the case in the trial court.

An important complication in this analysis is the issue of recidivism. Because treatment was assigned only to arrestees when they attended their first appearance hearing, arrestees who recidivated received the treatment assignment multiple times while nonrecidivating arrestees received it only once. Moreover, the treatment status of a previous case could affect *whether* an individual recidivated, and so the initial and subsequent treatment assignments could affect how frequently arrestees reappeared in the data. Therefore, assessing the criminal justice impacts of providing the PSA required accounting for selective eligibility.

Previous analyses have avoided this issue by focusing on only each individual's first case, limiting their scope to first-time arrests (Imai et al. (2023), Ben-Michael et al. (2025a,b)). However, in our extended dataset of 4,445 arrestees from Dane County, 968 had at least two cases and 315 had at least three cases. Note that individuals could be arrested multiple times for different reasons, each varying in how they related to our three main outcomes of interest. For example, an arrestee who FTA'd could be rearrested for committing a new crime.

Table 1 shows the distribution of the three outcomes for the first, second, and third arrests, where the corresponding percentage is in parentheses. The observed heterogeneity in the outcome distribution across subgroups highlights the diversity across arrestees with different cases. Therefore, restricting the analysis to only first cases jeopardizes the external validity of the experimental results to the overall population and overlooks the heterogeneity across different cases of the arrestees. This heterogeneity is crucial for evaluating the effects of PSA provision and judges' decisions, because arrestees with different numbers of cases may exhibit different behaviors. The issue of recidivism also motivates an exploration of arrestees' behavior over time under specific PSA provision or bail decision strategies. For instance, one might be interested in the total number of crimes an arrestee would commit if a PSA report

TABLE 1

The distribution of negative outcomes for the first, second, and third cases in the treatment and control groups: failure to appear (FTA), new criminal activity (NCA), and new violent criminal activity (NVCA). The table shows the number of cases in each category with the corresponding percentage in parentheses. The percentages are computed with respect to the total number of cases for each arrest under each treatment condition

	PSA (Treatment)				No PSA (Control)			
	FTA	NCA	NVCA	Group total	FTA	NCA	NVCA	Group total
First arrest	462 (20.1)	700 (31.7)	188 (8.5)	2,210 (100)	471 (21.1)	653 (29.2)	184 (8.2)	2,235 (100)
Second arrest	143 (29.3)	210 (43.0)	63 (12.9)	488 (100)	124 (25.8)	205 (42.7)	57 (11.9)	480 (100)
Third arrest	39 (23.8)	87 (53.0)	23 (14.0)	164 (100)	43 (28.5)	68 (45.0)	20 (13.2)	151 (100)

were provided to the judge for each of their cases, relative to the number of crimes if a PSA report were never provided. In the remainder of this paper, we develop a methodological framework to define and estimate causal quantities with selective eligibility.

3. Causal framework under selective eligibility. We now present the proposed methodology for longitudinal causal inference with selective eligibility. We first present assumptions and introduce causal estimands under this setting. We then derive identification formulae and develop efficient estimators.

3.1. Notation and assumptions. We consider a simple random sample of n units from a target population. Suppose that each unit can become eligible for the treatment up to a total of T times. We use S_{it} to indicate whether unit i is eligible for the t th treatment where $i = 1, 2, \dots, n$ and $t = 1, 2, \dots, T$. We assume that all units are eligible for the first-period treatment, that is, $S_{i1} = 1$ for all i . Furthermore, a unit can only become eligible for the $(t + 1)$ th treatment if it was previously eligible for the t th treatment, that is, $S_{it} \geq S_{i,t+1}$ for all $i = 1, 2, \dots, n$ and $t = 1, 2, \dots, T - 1$.

We use Z_{it} to represent the indicator variable for the t th treatment assignment for unit i . This variable is defined only when the unit becomes eligible for the t th treatment, that is, $S_{it} = 1$. We use $\bar{Z}_{it} = (Z_{i1}, \dots, Z_{it})$ to denote the treatment history up to the t th treatment for unit i , and $\bar{z}_t = (z_1, \dots, z_t)$ to denote a treatment sequence from time 1 to time t . The outcome is measured after each treatment. We let Y_{it} denote the outcome for unit i observed after the t th treatment. Similar to the treatment, the outcome Y_{it} is not well defined for units who do not become eligible for the t th treatment, that is, $S_{it} = 0$.

We adopt the potential outcomes framework under the stable unit treatment value assumption (SUTVA; Rubin (1980)). This implies that the treatment status of one unit does not affect the eligibility and outcome of another unit. However, we allow for within-unit carryover effects. That is, for a given unit, the treatment status at any time period may affect its own future eligibility and outcome. Let $S_{it}(\bar{z}_T)$ and $Y_{it}(\bar{z}_T)$ be the potential values of eligibility and outcome, given a treatment sequence $\bar{z}_T = (z_1, \dots, z_T)$, respectively. Then the observed eligibility and outcome are given by $S_{it} = S_{it}(\bar{Z}_{iT})$ and $Y_{it} = Y_{it}(\bar{Z}_{iT})$.

We assume that future treatment does not affect current eligibility and current outcome.

ASSUMPTION 1 (No Anticipation). For any $i = 1, 2, \dots, n$ and $t = 1, 2, \dots, T$, $Y_{it}(\bar{z}_T) = Y_{it}(\bar{z}'_T)$ if $\bar{z}_t = \bar{z}'_t$, and $S_{it}(\bar{z}_T) = S_{it}(\bar{z}'_T)$ if $\bar{z}_{t-1} = \bar{z}'_{t-1}$.




FIG. 1. The temporal ordering of the eligibility indicator S_{it} , treatment assignment Z_{it} , and potential outcome $Y_{it}(\bar{z}_t)$ in an example with two time periods. At each time point, eligibility is measured first. Treatment and outcome variables are only defined when a unit becomes eligible.

Assumption 1 is common in causal inference with time series and panel data (e.g., [Bojinov and Shephard \(2019\)](#), [Ben-Michael, Feller and Rothstein \(2022\)](#), [Papadogeorgou et al. \(2022\)](#), [Jiang, Chen and Ding \(2023\)](#)). The assumption is plausible as long as the units are not aware of their future treatment. Under Assumption 1 we can simplify $Y_{it}(\bar{z}_T)$ as $Y_{it}(\bar{z}_t)$ and $S_{it}(\bar{z}_T)$ as $S_{it}(\bar{z}_{t-1})$.

Figure 1 illustrates the temporal ordering of the eligibility, treatment, and outcome variables with an example of two time periods. In addition, we observe a set of covariates $\mathbf{X}_{it}$ for unit i at time t . These covariates may include both time-invariant and time-varying covariates, including past outcomes. We use $\bar{\mathbf{X}}_{it} = \{\mathbf{X}_{1t}, \dots, \mathbf{X}_{it}\}$ to denote the covariate history up to the t th treatment. Similar to the treatment, the covariate history $\bar{\mathbf{X}}_{it}$ is defined only for units eligible at time t .

The proposed methodology is applicable to both RCTs and observational studies. We generalize the sequential ignorability assumption to the settings with selective eligibility.

ASSUMPTION 2 (Sequential ignorability with selective eligibility). For any $i = 1, 2, \dots, n$ and $t = 1, 2, \dots, T$,

$$Z_{it} \perp\!\!\!\perp \{Y_{it}(\bar{z}_t), S_{i,t+1}(\bar{z}_t), \dots, S_{iT}(\bar{z}_{T-1})\} | \bar{Z}_{i,t-1} = \bar{z}_{t-1}, S_{it} = 1, \bar{\mathbf{X}}_{it} = \bar{\mathbf{x}}_t,$$

for all $\bar{z}_{t-1}, \bar{z}_t, \bar{z}_{t+1}, \dots, \bar{z}_{T-1}$, and $\bar{\mathbf{x}}_t$.

The assumption implies that, for eligible units, the treatment assignment is conditionally independent of the current potential outcome and the potential values of all future eligibility indicator variables, given the treatment and covariate histories. In addition, we generalize the standard overlap condition to the current setting.

ASSUMPTION 3 (Overlap condition with selective eligibility). For any $\bar{z}_{t-1}$ and $\bar{\mathbf{x}}_t$, $0 < \Pr(Z_{it} = 1 | S_{it} = 1, \bar{Z}_{i,t-1} = \bar{z}_{t-1}, \bar{\mathbf{X}}_{it} = \bar{\mathbf{x}}_t) < 1$.

Assumption 3 requires that eligible units at each time t have a nonzero probability of being assigned to both treatment and control.

We do not impose sequential ignorability or an overlap condition on the eligibility indicator S_{it} . Thus, selective eligibility can either be affected by unobserved variables or be a deterministic function of past covariates, treatments, and outcomes—a scenario commonly encountered in clinical studies. In general, observational studies face additional difficulties in satisfying Assumption 3, particularly when the covariate is high dimensional (e.g., [D'Amour et al. \(2021\)](#)). Longitudinal causal inference with selective eligibility also leads to challenges in meeting this assumption, as the dimension of the treatment history grows with time and the number of eligible units may become small for certain treatment histories.

3.2. *Causal estimands.* We first consider a causal effect of the treatment at each time t given a treatment history $\bar{z}_{t-1}$. Recall that the outcome $Y_{it}(\bar{z}_{t-1}, z)$ is defined only for eligible units at time t . Therefore, we introduce the following average eligible treatment effect (ETE):

$$(1) \quad \tau_t(\bar{z}_{t-1}) = \mathbb{E}\{Y_{it}(\bar{z}_{t-1}, 1) - Y_{it}(\bar{z}_{t-1}, 0) | S_{it}(\bar{z}_{t-1}) = 1\}.$$

The ETE represents the average effect of the t th treatment among units who are eligible at time t when the treatment history is fixed at $\bar{Z}_{i,t-1} = \bar{z}_{t-1}$. Due to the potential carryover effects resulting from past treatments, the ETE depends on the entire vector of previous treatments $\bar{z}_{t-1}$, yielding a separate causal quantity for each of 2^{t-1} different values of $\bar{z}_{t-1}$.

We also consider the expected total outcome (ETO) across time periods,

$$\theta(\bar{z}_T) = \sum_{t=1}^T \mathbb{E}\{Y_{it}(\bar{z}_t) S_{it}(\bar{z}_{t-1})\},$$

which represents the sum of potential outcomes among eligible units averaged across units under a treatment sequence $\bar{Z}_T = \bar{z}_T$. This estimand captures the total amount of outcomes when a unit is eligible. For binary or count outcomes, it corresponds to the expected number of events accumulated over eligible time points. For continuous outcomes, such as monetary loss or duration of detention, it represents the expected total cost incurred while the unit remains eligible. The magnitude of this estimand depends on both the frequency of the outcome and the duration of eligibility. As such, it may be small if ineligibility tends to occur early or if the outcome is rare or infrequent.

We can generalize the ETO for various treatment strategies. Let ξ denote a treatment strategy, which is defined as a probability distribution over the sequence of treatments. Thus, ξ may represent either a deterministic or stochastic intervention. We use $\bar{Z}_{iT}^\xi = (Z_{i1}^\xi, \dots, Z_{iT}^\xi)$ to denote a sequence of interventions on unit i that is generated according to the treatment strategy ξ for the duration of T time periods. At each time period t , the treatment strategy sequentially generates Z_{it}^ξ from a Bernoulli distribution conditional on the treatment history $\bar{Z}_{i,t-1}^\xi = (Z_{i1}^\xi, \dots, Z_{i,t-1}^\xi)$ for $t = 2, \dots, T$. Therefore, the treatment strategy can be characterized by the parameters of these Bernoulli distributions, denoted by a vector of probabilities $(\xi_1(z_1), \xi_2(\bar{z}_2), \dots, \xi_T(\bar{z}_T))$, where $\xi_t(\bar{z}_t)$ denotes the probability that the stochastic intervention assigns $Z_{it}^\xi = z_t$ conditional on $\bar{Z}_{i,t-1}^\xi = \bar{z}_{t-1}$.

Due to selective eligibility, the stochastic intervention is administered only for treatment-eligible units; that is, the intervention Z_{it}^ξ is generated for unit i only if $S_{it}(\bar{Z}_{i,t-1}^\xi) = 1$. This implies that we can characterize a treatment strategy as the following factorized probability:

$$\xi_t(\bar{z}_t) = \Pr(Z_{i1}^\xi = z_1) \prod_{t'=1}^t \Pr(Z_{it'}^\xi = z_{t'} | \bar{Z}_{i,t'-1}^\xi = \bar{z}_{t'-1}, S_{it'} = 1).$$

Note that a deterministic treatment strategy is a special case of this stochastic intervention where the probability is one for a specific sequence $\bar{z}_t$ and is zero otherwise. Given the above definition, we can define the ETO under a treatment strategy ξ ,

$$(2) \quad \begin{aligned} \theta(\xi) &= \sum_{t=1}^T \mathbb{E}\{Y_{it}(\bar{Z}_{it}^\xi) S_{it}(\bar{Z}_{i,t-1}^\xi)\} \\ &= \sum_{t=1}^T \sum_{\bar{z}_t} \mathbb{E}\{Y_{it}(\bar{z}_t) | S_{it}(\bar{z}_{t-1}) = 1\} \Pr\{S_{it}(\bar{z}_{t-1}) = 1\} \xi_t(\bar{z}_t). \end{aligned}$$

The ETE and ETO are related to the treatment effects defined by [Young et al. \(2020\)](#), [Rytgaard, Gerds and van der Laan \(2022\)](#), and [Díaz, Hoffman and Hejazi \(2024\)](#). These studies

focus on longitudinal settings in which censoring or truncation events render subsequent outcomes undefined. However, there are two key differences. First, in these studies the censoring or truncation process occurs after treatment and hence may be affected by it. In contrast, in our context the eligibility indicator is determined prior to treatment assignment at each time point and is, therefore, not influenced by the treatment. Second, the ETE and ETO are defined explicitly among the eligible units at each time—that is, those for whom both treatment and outcome are well defined. In contrast, the aforementioned studies define and identify causal effects for the entire population under counterfactual scenarios in which the censoring or truncation process follows a different intervention regime. As a result, their identification strategies rely on sequential ignorability assumptions for the censoring or truncation process, which are not required in our approach.

3.3. Nonparametric identification. We now turn to the nonparametric identification of the causal estimands defined above. We focus on the identification of $\mathbb{E}\{Y_{it}(\bar{z}_t)|S_{it}(\bar{z}_{t-1}) = 1\}$ and $\Pr\{S_{it}(\bar{z}_{t-1}) = 1\}$, which in turn implies the identification of $\tau_t(\bar{z}_{t-1})$ and $\theta(\xi)$ via equations (1) and (2).

We begin by introducing the following quantities to simplify the presentation:

$$\begin{aligned}\mu_t(\bar{z}_t, \bar{\mathbf{x}}_t) &:= \mathbb{E}(Y_{it}|S_{it} = 1, \bar{\mathbf{Z}}_{it} = \bar{z}_t, \bar{\mathbf{X}}_{it} = \bar{\mathbf{x}}_t), \\ p_t(\bar{z}_{t-1}, \bar{\mathbf{x}}_{t-1}) &:= \Pr(S_{it} = 1|S_{i,t-1} = 1, \bar{\mathbf{Z}}_{i,t-1} = \bar{z}_{t-1}, \bar{\mathbf{X}}_{i,t-1} = \bar{\mathbf{x}}_{t-1}), \\ w_t(\bar{z}_t, \bar{\mathbf{x}}_t) &:= \Pr(Z_{it} = z_t|S_{it} = 1, \bar{\mathbf{Z}}_{i,t-1} = \bar{z}_{t-1}, \bar{\mathbf{X}}_{it} = \bar{\mathbf{x}}_t), \\ l_t(\bar{z}_{t-1}, \bar{\mathbf{x}}_t) &:= \Pr(X_{it} = \mathbf{x}_t|S_{it} = 1, \bar{\mathbf{Z}}_{i,t-1} = \bar{z}_{t-1}, \bar{\mathbf{X}}_{i,t-1} = \bar{\mathbf{x}}_{t-1}),\end{aligned}$$

which represent the conditional mean outcome, eligibility probability, propensity score, and covariate distribution, respectively, for eligible units at each time t , given the treatment and covariate histories up to time $t - 1$. When $t = 1$, we have $p_t(\bar{z}_{t-1}, \bar{\mathbf{x}}_{t-1}) = \Pr(S_{i1} = 1) = 1$, $l(\bar{z}_{t-1}, \bar{\mathbf{x}}_t) = \Pr(X_{i1} = \mathbf{x}_1)$, and $w_t(\bar{z}_t, \bar{\mathbf{x}}_t) = \Pr(Z_{i1} = z_1|X_{i1} = \mathbf{x}_1)$.

The following theorem provides two different identification formulae, based on different components of the observed data distribution, for $\mathbb{E}\{Y_{it}(\bar{z}_{t-1}, z_t)|S_{it}(\bar{z}_{t-1}) = 1\}$ and $\Pr\{S_{it}(\bar{z}_{t-1}) = 1\}$.

THEOREM 1 (Nonparametric Identification). *Under Assumptions 1–3, the following identification formulae hold for $\mathbb{E}\{Y_{it}(\bar{z}_{t-1}, z_t)|S_{it}(\bar{z}_{t-1}) = 1\}$ and $\Pr\{S_{it}(\bar{z}_{t-1}) = 1\}$:*

(a) *Formula based on the conditional mean outcome and eligibility probability:*

$$\begin{aligned}\mathbb{E}\{Y_{it}(\bar{z}_t)|S_{it}(\bar{z}_{t-1}) = 1\} &= \frac{\int \mu_t(\bar{z}_t, \bar{\mathbf{x}}_t) \prod_{t'=2}^t p_{t'}(\bar{z}_{t'-1}, \bar{\mathbf{x}}_{t'-1}) \prod_{t'=1}^{t-1} l_{t'}(\bar{z}_{t'-1}, \bar{\mathbf{x}}_{t'}) d\bar{\mathbf{x}}_t}{\int \prod_{t'=2}^t p_{t'}(\bar{z}_{t'-1}, \bar{\mathbf{x}}_{t'-1}) \prod_{t'=1}^{t-1} l_{t'}(\bar{z}_{t'-1}, \bar{\mathbf{x}}_{t'}) d\bar{\mathbf{x}}_{t-1}}, \\ \Pr\{S_{it}(\bar{z}_{t-1}) = 1\} &= \int \prod_{t'=2}^t p_{t'}(\bar{z}_{t'-1}, \bar{\mathbf{x}}_{t'-1}) \prod_{t'=1}^{t-1} l_{t'}(\bar{z}_{t'-1}, \bar{\mathbf{x}}_{t'}) d\bar{\mathbf{x}}_{t-1},\end{aligned}$$

where integrals are replaced by summation for discrete covariates.

(b) *Formula based on the propensity score:*

$$\begin{aligned}\mathbb{E}\{Y_{it}(\bar{z}_{t-1}, z_t)|S_{it}(\bar{z}_{t-1}) = 1\} &= \frac{\mathbb{E}[Y_{it}\mathbf{1}\{\bar{\mathbf{Z}}_{it} = \bar{z}_t, S_{it} = 1\}/\pi_t(\bar{z}_t, \bar{\mathbf{X}}_{it})]}{\mathbb{E}[\mathbf{1}\{\bar{\mathbf{Z}}_{i,t-1} = \bar{z}_{t-1}, S_{it} = 1\}/\pi_{t-1}(\bar{z}_{t-1}, \bar{\mathbf{X}}_{i,t-1})]}, \\ \Pr\{S_{it}(\bar{z}_{t-1}) = 1\} &= \mathbb{E}\left[\frac{\mathbf{1}\{\bar{\mathbf{Z}}_{i,t-1} = \bar{z}_{t-1}, S_{it} = 1\}}{\pi_{t-1}(\bar{z}_{t-1}, \bar{\mathbf{X}}_{i,t-1})}\right],\end{aligned}$$

where $\pi_t(\bar{z}_t, \bar{\mathbf{X}}_{it}) = \prod_{t'=1}^t w_{t'}(\bar{z}_{t'}, \bar{\mathbf{X}}_{it'})$.

Theorem 1 does not require the sequential ignorability assumption for S_{it} . As a result, the identification formulas remain valid, even when the eligibility status of each unit depends on past outcomes or unobserved variables. This is because the estimands are defined specifically for the eligible subpopulation and do not involve extrapolation to ineligible units.

We provide some intuition for the results of Theorem 1 by focusing on $\mathbb{E}\{Y_{it}(\bar{z}_{t-1}, z_t) | S_{it}(\bar{z}_{t-1}) = 1\}$. The first identification formula, given in Theorem 1(a), is based on the models for the conditional mean outcome and eligibility probability. For example, when $t = 2$, this formula can be simplified as

$$\mathbb{E}\{Y_{i2}(\bar{z}_2) | S_{i2}(z_1) = 1\} = \frac{\mathbb{E}[\mathbb{E}\{\mu_2(\bar{z}_2, \bar{X}_{i2}) | z_1, \mathbf{X}_{i1}\} p_2(z_1, \mathbf{X}_{i1})]}{\mathbb{E}\{p_2(z_1, \mathbf{X}_{i1})\}},$$

where the covariates are integrated out. Therefore, we can view this formula as a g-formula for longitudinal causal inference with selective eligibility. In contrast, the first equation given in Theorem 1(b) presents the IPW formula, relying only on the propensity score at each time, given the treatment and covariate histories $w_t(\bar{z}_t, \bar{\mathbf{x}}_t)$. In both formulae the numerator equals $\mathbb{E}\{Y_{it}(\bar{z}_{t-1}, z_t) S_{it}(\bar{z}_{t-1})\}$, and the denominator equals $\Pr\{S_{it}(\bar{z}_{t-1}) = 1\}$.

As noted earlier, Theorem 1 directly implies the identification of $\tau_t(\bar{z}_{t-1})$ and $\theta(\xi)$. The identification of $\tau_t(\bar{z}_{t-1})$ requires the overlap condition (Assumption 3) to hold only for the corresponding treatment history $\bar{\mathbf{Z}}_{i,t-1} = \bar{z}_{t-1}$. Similarly, the identification of $\theta(\xi)$ only requires the overlap condition to hold for all $\bar{z}_t$ with $\xi(\bar{z}_t) \neq 0$.

The identification formulae in Theorem 1 have simpler forms when all the covariates are time invariant. We present this special case as a corollary.

COROLLARY 1 (Identification Formulae without Time-varying Covariates). *Under Assumptions 1—3 with $X_{it} = X_i$ for all i and t , the identification formulae in Theorem 1(a) simplify as*

$$\begin{aligned} \mathbb{E}\{Y_{it}(\bar{z}_t) | S_{it}(\bar{z}_{t-1}) = 1\} &= \frac{\mathbb{E}\{\mu_t(\bar{z}_t, \mathbf{X}_i) \prod_{t'=2}^t p_{t'}(\bar{z}_{t'-1}, \mathbf{X}_i)\}}{\mathbb{E}\{\prod_{t'=2}^t p_{t'}(\bar{z}_{t'-1}, \mathbf{X}_i)\}}, \\ \Pr\{S_{it}(\bar{z}_{t-1}) = 1\} &= \mathbb{E}\left\{\prod_{t'=2}^t p_{t'}(\bar{z}_{t'-1}, \mathbf{X}_i)\right\}. \end{aligned}$$

Unlike the formulae in Theorem 1(a), without time-varying covariates, the formulae in Corollary 1 no longer include the covariate distribution terms $l_t(\bar{z}_{t-1}, \bar{\mathbf{x}}_t)$. In addition, the formulae in Theorem 1(b) do not change, except that $\bar{X}_{it}$ is replaced with X_i .

3.4. Efficient influence functions. The identification formulae in Theorem 1 imply the existence of multiple estimators for the causal estimands. To combine them in a principled way, we derive the efficient influence functions (EIFs) of the estimands, which in turn yield estimators with certain robustness and optimality properties. These estimators can achieve parametric convergence rates even when nuisance functions are estimated at slower rates through nonparametric estimation methods (see, e.g., Bickel et al. (1993), Tsiatis (2006), Kennedy (2016)).

We begin by introducing the following two regression functions that are building blocks of the EIFs and deriving their identification formulae:

$$\begin{aligned} m_{Y_{it}}(\bar{z}_{t'}, \bar{X}_{it'}) &:= \mathbb{E}\{Y_{it}(\bar{z}_t) S_{it}(\bar{z}_{t-1}) | S_{it'} = 1, \bar{\mathbf{Z}}_{it'} = \bar{z}_{t'}, \bar{\mathbf{X}}_{it'}\}, \quad t' \leq t, \\ m_{S_{it}}(\bar{z}_{t'}, \bar{X}_{it'}) &:= \mathbb{E}\{S_{it}(\bar{z}_{t-1}) | S_{it'} = 1, \bar{\mathbf{Z}}_{it'} = \bar{z}_{t'}, \bar{\mathbf{X}}_{it'}\}, \quad t' \leq t-1. \end{aligned}$$

When $t' = 0$, we have $m_{Y_{it}} = \mathbb{E}\{Y_{it}(\bar{z}_t) S_{it}(\bar{z}_{t-1})\}$ and $m_{S_{it}} = \mathbb{E}\{S_{it}(\bar{z}_{t-1})\}$.

PROPOSITION 1 (Identification of Regression Functions). *Under Assumptions 1–3, we have the following identification formulae:*

$$m_{Y_t S_t}(\bar{z}_{t'}, \bar{X}_{it'}) = \begin{cases} \mathbb{E}(Y_{it} | S_{it} = 1, \bar{Z}_{it} = \bar{z}_{it}, \bar{X}_{it}) & t' = t, \\ \mathbb{E}\{S_{i,t'+1} m_{Y_t S_t}(\bar{z}_{t'+1}, \bar{X}_{i,t'+1}) | S_{it'} = 1, \bar{Z}_{it'} = \bar{z}_{t'}, \bar{X}_{it'}\} & t' < t. \end{cases}$$

$$m_{S_t}(\bar{z}_{t'}, \bar{X}_{it'}) = \begin{cases} \mathbb{E}(S_{it} | S_{i,t-1} = 1, \bar{Z}_{i,t-1} = \bar{z}_{t-1}, \bar{X}_{i,t-1}) & t' = t - 1, \\ \mathbb{E}\{S_{i,t'+1} m_{S_t}(\bar{z}_{t'+1}, \bar{X}_{i,t'+1}) | S_{it'} = 1, \bar{Z}_{it'} = \bar{z}_{t'}, \bar{X}_{it'}\} & t' < t - 1. \end{cases}$$

Proposition 1 shows how to obtain the identification formula for $m_{Y_t S_t}(\bar{z}_{t'}, \bar{X}_{it'})$ for any $t' \leq t$. Specifically, we first obtain the formula for $m_{Y_t S_t}(\bar{z}_t, \bar{X}_{it})$ (i.e., $t' = t$) and then sequentially derive the formula for $m_{Y_t S_t}(\bar{z}_{t'}, \bar{X}_{it'})$ for $t' = t - 1, t - 2, \dots, 1, 0$. Similarly, we can obtain the identification formulae for $m_{S_t}(\bar{z}_{t'}, \bar{X}_{it'})$ for all $t' \leq t - 1$. When $t' = 0$, the identification formula for m_{S_t} coincides with that of $\Pr\{S_{it}(\bar{z}_{t-1}) = 1\}$, while the identification formula for the ratio of $m_{Y_t S_t}$ and m_{S_t} equals the formula for $\mathbb{E}\{Y_{it}(\bar{z}_t) | S_{it}(\bar{z}_{t-1}) = 1\}$. Both formulae are given in Theorem 1(a).

Given Proposition 1, the following theorem presents the EIFs of $\tau_t(\bar{z}_{t-1})$ and $\theta(\xi)$.

THEOREM 2 (Efficient Influence Functions). *Under Assumptions 1–3, the EIFs of $\tau_t(\bar{z}_{t-1})$ and $\theta(\xi)$ are given by*

$$(3) \quad \phi_\tau = \frac{\phi_N(\bar{z}_{t-1}, 1) - \phi_N(\bar{z}_{t-1}, 0) - \tau_t(\bar{z}_{t-1}) \cdot \phi_D(\bar{z}_{t-1})}{\Pr\{S_{it}(\bar{z}_{t-1}) = 1\}},$$

$$(4) \quad \phi_\theta = \sum_{t=1}^T \sum_{\bar{z}_t} \phi_N(\bar{z}_t) \xi_t(\bar{z}_t) - \theta(\xi),$$

respectively, where

$$\begin{aligned} \phi_N(z_t) &= m_{Y_t S_t}(z_1, X_{i1}) + \frac{\mathbf{1}\{\bar{Z}_{it} = \bar{z}_t\} S_{it} \{Y_{it} - m_{Y_t S_t}(\bar{z}_{it}, \bar{X}_{it})\}}{\pi_t(\bar{z}_t, \bar{X}_{it})} \\ &\quad + \sum_{t'=2}^t \frac{\mathbf{1}\{\bar{Z}_{i,t'-1} = \bar{z}_{t'-1}\} \{S_{it'} m_{Y_t S_t}(\bar{z}_{t'}, \bar{X}_{it'}) - S_{i,t'-1} m_{Y_t S_t}(\bar{z}_{t'-1}, \bar{X}_{i,t'-1})\}}{\pi_{t'-1}(\bar{z}_{t'-1}, \bar{X}_{i,t'-1})}, \\ \phi_D(z_{t-1}) &= m_{S_t}(z_1, X_{i1}) + \frac{\mathbf{1}\{\bar{Z}_{i,t-1} = \bar{z}_{t-1}\} \{S_{it} - S_{i,t-1} m_{S_t}(\bar{z}_{t-1}, \bar{X}_{i,t-1})\}}{\pi_{t-1}(\bar{z}_{t-1}, \bar{X}_{i,t-1})} \\ &\quad + \sum_{t'=2}^{t-1} \frac{\mathbf{1}\{\bar{Z}_{i,t'-1} = \bar{z}_{t'-1}\} \{S_{it'} m_{S_t}(\bar{z}_{t'}, \bar{X}_{it'}) - S_{i,t'-1} m_{S_t}(\bar{z}_{t'-1}, \bar{X}_{i,t'-1})\}}{\pi_{t'-1}(\bar{z}_{t'-1}, \bar{X}_{i,t'-1})}. \end{aligned}$$

The identification formulae of $m_{Y_t S_t}(\bar{z}_{t'}, \bar{X}_{it'})$ and $m_{S_t}(\bar{z}_{t'}, \bar{X}_{it'})$ are given in Proposition 1.

We can show that $\phi_N(\bar{z}_t) - m_{Y_t S_t}$ and $\phi_D(\bar{z}_{t-1}) - m_{S_t}$ are the EIFs of $\mathbb{E}\{Y_{it}(\bar{z}_{t-1}, z_t) S_{it}(\bar{z}_{t-1})\}$ and $\Pr\{S_{it}(\bar{z}_{t-1}) = 1\}$, respectively. The EIFs do not include terms corresponding to the perturbation of treatment distribution. Therefore, like the case of the average treatment effect estimation, knowing the propensity score does not affect the EIFs (Hahn (1998)).

Theorem 2 generalizes the EIFs under the standard longitudinal causal inference setting (van der Laan and Gruber (2012), Kennedy (2019), Díaz et al. (2023)), where all units are eligible for all treatments. In this case, $\tau_t(\bar{z}_{t-1})$ reduces to the effect of assigning treatment at time t versus control at time t , given the prior treatment history $\bar{z}_t - 1$. Correspondingly, ϕ_τ

reduces to the EIF in unconfounded longitudinal studies. In addition, if there is only one time period $T = 1$, $\tau_t(\bar{z}_{t-1})$ reduces to the average treatment effect in a single time point, and ϕ_τ becomes the classic doubly robust formula (Hahn (1998)).

Given these EIFs, we can solve $\mathbb{E}(\phi_\tau) = 0$ and $\mathbb{E}(\phi_\theta) = 0$ to obtain the identification formulae for $\tau_t(\bar{z}_{t-1})$ and $\theta(\xi)$, which are given in the following corollary.

COROLLARY 2 (Identification formulae based on the efficient influence functions). *Under Assumptions 1–3,*

$$\tau_t(\bar{z}_{t-1}) = \frac{\mathbb{E}\{\phi_N(\bar{z}_{t-1}, 1) - \phi_N(\bar{z}_{t-1}, 0)\}}{\mathbb{E}\{\phi_D(\bar{z}_{t-1})\}},$$

$$\theta(\xi) = \sum_{t=1}^T \sum_{\bar{z}_t} \mathbb{E}\{\phi_N(\bar{z}_t)\} \xi_t(\bar{z}_t).$$

When all the covariates are time invariant, the EIFs in Theorem 2 and identification formulae in Corollary 2 can also be simplified by replacing the formulae for $m_{Y_t S_t}(\bar{z}_{t'}, \bar{X}_{it'})$ and $m_{S_t}(\bar{z}_{t'}, \bar{X}_{it'})$ in Proposition 1 with the following:

$$(5) \quad m_{Y_t S_t}(\bar{z}_{t'}, \bar{X}_{it'}) = \mu_t(\bar{z}_t, \mathbf{X}_i) \prod_{s=t'}^t p_s(\bar{z}_{s-1}, \mathbf{X}_i),$$

$$(6) \quad m_{S_t}(\bar{z}_{t'}, \bar{X}_{it'}) = \prod_{s=t'}^t p_s(\bar{z}_{s-1}, \mathbf{X}_i),$$

where $\mathbf{X}_i$ is the set of time-invariant covariates. Unlike those shown in Theorem 2, these formulae depend only on the conditional outcome mean and eligibility probabilities. As a result, they do not need to be iteratively defined.

3.5. Estimation and inference. Using the EIFs derived above, we construct estimators that have a double robustness property and attain parametric convergence rates, even when nuisance functions are estimated using flexible machine learning algorithms. We also introduce alternative simple estimators that do not require models for the outcomes and eligibility probabilities.

3.5.1. Estimators based on the EIFs. Corollary 2 motivates estimating the causal quantities by replacing the nuisance functions with their estimates. We denote the resulting estimators as $\hat{\tau}_t(\bar{z}_{t-1})$ for $\tau_t(\bar{z}_{t-1})$ and $\hat{\theta}(\xi)$ for $\theta(\xi)$, respectively. Detailed expressions for these estimators are provided in the Supplementary Material (Jiang et al. (2026)).

We discuss how to estimate the nuisance functions. Because $\pi_t(\bar{z}_t, \bar{X}_{it}) = \prod_{s=1}^t w_s(\bar{z}_s, \bar{X}_{is})$, we estimate $w_s(\bar{z}_s, \bar{X}_{is})$ for each s by modeling Z_{is} , given $(\bar{Z}_{i,s-1}, \bar{X}_{is})$ for units with $S_{is} = 1$, and then multiplying these estimates across s to obtain $\hat{\pi}_t(\bar{z}_t, \bar{X}_{it})$. For $\hat{m}_{Y_t S_t}(\bar{z}_{t'}, \bar{X}_{it'})$, we propose the following recursive modeling approach based on Proposition 1(a):

1. Fit a model of Y_{it} , given $\bar{Z}_{it}$ and $\bar{X}_{it}$, for units with $S_{it} = 1$ to obtain an estimate of $m_{Y_t S_t}(\bar{z}_t, \bar{X}_{it}) = \mathbb{E}(Y_{it} | S_{it} = 1, \bar{Z}_{it} = \bar{z}_t, \bar{X}_{it})$, denoted by $\hat{m}_{Y_t S_t}(\bar{z}_t, \bar{X}_{it})$.

2. Fit a model of $S_{i,t'+1} m_{Y_t S_t}(\bar{z}_{t'+1}, \bar{X}_{i,t'+1})$, given $\bar{Z}_{it'}$ and $\bar{X}_{it'}$, for units with $S_{it'} = 1$, using $\hat{m}_{Y_t S_t}(\bar{z}_{t'+1}, \bar{X}_{i,t'+1})$, to recursively obtain an estimate of $m_{Y_t S_t}(\bar{z}_{t'}, \bar{X}_{it'}) = \mathbb{E}\{S_{i,t'+1} m_{Y_t S_t}(\bar{z}_{t'+1}, \bar{X}_{i,t'+1}) | S_{it'} = 1, \bar{Z}_{it'} = \bar{z}_{t'}, \bar{X}_{it'}\}$ for $t' = t-1, t-2, \dots, 1$.

The model in Step 2 may be difficult to specify since $S_{i,t'+1}m_{Y_t S_t}(\bar{z}_{t'+1}, \bar{X}_{i,t'+1}) = 0$ for units with $S_{i,t'+1} = 0$. We could alternatively write

$$(7) \quad \begin{aligned} & \mathbb{E}\{S_{i,t'+1}m_{Y_t S_t}(\bar{z}_{t'+1}, \bar{X}_{i,t'+1})|S_{it'} = 1, \bar{Z}_{it'} = \bar{z}_{t'}, \bar{X}_{it'}\} \\ &= \mathbb{E}\{m_{Y_t S_t}(\bar{z}_{t'+1}, \bar{X}_{i,t'+1})|S_{i,t'+1} = 1, \bar{Z}_{it'} = \bar{z}_{t'}, \bar{X}_{it'}\}p_{t'+1}(\bar{z}_{t'}, \bar{X}_{it'}). \end{aligned}$$

Then the estimate of this term can be obtained by fitting a model of $\hat{m}_{Y_t S_t}(\bar{z}_{t'+1}, \bar{X}_{i,t'+1})$, given $\bar{Z}_{it'}$ and $\bar{X}_{it'}$ for units with $S_{it'} = 1$, and a model of $S_{it'}$, given $\bar{Z}_{i,t'-1}$, $\bar{X}_{i,t'-1}$ for units with $S_{i,t'-1} = 1$.

The recursive estimation procedure involves a sequence of conditional regressions, which may introduce model incompatibility if all models are parametric. To mitigate this risk and enhance robustness, we recommend using flexible machine learning methods for each step in the recursion.

For $\hat{m}_{S_t}(\bar{z}_{t'}, \bar{X}_{it'})$, we propose a similar recursive approach based on Proposition 1(b):

1. Fit a model of S_{it} , given $\bar{Z}_{i,t-1}$ and $\bar{X}_{i,t-1}$, for units with $S_{i,t-1} = 1$ to obtain an estimate of $m_{S_t}(\bar{z}_{t-1}, \bar{X}_{i,t-1}) = \mathbb{E}(S_{it}|S_{i,t-1} = 1, \bar{Z}_{i,t-1} = \bar{z}_{t-1}, \bar{X}_{i,t-1})$, denoted by $\hat{m}_{S_t}(\bar{z}_{t-1}, \bar{X}_{i,t-1})$.
2. Fit a model of $S_{i,t'+1}m_{S_t}(\bar{z}_{t'+1}, \bar{X}_{i,t'+1})$, given $\bar{Z}_{it'}$, $\bar{X}_{it'}$ for units with $S_{it'}$, using $\hat{m}_{S_t}(\bar{z}_{t'+1}, \bar{X}_{i,t'+1})$, to recursively obtain an estimate of $m_{S_t}(\bar{z}_{t'}, \bar{X}_{it'}) = \mathbb{E}\{S_{i,t'+1}m_{S_t}(\bar{z}_{t'+1}, \bar{X}_{i,t'+1})|S_{it'} = 1, \bar{Z}_{it'} = \bar{z}_{t'}, \bar{X}_{it'}\}$ for $t' = t-2, t-3, \dots, 1$.

Similar to the above, for Step 2, an alternative estimation strategy is to multiply the estimates of $\mathbb{E}\{m_{S_t}(\bar{z}_{t'+1}, \bar{X}_{i,t'+1})|S_{it'} = 1, \bar{Z}_{it'} = \bar{z}_{t'}, \bar{X}_{it'}\}$ and $p_{t'+1}(\bar{z}_{t'}, \bar{X}_{it'})$.

The EIF-based estimation relies on the models for the conditional mean outcome, eligibility probability, and propensity score. However, we will show that the consistency of the proposed estimators for $\tau_t(\bar{z}_{t-1})$ and $\theta(\xi)$ does not require the correct specification of all models. The following theorem establishes the consistency of the proposed estimators under either correct specification of the conditional mean and eligibility probability models, or correct specification of the propensity score model.

THEOREM 3 (Double Robustness). *Let $\tilde{m}_{Y_t S_t}(\bar{z}_{t'}, \bar{x}_{t'})$, $\tilde{m}_{S_t}(\bar{z}_{t'}, \bar{x}_{t'})$, and $\tilde{\pi}_{t'}(\bar{z}_{t'}, \bar{x}_{t'})$ denote the probability limit of $\hat{m}_{Y_t S_t}(\bar{z}_{t'}, \bar{x}_{t'})$, $\hat{m}_{S_t}(\bar{z}_{t'}, \bar{x}_{t'})$, and $\hat{\pi}_{t'}(\bar{z}_{t'}, \bar{x}_{t'})$, respectively. Then $\hat{\tau}_t(\bar{z}_{t-1})$ is consistent for $\tau_t(\bar{z}_{t-1})$ if either of the following two equalities hold:*

$$\begin{aligned} (\tilde{m}_{Y_t S_t}(\bar{z}_{t'}, \bar{x}_{t'}), \tilde{m}_{S_t}(\bar{z}_{t'}, \bar{x}_{t'})) &= (m_{Y_t S_t}(\bar{z}_{t'}, \bar{x}_{t'}), m_{S_t}(\bar{z}_{t'}, \bar{x}_{t'})) \quad \text{for all } t', \\ \tilde{\pi}_{t'}(\bar{z}_{t'}, \bar{x}_{t'}) &= \pi_{t'}(\bar{z}_{t'}, \bar{x}_{t'}) \quad \text{for all } t'. \end{aligned}$$

Similarly, $\hat{\theta}(\xi)$ is consistent for $\theta(\xi)$ if either of the following two equalities holds:

$$\tilde{m}_{Y_t S_t}(\bar{z}_{t'}, \bar{x}_{t'}) = m_{Y_t S_t}(\bar{z}_{t'}, \bar{x}_{t'}), \quad \tilde{\pi}_{t'}(\bar{z}_{t'}, \bar{x}_{t'}) = \pi_{t'}(\bar{z}_{t'}, \bar{x}_{t'}) \quad \text{for all } t'.$$

The proof shows that the bias of $\hat{\phi}_N(\bar{z}_t)$ for estimating $\mathbb{E}\{Y_{it}(\bar{z}_t)S_{it}(\bar{z}_{t-1})\}$ is given by

$$\begin{aligned} & \sum_{t'=1}^t \int \{m_{Y_t S_t}(\bar{z}_{t'}, \bar{x}_{t'}) - \tilde{m}_{Y_t S_t}(\bar{z}_{t'}, \bar{x}_{t'})\} \left\{ \frac{\pi_{t'}(\bar{z}_{t'}, \bar{x}_{t'})}{\tilde{\pi}_{t'}(\bar{z}_{t'}, \bar{x}_{t'})} - \frac{\pi_{t'-1}(\bar{z}_{t'-1}, \bar{x}_{t'-1})}{\tilde{\pi}_{t'-1}(\bar{z}_{t'-1}, \bar{x}_{t'-1})} \right\} \\ & \times \prod_{s=2}^{t'} p_s(\bar{z}_{s-1}, \bar{x}_{s-1}) \prod_{s=1}^{t'} l_s(\bar{z}_{s-1}, \bar{x}_s) d\bar{x}_{t'}, \end{aligned}$$

while the bias of $\hat{\phi}_D(\mathbf{z}_{t-1})$ for estimating $\Pr\{S_{it}(\mathbf{z}_{t-1}) = 1\}$ has the following expression:

$$\begin{aligned} & \sum_{t'=1}^{t-1} \int \{m_{S_t}(\bar{\mathbf{z}}_{t'}, \bar{\mathbf{x}}_{t'}) - \tilde{m}_{S_t}(\bar{\mathbf{z}}_{t'}, \bar{\mathbf{x}}_{t'})\} \left\{ \frac{\pi_{t'}(\bar{\mathbf{z}}_{t'}, \bar{\mathbf{x}}_{t'})}{\tilde{\pi}_{t'}(\bar{\mathbf{z}}_{t'}, \bar{\mathbf{x}}_{t'})} - \frac{\pi_{t'-1}(\bar{\mathbf{z}}_{t'-1}, \bar{\mathbf{x}}_{t'-1})}{\tilde{\pi}_{t'-1}(\bar{\mathbf{z}}_{t'-1}, \bar{\mathbf{x}}_{t'-1})} \right\} \\ & \quad \times \prod_{s=2}^{t'} p_s(\bar{\mathbf{z}}_{s-1}, \bar{\mathbf{x}}_{s-1}) \prod_{s=1}^{t'} l_s(\bar{\mathbf{z}}_{s-1}, \bar{\mathbf{x}}_s) d\bar{\mathbf{x}}_{t'}. \end{aligned}$$

The double robustness follows immediately from these bias expressions.

Finally, if all models are correctly specified, then the estimators are asymptotically Normal and semiparametrically efficient under certain regularity conditions.

THEOREM 4 (Asymptotic Normality and Efficiency). *Suppose that in addition to Assumptions 1–3, the following conditions hold:*

(a) $\tilde{m}_{Y_t S_t}(\bar{\mathbf{z}}_{t'}, \bar{\mathbf{x}}_{t'}) = m_{Y_t S_t}(\bar{\mathbf{z}}_{t'}, \bar{\mathbf{x}}_{t'})$, $\tilde{m}_{S_t}(\bar{\mathbf{z}}_{t'}, \bar{\mathbf{x}}_{t'}) = m_{S_t}(\bar{\mathbf{z}}_{t'}, \bar{\mathbf{x}}_{t'})$, and $\tilde{\pi}_{t'}(\bar{\mathbf{z}}_{t'}, \bar{\mathbf{x}}_{t'}) = \pi_{t'}(\bar{\mathbf{z}}_{t'}, \bar{\mathbf{x}}_{t'})$ for all t' .

(b) *The estimated functions $\hat{m}_{Y_t S_t}(\bar{\mathbf{z}}_{t'}, \bar{\mathbf{x}}_{t'})$, $\hat{m}_{S_t}(\bar{\mathbf{z}}_{t'}, \bar{\mathbf{x}}_{t'})$, and $\hat{\pi}_{t'}(\bar{\mathbf{z}}_{t'}, \bar{\mathbf{x}}_{t'})$ are in a Donsker class.*

(c) $\delta < \pi_t(\bar{\mathbf{z}}_t, \bar{\mathbf{x}}_t) < 1 - \delta$ and $\delta < \hat{\pi}_t(\bar{\mathbf{z}}_t, \bar{\mathbf{x}}_t) < 1 - \delta$ for all $\bar{\mathbf{x}}_t$ and some $\delta \in (0, \frac{1}{2})$.

(d) *For all t' ,*

$$\|m_{Y_t S_t}(\bar{\mathbf{z}}_{t'}, \bar{\mathbf{X}}_{it'}) - \tilde{m}_{Y_t S_t}(\bar{\mathbf{z}}_{t'}, \bar{\mathbf{X}}_{it'})\|_{2,s} \cdot \|\hat{w}_{t'}(\bar{\mathbf{z}}_{t'}, \bar{\mathbf{X}}_{it'}) - w_t(\bar{\mathbf{z}}_{t'}, \bar{\mathbf{X}}_{it'})\|_{2,s} = o_P(n^{-1/2}),$$

$$\|m_{S_t}(\bar{\mathbf{z}}_{t'}, \bar{\mathbf{X}}_{it'}) - \tilde{m}_{S_t}(\bar{\mathbf{z}}_{t'}, \bar{\mathbf{X}}_{it'})\|_{2,s} \cdot \|\hat{w}_{t'}(\bar{\mathbf{z}}_{t'}, \bar{\mathbf{X}}_{it'}) - w_t(\bar{\mathbf{z}}_{t'}, \bar{\mathbf{X}}_{it'})\|_{2,s} = o_P(n^{-1/2}),$$

where $\|\cdot\|_{2,s}$ is defined as follows for any function of $\bar{\mathbf{X}}_{it'}$, $f(\bar{\mathbf{X}}_{it'})$:

$$\|f(\bar{\mathbf{X}}_{it'})\|_{2,s} = \int f^2(\bar{\mathbf{x}}_{t'}) \prod_{s=2}^{t'} p_s(\bar{\mathbf{z}}_{s-1}, \bar{\mathbf{x}}_{s-1}) \prod_{s=1}^{t'} l_s(\bar{\mathbf{z}}_{s-1}, \bar{\mathbf{x}}_s) d\bar{\mathbf{x}}_{t'}.$$

Then $\{\hat{\phi}_N(\bar{\mathbf{z}}_{t-1}, 1) - \hat{\phi}_N(\bar{\mathbf{z}}_{t-1}, 0)\}/\hat{\phi}_D(\bar{\mathbf{z}}_{t-1})$ and $\sum_{t=1}^T \sum_{\bar{\mathbf{z}}_t} \hat{\phi}_N(\bar{\mathbf{z}}_t) \xi(\bar{\mathbf{z}}_t)$ are asymptotically Normal and achieve the semiparametric efficiency bound.

Theorem 4 extends double machine learning (e.g., Chernozhukov et al. (2018)) to longitudinal settings with selective eligibility. Condition (a) requires the consistency of all nuisance models; Condition (b) imposes constraints on the complexity of the spaces of their estimators, which can be relaxed by cross-fitting; Condition (c) requires the overlap condition to hold for both the true and estimated treatment probabilities, enabling bounding the estimation errors; Condition (d) requires the nuisance estimators to converge at a fast enough rate, though it can be slower than the $n^{1/2}$ parametric convergence rates.

3.5.2. Alternative simple estimators. The above estimators are efficient but require complex modeling. In practice, researchers may prefer a simpler estimation strategy, even at the expense of some efficiency. We consider two alternative estimation strategies for $\mathbb{E}\{Y_{it}(\bar{\mathbf{z}}_t)|S_{it}(\bar{\mathbf{z}}_{t-1}) = 1\}$ and $\Pr\{S_{it}(\bar{\mathbf{z}}_{t-1}) = 1\}$ using the identification results shown in Theorem 1.

First, based on the identification formula of Theorem 1(a), we obtain the estimates $\hat{m}_{Y_t S_t}(\bar{\mathbf{z}}_t, \bar{\mathbf{X}}_{it})$ and $\hat{m}_{S_t}(\bar{\mathbf{z}}_t, \bar{\mathbf{X}}_{it})$ and then compute the following estimators:

$$(8) \quad \hat{\mathbb{E}}\{Y_{it}(\bar{\mathbf{z}}_t)|S_{it}(\bar{\mathbf{z}}_{t-1}) = 1\} = \frac{\sum_{i=1}^n \hat{m}_{Y_t S_t}(\bar{\mathbf{z}}_t, \bar{\mathbf{X}}_{it})}{\sum_{i=1}^n \hat{m}_{S_t}(\bar{\mathbf{z}}_t, \bar{\mathbf{X}}_{it})},$$

$$(9) \quad \hat{\Pr}\{S_{it}(\bar{\mathbf{z}}_{t-1}) = 1\} = \frac{1}{n} \sum_{i=1}^n \hat{m}_{S_t}(\bar{\mathbf{z}}_t, \bar{\mathbf{X}}_{it}).$$

This estimation strategy relies only on the outcome mean and eligibility probability models.

Next, based on Theorem 1(b), we estimate $\hat{\pi}_{t-1}(\bar{z}_{t-1}, \bar{X}_{i,t-1})$ by modeling the treatment assignment. Using these estimates, we obtain the following estimators:

$$(10) \quad \hat{\mathbb{E}}\{Y_{it}(\bar{z}_t)|S_{it}(\bar{z}_{t-1}) = 1\} = \frac{\sum_{i=1}^n Y_{it} \mathbf{1}\{\bar{Z}_{it} = \bar{z}_t, S_{it} = 1\} / \hat{\pi}_{t-1}(\bar{z}_{t-1}, \bar{X}_{i,t-1})}{\sum_{i=1}^n \mathbf{1}\{\bar{Z}_{i,t-1} = \bar{z}_{t-1}, S_{it} = 1\} / \hat{\pi}_{t-1}(\bar{z}_{t-1}, \bar{X}_{i,t-1})},$$

$$(11) \quad \hat{\Pr}\{S_{it}(\bar{z}_{t-1}) = 1\} = \sum_{i=1}^n \frac{\mathbf{1}\{\bar{Z}_{i,t-1} = \bar{z}_{t-1}, S_{it} = 1\}}{\hat{\pi}_{t-1}(\bar{z}_{t-1}, \bar{X}_{i,t-1})}.$$

This estimation strategy relies only on the modeling of the propensity score. However, it may be unstable if some of the estimated propensity scores are close to zero or one.

Finally, the recursive modeling approach becomes unnecessary when the covariates are time invariant. From equations (5) and (6), we can plug in the estimated outcome mean and the eligibility probabilities in the identification formula and replace the expectation with the empirical average to obtain the estimators of $\hat{m}_{Y,S_i}(\bar{z}_{i,t'}, \bar{X}_{i,t'})$ and $\hat{m}_{S_i}(\bar{z}_{i,t'}, \bar{X}_{i,t'})$.

4. Simulation study.

4.1. *Setup.* We conduct a proof-of-concept simulation designed to verify the implementation of the proposed method and illustrate its key theoretical properties. We first generate the potential values of outcome and eligibility indicators, and then compute their observed values after generating the treatment:

1. Generate four independent standard Normal random variables, $\mathbf{X}_i = (X_{i1}, \dots, X_{i4})$.
2. Generate $Y_{i1}(z_1)$ from a Normal linear model:

$$\mathbb{E}\{Y_{i1}(z_1)|\mathbf{X}_i\} = -1 + z_1 + 0.5X_{i1} - X_{i3}.$$

3. Generate $S_{i2}(z_1)$ from the following model:

$$\Pr\{S_{i2}(z_1) = 1|\mathbf{X}_i\} = \begin{cases} \text{logit}^{-1}(1 + 0.5z_1 + 0.5X_{i2} - 0.5X_{i3} - X_{i4} - \delta Y_{i1}(z_1)) & \text{if } -1 < X_{i1} < 1, \\ 0 & \text{if } z_1 = 0 \text{ and } X_{i1} \leq -1, \\ 1 & \text{if } z_1 = 1 \text{ and } X_{i1} \geq 1. \end{cases}$$

4. Generate $Y_{i2}(z_1, z_2)$ and $S_{i3}(z_1, z_2)$ from Normal linear and logistic models conditional on $S_{i2}(z_1) = 1$ while setting $S_{i3}(z_1, z_2) = 0$ for units with $S_{i2}(z_1) = 0$,

$$\begin{aligned} &\mathbb{E}\{Y_{i2}(z_1, z_2)|S_{i2}(z_1) = 1, \mathbf{X}_i, Y_{i1}(z_1)\} \\ &= 0.5 - 0.5z_1 - 0.5z_1z_2 + X_{i2} - 0.5X_{i4} + 0.5z_1X_{i3} - \delta Y_{i1}(z_1) + 2\delta z_2Y_{i1}(z_1), \end{aligned}$$

$$\begin{aligned} &\Pr\{S_{i3}(z_1, z_2) = 1|S_{i2}(z_1) = 1, \mathbf{X}_i\} \\ &= \text{logit}^{-1}(1 - 0.5z_1 + z_2 + 0.5z_2X_{i2} - X_{i3} + \delta Y_{i2}(z_1, z_2)). \end{aligned}$$

5. Generate $Y_{i3}(z_1, z_2, z_3)$ from a Normal linear model conditional on $S_{i3}(z_1, z_2) = 1$:

$$\begin{aligned} &\mathbb{E}\{Y_{i3}(z_1, z_2, z_3)|S_{i3}(z_1, z_2) = 1, \mathbf{X}_i, Y_{i1}(z_1), Y_{i2}(z_1, z_2)\} \\ &= -1 - 0.25z_1 - 0.5z_2 - 0.5z_2z_3 + X_{i1} - 0.5X_{i3} - 2\delta Y_{i1}(z_1) - 0.5\delta z_2Y_{i1}(z_1) \\ &\quad + 2\delta Y_{i2}(z_1, z_2) + \delta z_3Y_{i2}(z_1, z_2) \end{aligned}$$

6. Generate Z_{i1} from a logistic model and compute the observed data Y_{i1} and S_{i2} :

$$\Pr\{Z_{i1} = 1|\mathbf{X}_i\} = \text{logit}^{-1}(0.2 + 0.2X_{i1} - 0.4X_{i2}).$$

7. Generate Z_{i2} from a logistic model conditional on $S_{i2} = 1$, and compute the observed data Y_{i2} and S_{i3} :

$$\Pr\{Z_{i2} = 1 | S_{i2} = 1, Z_{i1} = z_1, \mathbf{X}_i, Y_{i1}\} = \text{logit}^{-1}(0.5 - 0.5z_1 + 0.5X_{i2} - 0.5X_{i4} + \delta Y_{i1}).$$

8. Generate Z_{i3} from a logistic model conditional on $S_{i3} = 1$, and compute Y_{i3} :

$$\begin{aligned} \Pr\{Z_{i3} = 1 | S_{i3} = 1, Z_{i1} = z_1, Z_{i2} = z_2, \mathbf{X}_i, Y_{i1}, Y_{i2}\} \\ = \text{logit}^{-1}(1 - 0.5z_1 - 0.5z_2 + 0.5X_{i1} - 0.5X_{i3} + \delta Y_{i1} - \delta Y_{i2}). \end{aligned}$$

In this data generating process, δ determines the degree to which the outcome in the preceding periods affects the current treatment assignment and outcome; we consider two scenarios: $\delta = 0$ and $\delta = 0.5$. When $\delta = 0$, Assumption 2 holds with time invariant confounders $\mathbf{X}_{it} = \mathbf{X}_i$ for $t = 1, 2, 3$. When $\delta = 0.5$, Assumption 2 holds with time varying covariates $\mathbf{X}_{i1} = \mathbf{X}_i$, $\mathbf{X}_{i2} = (\mathbf{X}_i, Y_{i1})$, and $\mathbf{X}_{i3} = (\mathbf{X}_i, Y_{i1}, Y_{i2})$. Moreover, the construction of $S_{i2}(z_1)$ in Step 3 implies that both the ignorability assumption and the overlap condition are violated for the eligibility indicator.

Under each scenario, we consider the outcome regression estimators based on equations (8) and (9), the IPW estimators based on equations (10) and (11), and the DR estimators. The causal quantities of interest include τ_1 , $\tau_2(z_1)$, and $\tau_3(z_1, z_2)$ for $z_1, z_2 = 0, 1$, and the expected cumulated outcomes under the treat-nobody strategy θ_0 and treat-everyone strategy θ_1 .

Following Kang and Schafer (2007) and Kennedy (2019), we use both parametric and nonparametric methods to estimate the nuisance functions with both correctly specified covariates $\mathbf{X}_i$ and misspecified covariates $\mathbf{X}_i^*$, which are constructed by transforming the true covariates as follows: $X_1^* = \exp(X_1/2)$, $X_2^* = \frac{X_2}{1+\exp(X_1)} + 10$, $X_3^* = (\frac{X_1 X_3}{25} + 0.6)^3$, $X_4^* = (X_1 + X_4 + 20)^2$.

In total, we have four different versions of the simulation. For parametric estimation we use linear models for the continuous variables and logistic models for binary variables. For nonparametric estimation we use the Super Learner ensemble (van der Laan, Polley and Hubbard (2007)) that combines parametric models, random forest, and generalized additive models.

4.2. Results. Figure 2 presents the absolute bias of each estimator based on 1,000 Monte Carlo simulations under the scenario with time-invariant confounders and a sample size of $n = 5,000$. Figure 3 shows the root mean square error (RMSE) under the same settings, while Figures S3 and S4 in the Appendix present the corresponding results with time-varying confounders.

Across all scenarios using nonparametric methods (“CorML” and “MisML”), the doubly robust (DR) estimator generally demonstrates comparable or superior performance relative to the outcome regression and IPW estimators. This advantage is especially pronounced for τ_1 , $\tau_2(1)$, and $\tau_2(0)$, where the DR estimator exhibits consistently low bias and RMSE. For $\tau_3(\cdot, \cdot)$, however, the DR estimator occasionally shows higher bias and RMSE than outcome regression, possibly due to increased variability in propensity score estimation when the number of eligible units is small. This challenge is further amplified under the data-generating process without overlap, where extreme propensity scores hinder accurate estimation. The IPW estimator performs the worst in most scenarios, with large bias and RMSE across all scenarios. Overall, these results underscore the practical difficulties of estimating later-period effects in longitudinal settings with selective eligibility.

In scenarios using parametric methods, outcome regression performs best when the covariates are correctly specified. This is expected, as outcome regression corresponds to the maximum likelihood estimator under correct specification. When covariates are misspecified,




FIG. 2. Absolute bias of each estimator across 1,000 replications with $\delta = 0$ and $n = 5,000$. Each plot presents the results for one target causal quantity. The labels on the horizontal axis—“CorML,” “CorP,” “MisML,” and “MisP”—stand for the nonparametric estimation with correctly specified covariates, parametric estimation with correctly specified covariates, nonparametric estimation with misspecified covariates, and parametric estimation with misspecified covariates, respectively. Within each group, the bars correspond to the regression, IPW, and doubly robust estimators.

both the DR and outcome regression estimators exhibit bias, and their relative performance varies across estimands. Across all parametric settings, the IPW estimator remains the least reliable, showing the largest biases and RMSEs.

We also evaluate the coverage rates of the confidence interval based on the proposed doubly robust estimator. Details of the confidence interval construction are provided in the Appendix. Table 2 presents the results. When covariates are correctly specified, the estimator achieves nominal coverage across all target parameters under both time-invariant and time-varying confounding scenarios. In contrast, when covariates are misspecified, we observe a decline in coverage rates for several parameters, particularly for the effects at time 3. This pattern suggests that the estimator is sensitive to model misspecification when estimating more granular subgroup effects.

To further illustrate the consequences of ignoring selective eligibility, we implement a naive estimation strategy that treats each case as an independent observation and applies a standard doubly robust estimator, without accounting for the longitudinal eligibility structure. Specifically, we estimate the expected total outcomes θ_0 , θ_1 , and their difference $\theta_1 - \theta_0$ using this case-level approach, with nuisance functions estimated nonparametrically. We present the results in the Appendix. The coverage rates of the resulting confidence intervals are essentially zero, highlighting the risk of severely misleading inference when eligibility is not properly accounted for in the analysis. Interestingly, coverage is somewhat better under the




FIG. 3. Root mean squared error (RMSE) of each estimator with $\delta = 0$ and $n = 5,000$. Each subplot presents the results for one target causal quantity. Within each subplot, the x-axis labels “CorML,” “CorP,” “MisML,” and “MisP” stand for the nonparametric estimation with correctly specified covariates, parametric estimation with correctly specified covariates, nonparametric estimation with misspecified covariates, and parametric estimation with misspecified covariates, respectively. Within each group, the bars correspond to the regression, IPW, and doubly robust estimators.

time-varying confounding scenario than in the time-invariant case, possibly because recursive fitting of the nuisance functions partially mitigates misspecification.

Additional simulation results in the Appendix include scenarios with binary outcomes generated from logistic models with both time-invariant and time-varying confounders. These results are qualitatively similar.

5. Analysis of the PSA experiment. We now turn to analyzing the PSA experiment described in Section 2, using the methodology from Section 3 to account for selective eligi-

TABLE 2
Coverage rates of the doubly robust estimator using nonparametric methods under different model specifications for $\delta = 0$ (time-invariant covariates) and $\delta = 0.5$ (time-varying covariates) under $n = 5,000$. “Cor” indicates correctly specified covariates, and “Mis” indicates misspecified covariates

δ		τ_1	$\tau_2(1)$	$\tau_2(0)$	$\tau_3(1, 1)$	$\tau_3(0, 1)$	$\tau_3(1, 0)$	$\tau_3(0, 0)$	θ_0	θ_1
$\delta = 0$	Cor	0.956	0.948	0.958	0.945	0.945	0.951	0.951	0.953	0.940
	Mis	0.954	0.929	0.932	0.685	0.767	0.815	0.877	0.907	0.886
$\delta = 0.5$	Cor	0.957	0.946	0.951	0.951	0.938	0.953	0.941	0.941	0.952
	Mis	0.952	0.956	0.935	0.831	0.830	0.918	0.900	0.941	0.921

bility. We limit our analyses to estimating the causal effects up to the third case $T = 3$, as the number of cases decreases to only 113 when considering the fourth case, making reliable estimation infeasible.

We consider three components. First, we evaluate whether the judge becomes more likely to follow the PSA recommendation in later cases for the same arrestee. Second, we assess the impact of PSA provision on each of the three negative outcomes: FTA, NCA, and NVCA. Third, we examine whether PSA provision improves the quality of the judge's decisions, as measured by the misclassification rate. For each component, we estimate both the ETes and the ETOs under two hypothetical policies: one that provides PSA recommendations to all cases and the other that never provides PSA recommendations for any case.

5.1. Setup. Let Z_{it} represent the randomized treatment assignment variable with $Z_{it} = 1$ implying that the judge received the PSA recommendation for arrestee i after the t th arrest for $i = 1, 2, \dots, n$ and $t = 1, 2, 3$. For each component of the application, we estimate several causal quantities. They include the average effect of the first PSA provision τ_1 , the average effects of the second PSA provision $\tau_2(z_1)$ for $z_1 = 0, 1$, and the average effects of the third PSA provision $\tau_3(z_1, z_2)$ for $z_1, z_2 = 0, 1$. In addition, we also estimate the expected number of outcomes under two policy-relevant scenarios: if the judge were to receive a PSA recommendation for every case, θ_1 , and if the judge were to receive a PSA recommendation for no cases, θ_0 . Finally, we estimate the difference between these two scenarios, that is, $\theta_1 - \theta_0$.

Because the PSA provision is randomized, Assumption 2 holds without conditioning on covariates. Nevertheless, for potential efficiency gains, we include age, gender, race, and the interaction between gender and race in our models. Because the propensity score model is always correctly specified, the DR estimator will be consistent regardless of whether the other nuisance functions are correctly specified.

5.2. Empirical evidence of selective eligibility. Before turning to outcome analysis, we examine the presence of selective eligibility in the data. For time $t = 2$, we nonparametrically estimate the conditional probabilities $\Pr(S_{i2}(1) = 1 | X_{i2})$ and $\Pr(S_{i2}(0) = 1 | X_{i2})$ for each unit. We then compute the unit-level differences between these two probabilities. A difference away from zero implies that treatment at time $t = 1$ influences the likelihood of eligibility at time $t = 2$, providing evidence of selective eligibility. For time $t = 3$, we focus on the subset of units that remained eligible for the second treatment. We compare $\Pr(S_{i3}(0, 1) = 1 | X_{i3}, S_{i2}(0) = 1)$ with $\Pr(S_{i3}(0, 0) = 1 | X_{i3}, S_{i2}(0) = 1)$, and $\Pr(S_{i3}(1, 1) = 1 | X_{i3}, S_{i2}(1) = 1)$ with $\Pr(S_{i3}(1, 0) = 1 | X_{i3}, S_{i2}(1) = 1)$ to assess whether eligibility at time $t = 3$ is affected by the second treatment.

Figure 4 displays the histograms of the estimated unit-level differences in eligibility probabilities under alternative treatment conditions. In each panel the vertical dashed line marks zero difference. We observe substantial variation in predicted differences across units, especially for $t = 2$, where the distribution is visibly spread and centered slightly below zero. This suggests that treatment at the first time point may reduce the likelihood of being eligible at time 2 (i.e., being rearrested) for some units. For time $t = 3$, the differences are smaller in magnitude but still show visible dispersion, confirming that eligibility is not independent of treatment history.

5.3. Effect on the agreement between the judge's decision and PSA recommendation. We first examine whether the provision of the PSA affects the likelihood that the judge's decision aligns with the PSA recommendation. For this analysis we define the outcome $Y_{it} = 1$ if the judge's release decision agrees with the PSA recommendation after the t th arrest, and $Y_{it} = 0$ otherwise. Figure 5 presents the estimated effects of PSA provision on this agreement outcome, along with 95% confidence intervals.




FIG. 4. *Distribution of differences in fitted eligibility probabilities under alternative treatment histories using nonparametric methods. Each panel displays the distribution of unit-level differences in the fitted conditional probability of being eligible ($S_{it} = 1$) under different treatment conditions. The y-axis shows counts of units in each bin.*

We find that PSA provision generally increases the probability of agreement between the judge’s decision and the PSA recommendation. The estimated effects are statistically significant for the first two arrest cases, indicating that PSA provision influences judicial decisions in the direction of the PSA recommendation at early stages. For the third arrest case, the estimated effects are imprecisely estimated due to the smaller number of eligible cases. While some point estimates, such as $\tau_3(1, 1)$, are positive and larger in magnitude, their confidence intervals include zero. We, therefore, do not interpret these estimates as evidence of a systematic effect at later arrest stages. Similarly, while $\tau_2(1)$ is slightly larger than $\tau_2(0)$, the difference is not statistically significant. Overall, the evidence supports an effect of PSA provision on agreement at earlier arrests, with limited and inconclusive evidence at later cases.




FIG. 5. *Estimated average effects of PSA provision on whether the judge’s decision agrees with the PSA recommendation using nonparametric models. Each point represents a treatment effect, defined as the difference in agreement probabilities under different treatment histories. The figure shows the estimated effect at time $t = 1$ on the top (τ_1), the two estimated effects at time $t = 2$ in the middle ($\tau_2(z_1)$), the four estimated effects at time $t = 3$ in the bottom ($\tau_3(z_1, z_2)$), and the expected total number of agreements between the judge’s decision and the PSA recommendation under the never-treat strategy (θ_0) and always-treat strategy (θ_1). The horizontal lines represent the 95% confidence intervals.*




FIG. 6. Estimated average effects of PSA provision on the three negative outcomes (FTA, NCA, and NVCA) based on nonparametric models. Each point represents a treatment effect, defined as the difference in probabilities of negative outcomes under different treatment histories. Each panel includes the estimated effect at time $t = 1$ on the top (τ_1), the two estimated effects at time $t = 2$ in the middle ($\tau_2(z_1)$), the four estimated effects at time $t = 3$ in the bottom ($\tau_3(z_1, z_2)$), and the expected number of each negative outcome under the never-treat strategy (θ_0) and the always-treat strategy (θ_1). The horizontal lines represent 95% confidence intervals.

5.4. *Effect on the negative outcomes.* Figure 6 shows the estimated effects of PSA provision on the occurrence of each negative outcome with their 95% confidence intervals. We find that none of the estimated effects is statistically significant, with the exception of the estimated difference in the expected number of NCAs under the always-provide and never-provide strategies (middle panel). This result suggests that providing the PSA to the judge for all cases may lead to a higher number of total NCAs, although we would want to replicate this analysis in other essentially identical field RCTs we have conducted before so concluding.

Similarly, several of the point estimates for ETEs on NCA and NVCA are positive, though their intervals are wide and include zero. These results suggest a possible pattern in which PSA provision does not reduce and may even slightly increase certain types of negative outcomes. This pattern is broadly consistent with the earlier findings based on the first arrest case reported in Imai et al. (2023). While we do not interpret these as evidence of a true effect, the consistent direction of some estimates may warrant further examination in future work.

5.5. *Effect on the misclassification rate of the judge’s decision.* Finally, we evaluate whether the PSA provision helps improve the quality of the judge’s decisions. Following the framework proposed in Ben-Michael et al. (2025a), we define the quality of the decision in terms of a misclassification rate that captures whether the judge’s release decision aligns with the risk level of the arrestee. The method uses a risk-based benchmark: whether the arrestee would have committed a negative outcome if released without bail. A decision is classified as misaligned if either: (i) the judge releases an arrestee without bail and this results in a negative outcome, or (ii) the judge assigns bail to an arrestee when there counterfactually would not have been a bad outcome had the arrestee been released without bail. The outcome variable used in this analysis is a compound measure capturing this misclassification, defined as $(1 - Y_{it})D_{it} - (1 - Y_{it})(1 - D_{it})$, where Y_{it} is the realized negative outcome and D_{it} is the judge’s release decision, where $D_{it} = 1$ corresponds to the cash bail decision and $D_{it} = 0$ corresponds to releasing without cash bail. We refer readers to Ben-Michael et al. (2025a) for the full construction of this measure.




FIG. 7. Estimated average effects of PSA provision on the misclassification rates of the judge's decision in terms of the three negative outcomes using nonparametric models. Each point represents a treatment effect, defined as the difference in misclassification rates under different treatment histories. The figure shows the estimated effect at time $t = 1$ on the top (τ_1), the two estimated effects at time $t = 2$ in the middle ($\tau_2(z_1)$), the four estimated effects at time $t = 3$ in the bottom ($\tau_3(z_1, z_2)$), and the overall difference in expected number of misclassifications per person under the never-treat strategy and always-treat strategy ($\theta_1 - \theta_0$). Because the individual quantities θ_0 and θ_1 do not have interpretable meanings in this context, we report only their difference. The horizontal lines represent the 95% confidence intervals.

Figure 7 displays the estimated effects of PSA provision on the misclassification rates of the judge's decisions in terms of each negative outcome. The estimated overall effect on the expected number of misclassifications per arrestee ($\theta_1 - \theta_0$) is positive and statistically significant for NCA, indicating that PSA provision may worsen classification performance for this outcome. For FTA and NVCA, the overall effects are close to zero. Other than that, most of the estimated effects are small and not statistically significant. Several point estimates for the third case are positive for NCA and negative for FTA. However, in all cases the confidence intervals are wide and include zero. These results largely align with earlier findings based on first-case-only analyses in Ben-Michael et al. (2025a).

6. Concluding remarks. In this paper we proposed a general causal framework for defining and estimating treatment effects in longitudinal studies with selective eligibility. We established identification results and developed estimation approaches for the causal quantities. The proposed methodology, therefore, generalizes existing methods for standard longitudinal studies to this new setting.

In our randomized evaluation of the PSA, we find that providing the PSA increases the rate of agreement between the algorithm's recommendation and the judge's decision, with potentially larger impacts on agreement when the PSA is repeatedly provided for the same arrestee across rearrests. We mostly find null impacts of the provision of the PSA on the three negative outcomes, as well as on the misclassification rate of the judge's decisions, though there is some evidence that the total number of new criminal activities and the associated total number of misclassifications would increase if the PSA were to always be provided to the judge relative to cases where it were never provided to the judge.

Our analysis uses t to denote periods rather than calendar time, which may lead to two important issues. First, both judicial decision-making and arrestee behavior may exhibit temporal trends. This mismatch can challenge the assumption that cases at the same value of t are identically distributed, potentially introducing bias if external conditions evolve over

calendar time. Second, the interpretation of our causal estimands depends on the observation window and the timing of events. A shorter or longer study period would result in a different eligible population at each time point, affecting both identification and generalizability. Addressing these issues would require modeling the data-generating process in continuous time and incorporating censoring mechanisms for eligibility status. We view this as an important limitation of the current framework and a valuable direction for future research.

There are several avenues for future work. While we focused on nonparametric identification and estimation, with a large value of t , the causal effect $\tau_t(z_{t-1})$ conditions on a long treatment history, increasing the likelihood of violating the overlap condition required for identification. In addition, the number of causal estimands grows exponentially with t . In such cases, incorporating modeling assumptions becomes necessary for addressing identification challenges and improving estimation. One possible approach is to apply structural nested mean models under selective eligibility (Robins (1992), Hernán and Robins (2020)). Our estimands can also be extended to condition on time-varying covariates. The identification strategy remains valid under this extension. However, efficient estimation becomes more challenging and typically requires additional technical development.

It is also of interest to extend the proposed framework to continuous time longitudinal studies with selective eligibility. The survival analysis literature has formulated the problem as semicompeting risks in which the primary outcome is the nonterminal event and the process of selective eligibility acts as a competing risk (Fine, Jiang and Chappell (2001)). However, most existing research does not involve time-varying treatments (Comment et al. (2025), Xu et al. (2022)). Our framework may be generalized to assess the effect of a time-varying treatment in the presence of semicompeting risks.

Acknowledgments. The authors would like to thank the anonymous referees, an Associate Editor, and the Editor for their constructive comments that improved the quality of this paper.

Funding. The first author was supported in part by the National Natural Science Foundation of China (Grants 12371285, 12292984).

The second author was supported in part by the National Institute of Standards and Technology (Grant 60NANB24D231) and the Carnegie Mellon University AI Measurement Science and Engineering Center (AIMSEC).

The third and fourth authors were supported in part by a grant from Arnold Ventures, formerly the Laura & John Arnold Foundation.

The fifth author was supported in part by NSF Grant SES-2051196.

SUPPLEMENTARY MATERIAL

Supplementary appendix (DOI: [10.1214/26-AOAS2201SUPPA](https://doi.org/10.1214/26-AOAS2201SUPPA); .pdf). Proofs of the results and additional results on the simulation and empirical analysis.

Code (DOI: [10.1214/26-AOAS2201SUPPB](https://doi.org/10.1214/26-AOAS2201SUPPB); .zip). Code for implementing the proposed method.

REFERENCES

- ANDERSSON, F., HOLZER, H. J., LANE, J. I., ROSENBLUM, D. and SMITH, J. (2024). Does federally funded job training work?: Nonexperimental estimates of WIA training impacts using longitudinal data on workers and firms. *J. Hum. Resour.* **59** 1244–1283.
- BADEN, L. R., EL SAHLY, H. M., ESSINK, B., KOTLOFF, K., FREY, S., NOVAK, R., DIEMERT, D., SPECTOR, S. A., ROUPHAEL, N. et al. (2021). Efficacy and safety of the mRNA-1273 SARS-CoV-2 vaccine. *N. Engl. J. Med.* **384** 403–416.

- BAER, B. R., STRAWDERMAN, R. L. and ERTEFAIE, A. (2023). Causal inference for the expected number of recurrent events in the presence of a terminal event. Technical Report, arXiv preprint. Available at [arXiv:2306.16571](https://arxiv.org/abs/2306.16571).
- BANG, H. and ROBINS, J. M. (2005). Doubly robust estimation in missing data and causal inference models. *Biometrics* **61** 962–972. [MR2216189 https://doi.org/10.1111/j.1541-0420.2005.00377.x](https://doi.org/10.1111/j.1541-0420.2005.00377.x)
- BEN-MICHAEL, E., FELLER, A. and ROTHSTEIN, J. (2022). Synthetic controls with staggered adoption. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **84** 351–381. [MR4412990](https://doi.org/10.1111/rssb.12990)
- BEN-MICHAEL, E., GREINER, D. J., HUANG, M., IMAI, K., JIANG, Z. and SHIN, S. (2025a). Does AI help humans make better decisions? A statistical evaluation framework for experimental and observational studies. *Proc. Natl. Acad. Sci. USA* **122** e2505106122. [MR5001409 https://doi.org/10.1073/pnas.2505106122](https://doi.org/10.1073/pnas.2505106122)
- BEN-MICHAEL, E., GREINER, D. J., IMAI, K. and JIANG, Z. (2025b). Safe policy learning through extrapolation: Application to pre-trial risk assessment. *J. Amer. Statist. Assoc.* **120** 1386–1399. [MR4973866 https://doi.org/10.1080/01621459.2025.2489135](https://doi.org/10.1080/01621459.2025.2489135)
- BICKEL, P. J., KLAASSEN, C. A. J., RITOV, Y. and WELLNER, J. A. (1993). *Efficient and Adaptive Estimation for Semiparametric Models. Johns Hopkins Series in the Mathematical Sciences*. Johns Hopkins Univ. Press, Baltimore, MD. [MR1245941](https://doi.org/10.1215/0013780X-1993-0001)
- BOJINOV, I. and SHEPHARD, N. (2019). Time series experiments and causal estimands: Exact randomization tests and trading. *J. Amer. Statist. Assoc.* **114** 1665–1682. [MR4047291 https://doi.org/10.1080/01621459.2018.1527225](https://doi.org/10.1080/01621459.2018.1527225)
- BORUVKA, A., ALMIRALL, D., WITKIEWITZ, K. and MURPHY, S. A. (2018). Assessing time-varying causal effect moderation in mobile health. *J. Amer. Statist. Assoc.* **113** 1112–1121. [MR3862343 https://doi.org/10.1080/01621459.2017.1305274](https://doi.org/10.1080/01621459.2017.1305274)
- CHERNOZHUKOV, V., CHETVERIKOV, D., DEMIRER, M., DUFLO, E., HANSEN, C., NEWEY, W. and ROBINS, J. (2018). Double/debiased machine learning for treatment and structural parameters. *Econom. J.* **21** C1–C68. [MR3769544 https://doi.org/10.1111/ectj.12097](https://doi.org/10.1111/ectj.12097)
- COMMENT, L., MEALLI, F., HANEUSE, S. and ZIGLER, C. M. (2025). Survivor average causal effects for continuous time: A principal stratification approach to causal inference with semicompeting risks. *Biom. J.* **67** e70041. [MR4875065 https://doi.org/10.1002/bimj.70041](https://doi.org/10.1002/bimj.70041)
- D’AMOUR, A., DING, P., FELLER, A., LEI, L. and SEKHON, J. (2021). Overlap in observational studies with high-dimensional covariates. *J. Econometrics* **221** 644–654. [MR4215042 https://doi.org/10.1016/j.jeconom.2019.10.014](https://doi.org/10.1016/j.jeconom.2019.10.014)
- DÍAZ, I., HOFFMAN, K. L. and HEJAZI, N. S. (2024). Causal survival analysis under competing risks using longitudinal modified treatment policies. *Lifetime Data Anal.* **30** 213–236. [MR4684790 https://doi.org/10.1007/s10985-023-09606-7](https://doi.org/10.1007/s10985-023-09606-7)
- DÍAZ, I., WILLIAMS, N., HOFFMAN, K. L. and SCHENCK, E. J. (2023). Nonparametric causal effects based on longitudinal modified treatment policies. *J. Amer. Statist. Assoc.* **118** 846–857. [MR4595460 https://doi.org/10.1080/01621459.2021.1955691](https://doi.org/10.1080/01621459.2021.1955691)
- DING, P., GENG, Z., YAN, W. and ZHOU, X.-H. (2011). Identifiability and estimation of causal effects by principal stratification with outcomes truncated by death. *J. Amer. Statist. Assoc.* **106** 1578–1591. [MR2896858 https://doi.org/10.1198/jasa.2011.tm10265](https://doi.org/10.1198/jasa.2011.tm10265)
- FINE, J. P., JIANG, H. and CHAPPELL, R. (2001). On semi-competing risks data. *Biometrika* **88** 907–919. [MR1872209 https://doi.org/10.1093/biomet/88.4.907](https://doi.org/10.1093/biomet/88.4.907)
- GREINER, J., HALEN, R., STUBENBERG, M. and GRIFFIN, C. L. (2020). Randomized Control Trial Evaluation of the Implementation of the PSA-DMF System in Dane County, WI. Technical Report, Access to Justice Lab, Harvard Law School.
- GROSSI, G., MARIANI, M., MATTEI, A. and MEALLI, F. (in press). Bayesian principal stratification with longitudinal data and truncation by death. *Econometrics and Statistics*.
- HAHN, J. (1998). On the role of the propensity score in efficient semiparametric estimation of average treatment effects. *Econometrica* **66** 315–331. [MR1612242 https://doi.org/10.2307/2998560](https://doi.org/10.2307/2998560)
- HERNÁN, M. A. and ROBINS, J. M. (2020). *Causal Inference*. CRC Press/CRC, Boca Raton, FL.
- IMAI, K. (2008). Sharp bounds on the causal effects in randomized experiments with “truncation-by-death”. *Statist. Probab. Lett.* **78** 144–149. [MR2382067 https://doi.org/10.1016/j.spl.2007.05.015](https://doi.org/10.1016/j.spl.2007.05.015)
- IMAI, K., JIANG, Z., GREINER, D. J., HALEN, R. and SHIN, S. (2023). Experimental evaluation of algorithm-assisted human decision-making: Application to pretrial public safety assessment. *J. Roy. Statist. Soc., Ser. A, Statist. Soc.* **186** 167–216. [MR4690118 https://doi.org/10.1093/jrsssa/qnad010](https://doi.org/10.1093/jrsssa/qnad010)
- JANVIN, M., YOUNG, J. G., RYALEN, P. C. and STENSRUD, M. J. (2024). Causal inference with recurrent and competing events. *Lifetime Data Anal.* **30** 59–118. [MR4684786 https://doi.org/10.1007/s10985-023-09594-8](https://doi.org/10.1007/s10985-023-09594-8)
- JIANG, Z., BEN-MICHAEL, E., GREINER, D. J., HALEN, R. and IMAI, K. (2026). Supplement to “Longitudinal causal inference with selective eligibility: Application to pretrial risk assessment.” <https://doi.org/10.1214/26-AOAS2201SUPPA>, <https://doi.org/10.1214/26-AOAS2201SUPPB>

- JIANG, Z., CHEN, S. and DING, P. (2023). An instrumental variable method for point processes: Generalized Wald estimation based on deconvolution. *Biometrika* **110** 989–1008. [MR4667436](#) <https://doi.org/10.1093/biomet/asad005>
- JOSEFSSON, M. and DANIELS, M. J. (2021). Bayesian semi-parametric G-computation for causal inference in a cohort study with MNAR dropout and death. *J. R. Stat. Soc., Ser. C, Appl. Stat.* **70** 398–414. [MR4226674](#) <https://doi.org/10.1111/rssc.12464>
- JOSEFSSON, M., DE LUNA, X., DANIELS, M. J. and NYBERG, L. (2016). Casual inference with longitudinal outcomes and non-ignorable dropout: Estimating the effect of living alone on cognitive decline. *J. R. Stat. Soc., Ser. C, Appl. Stat.* **65** 131–144. [MR3438242](#) <https://doi.org/10.1111/rssc.12110>
- KANG, J. D. Y. and SCHAFER, J. L. (2007). Demystifying double robustness: A comparison of alternative strategies for estimating a population mean from incomplete data. *Statist. Sci.* **22** 523–539. [MR2420458](#) <https://doi.org/10.1214/07-STS227>
- KENNEDY, E. H. (2016). Semiparametric theory and empirical processes in causal inference. In *Statistical Causal Inferences and Their Applications in Public Health Research. ICSA Book Ser. Stat.* 141–167. Springer, Cham. [MR3617956](#)
- KENNEDY, E. H. (2019). Nonparametric causal effects based on incremental propensity score interventions. *J. Amer. Statist. Assoc.* **114** 645–656. [MR3963169](#) <https://doi.org/10.1080/01621459.2017.1422737>
- KURLAND, B. F., JOHNSON, L. L., EGLESTON, B. L. and DIEHR, P. H. (2009). Longitudinal data with follow-up truncated by death: Match the analysis method to research aims. *Statist. Sci.* **24** 211–222. [MR2655850](#) <https://doi.org/10.1214/09-STS293>
- LEE, K. and DANIELS, M. J. (2013). Causal inference for bivariate longitudinal quality of life data in presence of death by using global odds ratios. *Stat. Med.* **32** 4275–4284. [MR3118354](#) <https://doi.org/10.1002/sim.5857>
- LEE, K., DANIELS, M. J. and SARGENT, D. J. (2010). Causal effects of treatments for informative missing data due to progression/death. *J. Amer. Statist. Assoc.* **105** 912–929. [MR2752589](#) <https://doi.org/10.1198/jasa.2010.ap08739>
- LOCHMAN, J. E. and WELLS, K. C. (2004). The coping power program for preadolescent aggressive boys and their parents: Outcome effects at the 1-year follow-up. *J. Consult. Clin. Psychol.* **72** 571.
- PAPADOGEORGOU, G., IMAI, K., LYALL, J. and LI, F. (2022). Causal inference with spatio-temporal data: Estimating the effects of airstrikes on insurgent violence in Iraq. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **84** 1969–1999. [MR4515563](#) <https://doi.org/10.1111/rssb.12548>
- RICHARDSON, T. S. and ROBINS, J. M. (2013). Single world intervention graphs (SWIGs): A unification of the counterfactual and graphical approaches to causality. Center for the statistics and the social sciences, University of Washington series. Working Paper 128, 2013.
- ROBINS, J. (1986). A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. *Math. Model.* **7** 1393–1512.
- ROBINS, J. (1992). Estimation of the time-dependent accelerated failure time model in the presence of confounding factors. *Biometrika* **79** 321–334. [MR1185134](#) <https://doi.org/10.1093/biomet/79.2.321>
- ROBINS, J. M. and HERNÁN, M. A. (2009). Estimation of the causal effects of time-varying exposures. In *Longitudinal Data Analysis. Chapman & Hall/CRC Handb. Mod. Stat. Methods* 553–599. CRC Press, Boca Raton, FL. [MR1500133](#)
- RUBIN, D. B. (1980). Comment on “Randomization analysis of experimental data: The Fisher randomization test”. *J. Amer. Statist. Assoc.* **75** 591–593.
- RUBIN, D. B. (2006). Causal inference through potential outcomes and principal stratification: Application to studies with “censoring” due to death. *Statist. Sci.* **21** 299–309. [MR2339125](#) <https://doi.org/10.1214/088342306000000114>
- RYTGAARD, H. C., GERDS, T. A. and VAN DER LAAN, M. J. (2022). Continuous-time targeted minimum loss-based estimation of intervention-specific mean outcomes. *Ann. Statist.* **50** 2469–2491. [MR4500616](#) <https://doi.org/10.1214/21-aos2114>
- RYTGAARD, H. C. W. and VAN DER LAAN, M. J. (2024). One-step targeted maximum likelihood estimation for targeting cause-specific absolute risks and survival curves. *Biometrika* **111** 129–145. [MR4704562](#) <https://doi.org/10.1093/biomet/asad033>
- SCHINDL, K., SHEN, S. and KENNEDY, E. H. (2024). Incremental effects for continuous exposures. Technical Report, arXiv preprint. Available at [arXiv:2409.11967](#).
- SHARDELL, M. and FERRUCCI, L. (2018). Joint mixed-effects models for causal inference with longitudinal data. *Stat. Med.* **37** 829–846. [MR3760452](#) <https://doi.org/10.1002/sim.7567>
- SHARDELL, M., HICKS, G. E. and FERRUCCI, L. (2015). Doubly robust estimation and causal inference in longitudinal studies with dropout and truncation by death. *Biostatistics* **16** 155–168. [MR3365418](#) <https://doi.org/10.1093/biostatistics/kxu032>
- TSIATIS, A. A. (2006). *Semiparametric Theory and Missing Data. Springer Series in Statistics*. Springer, New York. [MR2233926](#)

- VAN DER LAAN, M. J. and GRUBER, S. (2012). Targeted minimum loss based estimation of causal effects of multiple time point interventions. *Int. J. Biostat.* **8** 9. [MR2923282](#) <https://doi.org/10.1515/1557-4679.1370>
- VAN DER LAAN, M. J., POLLEY, E. C. and HUBBARD, A. E. (2007). Super learner. *Stat. Appl. Genet. Mol. Biol.* **6** 25. [MR2349918](#) <https://doi.org/10.2202/1544-6115.1309>
- WANG, L., RICHARDSON, T. S. and ZHOU, X.-H. (2017). Causal analysis of ordinal treatments and binary outcomes under truncation by death. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **79** 719–735. [MR3641404](#) <https://doi.org/10.1111/rssb.12188>
- WU, X., WEINBERGER, K. R., WELLENIUS, G. A., DOMINICI, F. and BRAUN, D. (2024). Assessing the causal effects of a stochastic intervention in time series data: Are heat alerts effective in preventing deaths and hospitalizations? *Biostatistics* **25** 57–79. [MR4678531](#) <https://doi.org/10.1093/biostatistics/kxad002>
- XU, Y., SCHARFSTEIN, D., MÜLLER, P. and DANIELS, M. (2022). A Bayesian nonparametric approach for evaluating the causal effect of treatment in randomized trials with semi-competing risks. *Biostatistics* **23** 34–49. [MR4366034](#) <https://doi.org/10.1093/biostatistics/kxaa008>
- YANG, F. and DING, P. (2018). Using survival information in truncation by death problems without the monotonicity assumption. *Biometrics* **74** 1232–1239. [MR3908141](#) <https://doi.org/10.1111/biom.12883>
- YOUNG, J. G., STENSRUD, M. J., TCHETGEN TCHETGEN, E. J. and HERNÁN, M. A. (2020). A causal framework for classical statistical estimands in failure-time settings with competing events. *Stat. Med.* **39** 1199–1236. [MR4075855](#) <https://doi.org/10.1002/sim.8471>
- ZHANG, J. L. and RUBIN, D. B. (2003). Estimation of causal effects via principal stratification when some outcomes are truncated by “death”. *J. Educ. Behav. Stat.* **28** 353–368.