




## Nonparametric Causal Effects Based on Longitudinal Modified Treatment Policies

Iván Díaz, Nicholas Williams, Katherine L. Hoffman & Edward J. Schenck


**To cite this article:** Iván Díaz, Nicholas Williams, Katherine L. Hoffman & Edward J. Schenck (2023) Nonparametric Causal Effects Based on Longitudinal Modified Treatment Policies, *Journal of the American Statistical Association*, 118:542, 846-857, DOI: [10.1080/01621459.2021.1955691](https://doi.org/10.1080/01621459.2021.1955691)

**To link to this article:** <https://doi.org/10.1080/01621459.2021.1955691>

 [View supplementary material](#) 

 [Published online: 01 Sep 2021.](#)

 [Submit your article to this journal](#) 

 [Article views: 2836](#)

 [View related articles](#) 

 [View Crossmark data](#) 

 [Citing articles: 19 View citing articles](#) 



# Nonparametric Causal Effects Based on Longitudinal Modified Treatment Policies

Iván Díaz<sup>a</sup> , Nicholas Williams<sup>a</sup>, Katherine L. Hoffman<sup>a</sup>, and Edward J. Schenck<sup>b</sup>

<sup>a</sup>Division of Biostatistics, Department of Population Health Sciences, Weill Cornell Medicine, New York; <sup>b</sup>Division of Pulmonary & Critical Care Medicine, Department of Medicine, Weill Cornell Medicine, New York

## ABSTRACT

Most causal inference methods consider counterfactual variables under interventions that set the exposure to a fixed value. With continuous or multi-valued treatments or exposures, such counterfactuals may be of little practical interest because no feasible intervention can be implemented that would bring them about. Longitudinal modified treatment policies (LMTPs) are a recently developed nonparametric alternative that yield effects of immediate practical relevance with an interpretation in terms of meaningful interventions such as reducing or increasing the exposure by a given amount. LMTPs also have the advantage that they can be designed to satisfy the positivity assumption required for causal inference. We present a novel sequential regression formula that identifies the LMTP causal effect, study properties of the LMTP statistical estimand such as the efficient influence function and the efficiency bound, and propose four different estimators. Two of our estimators are efficient, and one is sequentially doubly robust in the sense that it is consistent if, for each time point, either an outcome regression or a treatment mechanism is consistently estimated. We perform numerical studies of the estimators, and present the results of our motivating study on hypoxemia and mortality in intubated Intensive Care Unit (ICU) patients. Software implementing our methods is provided in the form of the open source R package `lmt_p` freely available on GitHub ([https://github.com/nt-williams/lmt\\_p](https://github.com/nt-williams/lmt_p)) and CRAN.

## ARTICLE HISTORY

Received June 2020  
Accepted July 2021

## KEYWORDS

Continuous exposures; Longitudinal data; Modified treatment policies; Sequential double robustness; Targeted minimum loss-based estimation

## 1. Introduction

Most modern causal inference methods use a potential outcome framework to define causal effects as contrasts between the outcome distribution under different hypothetical counterfactual worlds. The defining characteristic of these counterfactual worlds is a series of interventions on the cause being evaluated. For example, causal effects for binary variables can be defined as the difference in mean outcome comparing a world where everyone receives an intervention versus a world where everyone receives a control—the so-called average treatment effect. For multi-valued exposures, most methods study the outcome expectation under a series of hypothetical worlds corresponding to assigning the same value of the exposure to all units in the population—the so-called dose-response curve. While informative, a dose-response analysis presents some challenges. First, for some exposures (e.g., physical activity) it is hard to conceive an intervention that would set them statically (e.g., make everyone exercise 30 min), even in principle. This presents a problem for dose-response effects because one cannot reasonably implement an intervention that would bring about such effects in the real world. Second, summarizing the infinite-dimensional dose-response curve often requires restrictive and arbitrary parametric assumptions. These assumptions, frequently encoded in marginal structural models, are typically wrong, and the models can be hard to interpret under model misspecification (Neugebauer and van der Laan 2007). Third, nonparametric approaches to estimation of the causal dose-response curve cannot achieve  $n^{1/2}$ -consistency,

because the parameter is not pathwise differentiable. For example, the estimators of Kennedy et al. (2017) converged at the  $(nh)^{1/2}$ -rate (where  $h$  is a bandwidth that decreases with sample size, giving the familiar  $n^{-2/5}$  rate under twice-differentiability). Fourth, the effects are not identified when some units have a zero chance to receive some of the exposure levels under consideration, a situation known as violations to the *positivity assumption*. Positivity violations are likely present with most continuous and multi-valued exposures, and are exacerbated when the exposures are measured at multiple time points.

As a solution to the above limitations, the causal inference literature has considered alternative definitions of causal effects, allowing for hypothetical worlds where the exposures can depend on characteristics of the unit (dynamic regimes), where the post-intervention distribution is a random draw from a given but possibly unknown distribution (stochastic interventions), and where the post-intervention exposure may be a random or deterministic function of the observed treatment value (e.g., Stock 1989; Robins, Hernán, and Siebert 2004). In this article, we adopt the latter approach. In this context, we have previously studied interventions that shift the exposure distribution (Díaz and van der Laan 2012) for single time point studies. Haneuse and Rotnitzky (2013) generalized our previous interventions and introduced the name *modified treatment policy (MTP)*, which we adopt in this article. MTPs can yield estimates with a familiar interpretation as the expected change in mean response for a given change in the exposure

of interest—without the need to impose unrealistic parametric assumptions such as linearity in the causal model. In our motivating example, we estimate the effect on 14-day survival of an increase of 50 units in the arterial partial pressure of oxygen to fraction of inspired oxygen (P/F) ratio, where the only patients with respiratory failure (P/F ratio < 300) receive the intervention.

For the longitudinal case, Robins, Hernán, and Siebert (2004) introduced *dynamic regimes depending on the natural value of treatment*, where the *natural value of treatment* is the value that treatment would take at time  $t$  if the intervention was discontinued right before time  $t$ . Richardson and Robins (2013) and Young, Hernán, and Robins (2014) formalized the effect definition and showed conditions under which the effects are identified by the extended g-formula of Robins, Hernán, and Siebert (2004). A key finding of these articles is that the sequential randomization assumption required for identifying effects that depend on the natural value of treatment is stronger than that required for identification of the effect of dynamic regimes. Although Richardson and Robins (2013) and Young, Hernán, and Robins (2014) used the same name *dynamic regimes depending on the natural value of treatment* to refer to the interventions, the definitions given in the papers are different. In this article we use the name *longitudinal modified treatment policies (LMTP)* to refer to the definition of Richardson and Robins (2013), as the interventions in this work are a natural extension of the definition of Haneuse and Rotnitzky (2013) for single time point studies.

We present a novel alternative expression of the extended g-formula in terms of sequential regressions. We show conditions under which the sequential regression formula allows the generalization of estimators for dynamic regimes to the case of LMTPs. We also introduce a stochastic intervention where instead of setting the exposure equal to the LMTP, we set it to a random draw from the LMTP distribution. We call this intervention LMTP stochastic intervention (LMTP-SI). Young, Hernán, and Robins (2014) showed that the sequential randomization assumption required for identification of LMTP-SI is identical to that required for identification of the effect of dynamic interventions. Thus, the LMTP functional can still be interpreted as a causal effect even under the standard identifiability assumptions required for dynamic regimes.

We propose four different estimators for the LMTP sequential regression functional. The first estimator is a simple extension of the inverse parametric probability weighted (IPW) of Young, Hernán, and Robins (2014); the second is a simple extension of the g-computation estimator presented by Robins, Hernán, and Siebert (2004) and Taubman et al. (2009). IPW and g-computation estimation require estimating nuisance parameters related to the conditional density of treatment and outcome regressions conditional on the history of a unit. When these nuisance parameters are estimated within pre-specified parametric models, the LMTP estimators are asymptotically Gaussian and the Delta method or the bootstrap are guaranteed to yield valid asymptotic  $p$ -values and confidence intervals. However, whenever the nuisance estimators are

data-adaptive, such as when model selection is performed, the asymptotic distribution of the IPW and g-computation estimators is generally unknown, making it difficult to perform hypothesis tests and quantify the uncertainty around the estimate. We thus develop two additional estimators that overcome this issue.

Our third and fourth estimation strategies have roots in semi-parametric estimation theory (e.g., von Mises 1947; Bickel et al. 1997; van der Vaart 1998; Robins et al. 2009), in the theory for doubly robust estimation using estimating equations (Robins, Rotnitzky, and Zhao 1994; Robins 2000; van der Laan and Robins 2003; Bang and Robins 2005), and in the targeted learning framework (van der Laan and Rubin 2006; van der Laan and Rose 2011, 2018). Central to this theory is the study of the efficient influence function (EIF) or canonical gradient, which characterizes the efficiency bound of the LMTP functional and allows the development of estimators under slow convergence rates for the nuisance parameters involved (Robins et al. 2009). We derive the EIF for LMTPs under the assumption that the LMTP does not depend on the data-generating mechanism. We show that this EIF has a similar structure to the EIF for dynamic interventions under a differentiability and invertibility condition on the LMTP. We then show that, under these conditions, the EIF allows for multiply robust estimation. This is a surprising fact that should not be expected for general interventions that depend on the natural value of treatment, since the parameter functional depends on both the outcome regression and the exposure mechanism. We exploit the similarity of the EIF to generalize some estimators for dynamic intervention available in the literature. In particular, we develop a targeted minimum loss-based estimator (TMLE), which is a natural extension of the estimator of van der Laan and Gruber (2012). We show that the TMLE is multiply robust consistent in the sense that it allows for consistent estimation under  $\tau + 1$  configurations of consistent estimation of the nuisance parameters, where  $\tau$  is the number of time points under consideration. We then develop a sequentially doubly robust (SDR) estimator based on multiply robust unbiased transformations by generalizing the estimators of Luedtke et al. (2017) and Rotnitzky, Robins, and Babino (2017), which are themselves related to ideas dating back to Rubin and van der Laan (2007) and Tchetgen Tchetgen (2009). Luedtke et al. (2017) referred to this estimator as *SDR via doubly robust transformations*; we adopt this name and shorten it to SDR for simplicity. The SDR estimator is expected to be consistent under  $2^\tau$  configurations of consistent estimation of the nuisance parameters. This is in complete analogy to the  $(\tau + 1)$ - and  $2^\tau$ -multiply robust consistency results discussed by Luedtke et al. (2017), Rotnitzky, Robins, and Babino (2017), and Molina et al. (2017) for the case of a dynamic intervention. We use cross-fitting to obtain  $n^{1/2}$ -convergence of our estimators while avoiding entropy conditions that may be violated by data adaptive estimators of the nuisance parameters (Zheng and van der Laan 2011; Chernozhukov et al. 2018). Finally, we note that since the LMTP parameter studied in this paper generalizes dynamic interventions for longitudinal studies, our methods and software can also be used for estimation of those parameters.

## 2. Notation and Definition of Causal Effects

Let  $Z_1, \dots, Z_n$  denote a sample of iid observations with  $Z = (L_1, A_1, L_2, A_2, \dots, L_\tau, A_\tau, Y) \sim P$ , where  $L_t$  denotes time-varying covariates,  $A_t$  denotes a vector of intervention variables such as treatment and/or censoring status, and  $Y = L_{\tau+1}$  denotes an outcome such as survival at the end of study follow-up. We let  $Pf = \int f(z)dP(z)$  for a given function  $f(z)$ . We use  $P_n$  to denote the empirical distribution of  $Z_1, \dots, Z_n$ , and assume  $P$  is an element of the nonparametric statistical model defined as all continuous densities on  $Z$  with respect to a dominating measure  $\nu$ . We let  $E$  denote the expectation with respect to  $P$ , that is,  $E\{f(Z)\} = \int f(z)dP(z)$ . We also let  $\|f\|^2$  denote the  $L_2(P)$  norm  $\int f^2(z)dP(z)$ . We use  $\bar{X}_t = (X_1, \dots, X_t)$  to denote the history of a variable, use  $\underline{X}_t = (X_t, \dots, X_\tau)$  to denote the future of a variable, and use  $H_t = (\bar{A}_{t-1}, \bar{L}_t)$  to denote the history of all variables up until just before  $A_t$ . For the complete history of a random variable, we simplify  $\bar{X}_\tau$  as  $\bar{X}$ . We let  $g_t(a_t | h_t)$  denote the probability density function of  $A_t$  conditional on  $H_t = h_t$ . We use calligraphic font to denote the support of a random variable, for example,  $\mathcal{A}_t$  denotes the support of  $A_t$ .

We formalize the definition of the causal effects using a nonparametric structural equation model (Pearl 2009). Specifically, for each time point  $t$ , we assume the existence of deterministic functions  $f_{L_t}$ ,  $f_{A_t}$ , and  $f_Y$  such that  $L_t = f_{L_t}(A_{t-1}, H_{t-1}, U_{L,t})$ ,  $A_t = f_{A_t}(H_t, U_{A,t})$ , and  $Y = f_Y(A_\tau, H_\tau, U_Y)$ . Here,  $U = (U_{L,t}, U_{A,t}, U_Y : t \in \{1, \dots, \tau\})$  is a vector of exogenous variables, with unrestricted joint distribution. Sufficient assumptions to identify the effects we discuss will be given in Section 3. LMTP effects can be defined in terms of hypothetical interventions where the equation  $A_t = f_{A_t}(H_t, U_{A,t})$  is removed from the structural model, and the exposure is assigned as a new random variable  $A_t^{\text{dl}}$ . An intervention that sets the exposures up to time  $t - 1$  to  $\bar{A}_{t-1}^{\text{dl}}$  generates counterfactual variables  $L_t(\bar{A}_{t-1}^{\text{dl}}) = f_{L_t}(A_{t-1}^{\text{dl}}, H_{t-1}(\bar{A}_{t-2}^{\text{dl}}), U_{L,t})$  and  $A_t(\bar{A}_{t-1}^{\text{dl}}) = f_{A_t}(H_t(\bar{A}_{t-1}^{\text{dl}}), U_{A,t})$ , where the counterfactual history is defined recursively as  $H_t(\bar{A}_{t-1}^{\text{dl}}) = (\bar{A}_{t-1}^{\text{dl}}, \bar{L}_t(\bar{A}_{t-1}^{\text{dl}}))$ . The variable  $A_t(\bar{A}_{t-1}^{\text{dl}})$  is referred to as the *natural value of treatment* (Richardson and Robins 2013; Young, Hernán, and Robins 2014), and represents the value of treatment that would have been observed at time  $t$  under an intervention carried out up until time  $t - 1$  but discontinued thereafter. An intervention where all the treatment variables up to  $t = \tau$  are intervened on generates a counterfactual outcome  $Y(\bar{A}^{\text{dl}}) = f_Y(A_\tau^{\text{dl}}, H_\tau(\bar{A}_{\tau-1}^{\text{dl}}), U_Y)$ . Causal effects will be defined in terms of the distribution of this counterfactual. We discuss two types of effects: longitudinal modified treatment policies (LMTPs), and LMTP stochastic interventions (LMTP-SI). Both effects are defined in terms of a user-given function  $\text{dl}(a_t, h_t)$  that maps a treatment value  $a_t$  and a history  $h_t$  into a new exposure value. For fixed values  $\bar{a}_t, \bar{h}_t$ , we recursively define  $a_t^{\text{dl}} = \text{dl}(a_t, h_t^{\text{dl}})$ , where  $h_t^{\text{dl}} = (\bar{a}_{t-1}^{\text{dl}}, \bar{h}_t)$ . The LMTPs that we study are thus defined as follows:

**Definition 1 (Longitudinal modified treatment policies (LMTP)).** An intervention  $A_t^{\text{dl}}$  is said to be a longitudinal modified treatment policy if it is defined as  $A_t^{\text{dl}} = \text{dl}(A_t(\bar{A}_{t-1}^{\text{dl}}), H_t(\bar{A}_{t-1}^{\text{dl}}))$  for a user-given function  $\text{dl}$ . LMTP causal effects are defined as

contrasts between the distribution of counterfactual outcomes  $Y(\bar{A}^{\text{dl}})$  under different functions  $\text{dl}$ .

This definition is a longitudinal generalization of the modified treatment policies defined by Haneuse and Rotnitzky (2013) for a single time point, because at  $t = 1$  the factual and natural values of treatment coincide. For  $\tau > 1$ , the modified treatment policies in Definition 1 are a particular case of the interventions that depend on the natural value of treatment first discussed by Robins, Hernán, and Siebert (2004) and formalized by Richardson and Robins (2013) and Young, Hernán, and Robins (2014).

To illustrate some of our methods and ideas, we will consider the following important examples of regimes  $\text{dl}$ .

**Example 1 (Threshold LMTP).** We revisit a problem posed by Taubman et al. (2009), in which we are interested in assessing the effect of exercising at least 30 minutes a day on the risk of coronary heart disease (CHD). Let  $A_t$  denote the minutes exercised by each study participant, let  $Y$  denote an indicator of CHD by the end of the study, and let  $L_t$  denote confounders such as comorbidities and lifestyle variables. Let  $\text{dl}(a_t, h_t) = \mathbb{1}(a_t \geq 30)a_t + \mathbb{1}(a_t < 30)30$ . At time point  $t = 1$ , we define the intervention as  $A_1^{\text{dl}} = \text{dl}(A_1, H_1)$ , which would set the number of minutes exercised to 30 if the participant exercised less than 30, and would leave it unchanged otherwise. At time points  $t > 1$ , we are interested in setting the amount of physical activity to 30 if the participant's natural value of treatment at time  $t$  is smaller than 30, and set it to the natural value of treatment otherwise. That is, we define  $A_t^{\text{dl}} = \text{dl}(A_t(\bar{A}_{t-1}^{\text{dl}}), H_t(\bar{A}_{t-1}^{\text{dl}}))$ .

**Example 2 (Shift LMTP).** Let  $A_t$  denote a continuous exposure, such as a drug dose or a physiological measurement such as P/F ratio that can be modified through intervention. To define this intervention, assume that  $A_t$  is supported as  $P(A_t \leq u_t(h_t) | H_t = h_t) = 1$  for some  $u_t$ . Then, for a user-given value  $\delta$ , we let

$$\text{dl}(a_t, h_t) = \begin{cases} a_t + \delta & \text{if } a_t \leq u_t(h_t) - \delta \\ a_t & \text{if } a_t > u_t(h_t) - \delta. \end{cases} \quad (1)$$

Then we define  $A_t^{\text{dl}} = \text{dl}(A_t(\bar{A}_{t-1}^{\text{dl}}), H_t(\bar{A}_{t-1}^{\text{dl}}))$ . This intervention was first introduced in the context of a single time point by Díaz and van der Laan (2012), and was further discussed by Díaz and van der Laan (2018) and Haneuse and Rotnitzky (2013), and by Díaz and Hejazi (2020) in the context of mediation. This intervention considers hypothetical worlds in which the natural exposure at time  $t$  is increased by a user-given value  $\delta$ , whenever such increase is feasible for a unit with history  $H_t(A_{t-1}^{\text{dl}})$ . In our motivating example, we assess the effect on mortality of an intervention that would increase a patient's P/F ratio by 50 units for patients with acute respiratory failure (P/F ratio  $< 300$ ). Alternatively, we could define a multiplicative shift as  $\text{dl}(a_t, h_t) = a_t \delta(h_t)$  for a shift function  $\delta(h_t)$  that may depend on the history  $h_t$ .

In some cases, an alternative intervention may be of interest. Consider for example the implementation of a public health intervention to encourage people to exercise more. In this case, one could be interested in a post-intervention exposure where the distribution of time exercise has shifted, but that does not necessarily set each individual's exposure to  $A_t^{\text{dl}} =$

$d(A_t(\bar{A}_{t-1}^{dl}), H_t(\bar{A}_{t-1}^{dl}))$  with  $d$  defined in Equation (1). This intervention is different from a modified treatment policy in that we are not necessarily interested in shifting an individual's exposure. Instead, we wonder what would have happened in a hypothetical world where each unit's exposure was a random draw from a shifted distribution, representing a behavior shift at the population level. A LMTP stochastic intervention is then formalized as:

**Definition 2 (LMTP stochastic intervention (LMTP-SI)).** Let  $G_t^{dl}$  denote the distribution of  $d(A_t, H_t)$  conditional on  $H_t$ . A longitudinal modified treatment policy stochastic intervention  $\bar{Q}^{dl}$  is defined as a sequence of random draws from  $G_t^{dl} : t = 1, \dots, \tau$ . LMTP-SI Causal effects are defined as contrasts between the distribution of counterfactual outcomes  $Y(\bar{Q}^{dl})$  under different functions  $d$ .

In the examples discussed so far we consider interventions where the function  $d(a_t, h_t)$  only depends on values of the exposure  $a_t$  and the history  $h_t$ . We note that the definition of Richardson and Robins (2013) allowed for random regimes by letting the functions  $d(a_t, h_t, \varepsilon_t)$  depend on a randomizer  $\varepsilon_t$ . It is easy to imagine such settings in practice. For example, Robins, Hernán, and Siebert (2004) proposed to study the effect on coronary heart disease of interventions where a random half of smokers quit smoking forever. The methods that we present next also allow for this type of random intervention. However, we require some assumptions on the randomizer  $\varepsilon_t$ , namely that (i) it is drawn independently across units and independently of  $U$ , and (ii) its distribution does not depend on  $P$ . Under these two assumptions, and with the aim of simplifying notation, we assume without loss of generality that the time varying vector  $L_t$  contains the randomizer  $\varepsilon_t$ . Another example of such a randomized intervention is given below.

**Example 3 (Incremental propensity score interventions).** Kennedy (2019) proposed an intervention for binary exposures where  $A_t^{dl}$  is a draw from a Bernoulli distribution with shifted propensity score given by

$$g_t^{dl}(1 | h_t) = \frac{\delta g_t(1 | h_t)}{\delta g_t(1 | h_t) + 1 - g_t(1 | h_t)}.$$

Let  $\varepsilon_t$  denote a random draw from a uniform distribution in the interval (0, 1). Define the intervention as  $d_t(a_t, h_t) = \mathbb{1}\{\varepsilon_t < g_t^{dl}(1 | h_t)\}$ .

In what follows we are concerned with identification and nonparametric estimation of the causal parameters

$$\theta_{lmtp} = E\{Y(\bar{A}^{dl})\}, \text{ and } \theta_{lmtp-si} = E\{Y(\bar{Q}^{dl})\},$$

where  $Y(\bar{A}^{dl})$  and  $Y(\bar{Q}^{dl})$  are defined above. The difference between LMTP and LMTP-SI is subtle but important, and leads to different assumptions required for identification, as we will see in the next section.

### 3. Identification of Causal Effects

The first step in developing estimators for  $\theta_{lmtp}$  and  $\theta_{lmtp-si}$  is to derive an identification result that allows us to write these

causal parameters as a function of only the distribution  $P$  of the observed data  $Z$ . Sufficient assumptions to identify  $\theta_{lmtp}$  and  $\theta_{lmtp-si}$  were first given by Richardson and Robins (2013) and Young, Hernán, and Robins (2014). We first present sufficient assumptions for identification under the assumed NPSEM, and then discuss their implications in several scenarios.

**Assumption 1 (Positivity).** If  $(a_t, h_t) \in \text{supp}\{A_t, H_t\}$  then  $(d(a_t, h_t), h_t) \in \text{supp}\{A_t, H_t\}$  for  $t \in \{1, \dots, \tau\}$ .

**Assumption 2 (Standard sequential randomization).**  $U_{A,t} \perp\!\!\!\perp \underline{U}_{L,t+1} | H_t$  for all  $t \in \{1, \dots, \tau\}$ .

**Assumption 3 (Strong sequential randomization).**  $U_{A,t} \perp\!\!\!\perp (\underline{U}_{L,t+1}, \underline{U}_{A,t+1}) | H_t$  for all  $t \in \{1, \dots, \tau\}$ .

**Assumption 1** is equivalent to the assumption presented in Young, Hernán, and Robins (2014), and simply states that the distribution of interest is supported in the data. Consider our motivating example of the effect of P/F ratio on survival on ICU patients. Under no loss-to-follow-up, this assumption states that if it is possible to find a patient with history  $h_t$  who has a P/F ratio of  $a_t$  at time  $t$ , then it is also possible to find a patient with history  $h_t$  who has a P/F ratio of  $d(a_t, h_t)$ . We note that this assumption may be enforced by definition of  $d$  if sufficient information is available about the conditional support of  $a_t$  conditional on  $h_t$ . Furthermore, if  $a_t$  is multivariate and includes missingness or censoring indicators, then **Assumption 1** also states that for every observed history  $h_t$  there is a probability greater than zero of observing a patient who is not lost-to-follow-up at time  $t$ . **Assumption 2** is standard for the identification of dynamic regimes, and is satisfied if all the common causes of the intervention variable  $A_t$  and  $L_s : s > t$  are measured and recorded in  $H_t$ . **Assumption 3** is stronger than **Assumption 2**, and is satisfied if all common causes of the intervention variable  $A_t$  and  $(A_s, L_s) : s > t$  are measured and recorded in  $H_t$ . **Assumption 3** is similar in nature to the independence assumption required by Richardson and Robins (2013) (see Theorem 31 in that reference).

We have the following identification theorem, which allows us to compute the parameters  $\theta_{lmtp}$  and  $\theta_{lmtp-si}$  as a function only of the observed data distribution.

**Theorem 1 (Identification of the effect of LMTPs).** Set  $m_{\tau+1} = Y$ . For  $t = \tau, \dots, 1$ , recursively define

$$m_t : (a_t, h_t) \mapsto E \left[ m_{t+1}(A_{t+1}^{dl}, H_{t+1}) | A_t = a_t, H_t = h_t \right], \tag{2}$$

and define  $\theta = E[m_1(A_1^{dl}, L_1)]$ . Then we have

- (i) Under **Assumptions 1** and **2**,  $\theta_{lmtp-si}$  is identified as  $\theta$ .
- (ii) Under **Assumptions 1** and **3**,  $\theta_{lmtp}$  is identified as  $\theta$ .

The identification expression given in **Theorem 1** in terms of sequential regressions is an alternative expression of the extended g-formula of Robins, Hernán, and Siebert (2004) and Richardson and Robins (2013). In addition to allowing for a variety of interesting interventions, this setup allows for a variety of data structures. In particular, it can handle loss-to-follow-up, survival analysis, and missing exposures as follows. Let  $A_t =$

$(A_{1,t}, A_{2,t})$ , where  $A_{1,t}$  denotes the exposure at time  $t$ , and  $A_{2,t}$  is equal to one if the unit remains uncensored at time  $t + 1$  and zero otherwise. Assume monotone loss-to-follow-up so that  $A_{2,t} = 0$  implies  $A_{2,k} = 0$  for all  $k > t$ , in which case  $(L_k, A_{1,k})$  and  $Y$  become degenerate for  $k > t$ . In this case we could be interested in a hypothetical world in which there is no loss-to-follow-up and the exposure  $A_{1,t}$  is shifted as in Equation (1) or any other intervention of interest. In particular, we can define  $A_t^{\text{dl}} = (\text{dl}\{A_{t,1}(\bar{A}_{t-1,1}^{\text{dl}}), H_t(\bar{A}_{t-1,1}^{\text{dl}})\}, 1)$ . Time-varying outcomes can be incorporated by letting  $L_t = (X_t, Y_{t-1})$ , where  $X_t$  denotes the time-varying covariates of interest and  $Y_t$  denotes the time-varying outcome, and letting  $Y = Y_\tau$ . If a prior time point  $k$  is of interest, then we let  $Y = Y_k$  and truncate the sequence at  $\tau = k$ . Time-to-event analysis may be performed by letting  $Y_t$  denote an indicator that a unit is event free at time  $t$ , and letting  $A_{2,t}$  denote an indicator that the unit is uncensored at time  $t + 1$ . This definition of the random variables imposes restrictions on the functions  $m_t$  and  $r_t$ , which must be taken into account in estimation. For example, we know by definition that  $m_t(a_t, h_t) = 1$  for  $h_t$  such that  $y_s = 1$  for any  $s < t$ .

#### 4. Optimality Theory

Thus, far we have derived a novel sequential regression formula that identifies the effect of an LMTP. We now turn our attention to a discussion of efficiency theory for its estimation in the nonparametric model. The *efficient influence function* (EIF) is a key object in semi-parametric estimation theory, as it characterizes the asymptotic behavior of all regular and efficient estimators (Bickel et al. 1997). Knowledge of the EIF has important practical implications. First, the EIF is often useful in constructing locally efficient estimators. Second, the EIF estimating equation often enjoys desirable properties such as multiple robustness, which allows for some components of the data distribution to be inconsistently estimated while preserving consistency of the estimator. Third, asymptotic analysis of estimators constructed using the EIF often yields second-order bias terms, which require slow convergence rates (e.g.,  $n^{-1/4}$ ) for the nuisance parameters involved, thereby enabling the use of flexible regression techniques in estimating these quantities. Before we proceed developing such theory, note that it is not possible to construct  $n^{1/2}$ -consistent estimators of  $\theta$  for the threshold intervention in Example 1 where  $\text{dl}(a_t, l_t) = a_t \mathbb{1}(a_t \leq \delta) + \delta \mathbb{1}(a_t > \delta)$ . This is because the parameter is not pathwise differentiable. Intuitively, inspection of the parameter definition for  $\tau = 1$  yields the reason for the lack of  $n^{1/2}$ -estimability

$$\begin{aligned} \theta &= E\{m(A \mathbb{1}(A \leq \delta) + \delta \mathbb{1}(A > \delta), L)\} \\ &= E\{m(A, L) \mathbb{1}(A \leq \delta)\} + E\{m(\delta, L) \mathbb{1}(A > \delta)\}. \end{aligned}$$

The term  $E\{m(\delta, L) \mathbb{1}(A > \delta)\}$  in this expression involves estimation of the causal effect of a static intervention setting a continuous exposure to  $A = \delta$ . Efficient estimation theory is not available for estimation of such parameters in the nonparametric model (Bickel et al. 1997), since all possible gradients of the pathwise derivative would necessarily involve a Dirac delta function at  $\delta$ . An alternative approach to overcome this issue is to redefine the regime  $\text{dl}$  so that the parameter becomes pathwise differentiable. Such approach is taken by Díaz and

van der Laan (2013a); the interested reader is encouraged to consult the original research article. In this article, we avoid this problem by only considering interventions that satisfy the following assumption, which is a straightforward generalization of the assumption of Haneuse and Rotnitzky (2013) for  $\tau = 1$ .

**Assumption 4** (Piecewise smooth invertibility for continuous exposures). For each  $h_t$ , assume that the support of  $A_t$  conditional on  $H_t = h_t$  may be partitioned into subintervals  $\mathcal{I}_{t,j}(h_t) : j = 1, \dots, J_t(h_t)$  such that  $\text{dl}(a_t, h_t)$  is equal to some  $\text{dl}_j(a_t, h_t)$  in  $\mathcal{I}_{t,j}(h_t)$  and  $\text{dl}_j(\cdot, h_t)$  has inverse function  $\text{bl}_j(\cdot, h_t)$  with derivative  $\text{bl}'_j(\cdot, h_t)$  with respect to  $a_t$ .

Define

$$g_t^{\text{dl}}(a_t | h_t) = \sum_{j=1}^{J_t(h_t)} \mathbb{1}_{t,j}\{\text{bl}_j(a_t, h_t), h_t\} g_t\{\text{bl}_j(a_t, h_t) | h_t\} |\text{bl}'_j(a_t, h_t)|, \quad (3)$$

where  $\mathbb{1}_{t,j}\{u, h_t\} = 1$  if  $u \in \mathcal{I}_{t,j}(h_t)$  and  $\mathbb{1}_{t,j}\{u, h_t\} = 0$  otherwise. Under Assumption 4, it is easy to show that the p.d.f. of  $A_t^{\text{dl}}$  conditional on the history  $h_t$  is  $g_t^{\text{dl}}(a_t | h_t)$ . In the case of equation (1) the post-intervention p.d.f. becomes

$$\begin{aligned} g_t^{\text{dl}}(a_t | h_t) &= g_t(a_t - \delta | h_t) \mathbb{1}\{a_t < u_t(h_t)\} \\ &\quad + g_t(a_t | h_t) \mathbb{1}\{a_t + \delta \geq u_t(h_t)\}, \end{aligned}$$

which shows that piecewise smoothness is sufficient to handle interventions such as Equation (1) which are not smooth in the whole range of the exposure.

For discrete exposure variables, we let

$$g_t^{\text{dl}}(a_t | h_t) = \sum_{s_t \in \mathcal{A}_t} \mathbb{1}\{\text{dl}(s_t, h_t) = a_t\} g_t(s_t | h_t).$$

Assumption 4 and expression (3) ensure that we can use the change of variable formula when computing integrals over  $\mathcal{A}_t$  for continuous exposures. This is useful for studying properties of the parameter and estimators we propose.

Efficiency theory in this paper focuses on functions  $\text{dl}$  that do not depend on  $P$  (recall that the function is deterministic but allowed to take a randomizer as argument). Assumption 4 together with this assumption will ensure that the efficient influence function of  $\theta$  for LMTPs has a structure similar to the influence function for the effect of dynamic regimes. This yields two important advantages for estimation. First, the structure of the EIF allows for multiply robust estimation, which is not generally possible for random regimes  $\text{dl}$  that depend on  $P$ . See Díaz and van der Laan (2013a) and Kennedy (2019) (Example 3) for examples. Second, this similarity will allow us to generalize existing estimators for dynamic regimes. In what follows it will be useful to define the density ratio

$$r_t(a_t, h_t) = \frac{g_t^{\text{dl}}(a_t | h_t)}{g_t(a_t | h_t)},$$

and the function

$$\begin{aligned} \phi_t : z \mapsto & \sum_{s=t}^{\tau} \left( \prod_{k=t}^s r_k(a_k, h_k) \right) \{m_{s+1}(a_{s+1}^{\text{dl}}, h_{s+1}) - m_s(a_s, h_s)\} \\ & + m_t(a_t^{\text{dl}}, h_t) \end{aligned} \quad (4)$$

for  $t = \tau, \dots, 1$ . When necessary, we use the notation  $\phi_t(z; \eta)$  or  $\phi_t(z; \underline{\eta}_t)$  to highlight the dependence of  $\phi_t$  on  $\underline{\eta}_t = (r_t, m_t, \dots, r_\tau, m_\tau)$ . We also use  $\eta$  to denote  $(r_1, m_1, \dots, r_\tau, m_\tau)$ , and define  $\phi_{\tau+1}(Z; \eta) = Y$ .

**Theorem 2 (Efficient Influence Function).** Assume one of:

1.  $A_t$  is a discrete random variable for all  $t$ , or
2.  $A_t$  is a continuous random variable and the modified treatment policy  $\mathfrak{d}$  satisfies [Assumption 4](#).

Assume that  $\mathfrak{d}$  does not depend on  $P$ . The efficient influence function for estimating  $\theta = E[m_1(A^{\mathfrak{d}}, L_1)]$  in the nonparametric model is given by  $\phi_1(Z) - \theta$ .

Note that if  $\tau = 1$ , then  $\phi_1(Z) - \theta$  reduces to the efficient influence function for a single time point intervention presented by Díaz and van der Laan (2012), equal to  $r(A, L)[Y - m(A, L)] + m(\mathfrak{d}(A, L), L) - \theta$ . For the case  $\tau = 1$ , the efficiency bound equals  $\text{Var}\{\phi_1(Z) - \theta\} = E[r^2(A, L)\text{Var}(Y | A, L)] + E[m(\mathfrak{d}(A, L), L) - \theta]^2$ , which depends on three features of the data distribution and intervention: (i) the conditional variance of the outcome  $\text{Var}(Y | A, L)$ , (ii) the amount of treatment effect heterogeneity  $E[m(\mathfrak{d}(A, L), L) - \theta]^2$ , and (iii) the extent to which the intervention  $\mathfrak{d}$  differs from the observed regime, as measured by the density ratio  $r(A, L)$ . Furthermore, interventions that exert large exposure density changes in areas of high outcome variability are expected to yield larger efficiency bounds, since  $r^2(A, L)$  and the variance  $\text{Var}(Y | A, L)$  are positively correlated in such cases.

In the following, we let  $\eta' = (r'_1, m'_1, \dots, r'_\tau, m'_\tau)$  denote some value of  $\eta$ . This value will typically represent the probability limit of a given estimator  $\hat{\eta}$ . The efficient influence function satisfies the following property, which will be crucial to establish consistency of some estimators of  $\theta$  under multiple robustness assumptions, and to construct estimators of  $\theta$  under slow convergence rates for estimation of  $r_t$  and  $m_t$ .

**Lemma 1 (First-order parameter approximation).** Let

$$C'_{t,s} = \prod_{r=t+1}^{s-1} r'_r(A_r, H_r).$$

For each  $t \in \{0, \dots, \tau - 1\}$ , and for any  $\eta'$ , define the second-order error term

$$\text{Rem}_t(a_t, h_t; \eta') = \sum_{s=t+1}^{\tau} E \left[ C'_{t,s} \{r'_s(A_s, H_s) - r_s(A_s, H_s)\} \{m'_s(A_s, H_s) - m_s(A_s, H_s)\} \mid A_t = a_t, H_t = h_t \right], \quad (5)$$

where for  $t = 0$  the conditioning event is the null set, and for  $t = \tau$  we let  $\text{Rem}_\tau(a_\tau, h_\tau; \eta') = 0$ . Under [Assumption 4](#) we have

$$m_t(a_t, h_t) = E[\phi_{t+1}(Z; \eta') \mid A_t = a_t, H_t = h_t] + \text{Rem}_t(a_t, h_t; \eta'). \quad (6)$$

This lemma is analogous to Lemma 1 in Luedtke et al. (2017) and Lemma 2 in Rotnitzky, Robins, and Babino (2017) for the standard g-formula for dynamic regimes. Furthermore,

this lemma shares important connections to the von Mises-type expansions used in some of the semi-parametric inference literature (e.g., von Mises 1947; van der Vaart 1998; Robins et al. 2009). For  $t = 0$ , inspection of this lemma teaches us that it is possible to construct a consistent estimator of  $\theta$  by averaging  $\phi_1(Z_i; \hat{\eta})$  across the sample, where  $\hat{\eta}$  is an estimator such that  $\text{Rem}_0(\hat{\eta}) = o_P(1)$ . The latter consistency can be achieved under the condition that for each  $t$ , either  $r_t$  or  $m_t$  can be estimated consistently. This robustness is an interesting property in light of the fact that the parameter  $\theta$  depends on both  $r_t$  and  $m_t$  for all  $t$ .

### 5. Estimation and Statistical Inference

In this section we assume that preliminary estimators  $\hat{r}_t$ , and  $\hat{m}_t$  are available. These estimators may be obtained from flexible regression techniques such as support vector machines, regression trees, boosting, neural networks, splines, or ensembles thereof (Breiman 1996; van der Laan, Polley, and Hubbard 2007). As previously discussed, the consistency of these estimators will determine the consistency of our estimators of the parameter  $\theta$ . In particular,  $m_t : t = 1, \dots, \tau$  may be estimated as follows. Start by running a preferred regression algorithm of  $m_{\tau+1,i} = Y_i$  on  $(A_{\tau,i}, H_{\tau,i})$ . Then evaluate the estimator  $\hat{m}_\tau$  at  $(A_{\tau,i}^{\mathfrak{d}}, H_{\tau,i})$ , that is, compute the prediction  $\hat{m}_\tau(A_{\tau,i}^{\mathfrak{d}}, H_{\tau,i})$ . Use this prediction as the pseudo-outcome in a regression on  $(A_{\tau-1,i}, H_{\tau-1,i})$ , to obtain an estimate  $\hat{m}_{\tau-1}$ . Compute the pseudo-outcome  $\hat{m}_{\tau-1}(A_{\tau-1,i}^{\mathfrak{d}}, H_{\tau-1,i})$  and iterate the process until obtaining an estimate  $\hat{m}_1$ . For this and other estimators presented below it will be necessary to specify a regression algorithm to estimate conditional expectations. In this paper we advocate for the use of data-adaptive regression methods, which avoid reliance on tenuous parametric assumptions that can invalidate the conclusions of an otherwise well-designed and conducted study. In particular, we propose to use an ensemble regression algorithm known as the Super Learner (van der Laan, Polley, and Hubbard 2007), which builds a convex combination of regression algorithms in a user-given library, with weights chosen to minimize the cross-validated prediction error. While estimation methods for conditional expectations abound in the statistics and machine learning literature, data-adaptive methods to estimate a multivariate density ratio such as  $r_t$  are scarce. In [Section 5.4](#) we present a method in which we recast the density ratio estimation problem as a classification problem based on  $2n$  observations. Once the problem is recast in this way, any classification method from the statistical learning literature (such as Super Learning) may be used to estimate the density ratio  $r_t$ .

We start this section by discussing two simple estimators: inverse probability weighting and g-computation or substitution estimators. These estimators cannot generally achieve  $n^{1/2}$ -consistency under data-adaptive estimation of the nuisance parameters in a nonparametric model. We then present a targeted minimum loss-based estimator which is locally efficient,  $n^{1/2}$ -consistent, and  $\tau + 1$ -multiply robust consistent, under assumptions. We then present a sequential regression estimator which has the same properties but has the additional advantage that it is  $2^\tau$ -multiply robust consistent.

### 5.1. Substitution and Inverse Probability Weighted Estimators

The substitution estimator simply uses the preliminary estimator of  $m_t$  described above along with the recursive definition in Equation (2)

$$\hat{\theta}_{\text{sub}} = \frac{1}{n} \sum_{i=1}^n \hat{m}_1(A_{1,i}^{\text{dl}}, L_{1,i}).$$

The IPW estimator is based on the observation that, under [Assumption 4](#), we have the following alternative representation of the parameter of interest (see [Lemma 2](#) in the supplementary materials):

$$\theta = \mathbb{E} \left[ \left( \prod_{t=1}^{\tau} r_t(A_t, H_t) \right) Y \right].$$

Thus, for estimators  $\hat{r}_t$  of  $r_t$ , we define

$$\hat{\theta}_{\text{ipw}} = \frac{1}{n} \sum_{i=1}^n \left( \prod_{t=1}^{\tau} \hat{r}_t(A_{t,i}, H_{t,i}) \right) Y_i.$$

This estimator is an extension of the estimator proposed by Young, Hernán, and Robins (2014), where we allow the use of machine learning methods to estimate the density ratio (see [Section 5.4](#)).

If  $m_t$  and  $r_t$  are estimated within prespecified parametric models, then, by the delta method, both  $\hat{\theta}_{\text{sub}}$  and  $\hat{\theta}_{\text{ipw}}$  are asymptotically linear. In addition, they are  $n^{1/2}$ -consistent if the models are correctly specified. The bootstrap or an influence function-based estimator may be used to construct asymptotically correct confidence intervals. However, if the time-varying variables are high-dimensional or there are too many time points such that smoothing is necessary, the required consistency of  $\hat{m}_t$  and  $\hat{r}_t$  will hardly be achievable within parametric models. This issue may be alleviated through the use of data-adaptive estimators. Unfortunately,  $n^{1/2}$ -consistency of  $\hat{\theta}_{\text{sub}}$  and  $\hat{\theta}_{\text{ipw}}$  will generally require that  $\hat{m}_t$  and  $\hat{r}_t$  are consistent in  $L_2(\mathbb{P})$ -norm at parametric rate, which is generally not possible when using data-adaptive estimation of high-dimensional regressions. Thus, the asymptotic distribution will generally be unknown, making it difficult to construct confidence intervals and hypothesis tests. In the following, we use the efficient influence function to propose two estimators that are  $n^{1/2}$ -consistent and efficient under weaker assumptions, requiring only  $n^{1/2}$ -convergence of the second-order regression bias term  $\text{Rem}_0(\hat{\eta})$ .

### 5.2. Targeted Minimum Loss-Based Estimator

We start by presenting a generalization of the estimator we proposed for the case of an LMTP in a single time point in Díaz and van der Laan (2018), which is also a generalization of the estimator proposed by van der Laan and Gruber (2012) for longitudinal dynamic regimes. Compared to the estimator we proposed in Díaz and van der Laan (2012), the proposal in Díaz and van der Laan (2018) has the advantage that it does not require a tilting model for the density  $g_t$ , which may be computationally intensive. In order to avoid imposing entropy conditions on the initial estimators, we use sample splitting and

cross-fitting (Klaassen 1987; Zheng and van der Laan 2011; Chernozhukov et al. 2018). Let  $\mathcal{V}_1, \dots, \mathcal{V}_J$  denote a random partition of the index set  $\{1, \dots, n\}$  into  $J$  prediction sets of approximately the same size. That is,  $\mathcal{V}_j \subset \{1, \dots, n\}$ ;  $\bigcup_{j=1}^J \mathcal{V}_j = \{1, \dots, n\}$ ; and  $\mathcal{V}_j \cap \mathcal{V}_{j'} = \emptyset$ . In addition, for each  $j$ , the associated training sample is given by  $\mathcal{T}_j = \{1, \dots, n\} \setminus \mathcal{V}_j$ . We let  $\hat{\eta}_j$  denote the estimator of  $\eta$  obtained by training the corresponding prediction algorithm using only data in the sample  $\mathcal{T}_j$ . Further, we let  $j(i)$  denote the index of the validation set which contains observation  $i$ .

The targeted minimum loss-based estimator  $\hat{\theta}_{\text{tmle}}$  is computed as a substitution estimator that uses an estimate  $\tilde{m}_{1,j(i)}$  carefully constructed to solve the cross-validated efficient influence function estimating equation  $\mathbb{P}_n\{\phi_1(\cdot, \tilde{\eta}_{j(\cdot)}) - \hat{\theta}_{\text{tmle}}\} = 0$ . The construction of  $\tilde{m}_{1,j(i)}$  is motivated by the observation that the efficient influence function of  $\theta$  can be expressed as a sum of terms of the form:

$$\left( \prod_{k=1}^t r_k(a_k, h_k) \right) \{m_{t+1}(a_{t+1}^{\text{dl}}, h_{t+1}) - m_t(a_t, h_t)\},$$

which take the form of score functions  $\omega(W)\{M - \mathbb{E}(M | W)\}$  for appropriately defined variables  $M$  and  $W$  and some weight function  $\omega$ . It is well known that if  $\mathbb{E}(M | W)$  is estimated within a weighted generalized linear model with canonical link that includes an intercept, then the weighted MLE estimate solves the score equation  $\sum_i \omega(W_i)\{M_i - \hat{\mathbb{E}}(M_i | W_i)\} = 0$ . TMLE uses this observation to iteratively tilt preliminary estimates of  $m_t$  toward a solution of the efficient influence function estimating equation. The algorithm is defined as follows:

Step 1. Initialize  $\tilde{\eta} = \hat{\eta}$  and  $\tilde{m}_{\tau+1,j(i)}(A_{\tau+1,i}^{\text{dl}}, H_{\tau+1,i}) = Y_i$ .

Step 2. For  $s = 1, \dots, \tau$ , compute the weights

$$\omega_{s,i} = \prod_{k=1}^s \hat{r}_{k,j(i)}(A_{k,i}, H_{k,i})$$

Step 3. For  $t = \tau, \dots, 1$ :

- Fit the generalized linear tilting model

$$\text{link } \tilde{m}_t^{\epsilon}(A_{t,i}, H_{t,i}) = \epsilon + \text{link } \tilde{m}_{t,j(i)}(A_{t,i}, H_{t,i})$$

where  $\text{link}(\cdot)$  is the canonical link. Here,  $\text{link} \tilde{m}_{t,j(i)}(a_t, h_t)$  is an offset variable (i.e., a variable with known parameter value equal to one). The parameter  $\epsilon$  may be estimated by running a generalized linear model of the pseudo-outcome  $\tilde{m}_{t+1,j(i)}(A_{t+1,i}^{\text{dl}}, H_{t+1,i})$  with only an intercept term, an offset term equal to  $\text{link} \tilde{m}_{t,j(i)}(A_{t,i}, H_{t,i})$ , and weights  $\omega_{t,i}$ , using all the data points in the sample. An outcome bounded in an interval  $[a, b]$  may be analyzed with logistic regression (i.e.,  $\text{link} = \text{logit}$ ) by mapping it to an outcome  $(0, 1)$  through the transformation  $(Y - a)/(b - a)(1 - 2\epsilon) + \epsilon$  for some small value  $\epsilon > 0$ . This approach has robustness advantages compared to fitting a linear model as it guarantees that the predictions in the next step remain within the outcome space (Gruber and van der Laan see 2010).



- Let  $\hat{\epsilon}$  denote the maximum likelihood estimate, and update the estimates as

$$\text{link}\tilde{m}_{t,j(i)}^{\hat{\epsilon}}(A_{t,i}, H_{t,i}) = \hat{\epsilon} + \text{link}\tilde{m}_{t,j(i)}(A_{t,i}, H_{t,i})$$

$$\text{link}\tilde{m}_{t,j(i)}^{\hat{\epsilon}}(A_{t,i}^{\text{dl}}, H_{t,i}) = \hat{\epsilon} + \text{link}\tilde{m}_{t,j(i)}(A_{t,i}^{\text{dl}}, H_{t,i}).$$

The above procedure with canonical link guarantees the following score equation is solved:

$$\frac{1}{n} \sum_{i=1}^n \left( \prod_{k=1}^t \hat{r}_{k,j(i)}(A_{k,i}, H_{k,i}) \right) \{ \tilde{m}_{t+1,j(i)}(A_{t+1,i}^{\text{dl}}, H_{t+1,i}) - \tilde{m}_{t,j(i)}^{\hat{\epsilon}}(A_{t,i}, H_{t,i}) \} = 0$$

- Update  $\tilde{m}_{t,j(i)} = \tilde{m}_{t,j(i)}^{\hat{\epsilon}}$ ,  $t = t - 1$ , and iterate.

Step 4. The TMLE is defined as

$$\hat{\theta}_{\text{tmle}} = \frac{1}{n} \sum_{i=1}^n \tilde{m}_{1,j(i)}(A_{1,i}^{\text{dl}}, L_{1,i}).$$

The iterative procedure and the score equation argument above guarantee that

$$\frac{1}{n} \sum_{i=1}^n \sum_{t=1}^{\tau} \left( \prod_{k=1}^t \hat{r}_{k,j(i)}(A_{k,i}, H_{k,i}) \right) \{ \tilde{m}_{t+1,j(i)}(A_{t+1,i}^{\text{dl}}, H_{t+1,i}) - \tilde{m}_{t,j(i)}(A_{t,i}, H_{t,i}) \} = 0,$$

and thus that  $P_n\{\phi_1^{\tilde{\eta}(c)} - \hat{\theta}_{\text{tmle}}\} = 0$ . This fact is crucial to prove the weak convergence result of  $\hat{\theta}_{\text{tmle}}$  in [Theorem 3](#), which is useful to construct confidence intervals and hypothesis tests.

**Theorem 3 (Weak convergence of TMLE).** Assume the conditions of [Theorem 2](#) hold. Assume that  $\sum_{t=1}^{\tau} \|\hat{r}_t - r_t\| \|\tilde{m}_t - m_t\| = o_p(n^{-1/2})$  and that  $P\{r_t(A_t, H_t) < c\} = P\{\hat{r}_t(A_t, H_t) < c\} = 1$  for some  $c < \infty$ . Then

$$\sqrt{n}(\hat{\theta}_{\text{tmle}} - \theta) \rightsquigarrow N(0, \sigma^2),$$

where  $\sigma^2 = \text{Var}\{\phi_1(Z; \eta)\}$  is the nonparametric efficiency bound.

Note that  $n^{1/2}$ -consistency of the TMLE requires consistent estimation of all nuisance parameters  $(r_t, m_t) : t = 1, \dots, \tau$  at the rates stated in the theorem. The delta method implies these rates would be trivially achieved if  $r_t$  and  $m_t$  were estimated within correctly prespecified parametric models. The required rates may also be achievable by many data-adaptive regression algorithms. For example, these rates would be satisfied if  $\|\hat{r}_t - r_t\| = o_p(n^{-1/4})$  and  $\|\tilde{m}_t - m_t\| = o_p(n^{-1/4})$ . For  $r_t$ , the  $n^{-1/4}$ -rate may be achievable by  $\ell_1$  regularization (Bickel et al. 2009), regression trees (Wager and Walther 2015), neural networks (Chen and White 1999), or the highly adaptive lasso (Benkeser and van der Laan 2016). For  $m_t$ , establishing the  $n^{1/4}$ -rate requires more careful analysis as the outcome is fitted from data. Specifically, it is possible to bound the regression error in terms of an error related to estimation of the outcome plus an error purely due to regression estimation. Methods to study this type of two-stage estimator can be found in Ai and Chen (2003), Rubin and van der Laan (2005), Foster and Syrgkanis (2019), and Kennedy (2020), among others.

Beyond the  $\sqrt{n}$ -consistency implied by [Theorem 3](#), the TMLE is also multiply robust consistent under  $\tau + 1$  configurations of consistency in estimation of the nuisance parameters  $r_t$  and  $m_t$ .

**Lemma 2 ( $\tau + 1$  multiply robust consistency of TMLE).** Assume the conditions of [Theorem 2](#) hold. Assume that there is a time  $k$  such that  $\|\tilde{m}_t - m_t\| = o_p(1)$  for all  $t > k$  and  $\|\hat{r}_t - r_t\| = o_p(1)$  for all  $t \leq k$ . Then we have  $\hat{\theta}_{\text{tmle}} = \theta + o_p(1)$ .

[Lemma 2](#) is a direct consequence of Equation (3) in the [supplementary materials](#) together with the expression for  $\text{Rem}_1(\eta)$  in Equation (5). Note that although  $m_t$  implicitly depends on  $g_{t+1}, \dots, g_{\tau}$ , it is possible to construct estimators that achieve the conditions of the lemma, as the parameterization (2) means  $m_t$  and  $g_{t+1}, \dots, g_{\tau}$  are in fact variation independent, that is, one can construct consistent estimators of  $m_t$  without relying on consistent estimators of  $g_{t+1}, \dots, g_{\tau}$ .

Inspection of  $\text{Rem}_1(\eta)$  teaches us that a result stronger than [Lemma 2](#) should be possible, that is, consistent estimation of  $\theta$  should be achievable under a weaker sequential doubly robust consistent (SDR) assumption that  $\|\tilde{m}_t - m_t\| = o_p(1) \vee \|\hat{r}_t - r_t\| = o_p(1)$  for all  $t$ . Note, however, that estimators of  $\theta$  (such as TMLE) that use the recursive definition in (2) to estimate  $m_t$  can only be expected to satisfy the assumption in the [Lemma 2](#) (and not  $2^{\tau}$ -multiple robustness) because  $m_t$  cannot generally be consistently estimated using the recursive definition (2) unless  $m_s$  is consistently estimated for all  $s > t$ . In the following we present a sequential regression estimator that overcomes this limitation by using expression (6) instead of Equation (2) to estimate  $m_t$ . This fact motivates our use of the name *sequentially doubly robust* since the estimator relies on the fact that each  $m_t$  can be consistently estimated even if  $m_s$  is inconsistently estimated for some  $s > t$ . This is in contrast to the more common name  *$2^{\tau}$ -multiply robust*, which only applies to the estimator of  $\theta$ .

### 5.3. Sequential Regression Estimator Using SDR Unbiased Transformations

In this section we use the multiply robust unbiased transformation in expression (6) to obtain an estimate of  $m_t$ . This estimator is an extension to LMTPs of estimators proposed by Luedtke et al. (2017) and Rotnitzky, Robins, and Babino (2017) for longitudinal dynamic regimes. We say that  $\phi_{t+1}$  is a multiply robust unbiased transformation for  $m_t$  due to the following proposition, which is a straightforward consequence of [Lemma 1](#).

**Proposition 1.** Let  $\eta'$  be such that either  $m'_s = m_s$  or  $r'_s = r_s$  for all  $s > t$ . Then we have

$$E[\phi_{t+1}(Z; \eta') \mid A_t = a_t, H_t = h_t] = m_t(a_t, h_t).$$

This lemma motivates the construction of the sequential regression estimator by iteratively regressing an estimate of the data transformation  $\phi_{t+1}(Z; \eta)$  on  $(A_t, H_t)$ , starting at  $\phi_{\tau+1}(Z; \eta) = Y$ . Similar ideas have been used by others to obtain estimates of various causal inference parameters (Buckley and James 1979; Rubin and van der Laan 2007; Díaz and

van der Laan 2013b; Kennedy et al. 2017; Luedtke et al. 2017). For preliminary cross-fitted estimates  $\hat{r}_{1,j(i)}, \dots, \hat{r}_{\tau,j(i)}$ , the estimator is defined as follows:

Step 1. Initialize  $\phi_{\tau+1}(Z_i; \check{\eta}_{\tau,j(i)}) = Y_i$  for  $i = 1, \dots, n$ .

Step 2. For  $t = \tau, \dots, 1$ :

- Compute the pseudo-outcome  $\check{Y}_{t+1,i} = \phi_{t+1}(Z_i; \check{\eta}_{t,j(i)})$  for all  $i = 1, \dots, n$ .
- For  $j = 1, \dots, J$ :
  - Regress  $\check{Y}_{t+1,i}$  on  $(A_{t,i}, H_{t,i})$  using any regression technique and using only data points  $i \in \mathcal{T}_j$ .
  - Let  $\check{m}_{t,j}$  denote the output, update  $\check{\eta}_{t,j} = (\hat{r}_{t,j}, \check{m}_{t,j}, \dots, \hat{r}_{\tau,j}, \check{m}_{\tau,j})$ , and iterate.

Step 3. Define the sequential estimator regression as

$$\hat{\theta}_{sr} = \frac{1}{n} \sum_{i=1}^n \phi_1(Z_i; \check{\eta}_{j(i)}).$$

To prove that the SDR estimator is sequentially doubly robust consistent, it will be useful to have the following alternative representation of  $\text{Rem}_t(a_t, h_t; \eta)$ . Define the data-dependent parameter

$$\check{m}_t^\dagger(a_t, h_t) = E[\phi_{t+1}(Z; \check{\eta}_t) \mid A_t = a_t, H_t = h_t],$$

where the outer expectation is only with respect of the distribution  $P$  of  $Z$  (i.e.,  $\check{\eta}$  is fixed). Equation (6) yields

$$m_t(a_t, h_t) = \check{m}_t^\dagger(a_t, h_t) + \text{Rem}_t(a_t, h_t; \check{\eta}). \quad (7)$$

An induction argument together with Equation (5) yields the lemma below (proved in the supplementary materials).

**Lemma 3.** Assume that  $P\{r_t(A_t, H_t) < c\} = P\{\hat{r}_t(A_t, H_t) < c\} = 1$  for some  $c < \infty$ . Then

$$\text{Rem}_0(\check{\eta}) = \sum_{t=1}^{\tau} O_P(\|\hat{r}_t - r_t\| \|\check{m}_t - \check{m}_t^\dagger\|). \quad (8)$$

In comparison to Equation (5), the representation of the remainder term in Equation (8) avoids iterative definitions of the regression error terms involved. Unlike  $\|\check{m}_t - m_t\|$ , which implicitly depends on all  $\|\check{m}_s - m_s\| : s > t$ , the error term  $\|\check{m}_t - \check{m}_t^\dagger\|$  depends only on the consistency of regression procedure used in 5.3 to estimate the outer expectation in expression (7). This representation is thus more useful to establish sequential doubly robust consistency. In particular, we have:

**Lemma 4 ( $2^\tau$ -multiply robust consistency of SDR estimator).** Assume the conditions of Theorem 2 hold. Assume that, for each time  $t$ , either  $\|\hat{r}_t - r_t\| = o_P(1)$  or  $\|\check{m}_t - \check{m}_t^\dagger\| = o_P(1)$ . Then we have  $\hat{\theta}_{sr} = \theta + o_P(1)$ .

The sequential regression estimator also satisfies a weak convergence result analogous to Theorem 3:

**Theorem 4 (Weak convergence of SDR estimator).** Assume the conditions of Theorem 2 hold. Assume that  $\sum_{t=1}^{\tau} \|\hat{r}_t - r_t\| \|\check{m}_t - \check{m}_t^\dagger\| = o_P(n^{-1/2})$  and that  $P\{r_t(A_t, H_t) < c\} = P\{\hat{r}_t(A_t, H_t) < c\} = 1$  for some  $c < \infty$ . Then

$$\sqrt{n}(\hat{\theta}_{sr} - \theta) \rightsquigarrow N(0, \sigma^2),$$

where  $\sigma^2 = \text{var}\{\phi_1(Z; \eta)\}$  is the nonparametric efficiency bound.

Like its TMLE counterpart, the above theorem may be used to construct asymptotically valid confidence intervals and hypothesis tests. For example, the standard error may be estimated as the empirical variance of  $\phi_1(Z; \check{\eta}_{j(i)})$ , and this standard error may be used to compute Wald-type confidence intervals.

We note that the rates required for  $n^{1/2}$ -consistency in Theorems 3 and 4 are the same for the TMLE and SDR, and in this sense the SDR estimator does not seem to confer asymptotic advantages with respect to the TMLE. Furthermore, the TMLE and SDR estimators proposed in this paper do not allow for the construction of multiply robust confidence regions or p-values. Much of the recent literature focused on solving this problem proposes the construction of estimators that remain  $n^{1/2}$ -consistent under multiply robust consistency assumptions on the nuisance estimators. The interested reader is referred to van der Laan (2014), Farrell (2015), Benkeser et al. (2016), Díaz and van der Laan (2017), Díaz (2019), Smucler, Rotnitzky, and Robins (2019), and references therein as examples of this literature. The SDR and TMLE estimators proposed here could possibly be adapted to satisfy this property through extension of the methods of Díaz (2019) to LMTPs.

Lemmas 2 and 4 only state the conditions for consistency and do not provide the convergence rate of the TMLE and SDR under misspecification of the nuisance parameters. For example, for  $\tau = 2$  with  $m_2$  inconsistently estimated we know the TMLE will always be inconsistent and the SDR estimator may be consistent if the other conditions of Lemma 4 hold. However, the rate of consistency of the SDR in this case is unknown. We conjecture that the additional robustness of the SDR conferred by Lemma 4 endows this estimator with better finite sample behavior compared to the TMLE. We explore this conjecture in our simulation studies in Section 1 of the supplementary materials.

An important possible drawback of the sequential regression estimator is that the pseudo-outcome  $\check{Y}_{t,i}$  may be outside the bounds of the original outcome, which may in turn yield a parameter estimate  $\hat{\theta}_{sr}$  out of bounds of the parameter space. This may be especially problematic if the intervention  $d$  is allowed to yield near violations of the positivity assumption, in which case the density ratio  $r_t$  may be highly variable. While this can be remedied using truncation, an extension of the sequential regression infinite-dimensional TMLE strategy proposed by Luedtke et al. (2017) for the case of a static intervention may offer a more principled solution.

#### 5.4. Density Ratio Estimation Via Classification

The estimators proposed in the previous sections require a preliminary estimator of the density ratio  $r_t$ . A possible strategy

to estimate this density ratio is to obtain estimates of the density  $g_t$  and plug them into the definition of  $g_t^{cl}$  to compute the ratio. This approach may fail for several reasons, the most important being that data-adaptive estimators for high-dimensional conditional densities are scarce in the machine and statistical learning literature. In this section we propose a different approach, in which we recast the problem of estimating the density ratio  $r_t$  in terms of a classification problem in an augmented dataset that contains  $2n$  observations (Qin 1998; Cheng et al. 2004).

For a fixed time point, consider an augmented dataset of size  $2n$  in which we have duplicated each observation. In this augmented dataset, one of the duplicated records gets assigned the actually observed exposure,  $A_t$ , and the other one gets assigned exposure under the intervention,  $A_t^{cl}$ . We also introduce an indicator variable  $\Lambda$ , which is equal to one if the duplicated observation corresponds to the treatment under intervention, and zero otherwise. The augmented dataset can be represented as follows:  $(H_{\lambda,i,t}, A_{\lambda,i,t}, \Lambda_{\lambda,i} : \lambda = 0, 1; i = 1, \dots, n)$ , where  $\Lambda_{\lambda,i} = \lambda$  indexes the duplicates,  $H_{\lambda,i,t} = H_{i,t}$  is the history variable and is equal for both duplicated records, and  $A_{\lambda,i,t} = \lambda \times A_{i,t}^{cl} + (1 - \lambda) \times A_{i,t}$  is the natural exposure level if  $\lambda = 0$ , and the intervened exposure level if  $\lambda = 1$ .

Denote the probability distribution of  $(H_t, A_t, \Lambda)$  in the augmented dataset by  $P^\lambda$ , and the corresponding density by  $P^\lambda$ . Define the following parameter of  $P^\lambda$ :

$$u_t^\lambda(a_t, h_t) = P^\lambda(\Lambda = 1 \mid A_t = a_t, H_t = h_t).$$

Then, we have the following relation between the density ratio  $r_t$  and the distribution  $P^\lambda$ :

$$\begin{aligned} r_t(a_t, h_t) &= \frac{P^\lambda(a_t, h_t \mid \Lambda = 1)}{P^\lambda(a_t, h_t \mid \Lambda = 0)} \\ &= \frac{P^\lambda(\Lambda = 1 \mid A_t = a_t, H_t = h_t)}{P^\lambda(\Lambda = 0 \mid A_t = a_t, H_t = h_t)} \\ &= \frac{u_t^\lambda(a_t, h_t)}{1 - u_t^\lambda(a_t, h_t)}, \end{aligned}$$

where the first equality follows by definition of  $r_t$  and the definition of conditional density, the second follows by Bayes rule and the observation that  $1/2 = P^\lambda(\Lambda = 1) = P^\lambda(\Lambda = 0)$ , and the last one by definition. Thus, estimation of the density ratio may be carried out by estimating  $u^\lambda$  in the augmented dataset via any classification method available in the machine and statistical learning literature (e.g., super learning). In order to preserve properties of the estimator such as the asymptotic normality presented in Theorems 3 and 4, cross-validation and cross-fitting should be performed as follows. Let  $\mathcal{V}_1, \dots, \mathcal{V}_J$  denote a random partition of the index set  $\{1, \dots, n\}$  as in Section 5.2, and let  $\mathcal{T}_j = \{1, \dots, n\} \setminus \mathcal{V}_j$ . The estimator  $\hat{r}_{t,j(i)}(A_{i,t}, H_{i,t})$  may be computed as  $\hat{u}_{t,j(i)}^\lambda(A_{i,t}, H_{i,t}) / \{1 - \hat{u}_{t,j(i)}^\lambda(A_{i,t}, H_{i,t})\}$  where  $\hat{u}_{t,j(i)}^\lambda$  is the result of training a classification algorithm using data  $(A_{\lambda,i,t}, H_{\lambda,i,t}, \Lambda_{\lambda,i} : \lambda = 0, 1; i \in \mathcal{T}_j(i))$ . This cross-validation scheme is necessary because the rows in the augmented dataset are not iid. Performing cross-validation on the index set  $\{i : 1, \dots, n\}$  rather than  $\{\lambda = 0, 1; i = 1, \dots, n\}$  ensures that certain independencies required for the proofs of Theorems 3 and 4 remain true.

## 6. Illustrative Application

To illustrate the proposed methods, we estimate the effect of an intervention on patients with acute respiratory failure on 14-day survival among 10,044 intubated ICU patients. The data used in this illustration is the Weill Cornell Critical care Database for Advanced Research (WC-CEDAR), a comprehensive data repository containing demographic, laboratory, procedure, diagnosis, medication, microbiology, and flow sheet data documented as part of standard care. I

The study cohort includes patients who are invasively mechanically ventilated. Study time begins at their time of intubation. The exposure of interest is the worst daily arterial partial pressure of oxygen to fraction of inspired oxygen (PaO<sub>2</sub> to Fio<sub>2</sub>, or P/F) ratio. The P/F ratio is a continuous measure of hypoxemia and is used to quantify the severity of acute respiratory failure. For example, the Berlin definition of Acute Respiratory Distress Syndrome (ARDS) uses a P/F ratio of  $< 100$  for its classification of severe ARDS,  $100 \leq$  P/F ratio  $< 200$  for moderate ARDS, and  $200 \leq$  P/F ratio  $< 300$  for mild. There are several ways physicians can intervene on an individual patient's P/F ratio, but the most direct modifications are through supplemental oxygen, for example, invasive and noninvasive forms of mechanical ventilation. When a patient is mechanically ventilated, various changes to device parameters can significantly affect the patient's P/F ratio. Among other interventions to increase patients' P/F ratios are neuromuscular blockade medications, physical techniques such as prone positioning, and treating the underlying condition e.g. administering loop diuretic, antibiotics or steroids. Some of these interventions have negative effects such as hypotension and renal failure. Our goal is to estimate the overall effect on survival from the time of invasive mechanical ventilation of an intervention that would increase the P/F ratio by 50 among patients with clinically defined acute respiratory failure (P/F ratio  $< 300$ ). An increase of 50 is both clinically feasible and clinically meaningful. Figure 1 shows an example of the pre- and postintervention distribution of the P/F ratio for day 1.

Baseline confounders include age, sex, race, number of Elixhauser comorbidities, and pneumonia status. Time-dependent




Figure 1. Pre- and postintervention distribution of the P/F ratio at day 1 in our illustrative application.

**Table 1.** Estimated 28-day ICU mortality under the observed P/F ratios and a treatment policy that increases P/F ratio by 50 units when below 300.

Estimator	No shift			Trt. Policy			Trt. Effect			P
	Estimate	SE	95% CI	Estimate	SE	95% CI	Estimate	SE	95% CI	
IPW	0.754	–	(–, –)	0.978	–	(–, –)	0.223	–	(–, –)	–
Sub.	0.877	–	(–, –)	0.891	–	(–, –)	0.014	–	(–, –)	–
TMLE	0.888	0.004	(0.880, 0.895)	0.896	0.009	(0.879, 0.914)	0.009	0.008	(–0.006, 0.024)	0.244
SDR	0.883	0.004	(0.875, 0.890)	0.898	0.009	(0.880, 0.915)	0.015	0.008	(–0.001, 0.030)	0.053

confounders include daily mechanical ventilation status (invasive or non-invasive) and the Sequential Organ Failure Assessment (SOFA) score with the pulmonary component removed. Patients are censored at their day of hospital discharge, as vital status was unknown after this point.

We estimate the effect of the LMTP using substitution, IPW, TMLE, and SDR estimation. The functions  $r_t$  and  $m_t$  are estimated using the Super Learner (van der Laan, Polley, and Hubbard 2007). We included multivariate adaptive regression splines (MARS), extreme gradient boosting, logistic regression with  $\ell_1$  regularization (LASSO), and simple logistic regression as our candidate learners. Importantly, the estimators of  $r_t$  and  $m_t$  condition on the complete history of all variables and do not make Markov assumptions. The average weights given to each learner at each of the 14 time points for  $r$  and  $m$  are shown in Figure 3 in the supplementary materials. Using SDR to adjust for right censoring, we estimated the 14-day survival under no intervention on P/F ratios to be 88.3% (95% CI 87.5, 89.0%). The estimate under an increase of 50 among hypoxic patients is estimated as 89.8% (95% CI 88.0, 91.5%). This intervention would increase 14-day survival by 1.5% (95% CI –0.1, 3%). All results, including those using the TMLE, substitution, and IPW estimators, can be found in Table 1.

While the substitution estimator as well as the TMLE and SDR estimators produced comparable results, the IPW estimates an increase of 22% in survival. The disagreement between the IPW and other estimators is due to the high variability of the density ratios, which make the IPW estimator highly variable. For example, the weights  $\prod_{t=1}^{\tau} \hat{r}_t(A_t, H_t)$  have a coefficient of variation of 5.5%. Here we note that this instability of the IPW does not mean that there are positivity violations in our application (the maximum value of  $\prod_{t=1}^{\tau} \hat{r}_t(A_t, H_t)$  is 97.7), but rather that the weights are highly variable and so is the IPW estimator. The distribution of the weights for the IPW estimator is presented in Figure 4 in the supplementary materials.

## Supplementary Materials

Web supplementary materials contain all theoretical results and additional results for the simulation study and application.

## ORCID

Iván Díaz  <http://orcid.org/0000-0001-9056-2047>

## References

- Ai, C., and Chen, X. (2003), “Efficient Estimation of Models With Conditional Moment Restrictions Containing Unknown Functions,” *Econometrica*, 71, 1795–1843. [853]
- Bang, H., and Robins, J. M. (2005), “Doubly Robust Estimation in Missing Data and Causal Inference Models,” *Biometrics*, 61, 962–973. [847]
- Benkeser, D., and van der Laan, M. (2016), “The Highly Adaptive Lasso Estimator,” In *2016 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*. IEEE; pp. 689–696. [853]
- Benkeser, D., Carone, M., van der Laan, M. J., and Gilbert, P. (2016), “Doubly-Robust Nonparametric Inference on the Average Treatment Effect,” Technical Report 356, U.C. Berkeley Division of Biostatistics Working Paper Series. [854]
- Bickel, P. J., Klaassen, C. A. J., Ritov, YA., and Wellner, J. A. (1997), *Efficient and Adaptive Estimation for Semiparametric Models*, New York: Springer-Verlag. [847,850]
- Bickel, P. J., Ritov, Y., and Tsybakov, A. B. (2009), “Simultaneous Analysis of Lasso and Dantzig Selector,” *The Annals of Statistics*, 37, 1705–1732. [853]
- Breiman, L. (1996), “Stacked Regressions,” *Machine Learning*, 24, 49–64. [851]
- Buckley, J., and James, I. (1979), “Linear Regression With Censored Data,” *Biometrika*, 66, 429–436. [853]
- Chen, X., and White, H. (1999), “Improved Rates and Asymptotic Normality for Nonparametric Neural Network Estimators,” *IEEE Transactions on Information Theory*, 45, 682–691. [853]
- Cheng, K. E., and Chu, C.K. “Semiparametric Density Estimation Under a Two-Sample Density Ratio Model,” *Bernoulli*, 10, 583–604, 2004. [855]
- Chernozhukov, V., Chetverikov, D., Demirer, M., Dufo, E., Hansen, C., Newey, W., and Robins, J. (2018), “Double/Debiased Machine Learning for Treatment and Structural Parameters,” *The Econometrics Journal*, 21, C1–C68. [847,852]
- Díaz, I. “Statistical Inference for Data-Adaptive Doubly Robust Estimators With Survival Outcomes,” *Statistics in Medicine*, 38, 2735–2748. [854]
- Díaz, I., and Hejazi, N. S. (2020), Causal mediation analysis for stochastic interventions. *Journal of the Royal Statistical Society, Series B*. Available at <https://rss.onlinelibrary.wiley.com/doi/abs/10.1111/rssb.12362>. [848]
- Díaz, I., and van der Laan, M. J. (2012), “Population Intervention Causal Effects Based on Stochastic Interventions,” *Biometrics*, 68, 541–549. [846,848,851,852]
- (2013a), “Assessing the Causal Effect of Policies: An Example Using Stochastic Interventions,” *The International Journal of Biostatistics*, 9, 161–174. [850]
- (2013b), “Targeted Data Adaptive Estimation of the Causal Dose-Response Curve,” *Journal of Causal Inference*, 1, 171–192. [854]
- (2017), “Doubly Robust Inference for Targeted Minimum Loss-Based Estimation in Randomized Trials With Missing Outcome Data,” *Statistics in Medicine*, 36, 3807–3819. [854]
- (2018), “Stochastic Treatment Regimes,” in *Targeted Learning in Data Science*, Cham: Springer; pp. 219–232. [848,852]
- Farrell, M. H. (2015), “Robust Inference on Average Treatment Effects With Possibly More Covariates Than Observations,” *Journal of Econometrics*, 189, 1–23. [854]
- Foster, D. J., and Syrgkanis, V. (2019), “Orthogonal Statistical Learning” arXiv:1901.09036. [853]
- Gruber, S., and van der Laan, M. J. (2010), “A Targeted Maximum Likelihood Estimator of a Causal Effect on a Bounded Continuous Outcome,” *The International Journal of Biostatistics*, 6. [852]
- Haneuse, S., and Rotnitzky, A. (2013), “Estimation of the Effect of Interventions That Modify the Received Treatment,” *Statistics in Medicine*. [846,847,848,850]
- Kennedy, E. H. (2019), “Nonparametric Causal Effects Based on Incremental Propensity Score Interventions,” *Journal of the American Statistical Association*, 114, 645–656. [849,850]

- (2020), “Optimal Doubly Robust Estimation of Heterogeneous Causal Effects,” arXiv:2004.14497. [853]
- Kennedy, E. H., Ma, Z., McHugh, M. D., and Small, D. S. (2017), “Non-Parametric Methods for Doubly Robust Estimation of Continuous Treatment Effects,” *Journal of the Royal Statistical Society, Series B*, 79, 1229–1245. [846,854]
- Klaassen, C. A. J. (1987), “Consistent Estimation of the Influence Function of Locally Asymptotically Linear Estimators,” *The Annals of Statistics*, 1548–1562. [852]
- Luedtke, A. R., Sofrygin, O., van der Laan, M. J., and Carone, M. (2017), “Sequential Double Robustness in Right-Censored Longitudinal Models,” arXiv:1705.02459. [847,851,853,854]
- Molina, J., Rotnitzky, A., Sued, M., and Robins, J. M. (2017), “Multiple Robustness in Factorized Likelihood Models,” *Biometrika*, 104, 561–581. [847]
- Neugebauer, R., and van der Laan, M. J. (2007), “Nonparametric causal Effects Based on Marginal Structural Models,” *Journal of Statistical Planning & Inference*, 137, 419 – 434. ISSN 0378-3758. DOI: 10.1016/j.jspi.2005.12.008. [846]
- Pearl, J. (2009), *Causality: Models, Reasoning, and Inference*. Cambridge: Cambridge University Press. [848]
- Qin, J. (1998), “Inferences for Case-Control and Semiparametric Two-Sample Density Ratio Models,” *Biometrika*, 85, 619–630. [855]
- 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. [847,848,849]
- Robins, J., Li, L., Tchetgen, E., and van der Vaart, A. W. (2009), “Quadratic Semiparametric Von Mises Calculus,” *Metrika*, 69, 227–247. [847,851]
- Robins, J. M., Hernán, M. A., and Siebert, U. (2004), “Effects of Multiple Interventions,” *Comparative Quantification of Health Risks: Global and Regional Burden of Disease Attributable to Selected Major Risk Factors*, 1, 2191–2230. [846,847,848,849]
- Robins, J. M. (2000), “Robust Estimation in Sequentially Ignorable Missing Data and Causal Inference Models,” in *Proceedings of the American Statistical Association*, p. 6–10. [847]
- Robins, J. M., Rotnitzky, A., and Zhao, L. P. (1994), “Estimation of Regression Coefficients When Some Regressors Are Not Always Observed,” *Journal of the American Statistical Association*, 89, 846–866. [847]
- Rotnitzky, A., Robins, J., and Babino, L. (2017), “On the Multiply Robust Estimation of the Mean of the g-functional” arXiv:1705.08582. [847,851,853]
- Rubin, D., and van der Laan, M. J. (2005), “A General Imputation Methodology for Nonparametric Regression With Censored Data.” Technical Report 194, U.C. Berkeley Division of Biostatistics Working Paper Series. [853]
- (2007), “A Doubly Robust Censoring Unbiased Transformation,” *The International Journal of Biostatistics*, 3. [847,853]
- Smucler, E., Rotnitzky, A., and Robins, J. M. (2019), “A Unifying Approach for Doubly-Robust  $\ell_1$  Regularized Estimation of Causal Contrasts,” arXiv: 1904.03737. [854]
- Stock, J. H. (1989), “Nonparametric Policy Analysis,” *Journal of the American Statistical Association*, 84, 567–575. [846]
- Taubman, S. L., Robins, J. M., Mittleman, M. A., and Hernán, M. A. (2009), “Intervening on Risk Factors for Coronary Heart Disease: An Application of the Parametric g-Formula,” *International Journal of Epidemiology*, 38, 1599–1611. [847,848]
- Tchetgen Tchetgen, E. J. (2009), “A Commentary on g. Molenberghs’s Review of Missing Data Methods,” *Drug Information Journal*, 43, 433–435. [847]
- van der Laan, M. J. (2014), “Targeted Estimation of Nuisance Parameters to Obtain Valid Statistical Inference,” *The International Journal of Biostatistics*, 10, 29–57. [854]
- van der Laan, M. J., and Gruber, S. (2012), “Targeted Minimum Loss Based Estimation of Causal Effects of Multiple Time Point Interventions,” *The International Journal of Biostatistics*, 8, 1–41. [847,852]
- van der Laan, M. J., and Robins, J. M. (2003), *Unified Methods for Censored Longitudinal Data and Causality*, New York: Springer. [847]
- van der Laan, M. J., and Rose, S. (2011), *Targeted Learning: Causal Inference for Observational and Experimental Data*, New York: Springer. [847]
- (2018), *Targeted Learning in Data Science: Causal Inference for Complex longitudinal Studies*, New York: Springer. [847]
- van der Laan, M. J., and Rubin, D. (2006), “Targeted Maximum Likelihood Learning,” *The International Journal of Biostatistics*, 2, 1–40. [847]
- van der Laan, M.J., Polley, E., and Hubbard, A. (2007), “Super Learner,” *Statistical Applications in Genetics & Molecular Biology*, 6. [851,856]
- van der Vaart, A. W. *Asymptotic Statistics*, Cambridge: Cambridge University Press, 1998. [847,851]
- von Mises, R. (1947), “On the Asymptotic Distribution of Differentiable Statistical Functions,” *The Annals of Mathematical Statistics*, 18, 309–348. [847,851]
- Wager, S., and Walther, G. “Adaptive Concentration of Regression Trees, With Application to Random Forests,” arXiv: 1503.06388, 2015. [853]
- Young, J. G., Hernán, M. A., and Robins, J. M. (2014), “Identification, Estimation and Approximation of Risk Under Interventions That Depend on the Natural Value of Treatment Using Observational Data,” *Epidemiologic Methods*, 3, 1–19. [847,848,849,852]
- Zheng, W., and van der Laan, M. J. (2011), “Cross-Validated Targeted Minimum-Loss-Based Estimation,” in *Targeted Learning*, eds. M. J. van der Laan, and S. Rose, New York: Springer-Verlag; pp. 459–474. [847,852]