Research Article

Iván Díaz*, Nicholas Williams, and Kara E. Rudolph

Efficient and flexible mediation analysis with time-varying mediators, treatments, and confounders

Abstract: Understanding the mechanisms of action of interventions is a major general goal of scientific inquiry. The collection of statistical methods that use data to achieve this goal is referred to as *mediation analysis*. Natural direct and indirect effects provide a definition of mediation that matches scientific intuition, but they are not identified in the presence of time-varying confounding. Interventional effects have been proposed as a solution to this problem, but existing estimation methods are limited to assuming simple (e.g., linear) and unrealistic relations between the mediators, treatments, and confounders. We present an identification result for interventional effects in a general longitudinal data structure that allows flexibility in the specification of treatment-outcome, treatment-mediator, and mediator-outcome relationships. Identification is achieved under the standard no-unmeasured-confounders and positivity assumptions. In this article, we study semi-parametric efficiency theory for the functional identifying the mediation parameter, including the non-parametric efficiency bound, and was used to propose non-parametrically efficient estimators. Implementation of our estimators only relies on the availability of regression algorithms, and the estimators in a general framework that allows the analyst to use arbitrary regression machinery were developed. The estimators are doubly robust, $\sqrt{n}$-consistent, asymptotically Gaussian, under slow convergence rates for the regression algorithms used. This allows the use of flexible machine learning for regression while permitting uncertainty quantification through confidence intervals and p-values. A free and open-source R package implementing the methods is available on GitHub. The proposed estimator to a motivating example from a trial of two medications for opioid-use disorder was applied, where we estimate the extent to which differences between the two treatments on risk of opioid use are mediated by craving symptoms.

Keywords: mediation analysis, longitudinal data, sequential double robustness, efficient estimation, machine learning

MSC 2020: 62D20

1 Introduction

Mediation analyses are paramount to scientific inquiry, as they provide a way of understanding the mechanisms through which effects operate [1]. For example, recent mediation analyses have helped to uncover the types of immune response that coronavirus disease 2019 vaccines trigger in order to prevent disease [2].

---

* Corresponding author: Iván Díaz, Department of Population Health, Division of Biostatistics, New York University Grossman School of Medicine, New York, United States, e-mail: ivan.diaz@nyu.edu
Nicholas Williams, Kara E. Rudolph: Department of Epidemiology, Mailman School of Public Health, Columbia University, New York, United States

Multiple methods have been developed for mediation analysis in the setting of a mediator measured at a single time point, using a counterfactual or potential outcome framework. For example, the *natural direct effect* is defined as the effect of exposure/treatment that would have been seen in a hypothetical world where the effect operating through the mediator is disabled. Conditions for identifiability for the natural direct effect and its indirect counterpart have been given in [3] and [4].

Though the definition of the natural (in)direct effects is scientifically interesting, their identification requires assumptions about the real world, which are guaranteed to not hold in many practical situations, such as with time-varying confounding (cf. recanting witness, [5]). This restriction limits the applicability of natural direct and indirect effects in practice, as intermediate confounders are expected to be present in many applications, for example, when the effect of an intervention operates through adherence [6]. Several methods have been proposed to do away with the cross-world independence assumption and/or the assumption of no intermediate confounders, for example, partial identification methods [7–9], so-called *separable effects* [7,10], and effects defined in terms of stochastic interventions on the exposure [11,12]. In this article, we tackle this problem focusing on mediation effects defined in terms of contrasts between counterfactuals in hypothetical worlds in which the treatment is set to some value deterministically, whereas the mediator is drawn stochastically from its counterfactual distribution under interventions on treatment [13–15]. Efficient non-parametric estimators that leverage machine learning to alleviate model misspecification bias have been recently proposed for these parameters [16,17], but they are limited to single time-point exposures. These effects have been called *interventional* effects [18].

Mediation analysis methodologies for time-varying exposures and mediators are limited. Non-parametric identification formulas for interventional effects in the case of mediators and treatments measured longitudinally, assuming that the time-dependent covariates are measured either before or after the mediator, but not both, are given in refs [18,19]. The authors propose estimation methods that rely on the unlikely ability to correctly specify parametric models for the distribution of the unobservable counterfactual outcomes. Similar interventional effects where the mediator is drawn from its counterfactual distribution conditional on all the past are proposed in ref. [20], together with non-parametric efficient estimators that rely on data-adaptive regression to alleviate model misspecification bias. However, the "direct effect" defined in ref. [20] does not capture the pathway from treatment through intermediate confounder to outcome and is thus not a direct effect in the sense that we are interested in this article. Similar methods for survival analysis and treatment at a single time point have also been proposed [21]. Marginal structural models for longitudinally measured mediators under treatment at a single time point and no loss-to-follow-up are proposed in ref. [22]. More recent work includes the development of a 4-way decomposition of immediate and delayed effects using an infinite time-horizon in a reinforcement learning framework [23], and other methods considering time-varying mediators but a single exposure at baseline [21,24–27].

In this article, we develop a general longitudinal causal mediation approach that fills several gaps in the aforementioned literature. The method we propose satisfies the following: (i) the direct effect is defined in terms of the effects operating through all the pathways that do not include the mediator at any time point, (ii) the methods allow for longitudinally measured mediators and treatments, with confounders possibly measured before and after treatment, (iii) the methods allow the use of data-adaptive regression to alleviate concerns of model misspecification bias, and (iv) the methods are endowed with formulas to construct efficient estimators and computation of valid standard errors and confidence intervals, even under the use of data-adaptive regression. One limitation of prior work remains: our proposed methods can only handle categorical mediators, and the computational complexity increases with the number of categories.

The remainder of this article is organized as follows. In Section 2, we introduce the parameters of interest as well as the identification result and some simple estimators; in Section 4, we discuss efficiency theory for the interventional mediation functional, presenting estimating equations and the efficiency bound in the non-parametric model; in Section 3, we discuss simple estimators based on inverse probability weighting and sequential regression, and in Section 5, we discuss the proposed efficient estimator as well as its asymptotic properties. In Section 6, we present the results of numerical studies, and finally, in Section 7, we present the results of an illustrative study on estimating the longitudinal effect of initiating treatment for

opioid-use disorder (OUD) with extended-release naltrexone (XR-NTX) vs buprenorphine-naloxone (BUP-NX) on risk of illicit opioid use during the fourth week of treatment that operates through the mediator of craving symptoms, using symptoms of depression and withdrawal as time-varying covariates. We include weekly measures for each of the first 4 weeks of treatment.

2 Notation and definition of interventional (in)direct effects

Let $X_1, \ldots, X_n$ denote a sample of i.i.d. observations with $X = (L_1, A_1, Z_1, M_1, L_2, \ldots, A_\tau, Z_\tau, M_\tau, L_{\tau+1} = Y) \sim \mathsf{P}$, where A_t denotes a vector of intervention variables such as treatment and/or loss-to-follow-up, Z_t denotes intermediate confounders, M_t denotes a mediator of interest, and L_t denotes the time-varying covariates. The outcome of interest is a variable $Y = L_{\tau+1}$ measured at the end of the study. We let $\mathsf{P}f = \int f(x)\mathrm{d}\mathsf{P}(x)$ for a given function $f(x)$. We use P_n to denote the empirical distribution of $X_1, \ldots, X_n$ and assume P is an element of the non-parametric statistical model defined as all continuous densities on X with respect to a dominating measure ν. We let E denote the expectation with respect to P, i.e., $\mathsf{E}\{f(X)\} = \int f(x)\mathrm{d}\mathsf{P}(x)$. We also let $\|f\|^2$ denote the $L_2(\mathsf{P})$ norm $\int f^2(x)\mathrm{d}\mathsf{P}(x)$. We use $\bar{W}_t = (W_1, \ldots, W_t)$ to denote the past history of a variable W, use $\underline{W}_t = (W_t, \ldots, W_\tau)$ to denote the future of a variable, and use $H_{A,t} = (\bar{L}_t, \bar{M}_{t-1}, \bar{Z}_{t-1}, \bar{A}_{t-1})$ to denote the history of all variables up until just before A_t. The random variables $H_{Z,t}$, $H_{M,t}$, and $H_{L,t}$ are defined similarly as $H_{Z,t} = (A_t, H_{A,t})$, $H_{M,t} = (Z_t, H_{Z,t})$, and $H_{L,t} = (M_{t-1}, H_{M,t-1})$. For the complete history and past of a random variable, we sometimes simplify $\bar{W}_\tau$ and $\underline{W}_1$ as $\bar{W}$. By convention, variables with an index $t \leq 0$ are defined as null, expectations conditioning on a null set are marginal, products of the type $\prod_{t=k}^{k-1} b_t$ and $\prod_{t=0}^{0} b_t$ are equal to one, and sums of the type $\sum_{t=k}^{k-1} b_t$ and $\sum_{t=0}^{0} b_t$ are equal to zero. We let $\mathsf{g}_{A,t}(a_t|h_{A,t})$ denote the probability mass function of A_t conditional on $H_{A,t} = h_{A,t}$ and assume A_t takes values on a finite set. The function $\mathsf{g}_{M,t}(m_t|h_{M,t})$ is defined similarly, and we also assume that M_t takes values on a finite set. The variables L_t and Z_t are allowed to take values on any set; i.e., they can be multivariate, continuous, etc.

We formalize the definition of the causal effects using a non-parametric structural equation model [28]. Specifically, for each time point t, we assume the existence of deterministic functions $f_{A,t}$, $f_{Z,t}$, $f_{M,t}$, and $f_{L,t}$ such that $A_t = f_{A,t}(H_{A,t}, U_{A,t})$, $Z_t = f_{Z,t}(H_{Z,t}, U_{Z,t})$, $M_t = f_{M,t}(H_{M,t}, U_{M,t})$, and $L_t = f_{L,t}(H_{L,t}, U_{L,t})$. Here, $U = (U_{A,t}, U_{Z,t}, U_{M,t}, U_{L,t}, U_Y : t \in \{1, \ldots, \tau\})$ is a vector of exogenous variables, with unrestricted joint distribution. This model can be expressed in the form of a directed acyclic graph (DAG) as in Figure 1. In this article, we are concerned with the definition and estimation of the causal effect of an intervention on $\bar{A}$ on Y, as well as its decomposition in terms of the effect through all paths involving the mediator $\bar{M}$ versus effects through all other mechanisms that do not involve any component of $\bar{M}$. Mediation effects will be defined in terms of hypothetical interventions where the equations $A_t = f_{A,t}(H_{A,t}, U_{A,t})$ and $M_t = f_{M,t}(H_{M,t}, U_{M,t})$ are removed from the structural model, and the treatment and mediator nodes are externally assigned as follows. Let $\bar{M}(\bar{a})$ denote the counterfactual mediator vector observed in a hypothetical world where $\bar{A} = \bar{a}$. For a value $\bar{m}$ in the range of $\bar{M}$, we also define a counterfactual outcome $Y(\bar{a}, \bar{m})$ as the outcome in a hypothetical world, where $\bar{A} = \bar{a}$ and $\bar{M} = \bar{m}$. These variables are defined in terms of our assumed NPSEM in the supplementary materials.
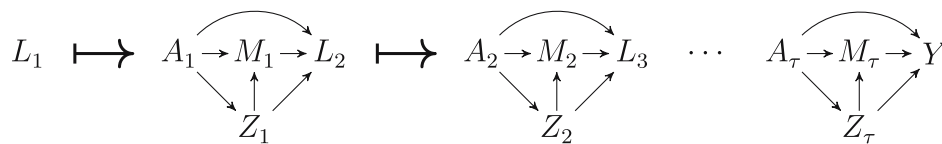

Figure 1: DAG. For simplicity, the symbol $\mapsto$ is used to indicate arrows from all nodes in its left to all nodes in its right.

Let $\bar{a}^\star = (a_1^\star, \ldots, a_\tau^\star)$ and $\bar{a}' = (a_1', \ldots, a_\tau')$ denote two user-specified values in the range of $\bar{A}$, and let $\bar{G}(\bar{a})$ denote a random draw from the distribution of the counterfactual variable $\bar{M}(\bar{a})$ conditional on $V \subseteq L_1$. We define the conditional interventional effect as $\mathsf{E}[Y(\bar{a}', \bar{G}(\bar{a}')) - Y(\bar{a}^\star, \bar{G}(\bar{a}^\star))]$ and decompose it as

$$\mathsf{E}[Y(\bar{a}', \bar{G}(a')) - Y(\bar{a}^\star, \bar{G}(\bar{a}^\star))] = \underbrace{\mathsf{E}[Y(\bar{a}', \bar{G}(\bar{a}')) - Y(\bar{a}', \bar{G}(\bar{a}^\star))]}_{\text{Indirect effect (through } \bar{M})} + \underbrace{\mathsf{E}[Y(\bar{a}', \bar{G}(\bar{a}^\star)) - Y(\bar{a}^\star, \bar{G}(\bar{a}^\star))]}_{\text{Direct effect (not through } \bar{M})}. \tag{1}$$

This is the definition of interventional effect in a longitudinal setting given by ref. [18]. In what follows, we focus on the identification and estimation of the parameters $\mathsf{E}[Y(\bar{a}', G(\bar{a}^\star))]$ for fixed $\bar{a}', \bar{a}^\star$, from which we can obtain the aforementioned effects. In a slight abuse of language, we will also refer to parameters of the type $\mathsf{E}[Y(\bar{a}', G(\bar{a}^\star))]$ as effects.

The aforementioned setup allows for the definition of causal effects for time-to-event outcomes subject to loss-to-follow-up or censoring as follows. Let $A_t = (A_{1,t}, A_{2,t})$, where $A_{1,t}$ denotes the exposure at time t, $A_{2,t}$ is equal to one if the unit remains uncensored at time $t + 1$ and zero otherwise, and $Y = L_{\tau+1}$ denotes the event-free status at the end of study follow-up. Assume monotone loss-to-follow-up so that $A_{2,t} = 0$ implies $A_{2,k} = 0$ for all $k > t$, in which case all the data for $k > t$ become degenerate. In this case, we could define the effects as above with $\bar{a}' = ((\bar{a}_{1,1}', 1), \ldots, (\bar{a}_{1,\tau}', 1))$ and $\bar{a}^\star = ((\bar{a}_{1,1}^\star, 1), \ldots, (\bar{a}_{1,\tau}^\star, 1))$, contrasting regimes where treatment at time t is set to $A_{1,t} = a_{1,t}'$ vs $A_{1,t} = a_{1,t}^\star$ while setting censoring status $A_{2,t} = 1$ as "not censored" for everyone. This definition of causal effect for time-to-event outcomes in terms of interventional effects bypasses some (but not all) problems that occur with natural effects due to the fact that a counterfactual longitudinal mediator may be truncated by death, which may render the counterfactual survival time undefined [29]. Furthermore, recent work has uncovered limitations of these effects as a tool to unveil mechanisms [30]. These and other issues related to the interpretation of these direct and indirect effects are discussed at length elsewhere. Interested readers are encouraged to consult the cited articles.

In what follows, we will use the notation $H_{A,t}'$, $H_{Z,t}'$, $H_{M,t}'$, and $H_{L,t}'$ to refer to the intervened histories of the variables under an intervention setting $\bar{A} = \bar{a}'$. For example, $H_{L,t}' = (\bar{L}_{t-1}, \bar{M}_{t-1}, \bar{Z}_{t-1}, \bar{A}_{t-1} = \bar{a}_{t-1}')$. The histories $H_{A,t}^\star$, $H_{Z,t}^\star$, $H_{M,t}^\star$, and $H_{L,t}^\star$ are defined analogously, e.g., $H_{L,t}^\star = (\bar{L}_{t-1}, \bar{M}_{t-1}, \bar{Z}_{t-1}, \bar{A}_{t-1} = \bar{a}_{t-1}^\star)$. The following assumptions will be sufficient to prove the identification of the parameter $\mathsf{E}[Y(\bar{a}', G(\bar{a}^\star))]$:

Assumption 1. (Conditional exchangeability of treatment and mediator) Assume:
(i) $A_t \perp\!\!\!\perp Y(\bar{a}', \bar{m}) | \bar{L}_t, \bar{M}_{t-1} = \bar{m}_{t-1}, \bar{Z}_{t-1}, \bar{A}_{t-1} = \bar{a}_{t-1}'$ for all $t, \bar{m}$,
(ii) $M_t \perp\!\!\!\perp Y(\bar{a}', \bar{m}) | \bar{L}_t, \bar{M}_{t-1} = \bar{m}_{t-1}, \bar{Z}_t, \bar{A}_t = \bar{a}_t'$ for all $t, \bar{m}$,
(iii) $A_t \perp\!\!\!\perp \underline{M}_t(\bar{a}^\star) | \bar{L}_t, \bar{M}_{t-1}, \bar{Z}_{t-1}, \bar{A}_{t-1} = \bar{a}_{t-1}^\star$ for all t.

Assumption 2. (Positivity of treatment and mediator assignment) Assume:
(i) $\prod_{s=1}^{t} \mathsf{p}(l_s | h_{L,s}') \prod_{s=1}^{t-1} \mathsf{p}(z_s | h_{Z,s}') > 0$ implies $\mathsf{g}_{A,t}(a_t' | h_{A,t}') > 0$ for all t,
(ii) $\prod_{s=1}^{t} \mathsf{p}(l_s | h_{L,s}') \mathsf{p}(z_s | h_{Z,s}') > 0$ implies $\mathsf{g}_{M,t}(m_t | h_{M,t}') > 0$ for all t and m_t,
(iii) $\prod_{s=1}^{t} \mathsf{p}(l_s | h_{L,s}^\star) \prod_{s=1}^{t-1} \mathsf{p}(m_s | h_{M,s}^\star) \mathsf{p}(z_s | h_{Z,s}^\star) > 0$ implies $\mathsf{g}_{A,t}(a_t^\star | h_{A,t}^\star) > 0$ for all t and all m_t.

Assumptions 1 (i), 1(ii), and 1(iii) together with the assumed DAG in Figure 1 state that $H_{A,t}$ contains all the common causes of A_t and Y, $H_{M,t}$ contains all the common causes of M_t and Y, and $H_{A,t}$ contains all the common causes of A_t and M_t, respectively. These are the standard assumptions of no-unmeasured confounders for the treatment-outcome, mediator-outcome, and treatment-mediator relations. Assumptions 2(i) and 2(ii) allow the identification of $\mathsf{E}[Y(\bar{a}', \bar{m}) | V = v]$ for all $\bar{m}$. Assumption 2(i) states that the value a_t' has positive probability conditional on values $(l_s, z_{s-1} : s \leq t)$ that have themselves a positive conditional probability of occurring. Assumption 2(ii) states that the value m_t has positive probability conditional on values $(l_s, z_s : s \leq t)$ that have a positive conditional probability of occurring. Assumption 2(iii) allows us to identify $\mathsf{P}(\bar{M}(a^\star) = \bar{m} | V = v)$ and states that $a_t^\star$ has a positive probability conditional on values $(l_s, m_{s-1}, z_{s-1} : s \leq t)$ that have positive conditional probability of occurring. Under these assumptions, we have the following identification result.

Theorem 1. (Identification) *Under Assumptions* 1 *and* 2, *the interventional effect* $\theta = \mathsf{E}[Y(\bar{a}', G(\bar{a}^\star))]$ *is identified as follows. Let*

$$\varphi(\bar{m}, v) = \int l_{\tau+1} \prod_{t=1}^{\tau} \mathsf{p}[l_{t+1}|h'_{L,t+1}]\mathsf{p}[z_t|h'_{Z,t}]\mathsf{p}(l_1|v)\mathrm{d}\nu(\bar{l}, \bar{z}),$$

$$\lambda(\bar{m}, v) = \int \prod_{t=1}^{\tau} \mathsf{p}[l_{t+1}, m_t, z_t|h^\star_{Z,t}]\mathsf{p}(l_1|v)\mathrm{d}\nu(\bar{l}, \bar{z}).$$

Then, θ *is identified as* $\theta = \mathsf{E}[\sum_{\bar{m}}\varphi(\bar{m}, V)\lambda(\bar{m}, V)].$

The aforementioned identification theorem result generalizes several identification formulas in the literature. When $\tau = 1$, $Z = \varnothing$, and $V = L_1$, this identification formula yields (by computing the corresponding contrasts) the identification formula for the natural direct and indirect effects as derived by ref. [4] under an additional cross-world counterfactual assumption. When $\tau = 1$ and $V = L_1$, and the confounder Z is present, this identification formula yields the identification formula for the interventional effect described by [15]. If $\tau > 1$, $V = L_1$, and $Z_t = \varnothing$ for all t, θ in Theorem 1 reduces to Formula (1) in ref. [18]. If $\tau > 1$ and $L_t = \varnothing$ for $t \le \tau$, θ in Theorem 1 reduces to the identification results for Figure 5 given in page 926 of [18]. Our result is a particular case of the identification results of ref. [31].

The choice of variable V in defining the effect should be guided by the total effect to decompose. For instance, if $V = L_1$ then it may be the case that the total effect in (1) will be closer to the average treatment effect $\mathsf{E}[Y(\bar{a}') - Y(\bar{a}^\star)]$ compared to $V \subset L_1$. In what follows, we assume for simplicity in the presentation that $V = \varnothing$, but the estimators presented can be easily adapted to other cases.

3 Inverse probability weighting and sequential regression

The identification formula in Theorem 1 is a complex functional of the distribution P of the observed data, and it is therefore hard to estimate. Consider, for example, a plug-in estimation strategy where the relevant components of the likelihood are estimated and then plugged in the identifying functional. This would involve finding estimates of the densities of L_t and Z_t conditional on the past data $H_{L,t}$ and $H_{Z,t}$. If L_t and Z_t are continuous or multivariate, or if $H_{L,t}$ and $H_{Z,t}$ are high-dimensional, these densities may be hard to estimate, complicating the construction of a plug-in estimator. Furthermore, the consistency of such an estimator would depend on our ability to consistently estimate the densities involved. In general, this would require the use of flexible estimators that allow automatic selection of non-linearities, interactions, etc. However, there is no statistical theory that allows the study of the sampling distribution (both asymptotic and finite-sample) of such estimator. This results in a lack of a theoretical foundation for the construction of confidence intervals, standard errors, p-values, and other quantities that are often of interest to quantify the uncertainty of statistical estimates.

To address this problem, an estimator based on the theory of doubly robust unbiased transformations [32], targeted minimum loss-based estimation [33,34], and double machine learning [35] was proposed. Concepts in these frameworks will be used to construct estimators that will allow us to use machine learning to address model misspecification bias, while also achieving asymptotic normality, allowing us to construct formulas for computation of standard errors, confidence intervals, and p-values. The asymptotically normal estimators will be constructed in Section 5. This section focuses on describing the two main building blocks: an inverse probability weighted estimator and a sequential regression estimator.

3.1 Inverse probability weighted estimation

We start with the description of our proposed estimators by observing that an alternative representation of the identifying functional allows inverse probability-weighted estimation without reliance on estimating the conditional density of L_t and Z_t. Specifically, define the following random variables:

$$G_{A,t}(H_{A,t})' = \frac{\mathbb{1}\{A_t = a_t'\}}{g_t(a_t'|H_{A,t})}, \quad G_{A,t}^\star(H_{A,t}) = \frac{\mathbb{1}\{A_t = a_t^\star\}}{g_t(a_t^\star|H_{A,t})}, \quad G_{M,t}(H_{M,t}, m_t) = \frac{\mathbb{1}\{M_t = m_t\}}{g_{M,t}(M_t|H_{M,t})},$$

as well as $K_{l,u}' = \prod_{r=l}^u G_{A,r}'$, $K_{l,u}^\star = \prod_{r=l}^u G_{A,r}^\star$, $H_{l,u} = \prod_{r=l}^u G_{M,r}$. Note that we have omitted the dependence on $H_{A,t}$, $H_{M,t}$, and m_t, and we will do so whenever there is no ambiguity. Then, the identifying functionals in Theorem 1 can also be written as $\varphi(\bar{m}) = E\{K_{1,\tau}' H_{1,\tau} Y\}$; $\lambda(\bar{m}) = E\{K_{1,\tau}^\star \mathbb{1}\{\bar{M} = \bar{m}\}\}$.

The aforementioned formulas reveal a simple way to construct an estimator of θ using re-weighting. First, models for the probability distributions of A_t and M_t could be constructed, e.g., using flexible regression or classification methods. Then, these models are used to compute estimates of the weights $K_{1,\tau}'$, $H_{1,\tau}$, and $K_{1,\tau}^\star$. Then, these weights are used to estimate φ and λ, where the outer expectations of the aforementioned display are estimated using empirical means of weighted outcomes Y and $\mathbb{1}\{\bar{M} = \bar{m}\}$. These estimators of φ and λ can then be plugged in the formula for θ. While this estimator is attractive for its simplicity, it has two major drawbacks. First, it relies on the ability to correctly estimate the treatment and mediator conditional distributions. If flexible estimators from the regression or classification literature (e.g., generalized linear models with ℓ_1 penalties, and random forests) are used to increase the chances of correct specification, there are no general theoretical results that allow the computation of standard errors of confidence intervals with correctness guarantees. Furthermore, this estimator is generally inefficient in the non-parametric model.

Due to the aforementioned issues, the use of the inverse probability-weighted estimator is generally discouraged in practice. However, the weights K and H will be fundamental to construct efficient and asymptotically normal estimators in Section 5, and for the construction of the efficiency bound in Section 4.

3.2 Sequential regression representation of the mediational g-computation formula

The construction of efficient and asymptotically normal estimators in Section 5 will rely on sequential regression procedures to estimate φ and λ. Those estimators are motivated by an alternative representation of θ in Theorem 1 in the form of sequential regressions, which we present in this section. This approach was first proposed by Bang and Robins [36] and has become standard in the estimation of total causal effects in longitudinal studies [34,37].

The sequential regression procedure computes the integrals in Theorem 1 through an alternative representation that iterates in time starting at the last time point of the study $t = \tau + 1$ and ending at $t = 0$. To start the iteration, set $Q_{Z,\tau+1} = Y$. For fixed values $\bar{a}'$ and $\bar{m}$, and for $t = \tau, \dots, 0$, recursively define the random variables

$$Q_{L,t}(H_{M,t}, \underline{m}_t) = E[Q_{Z,t+1}(H_{A,t+1}, \underline{m}_{t+1})|M_t = m_t, H_{M,t}], \tag{2}$$

$$Q_{Z,t}(\bar{H}_{A,t}, \underline{m}_t) = E[Q_{L,t}(H_{M,t}, \underline{m}_t)|A_t = a_t', H_{A,t}], \tag{3}$$

where we note that $Q_{L,t}$ is a random variable through its dependence on $H_{M,t}$, but that $\underline{m}_t$ is a non-random user-given value. Likewise, $Q_{Z,t}$ is a random variable only through its dependence on $H_{A,t}$. To simplify notation, we will sometimes omit the dependence of the aforementioned functions on $H_{A,t}$, $H_{M,t}$, and $\bar{m}$. In the proof of Proposition 1 (available in the Supplementary Materials), we show that $\varphi(\bar{m}) = Q_{L,0}(\bar{m})$. The counterfactual distribution $\lambda(\bar{m})$ may be identified as follows. Let $Q_{M,\tau+1} = 1$. For $t = \tau, \dots, 0$, recursively define

$$Q_{M,t}(H_{A,t}, \underline{m}_t) = E[\mathbb{1}\{M_t = m_t\}Q_{M,t+1}(H_{A,t+1}, \underline{m}_{t+1})|A_t = a_t^\star, H_{A,t}].$$

Then, we have $\lambda(\bar{m}) = Q_{M,0}(\bar{m})$. This leads to the following alternative expression for the mediational g-computation formula in terms of sequential regressions.

Proposition 1. (Sequential regression representation of the longitudinal mediation g-formula) *For θ defined in Theorem* 1, *we have $\theta = \sum_{\bar{m}} Q_{L,0}(\bar{m}) Q_{M,0}(\bar{m})$.*

We illustrate the aforementioned result for the case $\tau = 2$ in the supplementary materials. This alternative expression of the longitudinal mediation formula allows the construction of an estimator by sequential regression, estimating the parameters $Q_{L,t}(\bar{m})$ and $Q_{Z,t}(\bar{m})$ sequentially for $t = \tau, \ldots, 0$ to obtain an estimate of $Q_{L,0}(\bar{m})$, and then estimating $Q_{M,t}(\bar{m})$ sequentially for $t = \tau, \ldots, 0$ to obtain an estimate of $Q_{M,0}(\bar{m})$. The algorithm requires the estimation of sequential regressions as in (2) and (3). There are at least two alternatives to estimate these sequential regressions. The first is to perform sequential regression separately for each $\bar{m}$ in the range of $\bar{M}$. The second is to construct a pooled dataset where we pool all values $\bar{m}$ in the range of $\bar{M}$ to obtain the estimates for all $\bar{m}$ from a single sequential regression procedure. In this article, we pursue the second approach. In this setup, the values $\bar{m}$ will play the role of another variable to be included in the regression algorithms. Specifically, note that estimating $Q_{L,0}(\bar{m})$ for each $\bar{m}$ would, in principle, require fitting $\prod_{k=1}^{\tau} J_k$ sets of τ sequential regressions as in formula (2), one set of sequential regressions for each one of the possible values of $\bar{m}$. This can be computationally prohibitive. For example, estimating $Q_{L,0}(\bar{m})$ in a study with $\tau = 5$ and a mediator taking on three different values involves $3^5 = 243$ sets of five sequential regressions, each in a dataset of size n. To alleviate computational complexity, we propose to fit regressions in pooled datasets constructed sequentially from $t = \tau$ to $t = 1$, where at time point t, each observation in the dataset at time point $t + 1$ is repeated J_t times. Under this data pooling approach, estimating $Q_{L,0}$ in a study with $\tau = 5$ and a mediator taking on three different values involves five sequential regressions, each in datasets of size $3n$, 3^2n, 3^3n, 3^4n, and 3^5n, respectively, where the predictor set decreases in size as the dataset increases in size. In this setting, smoothing on the mediator values $\bar{m}$ may be carried out through any regression procedure such as a parametric model or a more flexible regression approach from the machine or statistical learning literature. A detailed algorithm with pseudo-code for computation of a sequential regression estimator is given in Algorithm S1 in the supplementary materials.

If the pooled sequential regressions are performed within *a priori* correctly specified parametric models, then estimators of θ based on Proposition 1 may be shown to be CAN, and the delta method or the non-parametric bootstrap may be used to construct confidence intervals. However, positing correct parametric models for the sequential regressions involved is generally unattainable *a priori*. This is especially difficult to do in the longitudinal setting. Even if it was known that the data Z followed a given parametric model, correctly specifying models for the sequential regressions such that they are congenial with the model requires specialized models and might be impossible [38]. Furthermore, in most cases, especially with a large number of variables, the data Z may not be accurately modeled using parametric models and may require data-adaptive regression (e.g., machine learning) tools that offer flexibility in modeling non-linearities and interactions are necessary to attain consistency of the sequential regressions and therefore consistency of the estimator of θ.

Under model selection or data-adaptive regression, the sampling distribution of the aforementioned sequential regression estimator is generally unknown, which hinders computation of confidence intervals and other uncertainty measures. In the next section, we discuss efficiency theory for estimation of θ, which will allow the use of data-adaptive regression techniques while also allowing the computation of valid (under assumptions) standard errors and confidence intervals.

4 Efficiency theory

The foundations of our proposed efficient estimation approach are in semi-parametric estimation theory [e.g., 39–42] and in the theory for doubly robust estimation of causal effects using sequential regression [e.g., 33,34,36,37,43–47]. Central to this theory is the study of the non-parametric efficient influence function (EIF) or canonical gradient, which characterizes the efficiency bound of the longitudinal mediation functional given in Theorem 1 and allows the development of estimators under slow convergence rates for the nuisance parameters involved [41]. Specifically, our estimators will involve finding von-Mises-type approximations for the parameters $Q_{L,t}(\bar{m})$, $Q_{Z,t}(\bar{m})$, and $Q_{M,t}(\bar{m})$, which can be intuitively understood as

first-order expansions with second-order error remainder terms. Because the errors in the expansion are second-order, this will mean that the resulting estimator of $\hat{\theta}$ will be consistent and asymptotically normal at rate $n^{1/2}$ as long as the second-order error terms converge to zero at rate $n^{1/2}$. This convergence rate would be satisfied, for example, if the all regression functions were used for estimation converge at rate $n^{1/4}$. We will elaborate on this discussion in Section 5 when we present the asymptotic normality theorems for the proposed estimators.

Let $\eta = \{\mathsf{G}'_{A,t}, \mathsf{G}^{\star}_{A,t}, \mathsf{G}_{M,t}, \mathsf{Q}_{Z,t}, \mathsf{Q}_{M,t}, \mathsf{Q}_{L,t} : t = 1, \ldots, \tau\}$. For $t = 0, \ldots, \tau$, define

$$\mathsf{D}_{L,t}(\underline{X}_t, \underline{m}_t) = \sum_{s=t}^{\tau} \mathsf{K}'_{t+1,s}\mathsf{H}_{t,s}\{\mathsf{Q}_{Z,s+1} - \mathsf{Q}_{L,s}\} + \sum_{s=t+1}^{\tau} \mathsf{K}'_{t+1,s}\mathsf{H}_{t,s-1}\{\mathsf{Q}_{L,s} - \mathsf{Q}_{Z,s}\} + \mathsf{Q}_{L,t}. \tag{4}$$

$$\mathsf{D}_{Z,t}(\underline{X}_t, \underline{m}_t) = \sum_{s=t}^{\tau} \mathsf{K}'_{t,s}\mathsf{H}_{t,s}\{\mathsf{Q}_{Z,s+1} - \mathsf{Q}_{L,s}\} + \sum_{s=t}^{\tau} \mathsf{K}'_{t,s}\mathsf{H}_{t,s-1}\{\mathsf{Q}_{L,s} - \mathsf{Q}_{Z,s}\} + \mathsf{Q}_{Z,t} \tag{5}$$

$$\mathsf{D}_{M,t}(\underline{X}_t, \underline{m}_t) = \sum_{s=t}^{\tau} \mathsf{K}^{\star}_{t,s}\left(\prod_{k=t}^{s-1}\mathbb{1}(M_k = m_k)\right)\{\mathbb{1}(M_s = m_s)\mathsf{Q}_{M,s+1} - \mathsf{Q}_{M,s}\} + \mathsf{Q}_{M,t}. \tag{6}$$

Whenever necessary, we make explicit the dependence of these functions on η using notation such as $\mathsf{D}_{Z,t}(\eta)$ or $\mathsf{D}_{Z,t}(\underline{X}_t, \underline{m}_t; \eta)$.

Lemma 1. (von-Mises-type approximation for $\mathsf{Q}_{Z,t}$, $\mathsf{Q}_{L,t}$, and $\mathsf{Q}_{M,t}$) *Let $\tilde{\eta}$ denote an arbitrary value of η. For second-order terms $\mathsf{R}_{L,t}(\eta, \tilde{\eta})$, $\mathsf{R}_{Z,t}(\eta, \tilde{\eta})$, and $\mathsf{R}_{M,t}(\eta, \tilde{\eta})$, we have the following first-order expansions:*

$$\mathsf{Q}_{L,t} = \mathsf{E}[\mathsf{D}_{Z,t+1}(\tilde{\eta})|M_t = m_t, H_{M,t}] + \mathsf{R}_{L,t}(\eta, \tilde{\eta}),$$
$$\mathsf{Q}_{Z,t} = \mathsf{E}[\mathsf{D}_{L,t}(\tilde{\eta})|A_t = a'_t, H_{A,t}] + \mathsf{R}_{Z,t}(\eta, \tilde{\eta}),$$

where we let $\mathsf{D}_{Z,\tau+1}(\tilde{\eta}) = Y$. For $H_{A,t}$ such that $\bar{M}_{t-1} = \bar{m}_{t-1}$, we also have

$$\mathsf{Q}_{M,t} = \mathsf{E}[\mathbb{1}\{M_t = m_t\}\mathsf{D}_{M,t+1}(\tilde{\eta})|A_t = a^{\star}_t, H_{A,t}] + \mathsf{R}_{M,t}(\eta, \tilde{\eta}),$$

where we let $\mathsf{D}_{M,\tau+1}(\tilde{\eta}) = 1$.

The terms $\mathsf{R}_{L,t}(\eta, \tilde{\eta})$, $\mathsf{R}_{Z,t}(\eta, \tilde{\eta})$, and $\mathsf{R}_{M,t}(\eta, \tilde{\eta})$ are the second-order error terms involving expectations of products of errors such as $(\tilde{\mathsf{G}}_{M,s} - \mathsf{G}_{M,s})(\tilde{\mathsf{Q}}_{L,s} - \mathsf{Q}_{L,s})$, $(\tilde{\mathsf{G}}'_{A,s} - \mathsf{G}'_{A,s})(\tilde{\mathsf{Q}}_{Z,s} - \mathsf{Q}_{Z,s})$, and $(\tilde{\mathsf{G}}^{\star}_{A,s} - \mathsf{G}^{\star}_{A,s})(\tilde{\mathsf{Q}}_{M,s} - \mathsf{Q}_{M,s})$, and their explicit form is given in the supplementary materials. This lemma and specifically the second-order form of these remainder terms have important implications in terms of estimation. Specifically, this lemma says that for any value $\tilde{\eta}$, which could represent an inconsistent estimator, regressing $\mathsf{D}_{Z,t+1}(\tilde{\eta})$ on $H_{M,t}$ among units with $M_t = m_t$ yields a consistent estimator of $\mathsf{Q}_{L,t}$, as long as $\mathsf{R}_{L,t}(\eta, \tilde{\eta})$ is small. Because $\mathsf{R}_{L,t}(\eta, \tilde{\eta})$ is a sum of products of errors, it may be reasonable to assume that this term is small for data-adaptive regression estimators of η. Specifically, in Section 5, consistency and asymptotic normality of the estimators will require that $\mathsf{R}_{L,t}(\eta, \hat{\eta})$ converges to zero in probability at rate $n^{-1/2}$ or faster. The second-order term structure of this remainder term means that this fast convergence rate can be achieved under slower convergence rates for each of the components of $\hat{\eta}$. For example, it will be achievable if all the components of $\hat{\eta}$ converge in probability to their correct limits at rate $n^{-1/4}$. This kind of rate is achievable by several data-adaptive regression methods, such as ℓ_1 regularization [48], regression trees [49], neural networks [50], or the highly adaptive lasso [51]. Analogous considerations apply to the estimation of $\mathsf{Q}_{Z,t}$ and $\mathsf{Q}_{M,t}$.

Remark 1. (Generalization to $V \neq \varnothing$) When $V = \varnothing$ (as we have assumed throughout), Lemma 1 can be used to prove that the function $\mathsf{D}_{Z,1}$ is the uncentered EIF of $\varphi(\bar{m})$. If V is discrete, EIF of $\varphi(\bar{m}, v)$ will have an additional term that accounts for conditioning on V. When V is continuous or high-dimensional, $\varphi(\bar{m}, v)$ is not pathwise differentiable, and EIF does not exist. Nonetheless, the lemma shows that $\mathsf{D}_{Z,1}$ is a doubly robust estimating function for $\varphi(\bar{m}, v)$ (meaning that the corresponding estimators will have second-order errors, see [32]), regardless of the dimension of V, which implies that estimators of θ that use this lemma for estimating $\varphi(\bar{m}, v)$ will satisfy an appropriately modified version of the asymptotic properties outlined below.

An application of the delta method along with Lemma 1 yields the following EIF for the estimation of θ in the non-parametric model (see the supplementary materials for a proof).

Theorem 2. (EIF for θ) *The EIF for θ in the non-parametric model is given by* $\mathsf{S}(X, \eta) = \sum_{\bar{m} \in \bar{\mathcal{M}}} [\{\mathsf{D}_{Z,1}(X, \bar{m}; \eta) - \varphi(\bar{m})\}\lambda(\bar{m}) + \{\mathsf{D}_{M,1}(X, \bar{m}; \eta) - \lambda(\bar{m})\}\varphi(\bar{m})]$.

This implies that the non-parametric efficiency bound for the estimation of θ is $\mathsf{Var}[\mathsf{S}(X, \eta)]$ and that an efficient estimator of θ will satisfy $\sqrt{n}(\hat{\theta} - \theta) = \frac{1}{\sqrt{n}}\sum_{i=1}^{n}\mathsf{S}(X_i, \eta) + o_{\mathsf{P}}(1)$, licensing the construction of Wald-type confidence intervals and hypothesis tests based on the central limit theorem. In the following section, we will describe an algorithm to construct such an estimator.

5 Efficient estimation using sequential doubly robust regression

Motivated by Lemma 1, we construct a sequentially doubly robust estimator as follows. Consider pseudo-outcomes $\hat{\mathsf{D}}_{Z,t+1}$, $\hat{\mathsf{D}}_{L,t}$, and $\hat{\mathsf{D}}_{M,t+1}$. The estimation algorithm proceeds sequentially from $t = \tau, \ldots, 1$ performing regression of these pseudo-outcomes in datasets $\mathcal{D}_t^+$ that pool over values of m_t. Specifics of this sequentially doubly robust estimator are presented in Algorithm 1.

Algorithm 1: Sequentially doubly robust estimator of θ

---

1      $\hat{\mathsf{D}}_{Z,\tau+1} \leftarrow Y$;

2      $\hat{\mathsf{D}}_{M,\tau+1} \leftarrow 1$;

3   **for** $t = \tau, \ldots, 1$ **do**

4          $\mathcal{D}_t^+ \leftarrow \mathcal{D}_{t+1}^+ \times \mathcal{M}_t$;

5
6          $\hat{\mathsf{Q}}_{L,t} \leftarrow \text{RegressAndPredict}(\text{outcome}=\hat{\mathsf{D}}_{Z,t+1}, \text{predictors} = (\underline{m}_t, H_{M,t}), \text{training set}$

7              $= \text{Subset}(\mathcal{D}_t^+, M_t = m_t), \text{prediction set} = \mathcal{D}_t^+)$;

8          $\hat{g}_t \leftarrow \text{RegressAndPredict}(A_t, H_{A,t}, \mathcal{D}, \mathcal{D}_t^+)$;

9
10         $\hat{g}_{M,t} \leftarrow \text{RegressAndPredict}(M_t, H_{M,t}, \mathcal{D}, \mathcal{D}_t^+)$;

11         $\hat{G}'_{A,t} \leftarrow \mathbb{1}\{A_t = a'_t\}/\hat{g}_t(A_t|H_{A,t})$;

12
13         $\hat{G}^{\star}_{A,t} \leftarrow \mathbb{1}\{A_t = a^{\star}_t\}/\hat{g}_t(A_t|H_{A,t})$;

14         $\hat{G}_{M,t} \leftarrow \mathbb{1}\{M_t = m_t\}/\hat{g}_{M,t}(M_t|H_{M,t})$;

15
         $\hat{\mathsf{D}}_{L,t} \leftarrow \mathsf{D}_{L,t}(\hat{\eta})$ using formula (4) and data $\mathcal{P}_{v,t}^+$;

         $\hat{\mathsf{Q}}_{Z,t} \leftarrow \text{RegressAndPredict}(\hat{\mathsf{D}}_{L,t}, (\underline{m}_t, H_{A,t}), \text{Subset}(\mathcal{D}_t^+, A_t = a'_t), \mathcal{D}_t^+)$;

         $\hat{\mathsf{Q}}_{M,t} \leftarrow \text{RegressAndPredict}(\mathbb{1}\{M_t = m_t\}\hat{\mathsf{D}}_{M,t+1}, (\underline{m}_t, H_{A,t}), \text{Subset}(\mathcal{D}_t^+, A_t = a^{\star}_t), \mathcal{D}_t^+)$;

         $\hat{\mathsf{D}}_{Z,t} \leftarrow \mathsf{D}_{Z,t}(\hat{\eta})$ using formula (5) and data $\mathcal{D}_t^+$;

         $\hat{\mathsf{D}}_{M,t} \leftarrow \mathsf{D}_{M,t}(\hat{\eta})$ using formula (6) and data $\mathcal{D}_t^+$;

16   **end**

17      $\hat{\varphi}_v(\bar{m}) \leftarrow \text{Mean}(\hat{\mathsf{D}}_{Z,1}(\bar{m}, L_1), \text{data} = \text{Subset}(\mathcal{D}_1^+, \bar{M} = \bar{m}))$;

18      $\hat{\lambda}_v(\bar{m}) \leftarrow \text{Mean}(\hat{\mathsf{D}}_{M,1}(\bar{m}, L_1), \text{data} = \text{Subset}(\mathcal{D}_1^+, \bar{M} = \bar{m}))$;

19      $\hat{\theta} \leftarrow \sum_{\bar{m}} \hat{\varphi}(\bar{m})\hat{\lambda}(\bar{m})$

Sequential doubly robust estimators decouple estimation of conditional expectations at time t from consistent estimation of sequences of estimators at time $t + 1, \ldots, \tau$, thereby achieving extra robustness to model misspecification. To introduce sequential double robustness, define the data-

dependent parameters $\check{Q}_{L,t} = E[D_{Z,t+1}(\hat{\eta})|M_t = m_t, H_{M,t}]$, $\check{Q}_{Z,t} = E[D_{L,t}(\hat{\eta})|A_t = a'_t, H_{A,t}]$, and $\check{Q}_{M,t} = E[\mathbb{1}\{M_t = m_t\}D_{M,t+1}(\hat{\eta})|A_t = a^*_t, H_{A,t}]$, where the expectation is with respect to the distribution of X, i.e., the estimator $\hat{\eta}$ is considered fixed. Then, we have the following result.

Lemma 2. (Sequential double robustness) *Assume that, at each time point t, we have*

(i) $\|\hat{G}'_{A,t} - G'_{A,t}\| = o_P(1)$ *or* $\|\hat{Q}_{Z,t} - \check{Q}_{Z,t}\| = o_P(1)$,

(ii) $\|\hat{G}_{M,t} - G_{M,t}\| = o_P(1)$ *or* $\|\hat{Q}_{L,t} - \check{Q}_{L,t}\| = o_P(1)$, *and*

(iii) $\|\hat{G}^*_{A,t} - G^*_{A,t}\| = o_P(1)$ *or* $\|\hat{Q}_{M,t} - \check{Q}_{M,t}\| = o_P(1)$,

then we have $\hat{\theta} = \theta + o_P(1)$, for $\hat{\theta}$ defined in Algorithm 1.

Here, we emphasize that the error terms $\|\hat{Q}_{Z,t} - \check{Q}_{Z,t}\|$, $\|\hat{Q}_{L,t} - \check{Q}_{L,t}\|$, and $\|\hat{Q}_{M,t} - \check{Q}_{M,t}\|$ only depend on the consistency of the regression procedure used at time point t and not on the consistency of estimators at any time point $s > t$. To illustrate why this is important, consider an alternative estimation strategy in which the parameters $Q_{L,t}$, $Q_{Z,t}$, and $Q_{M,t}$ are estimated directly (i.e., based on (2) and (3)) using flexible regression techniques, and the estimators of flexible regression techniques, and the estimators of $\varphi(\bar{m})$ and $\lambda(\bar{m})$ are constructed using a one-step estimator [42] by taking the empirical mean of $D_{L,1}(X, \bar{m}; \hat{\eta})$ and $D_{M,1}(X, \bar{m}; \hat{\eta})$. Note that the functions $Q_{L,t}$, $Q_{Z,t}$, and $Q_{M,t}$ are the functions of the sequences $\{Q_{L,t+1}, \dots, Q_{L,\tau}\}$, $\{Q_{Z,t+1}, \dots, Q_{Z,\tau}\}$, and $\{Q_{M,t+1}, \dots, Q_{M,\tau}\}$. It therefore appears that consistent estimation of these parameters at time t requires consistent estimation of these future sequences. In fact, an analysis of such a one-step estimator would yield error terms of the form $\|\hat{Q}_{Z,t} - Q_{Z,t}\|$. Unlike our error terms (which are stated in terms of convergence to $\check{Q}_{Z,t}$), convergence of these error terms would require consistency of $\hat{Q}_{L,s}$ and $\hat{Q}_{Z,s+1}$ for all $s \geq t$. The sequential regression construction using pseudo-outcomes avoids this problem. This lemma is the result of an application of Lemma 4 in [52]; more discussion on sequential double robustness may be found in [34,37,53].

5.1 Cross-fitting and asymptotic normality

Although the estimator of Algorithm 1 is sequentially doubly robust, thereby offering protection against model misspecification, it is not generally asymptotically normal nor efficient. Proving asymptotic normality and efficiency of this estimator would require the use of regression estimators that output functions in Donsker classes, which can be roughly understood as functions with bounded complexity [54]. This would importantly limit the ability to use flexible regression estimators that can further aid to mitigate model misspecification bias.

To solve the aforementioned problem, we will propose a cross-fitted sequentially doubly robust estimator, which helps to avoid assumptions on the complexity of the regression estimators used [35,55–57]. Cross-fitting will allow us to fully exploit the EIF given in Theorem 2 to construct estimators that are efficient and CAN under the use of flexible regression algorithms from the statistical and machine learning literature. These regression estimators are often better than parametric models at capturing the true functional form of the regression functions, therefore making the error terms $R_{L,0}(\hat{\eta}, \eta)$, $R_{Z,0}(\hat{\eta}, \eta)$, and $R_{M,0}(\hat{\eta}, \eta)$ small.

The cross-fitting algorithm is described in Algorithm S2, and the cross-fitted sequential regression algorithm using pooled datasets are described in Algorithm S3 in the supplementary materials. An R package implementing the algorithm is available on GitHub, and simulations testing the correctness of our implementation are given in the supplementary materials

Our proposed cross-fitted estimation algorithm satisfies sequential double robustness in the sense of Lemma 2. Furthermore, it satisfies the following weak convergence result.

Theorem 3. (Weak convergence of $\hat{\theta}$) *Assume that, at each time point t, we have*

(i) $\|\hat{G}'_{A,t} - G'_{A,t}\| \, \|\hat{Q}_{Z,t} - \check{Q}_{Z,t}\| = o_P(n^{-1/2})$,

(ii) $\|\hat{G}_{M,t} - G_{M,t}\| \, \|\hat{Q}_{L,t} - \check{Q}_{L,t}\| = o_P(n^{-1/2})$,

(iii) $\|\hat{G}^{\star}_{A,t} - G^{\star}_{A,t}\| \, \|\hat{Q}_{M,t} - \check{Q}_{M,t}\| = o_P(n^{-1/2})$, *and*

(iv) $P(G'_{A,t} < c) = P(G^{\star}_{A,t} < c) = P(G_{M,t} < c) = 1$ *for some* $c < \infty$.

Then, we have $\sqrt{n}(\hat{\theta} - \theta) \rightsquigarrow N(0, \sigma^2)$, *where* $\sigma^2 = \text{Var}[S(X; \eta)]$ *is the non-parametric efficiency bound.*

Note that the weak convergence of $\hat{\theta}$ requires a consistent estimation of the sequential regression functions at the rates stated in the theorem. Data-adaptive regression methods avoid reliance on parametric assumptions to achieve consistent estimation. Theorem 3 allows the construction of confidence intervals as $\hat{\theta} \pm z_{\alpha/2}\hat{\sigma}/\sqrt{n}$, where $\hat{\sigma}^2$ is the empirical variance of $S(X; \hat{\eta})$ and z_α is the quantile of a standard normal distribution.

The main difference between the construction in Algorithm 1 and Algorithm A3 is the introduction of sample splitting and the introduction of formulas to compute standard errors. Sample splitting requires careful handling of the difference training sets (denoted with a $\mathcal{T}$ in the algorithm) and prediction sets (denoted with a $\mathcal{P}$ in the algorithm) to ensure that the data for unit i are not used in the training set of regression functions used to make predictions for that unit. This allows the analysis of the estimators by conditioning on the training sample, and then "averaging" across training samples to obtain the final result.

6 Numerical studies

In this section, we perform a numerical simulation study with three main narrow but important objectives:
(1) To test the implementation of the algorithms in the R package lcm and demonstrate that the package produces correct answer in the sense that the estimator performance matches what is expected from theory.
(2) To provide an illustration of the performance of the estimator under the practical use case where machine learning is used for the estimation of the nuisance quantities.
(3) To illustrate the sequential double robustness of the estimator, which allows the algorithm to recover from inconsistent estimation of the outcome regression.

We achieve these two aims by constructing a Monte Carlo simulation as follows. First, we generate datasets with $\tau = 2$ time points from a distribution with binary A_t and Y, and where L_t, Z_t, and M_t are the categorical variables that can take on three levels each. All probability distributions are the functions of the type

$$P(W = k | H_W = h_W) = 0.8 \times \frac{\exp(h^t_{2,W}\beta_k)}{\sum_{k=1}^{K} \exp(h^t_{2,W}\beta_k)} + 0.1,$$

where W is one of A_t, Z_t, M_t, L_t, and Y, and $h^t_{2,W}$ is a vector with all main terms and two-way interactions for h^t_W. For A_t and Y, we set $K = 2$, and for L_t, Z_t, and M_t, we set $K = 3$. The coefficients β_k used in the simulation may be found on GitHub along with all other simulation codes.

In principle, estimation could be carried out using a non-parametric maximum-likelihood estimator (NPMLE) because the data are categorical. However, the data-generating mechanism consists of $3^6 \times 2^3 = 5,832$ categories, which would imply prohibitive sample sizes to obtain an NPMLE with good properties. We therefore use our developed estimators to compute estimates of the direct and indirect effects with $\bar{a}' = 1$ and $\bar{a}^\star = 0$, using flexible regression methods. We generated 1,000 datasets of sizes $n \in \{500, 1,000, 5,000\}$ and applied the estimators proposed. We then approximated quantities such as bias and coverage probabilities as empirical quantities across simulated datasets.

To achieve the first objective of this numerical study, all of the estimators of nuisance parameters in η consist of saturated logistic regression with an ℓ_1 regularization penalty, where the regularization parameter

Table 1: Simulation results for HAL estimation of nuisance parameters η

n	Estimand	|Bias|	$\sqrt{n} \times$ |Bias|	$n \times$ Var.	$n \times$ MSE	95% CI coverage
500	Direct	0.004	0.080	2.912	2.908	0.96
	Indirect	0.004	0.098	0.180	0.189	0.99
	Total	0.008	0.178	3.165	3.186	0.95
1,000	Direct	0.002	0.070	2.340	2.337	0.97
	Indirect	0.001	0.042	0.111	0.113	0.99
	Total	0.001	0.028	2.488	2.481	0.97
5,000	Direct	0.000	0.002	2.129	2.122	0.95
	Indirect	0.001	0.056	0.059	0.061	0.95
	Total	0.001	0.058	2.204	2.200	0.94

is chosen with cross-validation. This estimator is a particular case of the highly adaptive lasso estimator (HAL [51]), which has been proven to converge at the rates required by Theorem 3. This implies that under this estimation strategy, a correct implementation of our estimator should yield bias that converges to zero at $\sqrt{n}$-rate, as well as confidence intervals with the nominal coverage. The results of this simulation are presented in Table 1.

The results in Table 1 illustrate that the bias, indeed, converges to zero at rate $\sqrt{n}$, as expected. Nonetheless, the coverage of the confidence intervals is larger than the nominal value. This is due to the use of a non-targeted estimator of the variance of the estimator; this phenomenon has been observed before for this type of estimator (e.g., [58,59]). Alternative variance estimators that alleviate this problem have been proposed for other contexts [60]; the extension of those estimators to our problem will be the subject of future work.

To achieve the second objective of illustrating the performance of the methods in a typical setting under the use of typical data-adaptive regression methods in the statistics and machine learning literature, we fit the regressions in η with an ensemble method known as the super-learner [61], which builds a combination of algorithms in a user-supplied list of algorithms. The weights in the combination are chosen using cross-validation, and the super-learner has been shown to have certain desirable oracle optimality properties [62]. As candidate learners, we chose two estimators: (i) a logistic or linear (depending on whether the outcome takes values on {0, 1} or not) regression model with only main terms and (ii) a tree ensemble built using gradient boosting machines with hyper-parameters tuned using cross-validation. The results are presented in Table 2 and show good performance of the estimators in terms of both coverage of the confidence intervals and bias.

Note that neither of the candidates in the super learner estimator used in this simulation contains the true data-generating mechanisms. However, the candidates in the library are the type of flexible estimators that we encourage users of our approach to use in practice. Therefore, this table serves as an illustration of estimator consistency in practical situations where the outcome sequential regressions may be unknown

Table 2: Simulation results for super-learner estimation of nuisance parameters η

n	Estimand	|Bias|	$\sqrt{n} \times$ |Bias|	$n \times$ Var.	$n \times$ MSE	95% CI coverage
500	Direct	0.007	0.159	2.441	2.458	0.97
	Indirect	0.000	0.001	0.223	0.223	0.98
	Total	0.007	0.157	2.405	2.422	0.98
1,000	Direct	0.003	0.104	2.394	2.397	0.94
	Indirect	0.001	0.019	0.136	0.136	0.99
	Total	0.003	0.085	2.264	2.264	0.95
5,000	Direct	0.003	0.185	2.071	2.098	0.95
	Indirect	0.000	0.025	0.076	0.077	0.96
	Total	0.002	0.160	2.040	2.059	0.95

but are estimated using flexible regression. We caution, however, that performance of confidence intervals can vary wildly under model misspecification, that our theoretical results do not guarantee correct coverage under general forms of model misspecification, and that our simulation results for coverage and $\sqrt{n}$-bias should not be taken as indicative of general performance.

To achieve the third objective of illustrating the sequential double robustness of the estimator, we fit the regressions for the outcome Y using a logistic regression with main terms (inconsistent), the models for A_2 and M_2 using HAL (consistent), the models for L_2 and Z_2 using HAL (consistent), and the models for A_1 and M_1 using logistic regression with main terms (inconsistent). The results in Table 3 illustrate that this scenario leads to estimators with low large-sample bias (possibly consistent), but higher $\sqrt{n}$-bias (possibly not $\sqrt{n}$-consistent), as predicted by theory. As with the results of Table 2, the correct approximately coverage observed in this simulation should not be expected in general misspecification situations.

7 Illustrative application

We applied our proposed estimators to a longitudinal mediation question from a comparative effectiveness trial of XR-NTX vs BUP-NX for the treatment of OUD [63]. Specifically, we were interested in estimating the extent to which differences in use of illicit opioids during the first month of treatment between the two medications was due to mediation by self-reported craving of opioids, among those completing the detoxification requirement and initiating treatment. This involved estimating the interventional direct effect of being treated with XR-NTX vs BUP-NX on risk of using illicit opioids during the first 4 weeks of treatment, not operating through differences in craving of opioids; and the interventional indirect effect of being treated with XR-NTX vs BUP-NX on risk of using illicit opioids during the first 4 weeks of treatment that did operate through differences in craving. Patients report less opioid use when on XR-NTX vs BUP-NX (as well as methadone) [64,65], but the reasons underlying this difference are not well understood. In this study, patients were randomized to receive XR-NTX or BUP-NX. At time of randomization, a large number (over 30) of baseline covariates, denoted L_1 (listed in the supplemental materials), were measured. Although patient assignment to XR-NTX vs BUP-NX was randomized, we are estimating effects only among those who initiated treatment. Treatment initiation is not randomized and likely depends on patient characteristics. Initiation is also more difficult for those assigned to XR-NTX, because it requires complete detoxification [63]. Our adjustment for an extensive set of possibly confounding variables (L_1) helps address lack of randomization in the exposure groups. We use $A_1 = 1$ to denote initiation with XR-NTX and $A_1 = 0$ to denote initiation with BUP-NX. The outcome of this study is opioid use as detected by weekly urine drug screen or timeline followback interview [66]. We use L_t to denote opioid use measured at week $t + 1$ for $t \in \{1, 2, 3\}$, which detects use in the several days prior (via urine drug screen) to week prior (via interview), and where t represents the number of weeks since randomization. We hypothesized that if XR-NTX reduces

Table 3: Simulation results under misspecification of the outcome regression

n	Estimand	|Bias|	$\sqrt{n} \times$ |Bias|	$n \times$ Var.	$n \times$ MSE	95% CI coverage
500	Direct	0.005	0.104	2.201	2.204	0.96
	Indirect	0.000	0.003	0.161	0.160	1.00
	Total	0.005	0.107	2.210	2.214	0.98
1,000	Direct	0.001	0.045	2.092	2.087	0.96
	Indirect	0.000	0.004	0.117	0.117	0.99
	Total	0.001	0.041	2.117	2.111	0.98
5,000	Direct	0.004	0.305	2.049	2.136	0.95
	Indirect	0.000	0.001	0.063	0.062	0.98
	Total	0.004	0.304	2.031	2.117	0.96

craving more than BUP-NX, that this could provide a partial explanation of the lower opioid use while on XR-NTX. We use M_t to denote craving during week $t + 1$ since randomization, for $t \in \{1, 2, 3\}$. There may also be differences in depressive symptoms [67] and withdrawal symptoms [68] between the two medications, so we incorporated measures of each (Hamilton depression scale, subjective opioid withdrawal scale; see [69,70]), as time-varying confounders. We use Z_t to denote these confounders measured during week $t + 1$ for $t \in \{1, 2, 3\}$.

In addition, patients could drop out of treatment or otherwise be lost to follow-up starting in week 3 after randomization. Importantly, the outcomes L_t are always observed, as it is assumed that a patient who has missing opioid-use data (most likely due to a missed visit) would have been positive for opioid use [71,72]. For $t \in \{2, 3\}$, we use $A_t = 1$ to denote that a patient's Z_t and M_t have been measured, and we let $A_t = 0$ otherwise. Thus, our intervention variable is given by a combination of treatment and censoring, $\bar{A} = \{A_1, A_2, A_3\}$. In summary, we can write our observed data for this example as $O = (L_1, A_1, Z_1, M_1, L_2, A_2, Z_2, M_2, L_3, A_3, Z_3, M_3, L_4)$.

Note that the assumption of a negative (or positive) outcome for patients loss-to-follow-up may not be sensible in all applications. In such cases, we encourage the use of a loss-to-follow-up indicator variable and an intervention on loss-to-follow-up as described in Section 2.

We estimated the interventional direct effect of initiating treatment with XR-NTX vs BUP-NX on illicit opioid use during week 4 of treatment, not operating through craving, under a hypothetical intervention to prevent dropout. That is, for $\bar{a}' = (1, 1, 1)$ and $\bar{a}^\star = (0, 1, 1)$, we estimated $\mathsf{E}(Y(\bar{a}', \bar{G}(\bar{a}^\star)) - Y(\bar{a}^\star, \bar{G}(\bar{a}^\star)))$. We also estimated the interventional indirect effect of initiating treatment with XR-NTX vs BUP-NX on illicit opioid use during week 4 of treatment operating through craving, had those who dropped out not dropped out. That is, we estimated $\mathsf{E}(Y(\bar{a}', \bar{G}(\bar{a}')) - Y(\bar{a}', \bar{G}(\bar{a}^\star)))$. We used an ensemble of machine learning algorithms in fitting the nuisance parameters. The ensemble included an intercept-only model, lasso, multiple additive regression splines, and extreme gradient boosting. The weights in the ensemble were chosen using super-learning [61]. We used cross-fitting with fivefolds.

For the interventional direct effect, we estimated that initiating treatment with XR-NTX vs BUP-NX, not operating through craving, would reduce risk of using illicit opioids during week 4 of treatment by 8.8 percentage points (risk difference: −0.088, 95% CI: −0.129, −0.048). For the interventional indirect effect, we estimated that initiating treatment with XR-NTX vs BUP-NX, operating through craving, would not meaningfully decrease risk of using illicit opioids during week 4 of treatment (risk difference: −0.004, 95% CI: −0.019, 0.011). Thus, we conclude that reductions in risk of illicit opioid use during treatment with XR-NTX vs BUP-NX are not due to differences that operate through the treatments' effects on craving.

8 Discussion

Our approach generalizes interventional causal effects to allow for high-dimensional time-dependent variables measured post- and pre-treatment. We present an estimation algorithm that leverages machine learning to alleviate misspecification bias while retaining statistical properties such as $\sqrt{n}$-consistency, non-parametric efficiency, and asymptotic normality.

While this approach allows great flexibility in the data structure and estimation method, some limitations remain. We assume that the mediator is categorical and computational tractability of our proposed estimator requires that it takes values on a small set. This limitation seems fundamental and hard to overcome within an interventional effect framework, as all estimators will require the estimation of the density of the counterfactual variable $\bar{M}(\bar{a})$. We know of no method that can do this non-parametrically in the case of a continuous or high-dimensional variable $\bar{M}$, although recent approaches on the estimation of counterfactual densities in single time-point settings are promising [73].

In addition, interventional effects have some limitations. First, they do not decompose the average treatment effect $E[Y(\bar{a}') - Y(\bar{a}^\star)]$. Second, the interventional indirect effect can be non-zero, even when there are no paths of the type $A \rightarrow M \rightarrow Y$ in the DAG in Figure 1 [30]. Solving this limitation would require

a different framework for mediation analysis. We proposed such a framework for single time-point exposures in ref. [74], and future work will include generalizing that framework to the case of time-varying exposures and mediators.

Conflict of interest: Prof. Iván Díaz is a member of the Editorial Board of the Journal of Causal Inference but was not involved in the review process of this article.

References

[1] Vander Weele TJ. Mediation and mechanism. European J Epidemiol. 2009;24(5):217–24.
[2] Gilbert PB, Montefiori DC, McDermott A, Fong Y, Benkeser DC, Deng W, et al. Immune correlates analysis of the mRNA-1273 COVID-19 vaccine efficacy trial. MedRxiv. 2021.
[3] Robins JM, Greenland S. Identifiability and exchangeability for direct and indirect effects. Epidemiology. 1992;3:143–55.
[4] Pearl J. Direct and indirect effects. In: Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence. UAI '01. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.; 2001. p. 411–20. http://dl.acm.org/citation.cfm?id=647235.720084.
[5] Avin C, Shpitser I, Pearl J. Identifiability of path-specific effects. In: IJCAI International Joint Conference on Artificial Intelligence. 2005. p. 357–63.
[6] Rudolph KE, Gimbrone C, Díaz I. Helped into harm: mediation of a housing voucher intervention on mental health and substance use in boys. Epidemiology. 2021;32(3):336–46.
[7] Robins JM, Richardson TS. Alternative graphical causal models and the identification of direct effects. Causality and Psychopathology: Finding the Determinants of Disorders and Their Cures. 2010. p. 103–58.
[8] Tchetgen Tchetgen EJ, Phiri K. Bounds for pure direct effect. Epidemiology (Cambridge, Mass). 2014;25(5):775.
[9] Miles CH, Kanki P, Meloni S, Tchetgen Tchetgen EJ. On partial identification of the pure direct effect. 2015. arXiv: http://arXiv.org/abs/arXiv:150901652.
[10] Stensrud MJ, Young JG, Didelez V, Robins JM, Hernán MA. Separable effects for causal inference in the presence of competing events. J Amer Stat Assoc. 2022;117(537):175–83.
[11] Díaz I, Hejazi NS. Causal mediation analysis for stochastic interventions. J R Stat Soc Ser B (Stat Methodol). 2020;82(3):661–83.
[12] Hejazi NS, Rudolph KE, van der Laan MJ, Díaz I. Nonparametric causal mediation analysis for stochastic interventional (in) direct effects. 2020. arXiv: http://arXiv.org/abs/arXiv:200906203.
[13] Petersen ML, Sinisi SE, van der Laan MJ. Estimation of direct causal effects. Epidemiology. 2006;17(3):276–84.
[14] van der Laan MJ, Petersen ML. Direct effect models. Int J Biostat. 2008;4(1).
[15] Vander Weele TJ, Vansteelandt S, Robins JM. Effect decomposition in the presence of an exposure-induced mediator-outcome confounder. Epidemiology (Cambridge, Mass). 2014;25(2):300.
[16] Zheng W, van der Laan MJ. Targeted maximum likelihood estimation of natural direct effects. Int J Biostat. 2012;8(1):1–40.
[17] Díaz I, Hejazi NS, Rudolph KE, van Der Laan MJ. Nonparametric efficient causal mediation with intermediate confounders. Biometrika. 2021;108(3):627–41.
[18] Vander Weele TJ, Tchetgen Tchetgen EJ. Mediation analysis with time varying exposures and mediators dl. J R Stat Soc Ser B Stat Methodol. 2017;79(3):917.
[19] Tai AS, Lin SH, Chu YC, Yu T, Puhan MA, Vander Weele T. Causal mediation analysis with multiple time-varying mediators. Epidemiology. 2022;34(1):8–19.
[20] Zheng W, van der Laan M. Longitudinal mediation analysis with time-varying mediators and exposures, with application to survival outcomes. J Causal Inference. 2017;5(2).
[21] Vansteelandt S, Linder M, Vandenberghe S, Steen J, Madsen J. Mediation analysis of time-to-event endpoints accounting for repeatedly measured mediators subject to time-varying confounding. Stat Med. 2019;38(24):4828–40.
[22] Mittinty MN, Vansteelandt S. Longitudinal mediation analysis using natural effect models. Amer J Epidemiol. 2020;189(11):1427–35.
[23] Ge L, Wang J, Shi C, Wu Z, Song R. A reinforcement learning framework for dynamic mediation analysis. 2023. arXiv: http://arXiv.org/abs/arXiv:230113348.
[24] Loh WW, Moerkerke B, Loeys T, Vansteelandt S. Nonlinear mediation analysis with high-dimensional mediators whose causal structure is unknown. Biometrics. 2022;78(1):46–59.

[25] Tai AS, Lin PH, Huang YT, Lin SH. Path-specific effects in the presence of a survival outcome and causally ordered multiple mediators with application to genomic data. Stat Meth Med Res. 2022;31(10):1916–33.

[26] Tanner KT, Sharples LD, Daniel RM, Keogh RH. Methods of analysis for survival outcomes with time-updated mediators, with application to longitudinal disease registry data. Stat Meth Med Res. 2022;31(10):1959–75.

[27] Tai AS, Lin SH. Complete effect decomposition for an arbitrary number of multiple ordered mediators with time-varying confounders: a method for generalized causal multi-mediation analysis. Stat Meth Med Res. 2023;32(1):100–17.

[28] Pearl J. Causality: models, reasoning, and inference. Cambridge: Cambridge University Press; 2000.

[29] Didelez V. Defining causal mediation with a longitudinal mediator and a survival outcome. Lifetime Data Anal.. 2019;25(4):593–610.

[30] Miles CH. On the causal interpretation of randomized interventional indirect effects. 2022. arXiv: http://arXiv.org/abs/ arXiv:220300245.

[31] Richardson TS, Robins JM. 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. 2013;128(30):2013.

[32] Rubin D, van der Laan MJ. A doubly robust censoring unbiased transformation. Int J Biostat. 2007;3(1).

[33] van der Laan MJ, Rose S. Targeted learning: causal inference for observational and experimental data. New York: Springer; 2011.

[34] Luedtke AR, Sofrygin O, van der Laan MJ, Carone M. Sequential double robustness in right-censored longitudinal models. 2017. arXiv: http://arXiv.org/abs/arXiv:170502459.

[35] Chernozhukov V, Chetverikov D, Demirer M, Duflo E, Hansen C, Newey W, et al. Double/debiased machine learning for treatment and structural parameters. Econometr J. 2018;21(1):C1–68.

[36] Bang H, Robins JM. Doubly robust estimation in missing data and causal inference models. Biometrics. 2005;61(4):962–73.

[37] Rotnitzky A, Robins J, Babino L. On the multiply robust estimation of the mean of the g-functional. 2017. arXiv: http:// arXiv.org/abs/arXiv:170508582.

[38] Dukes O, Martinussen T, Tchetgen Tchetgen EJ, Vansteelandt S. On doubly robust estimation of the hazard difference. Biometrics. 2019;75(1):100–9.

[39] vonMises R. On the asymptotic distribution of differentiable statistical functions. Ann Math Stat. 1947;18(3):309–48.

[40] van der Vaart AW. Asymptotic statistics. Cambridge: Cambridge University Press; 1998.

[41] Robins J, Li L, Tchetgen Tchetgen E, van der Vaart AW. Quadratic semiparametric von mises calculus. Metrika. 2009;69(2–3):227–47.

[42] Bickel PJ, Klaassen CA, Ritov Y, Wellner JA. Efficient and adaptive estimation for semiparametric models. Berlin: Springer-Verlag; 1997.

[43] Robins JM. Robust estimation in sequentially ignorable missing data and causal inference models. In: Proceedings of the American Statistical Association. 2000.

[44] Robins JM, Rotnitzky A, Zhao LP. Estimation of regression coefficients when some regressors are not always observed. J Amer Stat Assoc. 1994 September;89(427):846–66.

[45] van der Laan MJ, Robins JM. Unified methods for censored longitudinal data and causality. New York: Springer; 2003.

[46] van der Laan MJ, Rubin D. Targeted maximum likelihood learning. Int J Biostat. 2006;2(1). https://doi.org/10.2202/1557-4679.1043.

[47] van der Laan MJ, Rose S. Targeted learning in data science: causal inference for complex longitudinal studies. New York: Springer; 2018.

[48] Bickel PJ, Ritov Y, Tsybakov AB. Simultaneous analysis of Lasso and Dantzig selector. Ann Stat. 2009;37(4):1705–32.

[49] Wager S, Walther G. Adaptive concentration of regression trees, with application to random forests. 2015. arXiv: http:// arXiv.org/abs/arXiv:150306388.

[50] Chen X, White H. Improved rates and asymptotic normality for nonparametric neural network estimators. IEEE Trans Inform Theory. 1999;45(2):682–91.

[51] Benkeser D, van der Laan M. The highly adaptive lasso estimator. In: 2016 IEEE International Conference on Data Science and Advanced Analytics (DSAA). IEEE; 2016. p. 689–96.

[52] Díaz I, Williams N, Hoffman KL, Schenck EJ. Nonparametric causal effects based on longitudinal modified treatment policies. J Amer Stat Assoc. 2021:1–16.

[53] Díaz I, Hoffman KL, Hejazi NS. Causal survival analysis under competing risks using longitudinal modified treatment policies. 2022. arXiv: http://arXiv.org/abs/arXiv:220203513.

[54] van der Vaart AW, Wellner JA. Weak convergence and empirical processes. New York: Springer-Verlag; 1996.

[55] Bickel PJ. On adaptive estimation. Ann Stat. 1982;10(3):647–71.

[56] Klaassen CA. Consistent estimation of the influence function of locally asymptotically linear estimators. Ann Stat. 1987;15(4):1548–62.

[57] Zheng W, van der Laan MJ. Cross-validated targeted minimum-loss-based estimation. In: Targeted learning. New York: Springer; 2011. p. 459–74.

[58] Benkeser D, Carone M, Laan MVD, Gilbert P. Doubly robust nonparametric inference on the average treatment effect. Biometrika. 2017;104(4):863–80.

[59] Balzer L, Ahern J, Galea S, van der Laan M. Estimating effects with rare outcomes and high dimensional covariates: knowledge is power. Epidemiol Meth. 2016;5(1):1–18.

[60] Tran L, Petersen M, Schwab J, van der Laan MJ. Robust variance estimation and inference for causal effect estimation. 2018. arXiv: http://arXiv.org/abs/arXiv:181003030.

[61] van der Laan MJ, Polley E, Hubbard A. Super learner. Stat Appl Genetics Molecular Biol. 2007;6(25):25.

[62] van der Laan MJ, Dudoit S, van der Vaart AW. The cross-validated adaptive epsilon-net estimator. In: Division of biostatistics. Berkeley: University of California; 2004.

[63] Lee JD, Nunes Jr EV, Novo P, Bachrach K, Bailey GL, Bhatt S, et al. Comparative effectiveness of extended-release naltrexone versus buprenorphine-naloxone for opioid relapse prevention (X: BOT): a multicentre, open-label, randomised controlled trial. Lancet. 2018;391(10118):309–18.

[64] Greiner MG, Shulman M, Choo TH, Scodes J, Pavlicova M, Campbell AN, et al. Naturalistic follow-up after a trial of medications for opioid use disorder: Medication status, opioid use, and relapse. J Substance Abuse Treatment. 2021;131:108447.

[65] Solli KK, Latif ZeH, Opheim A, Krajci P, Sharma-Haase K, Benth JS, et al. Effectiveness, safety and feasibility of extended-release naltrexone for opioid dependence: a 9-month follow-up to a 3-month randomized trial. Addiction. 2018;113(10):1840–9.

[66] Sobell L, Sobell M. Alcohol timeline followback users' manual. Toronto, Canada: Addiction Research Foundation; 1995.

[67] Rudolph KE, Díaz I, Hejazi NS, van der Laan MJ, Luo SX, Shulman M, et al. Explaining differential effects of medication for opioid use disorder using a novel approach incorporating mediating variables. Addiction. 2021;116(8):2094–103.

[68] SAMHSA. Medications for opioid use disorder for healthcare and addiction professionals, policymakers, patients, and families: treatment improvement protocol TIP 63. 2021. https://store.samhsa.gov/sites/default/files/SAMHSA_Digital_Download/PEP21-02-01-002.pdf.

[69] Hamilton M. The Hamilton depression scale-accelerator or break on antidepressant drug discovery. Psychiatry. 1960;23:56–62.

[70] Cooper ZD, Johnson KW, Pavlicova M, Glass A, Vosburg SK, Sullivan MA, et al. The effects of ibudilast, a glial activation inhibitor, on opioid withdrawal symptoms in opioid-dependent volunteers. Addiction Biol. 2016;21(4):895–903.

[71] Hser YI, Evans E, Huang D, Weiss R, Saxon A, Carroll KM, et al. Long-term outcomes after randomization to buprenorphine/naloxone versus methadone in a multi-site trial. Addiction. 2016;111(4):695–705.

[72] Weiss RD, Potter JS, Griffin ML, Provost SE, Fitzmaurice GM, McDermott KA, et al. Long-term outcomes from the national drug abuse treatment clinical trials network prescription opioid addiction treatment study. Drug Alcohol Depend. 2015;150:112–9.

[73] Kennedy EH, Balakrishnan S, Wasserman L. Semiparametric counterfactual density estimation. 2021. arXiv: http://arXiv.org/abs/arXiv:210212034.

[74] Díaz I. Causal influence, causal effects, and path analysis in the presence of intermediate confounding. 2022. arXiv: http://arXiv.org/abs/arXiv:220508000.