Local Longitudinal Modified Treatment Policies

Herbert Susmann¹ and Iván Díaz¹

¹Division of Biostatistics, Department of Population Health, New York University Grossman School of Medicine, USA

May 13, 2024

Abstract

Longitudinal Modified Treatment Policies (LMTPs) provide a framework for defining a broad class of causal target parameters for continuous and categorical exposures. We propose Local LMTPs, a generalization of LMTPs to settings where the target parameter is conditional on subsets of units defined by the treatment or exposure. Such parameters have wide scientific relevance, with well-known parameters such as the Average Treatment Effect on the Treated (ATT) falling within the class. We provide a formal causal identification result that expresses the Local LMTP parameter in terms of sequential regressions, and derive the efficient influence function of the parameter which defines its semi-parametric and local asymptotic minimax efficiency bound. Efficient semi-parametric inference of Local LMTP parameters requires estimating the ratios of functions of complex conditional probabilities (or densities). We propose an estimator for Local LMTP parameters that directly estimates these required ratios via empirical loss minimization, drawing on the theory of Riesz representers. The estimator is implemented using a combination of ensemble machine learning algorithms and deep neural networks, and evaluated via simulation studies. We illustrate in simulation that estimation of the density ratios using Riesz representation might provide more stable estimators in finite samples in the presence of empirical violations of the overlap/positivity assumption.

1 Introduction

Many causal estimands of scientific interest are defined in terms of contrasts of the marginal means of counterfactual outcomes under different interventions averaged over a population. In many cases, it is also of interest to understand the effect of an intervention only among those who received it (Heckman et al., 2001). For instance, in cross-sectional settings with a binary treatment the well-known Average Treatment Effect (ATE) is the expected difference in counterfactual outcomes under treatment and control averaged over the entire population, and the Average Treatment Effect on the Treated (ATT) averages only over the subpopulation that received the treatment.

As causal effects analogous to the ATE can be defined in more complex data structures and for more complex interventions, so can the ATT be generalized. In particular, we work within the framework of Longitudinal Modified Treatment Policies (LMTPs), which cover a broad class of causal effects for continuous, binary, and time to event outcomes (Díaz and van der Laan, 2012; Haneuse and Rotnitzky, 2013; Díaz et al., 2023; Hoffman et al., 2023; Díaz et al., 2024). LMTPs are defined as the population expected counterfactual outcome under an intervention defined as a modified treatment policy (MTP). An MTP intervention changes the natural value of the exposure (Young et al., 2014) (that is, the exposure an individual would have received under no intervention) according to a fixed function. For example, an MTP for a continuous exposure could be defined as a function that shifts the natural value of the exposure upwards by a fixed amount. The novelty of MTPs compared to dynamic treatment rules, which define treatment as a function of timevarying covariates and prior treatment status, is that MTPs are allowed to depend on the natural value of the exposure. LMTPs can handle continuous, categorical, multivariate, and time-to-event exposures in both cross-sectional and longitudinal data structures. We propose a novel causal parameter, referred to as a Local LMTP, defined as the expected counterfactual outcome under an MTP intervention among any subset of the population defined in terms of exposure status.

To illustrate the difference between LMTP and Local LMTP parameters, suppose we wish to estimate the effect of exercise on health outcomes in an observational study. The exposure is defined as the number of hours spent per week engaging in aerobic exercise, measured weekly over a 3 month period. The outcome is blood pressure measured at the end of the study period. The intervention is defined as an MTP in which 30 additional minutes of exercise are added to each individuals weekly exercise total. The LMTP causal estimand is defined as the population average blood pressure under the counterfactual of all individuals following the MTP. However, it may be of significant scientific interest to understand how the MTP affects individuals who otherwise would not have exercised at all. As such, a Local LMTP parameter could be defined as the average blood pressure under the MTP among the subpopulation who did not exercise.

We show that the Local LMTP parameter is causally identifiable under similar conditions as the LMTP parameter. The result expresses the extended g-formula in terms of sequential regressions, similar to the result for the LMTP parameter (Díaz et al., 2023) and for the ATE of a binary treatment (Luedtke et al., 2017; Bang and Robins, 2005).

We propose several estimators of the Local LMTP parameter. Substitution and inverse probability weighted (IPW) estimators can be used, however their performance depends on consistent estimation of the corresponding nuisance parameters, and their sampling distribution is generally unknown when data-adaptive methods (e.g., model selection) are used for estimating the nuisance parameters. As such, we also propose a multiply-robust estimator using Targeted Minimum Loss-Based Estimation (TMLE). The development of the TMLE estimator is based on deriving the Efficient Influence Function (EIF) of the target parameter, which not only allows us to construct an estimator that can accommodate data-adaptive regression methods, but also defines the efficiency bound for estimating the parameter in a non-parametric model.

The IPW and TMLE estimators require estimates of cumulated inverse probability (or inverse probability density) weights, which can be challenging to estimate in practice if the exposure is continuous, and can lead to empirical positivity violations if the exposure is continuous or categorical. One common approach to estimating inverse probability weights is to first estimate the required conditional probability (density) and then invert them to form weights. However, this can lead to significant instability in the resulting estimators, especially for high-dimensional or continuous treatments and in longitudinal data structures where inverse probabilities (densities) are cumulated. Díaz et al. (2023) used an approach based on a procedure for estimating density ratios, but this approach does not neatly generalize to Local LMTPs and can also suffer from instability

in the resulting estimators. As an alternative, we propose a strategy based on the Riesz representers that estimates the weights directly through empirical loss minimization (Chernozhukov et al., 2022a). This approach avoids the instabilities inherent when weights are formed by inverting and multiplying probabilities (densities). Furthermore, we demonstrate how ensemble learning can be applied such that a variety of flexible machine learning methods can be used to directly estimate the weights. The good performance of this approach is of wider general interest beyond estimating LMTPs and applies to any estimator that requires inverse probability weights, and especially to longitudinal settings where estimators require cumulating many inverse probabilities.

The main contributions of this work are twofold. First, we define the novel Local LMTP parameter, establish causal identification results, and propose a semi-parametric efficient estimator based on TMLE. Second, we estimate the inverse probability weights using empirical loss minimization based on Riesz representers, and show that these weight estimators can lead to improved finite sample performance of the TMLE. The rest of the manuscript unfolds as follows: in the remainder of the introduction we discuss related prior work. In Section 2 we rigorously define the Local LMTP parameter within a causal structural equation model and provide the causal identification result. In Section 3 we analyze the Local LMTP statistical parameter in the framework of semi-parametric efficiency theory and derive its EIF. In Section 4 we present an estimation strategy based on TMLE and direct estimation of the cumulated probability weights. The performance of the methods are investigated via simulation studies in Section 5.

1.1 Prior Work

The cross-sectional ATT has been extensively studied in the case of binary exposures (Heckman, 1995; Hahn, 1998; Shpitser and Pearl, 2009; Leacy and Stuart, 2014; Wang et al., 2017; Matsouaka et al., 2023). Doubly-robust estimators have been proposed based on TMLE (Hubbard et al., 2011) and augmented inverse probability weighting (Moodie et al., 2018). The ATT has also been extended to the case of categorical exposures (VanderWeele and Hernan, 2013). In the longitudinal setting, parameters analogous to the ATT have been studied in the context of instrumental variable designs (Tchetgen Tchetgen and Vansteelandt, 2013; Liu et al., 2015) and in a more general setting using Marginal Structural Models (Schomaker and Baumann, 2023).

LMTPs grew out of earlier work on defining and estimating causal effects in longitudinal settings and with complex interventions and non-binary exposures. Díaz and van der Laan (2012) proposed doubly robust estimators for causal effects defined via shift interventions for single time points. This work was further developed by (Haneuse and Rotnitzky, 2013), who introduced the term *modified treatment policy*. From another angle, there was research on dynamic treatment regimes in longitudinal settings (Robins et al., 2004; Richardson and Robins, 2013; Young et al., 2014), which also accommodates interventions that depend on the natural value of treatment or are stochastic. LMTPs synthesize these lines of research into a general framework for defining and estimating longitudinal causal effects based on MTPs (Díaz et al., 2023).

The goal of the present work is to bridge together the above strands of literature, proposing an end-to-end methodology for the definition, identification, and estimation of the effects of modified treatment policies within strata of the population defined by values of the exposure.

The definition of our causal parameters relies on the structural causal models of Pearl (2009), but analogous definitions could have been achieved under a potential outcome framework (Richardson and Robins, 2013). The analysis of the properties of the statistical identification formula, as

well as the development of estimators, builds on a long line of research in general semi-parametric efficiency theory (Bickel, 1982; van der Vaart and Wellner, 1996), and more recent results for longitudinal settings (Luedtke et al., 2017; Rotnitzky et al., 2017). The TMLE framework for constructing efficient estimators of parameters within non-parametric models is described thoroughly in van der Laan and Rubin (2006); van der Laan and Rose (2011). Our approach for estimating cumulated inverse probabilities is inspired by a recent research related to the use of Riesz representers in statistical estimation (Chernozhukov et al., 2021, 2022a,b, 2023).

2 Causal Effects

The longitudinal data structure for Local LMTPs is the same as for LMTPs (Díaz et al., 2023), which we now review. Let $Z = (L_1, A_1, L_2, A_2, \ldots, L_{\tau}, A_{\tau}, Y)$ be a random variable where, for $t \in \{1, \ldots, \tau\}$, L_t is a vector of time-varying covariates, A_t a vector of categorical, multivariate, or continuous exposure variables, and Y a binary or continuous outcome. Suppose we have a sample Z_1, \ldots, Z_n of of i.i.d. draws $Z \sim P_0$, where P_0 falls in a non-parametric statistical model \mathcal{M} . Denote the history and future of a random variable as $\bar{X}_t = (X_1, \ldots, X_t)$ and $\underline{X}_t = (X_t, \ldots, X_{\tau})$. For succinctness, we will write \bar{X} and \underline{X} to denote the complete history and future of a random variable (\bar{X}_{τ} and \underline{X}_{τ} , respectively). Let $H_t = (\bar{A}_{t-1}, \bar{L}_t)$ be the history of all variables until A_t . Let $g_t(a_t, h_t)$ be the probability density or probability mass function of A_t conditional on $H_t = h_t$, evaluated at a_t . We use $\mathrm{supp}\{\cdot \mid \mathcal{B}\}$ to denote the support of a random variable conditional on the set \mathcal{B} .

The causal model is formalized via a non-parametric structural equation model (Pearl, 2009). We assume the observed data are generated according to the following deterministic functions, for $\{1, \ldots, \tau\}$:

$$L_{t} = f_{L_{t}}(A_{t-1}, H_{t-1}, U_{L,t})$$

$$A_{t} = f_{A_{t}}(H_{t}, U_{A,t}),$$

$$Y = f_{Y}(A_{\tau}, H_{\tau}, U_{Y}),$$

where $U = (U_{L,t}, U_{A,t}, U_Y : t \in \{1, ..., \tau\})$ is a set of exogenous variables. Note that the model implies a particular time-ordering of the variables: each L_t happens before the corresponding A_t , and Y occurs last. Interventions are defined by replacing A_t with a new random variable A_t^d (we will give examples of the construction of this random variable in what follows). An intervention \bar{A}_{t-1}^d on all exposures from time 1 to time t-1 induces counterfactual variables $L_t(\bar{A}_{t-1}^d) = f_{L_t}(A_{t-1}^d, H_{t-1}(\bar{A}_{t-2}^d, U_{L,t}))$ and $A_t(\bar{A}_{t-1}^d) = f_{A_t}(H_t(\bar{A}_{t-1}^d), U_{A,t}))$, with counterfactual history defined as $H_t(\bar{A}_{t-1}^d) = (\bar{A}_{t-1}^d, \bar{L}_t(\bar{A}_{t-1}^d))$. The counterfactual variable $A_t(\bar{A}_{t-1}^d) =$ $f_{A_t}(H_t(\bar{A}_{t-1}^d), U_{A,t})$ is referred to as the *natural value of treatment* (Young et al., 2014), and is interpreted as the value treatment would have taken had the intervention been implemented but discontinued right after time t. Intervening on all treatment variables up to time $t = \tau$ induces the counterfactual outcome $Y(\bar{A}^d) = f_Y(A_{\tau}^d, H_{\tau}(\bar{A}_{\tau-1}^d), U_Y)$. LMTPs are a particular type of intervention in which A_t^d is defined as a function of the natural value of treatment at time t and the complete counterfactual history.

Definition 1. (*Díaz et al.*, 2023, *Definition 1*) The intervention A_t^d is an LMTP if it has a representation $A_t^d = d(A_t(\bar{A}_{t-1}^d), H_t(\bar{A}_{t-1}^d))$ for an arbitrary function d.

Shift interventions, described below, are a popular example of an LMTP originally proposed in Díaz and van der Laan (2012) and discussed further in Díaz and van der Laan (2018), Haneuse and Rotnitzky (2013), and Hoffman et al. (2023).

Example 1 (Shift LMTP). Suppose there exists some u_t such that $P(A_t > u_t | H_t = h_t) = 1$ for all $t \in \{1, ..., \tau\}$. For some fixed δ , define the intervention as

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

The literature on MTPs has so far only considered identification and estimation of parameters of the type $E[Y(\bar{A}^d)]$ which average counterfactual outcomes for all units in the population. In what follows, we present identification and estimation theory for parameters that condition on a subset of the treatment space, together with motivating examples where such parameters are of scientific relevance.

2.1 Definition of Causal Effects

In this section we formalize the causal effect associated with an LMTP conditional on exposure status. Let $\overline{B} \subset \overline{A}$ be a subset of the space of all possible longitudinal treatment assignments. Define the counterfactual Local LMTP outcome average as

$$\theta^* = \mathbb{E}\left\{Y(\bar{A}^d) \mid \bar{A} \in \bar{\mathcal{B}}\right\}.$$

Note that the parameter reduces to the traditional LMTP parameter if $\overline{B} = \overline{A}$ (that is, if the parameter conditions on the space of all possible treatment assignments). Next, we give examples of cross-sectional and longitudinal Local LMTP parameters of scientific interest.

Example 2 (Cross-sectional average treatment effect on the treated (ATT)). Assume the treatment A is binary, such as a medication. It is often of interest to estimate the average treatment effect on the treated, given by E[Y(1) - Y(0) | A = 1]. The average outcome under treatment for the treated population, E[Y(1) | A = 1] = E[Y | A = 1], can be easily identified and estimated. The average outcome under control for the treated population, E[Y(0) | A = 1], is a true counterfactual quantity and requires extra work. In the form of a local LMTP, letting d(A, W) = 0 and $\mathcal{B} = \{1\}$, this quantity is given by $\theta^* = E[Y(d) | A \in \mathcal{B}]$.

Example 3 (Longitudinal policy-relevant effects). For continuous or numerical treatments, it is often of scientific interest to investigate what would have been the outcome in a counterfactual world where treatment would have been increased by some user-given amount. For example, let A denote the particulate matter PM2.5 that a given individual is exposed to. The Environmental Protection Agency sets the National Ambient Air Quality Standards for Particulate Matter to "protect millions of Americans from harmful and costly health impacts, such as heart attacks and premature death" Environmental Protection Agency (2023). Current standards set the PM2.5 limit to 9 micrograms per cubic meter. Thus, any policy-relevant causal effect would have to take into account this standard. For instance, one can be interested in the effect of reducing PM2.5 on health outcomes by 10% for geographical areas which are non-compliant with the standard at time t, i.e., one may be interested in estimating $\theta^* = E[Y(d) \mid \overline{A} \in \overline{B}]$, where $\mathcal{B}_t = \{A_t : A_t > b\}$, $d(a_t, h_t) = \delta \times a_t$ with b = 9, $\delta = 0.9$, and Y is a health outcome of interest, for example myocardial infarction.

2.2 Causal Identification

Causal identification of the Local LMTP parameter θ^* can be achieved under the following two assumptions:

A1 (Positivity). For all $t \in \{1, \ldots, \tau\}$, if $(a_t, h_t) \in \text{supp}\{A_t, H_t \mid A_t \in \mathcal{B}_t\}$ then $(d(a_t, h_t), h_t) \in \text{supp}\{A_t, H_t\}$.

A2 (Strong sequential randomization). For all $t \in \{1, \ldots, \tau\}$, $U_{A,t} \perp (\underline{U}_{L,t+1}, \underline{U}_{A,t+1})|H_t$.

Positivity assumptions such as A1 are a standard requirement for many causal parameters (Petersen et al., 2012; van der Laan and Rose, 2011). Assumption A1 is weaker than that required for population-averaged LMTP outcomes, as here positivity is only required conditional on \mathcal{B} . This has important implications in applications, as different choices of the set \mathcal{B} can lead to assumptions with varying degrees of plausibility (e.g., see Example 2 below for the ATT). The strong sequential randomization assumption requires that A_t , L_{t+1} , and A_{t+1} are unconfounded (that is, all common causes of A_t , L_{t+1} , and A_{t+1} are included in H_t). This is a stronger assumption that what is required to identify dynamic treatment rules that do not depend on the natural value of treatment. See Díaz et al. (2023) for an in-depth discussion of these identification assumptions, and Richardson and Robins (2013); Young et al. (2014) for equivalent assumptions in alternative causal models. Below we explain the identification assumptions for each of the running examples.

Example 2 (continued). Assumption A1 requires that for any covariate value l such that P(A = 1 | L = l) > 0, then it must hold that P(A = 0 | L = l) > 0. This is weaker than the positivity assumption required for identification of the ATE, which states that for any covariate value l such that P(L = l) > 0, then it must hold that 0 < P(A = 1 | L = l) < 1. Assumption A2 requires that $U_A \perp U_Y | L$; that is, there are no unmeasured common causes of A and Y.

Example 3 (continued). Assumption A1 requires that for all time points, for any history h_t such that $P(A_t = a^* | H_t = h_t, A_t > b) > 0$, then it must hold that $P(A_t = \delta a^* | H_t = h_t) > 0$. That is, at all time points if a geographical area has a positive probability of having a PM2.5 value of a^* falling above the threshold given its longitudinal history, then it must also have a positive probability of having a PM2.5 value of δa^* . Assumption A2 requires that at each t, the longitudinal history H_t includes all common causes of A_t (above threshold PM2.5 at time t) and (A_s, L_s) , s > t (above threshold PM2.5 after time t and covariates and outcome after time t).

Under these assumptions, we establish an identification result for the Local LMTP parameter in terms of sequential regressions.

Theorem 1. Let $m_{\tau+1} = Y$, $A_{\tau+1} = 1$, $\mathcal{B}_{\tau+1} = \{1\}$. In a slight abuse of notation, let $A_t^d = d(A_t, H_t)$. Recursively define for $t = \tau, \ldots, 1$ the parameters

$$m_t: (a_t, h_t) \mapsto \mathbb{E}\left[m_{t+1}(A_{t+1}^d, H_{t+1}) | A_t = a_t, H_t = h_t, \underline{A}_{t+1} \in \underline{\mathcal{B}}_{t+1}\right],$$

and let $\theta_t = \mathbb{E}\left[m_t(A_t^d, L_t)|\underline{A}_t \in \underline{\mathcal{B}}_t\right]$. Under A1 and A2, the Local LMTP parameter is identified as $\theta_1 = \mathbb{E}\left[m_1(A_1^d, L_1)|\overline{A} \in \overline{\mathcal{B}}\right]$. For convenience, we will write $\theta = \theta_1$. While the parameters θ_t for t > 1 are not immediately relevant in terms of defining the main causal effect, they play an important role later in estimation. In the cross-sectional setting the identification result simplifies greatly, as shown in the running example for the ATT.

Example 2 (continued). Recall the counterfactual causal parameter is given by $\theta^* = \mathsf{E}[Y(0) | A = 1]$. The identification result implies this parameter is identified by $\theta = \mathsf{E}[\mathsf{E}[Y | W, A = 0] | A = 1]$.

Example 3 (continued). First, define recursively the parameters

$$m_t: (a_t, h_t) \mapsto \mathsf{E}[m_{t+1}(\delta A_{t+1}, H_{t+1})|A_t = a_t, H_t = h_t, \underline{A}_{t+1} \in \mathcal{B}_{t+1}],$$

and recalling that $m_{\tau+1} = Y$. The identification result implies that $\theta = \mathsf{E}[m_1(\delta A_1, L_1) \mid A_1 > b]$.

While we will discuss estimation in depth later, to gain intuition about the identification result for longitudinal settings note that it implies a simple estimation strategy based on sequential regressions. In particular, m_{τ} can be estimated as a regression of Y on A_{τ} and H_{τ} . This regression can then be used to obtain predictions under the hypothetical modified treatment policy $A_{\tau}^{d} = d(A_{\tau}, H_{\tau})$, which can in turn be regressed on $A_{\tau-1}, H_{\tau-1}$ among individuals with $\bar{A}_{\tau} \in \mathcal{B}_{\tau}$ to yield an estimate of $m_{\tau-1}$. This procedure can be iterated to t = 1, which is averaged over individuals with $\bar{A} \in \mathcal{B}$ to yield an point estimate of θ .

The parameter θ also admits an inverse probability weighted representation for a subset of MTPs that are sufficiently smooth, as stated in the following assumption.

A3 (Piecewise smooth invertibility for continuous exposures (Díaz et al., 2023)). For all $t \in \{1, ..., \tau\}$, and for all h_t , assume that the support of A_t conditional on $H_t = h_t$ is partitionable into subintervals $\mathcal{I}_{t,j}$: $j = 1, ..., J_t(h_t)$ such that $d(a_t, h_t)$ is equal to some $d_j(a_t, h_t)$ and $d_j(\cdot, h_t)$ has an inverse function that is differentiable with respect to a_t .

Next, define

$$g_{t,\bar{\mathcal{B}}}^{d}(a,h) = \int_{\mathcal{B}_{t}^{d}(a,h)} g_{t}(a',h) d\nu(a'),$$

where $\mathcal{B}_t^d(a,h) = \{a' \in \mathcal{B}_t : a = d(a',h)\}$ and letting $g_{t,\emptyset}^d(a,h) = 0$. Next, define the density ratio at time t as

$$r_t(a_t, h_t) = \frac{g_{t,\bar{\mathcal{B}}}^d(a_t, h_t)}{g_t(a_t, h_t)}.$$

Finally, let $\alpha_t(\bar{a}_t, \bar{h}_t) = \prod_{k=1}^t r_k(a_k, h_k)$ be the cumulative density ratios up to time t.

Theorem 2. Under A1, A2, and A3 the Local LMTP parameter is identified as

$$\theta = \frac{1}{P(\bar{A} \in \bar{\mathcal{B}})} \mathsf{E} \left[\alpha_{\tau}(\bar{A}_{\tau}, \bar{H}_{\tau}) Y \right].$$

This alternative identification result implies a simple inverse probability weighting estimator of θ by taking the mean of the observed outcomes weighted by the estimated cumulative density ratios α_{τ} . In the cross-sectional case for the ATT, the weighting simplifies as seen below. **Example 2** (continued). The inverse probability weighted representation of the ATT parameter is given by

$$\theta = \frac{1}{P(A=1)} \mathsf{E} \left[\frac{\mathbb{I}[A=0]P(A=1|L)}{P(A=0|L)} Y \right]$$

3 Non-parametric Efficiency Theory

In this section we investigate the semi-parametric properties of the Local LMTP parameter θ . These results draw on a long literature in semi-parametric efficiency theory; see Begun et al. (1983); Bickel et al. (1997); van der Vaart (1998), among many others. Kennedy (2016) provides a useful review of the relevant theory specialized to causal inference applications. Our main result is to derive the *efficient influence function* (EIF) of θ , a key object in the semi-parametric analysis of the parameter for several reasons. First, the variance of the EIF defines the efficiency bound for estimating θ in the non-parametric model \mathcal{M} (Bickel et al., 1997). Second, knowledge of the EIF is also crucial for developing non-parametric estimators of θ and deriving important properties such as their asymptotic sampling distribution. Developing estimators using the EIF flows from the following expansion, sometimes referred to as the *von-Mises expansion* (Robins et al., 2009; von Mises, 1947): for any $P \in \mathcal{M}$,

$$\theta(P) = \theta(P_0) - \mathsf{E}_{P_0}\{\mathsf{D}(O; P)\} + R(P, P_0), \tag{1}$$

where $R(P, P_0)$ is a second-order term of products of differences between functionals of P and P_0 . Plugging in an estimate \hat{P} for P_0 , and assuming that the second order estimation error $R(\hat{P}, P_0)$ is small enough, one then obtains an approximation to the bias of a plug-in estimator $\theta(\hat{P})$, namely $\theta(\hat{P}) - \theta(P_0) \approx -E_{P_0} \{D(O; \hat{P})\}$. A biased corrected estimator may then be constructed by subtracting an estimate of this bias from the plug-in estimator (i.e., the so-called one-step estimator) (Pfanzagl and Wefelmeyer, 1985; Emery et al., 2000), or constructing an estimator \hat{P} such that this bias converges to zero (i.e., targeted minimum loss-based estimation) (van der Laan and Rose, 2011).

The EIF and von-Mises expansion can also be used to derive important properties of nonparametric estimators. The form of the EIF may imply multiple-robustness properties in which only combinations of the nuisance parameters need to be estimated consistently for the estimator of θ to be consistent. In addition, careful analysis of the form of the remainder term $R(P, P_0)$ can reveal additional useful properties of non-parametric estimators of θ . For common parameters of interest such as the Average Treatment Effect, the remainder term has a product structure that implies nuisance parameters can be estimated at slow $n^{-1/4}$ rates, allowing for the use of flexible machine-learning algorithms in estimation.

Theorem 3 (Efficient influence function). Assume that d does not depend on P_0 . Let

$$\phi_t : z \mapsto \sum_{s=t}^{\tau} \left(\prod_{k=t}^s r_k(a_k, h_k) \right) \mathbb{1}[\underline{a}_{t+1} \in \underline{\mathcal{B}}_{t+1}] \left\{ m_{s+1}(a_{s+1}^d, h_{s+1}) - m_s(a_s, h_s) \right\} + \mathbb{1}[\underline{a}_t \in \underline{\mathcal{B}}_t] m_t(a_t^d, h_t).$$

The parameter θ_1 is pathwise differentiable and its EIF is given by

$$\mathsf{D}(Z;P) = \frac{1}{P(\bar{A} \in \bar{\mathcal{B}})} \left[\phi_1(Z) - \mathbb{1}[\bar{A} \in \bar{\mathcal{B}}] \theta_1 \right].$$

The EIF of the Local LMTP parameter simplifies to the EIF of the LMTP parameter when \bar{B} is set to the space of all possible longitudinal treatments. The non-parametric efficiency bound for estimators of θ is defined as the variance of the EIF: $E_{P_0}[D(Z; P_0)^2]$.

Next, we apply Theorem 3 to derive the EIFs of the parameters in the running examples.

Example 1 (continued). The efficient influence function for the ATT parameter θ is

$$\mathsf{D}(Z;\mathsf{P}) = \frac{\mathbb{1}\{A=0\}}{\mathsf{P}(A=1)} \frac{\mathsf{g}(1\mid W)}{\mathsf{g}(0\mid W)} \{Y - \mathsf{m}(A, W)\} + \frac{\mathbb{1}\{A=1\}}{\mathsf{P}(A=1)} \{\mathsf{m}(0, W) - \theta(\mathsf{P})\},\$$

which coincides with the efficient influence function given by Hubbard et al. (2011) for the ATT.

Example 2 (continued). The efficient influence function is given by

$$D(Z; \mathsf{P}) = \frac{1}{P(\bar{A} \in \bar{\mathcal{B}})} \left[\sum_{t=1}^{\tau} \left(\prod_{s=1}^{t} \frac{\mathbbm{1}\{A_s > \delta b\}}{\mathsf{P}(A_s > b)} \frac{\mathsf{g}_s(\delta^{-1}A_s \mid H_s)}{\mathsf{g}(A \mid W)} \right) \{m_{t+1}(\delta A_{t+1}, H_{t+1}) - \mathsf{m}_s(A_t, H_t)\} + \mathbbm{1}\{\bar{A} \in \bar{\mathcal{B}}\}\{\mathsf{m}_1(\delta A, H_1) - \theta(\mathsf{P})\} \right].$$

To complete the analysis, we give the form of the remainder term $R(P, P_0)$ of the von-Mises expansion (1). For all $t \in \{1, ..., \tau\}$ let r'_t , m'_t be the density ratios and outcome regressions corresponding to P and r_t , m_t the equivalent under P_0 . Let

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

Then using results from Appendix A.3, the remainder term $R(P, P_0)$ is given by

$$R(P, P_0) = \sum_{t=1}^{\tau} \mathbb{E}_{P_0} \left[C'_{1,t} \mathbb{1}[\underline{A}_{t+1} \in \underline{\mathcal{B}}_{t+1}] \times \left\{ r'_t(A_t, H_t) - r_t(A_t, H_t) \right\} \left\{ m'_t(A_t, H_t) - m_t(A_t, H_t) \right\} | A_t = a_t, H_t = H_t, \underline{A}_{t+1} \in \underline{\mathcal{B}}_{t+1} \right]$$

The form of the second-order remainder term suggests the possibility of forming multiply-robust estimators of θ .

4 Estimation

As foreshadowed in the previous section, the von-Mises expansion (1) suggests strategies for estimating θ based on forming an estimate \hat{P} of P_0 that ensures the bias term $\mathsf{E}_{P_0}[D(O; P)]$ and the second-order remainder term $R(\hat{P}, P_0)$ go to zero. Fortunately, estimating the entirety of P_0 is not necessary: it suffices to estimate only the parts of P_0 relevant to θ and its EIF D. In our case, the relevant parts are the regression parameters m_t and the cumulative ratios α_t . Collectively we refer to these nuisance parameters as $\eta = (m_t, \alpha_t : t \in \{1, \ldots, \tau\})$. We first discuss strategies for estimating the nuisance parameters m_t and α_t , and then show how they can be used to form estimators of θ .

The parameter $m_t(a_t, h_t)$ can be estimated using flexible regression techniques from the machine learning literature. However, estimation of the cumulative ratio α_t is more challenging. A first option is to estimate directly the densities $g_t(a_t \mid h_t)$ and plug in that estimate into the definition of α_t . While that option may be feasible for categorical A_t and for few time points, implementing for continuous or multivariate A_t would require estimation of conditional densities and computation of numerical integrals which may be challenging and computationally intensive. As an alternative, we estimate the cumulated density ratios directly via empirical loss minimization. Our approach is based on the theory of Riesz representers. Note that $\theta_t(Z) = \mathbb{E}[m_t(A_t^d, L_t) \mid \underline{A}_t \in \underline{\mathcal{B}}_t]$ is a linear functional of m_t . Therefore, by the Riesz representation theorem, there exists a squareintegrable function α_t such that, for all f with $\mathbb{E}[f(A_t, H_t)^2] < \infty$,

$$\mathbb{E}[f(A_t^d, H_t)] = \mathbb{E}[\alpha_t(A_t, H_t)f(A_t, H_t)].$$
⁽²⁾

The function α_t is referred to as a "Riesz representer". For the Local LMTP parameter, the Riesz representer is precisely the cumulative product of density ratios α_t , which can be written as the solution to an optimization problem over a candidate space \mathcal{A} :

$$\begin{split} \alpha_t &= \operatorname*{argmin}_{\alpha^* \in \mathcal{A}} \mathbb{E} \left[(\alpha^* (A_t, H_t) - \alpha_t (A_t, H_t))^2 \right] \\ &= \operatorname*{argmin}_{\alpha^* \in \mathcal{A}} \mathbb{E} \left[\alpha^* (A_t, H_t)^2 - 2\alpha^* (A_t, H_t) \alpha_t (A_t, H_t) + \alpha_t (A_t, H_t)^2 \right] \\ &= \operatorname*{argmin}_{\alpha^* \in \mathcal{A}} \mathbb{E} \left[\alpha^* (A_t, H_t)^2 - 2\alpha^* (A_t, H_t) \alpha_t (A_t, H_t) \right] \\ &= \operatorname*{argmin}_{\alpha^* \in \mathcal{A}} \mathbb{E} \left[\alpha^* (A_t, H_t)^2 - 2\alpha^* (A_t^d, H_t) \right] \end{split}$$
by (2)

To estimate α_t we can solve the corresponding empirical minimization problem:

$$\hat{\alpha}_t = \operatorname*{argmin}_{\alpha^* \in \mathcal{A}} \mathbb{E}_n \left[\alpha^* (A_t, H_t)^2 - 2\alpha^* (A_t^d, H_t) \right].$$
(3)

Crucially, estimating α_t in this manner avoids cumulating inverse probabilies (or densities) as is necessary for the the plug-in estimator. The choice of candidate space \mathcal{A} implies different options for practically solving the optimization problem. Methods based on flexible splines, random forests, or neural networks allow for rich candidate spaces \mathcal{A} Following (Chernozhukov et al., 2022a), we use a deep neural network to estimate α_t . Contrary to (Chernozhukov et al., 2022a) we chose to use cross-fitting when estimating α_t . Without cross-fitting, technical conditions are required to control the complexity of the estimators, such as Donsker (van der Laan and Rose, 2011) or critical radius assumptions (Chernozhukov et al., 2021; Wainwright, 2019). Cross-fitting is a commonly used method that obviates the need for such assumptions (Zheng and van der Laan, 2011; Chernozhukov et al., 2018). Therefore, with cross-fitting our method can be applied with any empirical loss minimization method used to estimate α_t .

With strategies in hand for estimating m_t and α_t , we now turn to forming estimators of θ .

4.1 Substitution and Inverse probability weighted estimators

A substitution estimator can be formed via the following algorithm:

- 1. Initialize $\hat{m}_{\tau+1,i}(A^{d}_{\tau+1,i}, H_{\tau+1,i}) = Y_i$.
- 2. For $t = \tau, ..., 1$:
 - Using any regression method, regress $\hat{m}_{t+1}(A_{t+1,i}^d, H_{t+1}, i)$ on (A_t, H_t) for all $i \in \{1, \ldots, n\}$ such that $\bar{A}_{\tau+1,i} \in \bar{\mathcal{B}}_{t+1}$.
 - Use the regression model to form predictions $\hat{m}_{t,i}(A_{t,i}^d, H_{t,i})$ for all $i \in \{1, \ldots, n\}$.
- 3. Form the substitution estimator as

$$\hat{\theta}_{sub} = \frac{1}{\sum_{i=1}^{n} \mathbb{1}[\bar{A}_i \in \bar{\mathcal{B}}]} \sum_{i=1}^{n} \mathbb{1}[\bar{A}_i \in \bar{\mathcal{B}}] \hat{m}_1(A_{1,i}^d, L_{1,i}).$$

Similarly, if we have an estimator $\hat{\alpha}_{\tau}$ of α_{τ} , for example via (3), an IPW estimator of θ_1 is given by

$$\hat{\theta}_{ipw} = \left(\frac{1}{n}\sum_{i=1}^{n}\mathbb{1}[\bar{A}_i \in \bar{\mathcal{B}}]\right)^{-1} \frac{1}{n}\sum_{i=1}^{n}\hat{\alpha}_{\tau}(A_{\tau,i}, H_{\tau,i})Y_i.$$

The substitution estimator will be consistent if all of the regressions \hat{m}_t are consistent. The IPW estimator is consistent if $\hat{\alpha}_{\tau}$ is estimated consistently.

4.2 Targeted minimum loss-based estimator

Targeted Minimum Loss-Based Estimation (TMLE) is a framework for constructing asymptotically efficient non-parametric plug-in estimators (van der Laan and Rose, 2011). The core idea is to solve the EIF estimating equation by carefully fluctuating initial estimates of the parts of P_0 relevant to the parameter of interest. The TMLE estimator for Local LMTPs is similar to the one given in Díaz et al. (2023) for LMTPs with modifications to incorporate the conditional structure of the Local LMTP parameter. Following their work, we use cross-fitting in order to avoid technical conditions on the complexity of the nuisance estimators (Klaassen, 1987; Zheng and van der Laan, 2011). Randomly partition the indexes $\{1, \ldots, n\}$ into J validation sets $\mathcal{V}_1, \ldots, \mathcal{V}_J$. For each $j \in \mathcal{J}$, the training sample is given by $\mathcal{J}_j = \{1, \ldots, n\} \setminus \mathcal{V}_j$. Let j(i) be the validation set containing index i. The goal of the TMLE algorithm is to form a set of updated estimates $\tilde{m}_{t,j(i)}$ such that the empirical EIF estimating equation is solved. Inspection of the form of the EIF shows that this will be the case when

$$\mathsf{E}_{n}\left[\sum_{t=1}^{\tau}\hat{\alpha}_{t}\mathbb{1}[\underline{A}_{t+1}\in\underline{\mathcal{B}}_{t+1}]\left\{\tilde{m}_{t+1}(A_{t+1}^{d},H_{t+1})-\tilde{m}_{t}(A_{t},H_{t})\right\}\right]\approx0.$$
(4)

TMLE ensures this is the case by fluctuating an initial estimate \hat{m}_t via a carefully chosen parametric submodel indexed by a parameter $\epsilon \in \mathbb{R}$ and loss function to form updated estimates \tilde{m}_t . The parametric submodel is chosen such that the gradient of the loss with respect to ϵ equals (4). When the parameter ϵ is estimated by minimizing the loss function, this ensures that at the minimizer the empirical mean of the gradient is approximately zero. As such, the updated estimates approximately solve the empirical EIF (4). For Local LMTPs, the fluctuation model is chosen to be a weighted generalized linear model with canonical link, an intercept parameter ϵ , and offset set to the initial estimates \hat{m}_t . The TMLE algorithm is as follows:

- 1. Initialize $\tilde{\eta} = \hat{\eta}$ and $\tilde{m}_{\tau+1,j(i)}(A^d_{\tau+1,i}, H_{\tau+1,i}) = Y_i$.
- 2. For $s = 1, \ldots, \tau$ compute weights

$$\omega_{s,i} = \mathbb{1}[\underline{A}_{t+1} \in \underline{\mathcal{B}}_{t+1}]\hat{\alpha}_{s,i}.$$

- 3. For $t = \tau, ..., 1$:
 - Find the maximum likelihood estimate $\hat{\epsilon}$ of ϵ under the model

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

with weights

• Update \tilde{m} as

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

4. Form the TMLE estimate as

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

The TMLE algorithm for Local LMTP parameters is similar to the TMLE for LMTPs Díaz et al. (2023). The key difference between the algorithms is in Step 2, where for Local LMTPs the weights are multiplied by an indicator of the future treatment assignment lying in the conditioning set. The statistical properties of TMLE flow from the fact that after this iterative procedure, (4) is satisfied. Using this, it is possible to show that the TMLE estimator is asymptotically normal and efficient.

Theorem 4. Assume that

$$\|\hat{\alpha}_1 - \alpha_1\| \|\tilde{m}_1 - m_1\| + \sum_{t=2}^{\tau} \left\| \frac{\hat{\alpha}_t}{\hat{\alpha}_{t-1}} - \frac{\alpha_t}{\alpha_{t-1}} \right\| \|\tilde{m}_t - m_t\| = o_P(n^{-1/2}).$$

Assume there exists some $c < \infty$ such that $P(\alpha_t < c) = 1$ and $P(\hat{\alpha}_t < c) = 1$. Then

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

where $\sigma^2 = \operatorname{Var}_{P_0}(\mathsf{D}(Z; P_0)).$

Satisfying the assumptions of Theorem 4 requires that all the nuisance parameters are estimated at sufficiently fast rates. Consistency of $\hat{\theta}_{tmle}$, however, can be achieved even if some of the nuisance parameters are not estimated consistently, as shown in the following theorem.

Lemma 1 (τ + 1 multiply robust consistency of TMLE). Assume there exists a $k \in \{1, ..., \tau - 1\}$ such that $\|\tilde{m}_t - m_t\| = o_p(1)$ for all t > k, $\|\hat{\alpha}_1 - \alpha_1\| = o_p(1)$, and $\|\frac{\hat{\alpha}_t}{\hat{\alpha}_{t-1}} - \frac{\alpha_t}{\alpha_{t-1}}\| = o_p(1)$ for all $1 < t \le k$. Then $\hat{\theta}_{tmle} - \theta = o_p(1)$.

It is also theoretically possible to construct estimator with stronger multiple robustness properties, similar to the sequentially double robust (SDR) estimator proposed for LMTPs (Díaz et al., 2023). This estimator would require cumulative density ratios of the form $\prod_{t=s}^{k} r_s$ for all combinations of $1 \le t < \tau$ and $t < s \le \tau$. For the LMTP SDR estimator, these ratios were estimated by plugging in estimates of r_t for $t \in \{1, \ldots, \tau\}$. It would also be possible to apply the Riesz Representer approach to estimate the cumulative ratios directly, although in practice it is computationally prohibitive for large τ as it requires solving (3) a total of $(\tau)(\tau - 1)/2$ times. In comparison, for the TMLE algorithm requires solving (3) only τ times.

5 Simulation studies

We investigate the finite-sample performance of the proposed estimators for the Local LMTP parameter through three simulation studies, each probing the estimators from a different angle. Simulation study 1 looks at the robustness of the TMLE estimator inconsistent nuisance parameter estimation. Simulation study 2 compares how the Riesz Representer estimation strategy compares to estimating individual density ratios. Simulation study 3 investigates how Riesz Representer approach fares under extreme practical positivity violations.

5.1 Simulation study 1

In the first simulation study we investigate the performance of the TMLE estimator for a Local LMTP parameter defined for a categorical exposure. The setup is adapted from the simulation study presented in Díaz et al. (2023). The data generating process is given by

$$\begin{split} L_1 &\sim \text{Categorical}(0.5, 0.25, 0.25),\\ A_1 | L_1 &\sim \text{Binomial}(5, \text{logit}^{-1}(-0.3L_1)), \end{split}$$

$$L_t | (\bar{A}_{t-1}, \bar{L}_{t-1}) \sim \text{Bernoulli}(\text{logit}^{-1}(-0.3L_{t-1} + 0.5A_{t-1})) \text{ for } t \in \{2, 3, 4\},$$

$$A_t | (\bar{A}_{t-1}, \bar{L}_t) \sim \text{Binomial}(5, \text{logit}^{-1}(-2.5 + A_{t-1} + 0.5L_t)) \text{ for } t \in \{2, 3, 4\},$$

$$Y | (\bar{A}_4, \bar{L}_4) \sim \text{Bernoulli}(\text{logit}^{-1}(-1 + 0.5A_4 - L_4)).$$

The LMTP is defined as

$$d(a_t, h_t) = \begin{cases} a_t - 1 & \text{if } a_t \ge 1, \\ a_t & \text{if } a_t < 1. \end{cases}$$

We define two Local LMTP parameters conditional on the final treatment assignment. That is, we set $\mathcal{B}^a = \{a\}$ for $a \in \{0, 1, \dots, 5\}$. 100 datasets for each sample size $N = \{250, 500, 1000\}$ were created by independently sampling from the data generating process. The nuisance estimators were estimated differently in four scenarios:

- 1. all nuisance parameters estimated consistently.
- 2. m_t estimated consistently for t > 2 and inconsistently otherwise; α_t consistently for $t \le 2$ and inconsistently otherwise.
- 3. m_t estimated consistently for t < 4 and inconsistently for t = 4; α_t estimated consistently for t = 4 and inconsistently otherwise.
- 4. all nuisance parameters estimated inconsistently.

Consistent estimation of m_t was achieved by Super Learning with SL.mean and SL.glm learners, where the generalized linear model was correctly specified. For α_t , we assume that the neural network algorithm was consistent for this DGP. For the inconsistent cases, only SL.mean was used for estimating m_t and α_t was estimated using only an intercept term and no covariates or treatment variables.

A subset of the results are shown in Table 1; complete results are available in the Appendix as Table 4. The TMLE estimator achieved near optimal coverage and near zero mean error and mean absolute error for nearly all combinations of sample size and conditioning sets in scenario (1), where all nuisance parameters are estimated consistently. In scenario (2), IPW appears to be inconsistent, as expected because α_4 is estimated inconsistently. As expected based on the $\tau + 1$ robustness result (Theorem 1), TMLE has good performance in terms of error metrics. Surprisingly, the substitution estimator also has good performance in this example. In scenario (3), TMLE also has good performance both in terms of empirical coverage and error metrics, which is unexpected as the scenario does not fulfill the requirements of Theorem 1. We hypothesize that estimating α_4 correctly, which is the cumulative densities from t = 1 to t = 4, allows TMLE to correct for the inconsistent estimation of m_4 . Finally, in scenario (4) all estimators have poor performance due to the inconsistent estimation of all nuisance parameters.

5.2 Simulation study 2

The second simulation study investigates how the Riesz representation approach for estimating the cumulated densities α_t compares to estimating each ratio r_t and then cumulating them to form

			95% Coverage	Ν	$MAE \times 100$		$ME \times 100$		
\mathcal{B}^{a}	Scenario	N	TMLE	IPW	Sub	TMLE	IPW	Sub	TMLE
$\{0\}$	1	250	96.0%	7.53	6.05	6.30	-3.26	-4.71	-0.77
		500	98.0%	4.55	3.31	3.83	-2.53	-2.26	-1.06
		1000	95.0%	3.14	2.65	2.72	-1.97	-1.39	-1.15
	2	250	56.0%	23.48	6.10	6.20	-23.48	-4.82	-4.92
		500	74.0%	22.84	3.20	3.19	-22.84	-2.29	-2.21
		1000	69.0%	23.13	2.56	2.54	-23.13	-1.36	-1.39
	3	250	96.0%	7.54	23.50	6.35	-3.27	-23.50	-0.90
		500	98.0%	4.48	22.86	4.00	-2.38	-22.86	-1.09
		1000	94.0%	3.10	23.19	2.79	-1.97	-23.19	-1.26
	4	250	0.0%	23.54	23.50	23.69	-23.54	-23.50	-23.69
		500	0.0%	22.83	22.86	22.92	-22.83	-22.86	-22.92
		1000	0.0%	23.13	23.19	23.24	-23.13	-23.19	-23.24

Table 1: Results of Simulation Study 1 showing empirical coverage of the 95% confidence intervals, Mean Absolute Error (MAE), and Mean Error (ME) for the inverse probability weighted estimator (IPW), substitution estimator (Sub), and Targeted minimum loss-based estimator (TMLE).

estimates of α_t . As such, we consider a relatively simple longitudinal data generating process and modified treatment policy because the main variable of interest in this simulation study is the number of time points. The data generating process is given by

$$L_1 \sim \text{Uniform}(0, 1),$$

$$A_1 | L_1 \sim \text{Bernoulli}(0.5),$$

$$L_t | (\bar{A}_{t-1}, \bar{L}_{t-1}) \sim \text{Normal}(-0.5L_{t-1} + A_{t-1}, 0.5^2) \text{ for } t \in \{2, \dots, \tau\},$$

$$A_t | L_t \sim \text{Bernoulli}(\text{logit}^{-1}(0.5 + 0.1L_t)),$$

$$Y | A_\tau, L_\tau \sim \text{Normal}(A_\tau + L_\tau, \sigma^2).$$

The modified treatment policy was defined as $d(a_t, h_t) = 1$. The LMTP parameter was not conditioned on treatment (or, equivalently, $\mathcal{B}_t = \{0, 1\}$ for all t). The outcome regression models were estimated using properly specified generalized linear models. The estimators of r_t were based on correctly specified logistic regressions which were then cumulated to form estimates of α_t (we refer to this as the "plugin" method, as estimates \hat{r}_t are plugged in to the definition of α_t).

Simulation results are shown in Table 3 and Figure 1. For sample sizes N = 500 and N = 1000, the RR estimator maintains near optimal empirical coverage for all τ , while the performance of the Plug-in estimator suffers as τ increases. A similar pattern is seen for the mean absolute error. The instability of the Plug-in estimator can be seen via the standard deviation of the estimated cumulated weights $\hat{\alpha}_{\tau}$ which are larger for the Plug-in estimator than the RR estimator.

5.3 Simulation study 3

The third simulation study investigates how the proposed estimation approach based on Riesz Representers performs under practical positivity violations. The data generating process for a

		95% Coverage		MAE	$E \times 100$	$\operatorname{sd}(\hat{\alpha}_{\tau})$		
N	au	RR	Plugin	RR	Plugin	RR	Plugin	
100	2	92%	95%	1.90	1.48	1.82	1.56	
	4	79%	98%	3.18	2.36	2.79	2.99	
	6	84%	86%	3.00	4.22	2.97	5.32	
	8	78%	81%	3.20	5.05	2.67	8.37	
	10	81%	72%	3.14	5.40	2.39	15.35	
	12	80%	70%	2.99	5.11	2.06	18.62	
500	2	97%	96%	0.66	0.66	1.79	1.49	
	4	92%	93%	1.39	1.06	3.17	2.69	
	6	90%	95%	1.69	1.50	4.12	4.52	
	8	94%	94%	1.85	2.64	4.59	7.47	
	10	92%	85%	1.84	4.11	4.50	11.54	
	12	94%	75%	1.52	4.61	4.47	18.41	
1000	2	97%	97%	0.45	0.40	1.76	1.48	
	4	96%	98%	0.76	0.57	3.17	2.64	
	6	93%	98%	1.17	1.02	4.35	4.38	
	8	94%	96%	1.25	1.58	4.94	7.10	
	10	89%	92%	1.55	2.68	5.20	11.50	
	12	93%	78%	1.43	4.19	5.15	17.42	

Table 2: Results of Simulation Study 2 comparing the performance of the TMLE estimator with Riesz representers α_t estimated using empirical loss minimization (RR) and via plug-in estimation (Plug-in). The estimators are compared by their empirical 95% confidence interval coverage, mean absolute error (MAE), and the mean standard deviation of the Riesz representers at time τ (sd($\hat{\alpha}_{\tau}$)).



Figure 1: Results of Simulation Study 2.

		95% (Coverage	$MAE \times 100$		sc	$l(\hat{\alpha}_1)$
β	N	RR	Plug-in	RR	Plug-in	RR	Plug-in
0	500	97%	97%	2.64	2.62	0.99	1.00
	1000	98%	97%	1.89	1.89	0.98	1.00
	5000	95%	95%	0.94	0.94	0.97	1.00
1	500	99%	98%	3.24	2.90	1.34	1.16
	1000	96%	96%	2.37	2.27	1.26	1.14
	5000	96%	95%	1.02	0.99	1.24	1.14
2	500	85%	92%	6.69	3.26	2.67	1.22
	1000	93%	96%	3.67	2.41	2.30	1.22
	5000	94%	94%	1.36	1.05	2.00	1.21
3	500	77%	94%	12.70	3.31	4.30	1.17
	1000	62%	93%	10.59	2.58	3.89	1.16
	5000	59%	93%	4.23	1.10	3.25	1.16

Table 3: Results of Simulation Study 3 comparing the performance of the TMLE estimator with Riesz representers α_t estimated using empirical loss minimization (RR) and via plug-in estimation (Plug-in). The estimators are compared by their empirical 95% confidence interval coverage, mean absolute error (MAE), and the mean standard deviation of the Riesz representers (sd($\hat{\alpha}_1$)).

generic observation (L, A, Y) is given by

 $L \sim \text{Normal}(0, 1),$ $A \sim \text{Binomial} \left(\text{logit}^{-1}(\beta L) \right),$ $Y \sim \text{Normal}(0.5L + A, 0.5^2).$

The parameter β controls the degree of covariate overlap between treatment (A = 1) and control (A = 0) groups. For example, when $\beta = 0$ then A is randomly assigned with probability 0.5, representing an experimental design. As β increases, the probability of treatment becomes increasingly associated with higher covariate values. The Local LMTP parameter is defined setting $\mathcal{B} = \{0, 1\}$, such that it reduces to the LMTP parameter.

100 simulation datasets were drawn from the data generating process for sample sizes $N \in \{500, 1000, 5000\}$. The TMLE estimator was applied twice to each simulation dataset, using alternatively the plugin and Riesz Representer strategies for estimating α_t . Nuisance parameters were estimated using well-specified generalized linear models. Thus, we note that the plugin TMLE has an advantage in that the density ratios r_t are estimated using well-specified models, while the Riesz Representer approach estimates the cumulated ratios directly using flexible neural networks.

This simulation study illustrates that the computational strategy of estimating the Riesz Representers directly cannot overcome extreme positivity violations by itself. This should not be surprising, as we have not introduced any additional structure to handle extrapolating the density ratios and outcome regression into regions where there is little or no overlap between treatment and control groups. The better performance of the TMLE based on plugin estimation of the cumulated densities is because it is based on well-specified parametric regressions which allows it to more accurately "fill-in" information in the low-overlap regions.

A Appendix

A.1 **Proof of Theorem 1**

Proof The proof follows closely that of (Díaz et al., 2023, Theorem 1). Let $A_{t+1}(\bar{a}_t)$ and $L_{t+1}(\bar{a}_t)$ be the counterfactual variables at time t + 1 when the exposure is set to $\bar{A}_t = \bar{a}_t$. Fix a time point t arbitrarily. For s = 1, let

$$A_{t+s}^{d,\dagger}(\bar{a}_t, \bar{l}_t) = d\left\{\bar{a}_t, A_{t+s}(\bar{a}_t), \bar{l}_t, L_{t+s}(\bar{a}_t)\right\},\$$

and for s>1 define $A_{t+s}^{d,\dagger}$ as

$$A_{t+s}^{d,\dagger}(\bar{a}_t, \bar{l}_t) = d\{\bar{a}_t, A_{t+1}(\bar{a}_t), A_{t+2}(\bar{a}_t, A_{t+1}^{d,\dagger}(\bar{a}_t)), \dots, A_{t+s}(\bar{a}_t, A_{t+1}^{d,\dagger}(\bar{a}_t), \dots, A_{t+s-1}^{d,\dagger}(\bar{a}_t)), \\ \bar{l}_t, L_{t+1}(\bar{a}_t), L_{t+2}(\bar{a}_t, A_{t+1}^{d,\dagger}(\bar{a}_t)), \dots, L_{t+s}(\bar{a}_t, A_{t+1}^{d,\dagger}(\bar{a}_t), \dots, A_{t+s-1}^{d,\dagger}(\bar{a}_t))\}.$$

In words, $A_{t+s}^{d,\dagger}$ is the modified outcome at time t + s when the MTP was applied up to time t, and then from t + 1 to t + s - 1 the treatment is set to the natural value of treatment. Similarly let

$$L_{t+s}^{\dagger}(\bar{a}_{t}) = L_{t+s}(\bar{a}_{t}, A_{t+s}^{d,\dagger}(\bar{a}_{t}), \dots, A_{t+s}^{d,\dagger}(\bar{a}_{2}))$$

and

$$Z_t(\bar{a}_t, \bar{l}_t) = f_Y(\bar{a}_t, \underline{A}_{t+1}^{d,\dagger}(\bar{a}_t, \bar{l}_t), \bar{l}_t, \underline{L}_{t+1}^{\dagger}(\bar{a}_t), U_Y)$$

Begin by rewriting the conditional expectation:

$$\mathbb{E}\left[Y(\bar{A}^{d}) \mid \bar{A} \in \bar{\mathcal{B}}\right]$$

= $\mathbb{E}\left[\frac{1}{P(\bar{A} \in \bar{\mathcal{B}})}Y(\bar{A}^{d})\mathbb{1}[\bar{A} \in \bar{\mathcal{B}}]\right]$
= $\mathbb{E}\left[\frac{1}{P(A_{1} \in \mathcal{B}_{1})P(\underline{A}_{2} \in \underline{\mathcal{B}}_{2})}Y(\bar{A}^{d})\mathbb{1}[A_{1} \in \mathcal{B}_{1}]\mathbb{1}[\underline{A}_{2} \in \underline{\mathcal{B}}_{2}]\right].$

Next, apply the law of iterated expectations and simplify:

$$= \int_{\mathcal{A}_1,\mathcal{L}_1} \frac{1}{P(A_1 \in \mathcal{B}_1)P(\underline{A}_2 \in \underline{\mathcal{B}}_2)} \mathbb{E} \left[Z_1(a_1^d, l_1)\mathbb{1}[\underline{A}_2 \in \underline{\mathcal{B}}_2] \mid A_1 = a_1, L_1 = l_1 \right] \mathbb{1}[A_1 \in \mathcal{B}_1] dP(a_1, l_1)$$
$$= \int_{\mathcal{A}_1,\mathcal{L}_1} \frac{1}{P(\underline{A}_2 \in \underline{\mathcal{B}}_2)} \mathbb{E} \left[Z_1(a_1^d, l_1)\mathbb{1}[\underline{A}_2 \in \underline{\mathcal{B}}_2] \mid A_1 = a_1, L_1 = l_1 \right] dP(a_1, l_1 \mid a_1 \in \mathcal{B}_1).$$

By Lemma 1 (below) and Assumptions A1 and A2,

$$= \int_{\mathcal{A}_1,\mathcal{L}_1} \frac{1}{P(\underline{A}_2 \in \underline{\mathcal{B}}_2)} \mathbb{E} \left[Z_1(a_1^d, l_1) \mathbb{1}[\underline{A}_2 \in \underline{\mathcal{B}}_2] \mid A_1 = a_1^d, L_1 = l_1 \right] dP(a_1, l_1 \mid a_1 \in \mathcal{B}_1)$$
$$= \int_{\mathcal{A}_1,\mathcal{L}_1} \mathbb{E} \left[Z_1(a_1^d, l_1) \mid A_1 = a_1^d, L_1 = l_1, \underline{A}_2 \in \underline{\mathcal{B}}_2 \right] dP(a_1, l_1 \mid a_1 \in \mathcal{B}_1)$$

Continue by applying the same steps to the inner expectation:

$$= \int_{\bar{\mathcal{A}}_2, \bar{\mathcal{L}}_2} \mathbb{E}\left[Z_1(a_1^d, l_1) \mid A_2 = a_2, L_2 = l_2, A_1 = a_1^d, L_1 = l_1, \underline{A}_2 \in \underline{\mathcal{B}}_2\right] dP(a_2, l_2 \mid a_2 \in \mathcal{B}_2) dP(a_1, l_1 \mid a_1 \in \mathcal{B}_1)$$

In the event $A_{t+1} = a_{t+1}, H_{t+1} = h_{t+1}^d, \underline{A}_t \in \underline{\mathcal{B}}_t$ then $Z_t(\bar{a}_t^d, \bar{l}_t) = Z_{t+1}(\bar{a}_{t+1}^d, \bar{l}_{t+1})$. Therefore:

$$= \int_{\bar{\mathcal{A}}_2, \bar{\mathcal{L}}_2} \mathbb{E}\left[Z_2(a_2^d, l_2) \mid A_2 = a_2, L_2 = l_2, A_1 = a_1^d, L_1 = l_1, \underline{A}_2 \in \underline{\mathcal{B}}_2\right] dP(a_2, l_2 \mid a_2 \in \mathcal{B}_2) dP(a_1, l_1 \mid a_1 \in \mathcal{B}_1)$$

Continue applying the previous steps until arriving at

$$= \int_{\bar{\mathcal{A}}_{\tau},\bar{\mathcal{L}}_{\tau}} \mathbb{E}\left[Z_{\tau}(\bar{a}_{\tau}^{d},\bar{l}_{\tau}) \mid A_{\tau} = a_{\tau}^{d}, H_{\tau} = h_{\tau}^{d}, \underline{A}_{\tau+1} \in \underline{\mathcal{B}}_{\tau+1} \right] \prod_{k=1}^{\tau} dP(a_{k},l_{k} \mid a_{k-1}^{\tau}h_{k-1}^{d}, a_{k} \in \mathcal{B}_{k}),$$

Finally, by applying the definition of Z_t :

$$= \int_{\bar{\mathcal{A}}_{\tau},\bar{\mathcal{L}}_{\tau}} \mathbb{E}\left[Y \mid A_{\tau} = a_{\tau}^{d}, H_{\tau} = h_{\tau}^{d}, \underline{A}_{\tau+1} \in \underline{\mathcal{B}}_{\tau+1}\right] \prod_{k=1}^{\tau} dP(a_{k}, l_{k} \mid a_{k-1}^{\tau} h_{k-1}^{d}, a_{k} \in \mathcal{B}_{k}),$$
(5)

Next, recursively apply the following starting with $t = \tau$ to arrive at the stated result:

$$\int_{\mathcal{A}_t, \mathcal{L}_t} m_t(a_t^d, h_t^d) dP(a_t, l_t \mid a_{t-1}^d, h_{t-1}^d, \underline{a}_t \in \underline{\mathcal{B}}_t) = \mathbb{E}\left[m(A_t^d, H_t) \mid A_{t-1} = a_{t-1}^d, H_{t-1} = h_{t-1}^d, \underline{A}_t \in \underline{\mathcal{B}}_t\right]$$

$$= m_{t-1}(a_{t-1}^d, h_{t-1}^d).$$

Lemma 2. Given Assumption A2, for all t it follows that $Z_t(\overline{a}_t, \overline{l}_t) \mathbb{1}[\underline{A}_{t+1} \in \underline{\mathcal{B}}_{t+1}] \perp A_t \mid H_t$.

Proof Under the assumed structural causal model, $Z_t(\bar{a}_t, \bar{l}_t)\mathbb{1}[\underline{A}_{t+1} \in \underline{\mathcal{B}}_{t+1}]$ is a deterministic function of $(\underline{U}_{L,t+1}, \underline{U}_{A,t+1})$.

A.2 Proof of Theorem 2

Proof Write

$$\mathsf{E}[r_t(A_t, H_t)m_t(A_t, H_t) \mid \underline{A}_{t+1} \in \underline{\mathcal{B}}_{t+1}]$$

$$= \mathsf{E}\left[\frac{g_{t,\overline{\mathcal{B}}}^d(A_t, H_t)}{g_t(A_t, H_t)}m_t(A_t, H_t) \mid \underline{A}_{t+1} \in \underline{\mathcal{B}}_{t+1}\right]$$

$$\begin{split} &= \mathsf{E}\left[\frac{\int_{\mathcal{B}_{t}^{d}(A_{t},H_{t})}g_{t}(a',H_{t})d\nu(a')}{g_{t}(A_{t},H_{t})}m_{t}(A_{t},H_{t})\mid\underline{A}_{t+1}\in\underline{\mathcal{B}}_{t+1}\right]\\ &= \mathsf{E}\left[\frac{\int\mathbbm{1}[A_{t}=d(a',H_{t}),a'\in\mathcal{B}_{t}]g_{t}(a',H_{t})d\nu(a')}{g_{t}(A_{t},H_{t})}m_{t}(A_{t},H_{t})\mid\underline{A}_{t+1}\in\underline{\mathcal{B}}_{t+1}\right]\\ &= \mathsf{E}\left[\int\frac{\int\mathbbm{1}[a=d(a',H_{t}),a'\in\mathcal{B}_{t}]g_{t}(a',H_{t})d\nu(a')}{g_{t}(a,H_{t})}m_{t}(a,H_{t})g_{t}(a,H_{t})d\nu(a)\mid\underline{A}_{t+1}\in\underline{\mathcal{B}}_{t+1}\right]\\ &= \mathsf{E}\left[\int\int\mathbbm{1}[a=d(a',H_{t}),a'\in\mathcal{B}_{t}]g_{t}(a',H_{t})d\nu(a')m_{t}(d(a',H_{t}),H_{t})d\nu(a)\mid\underline{A}_{t+1}\in\underline{\mathcal{B}}_{t+1}\right]\\ &= \mathsf{E}\left[m_{t}(A_{t}^{d},H_{t})d\nu(a)\mid\underline{A}_{t+1}\in\underline{\mathcal{B}}_{t+1}\right]. \end{split}$$

Applying this recursively for $t = \tau, \ldots, 1$ to (5) yields the result.

A.3 First-order expansion

Let $\eta' = (r'_1, m'_1, \dots, r'_{\tau}, m'_{\tau})$. Let

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

Define the second order term for $t = \{0, \dots, \tau - 1\}$ as

$$\operatorname{Rem}_{t}(a_{t}, h_{t}; \eta') =$$

$$\sum_{s=t+1}^{\tau} \left(\mathbb{E}[C_{t,s}' \mathbb{1}[\underline{A}_{s+1} \in \underline{\mathcal{B}}_{s+1}] \times \left\{ r_{s}'(A_{s}, H_{s}) - r_{s}(A_{s}, H_{s}) \right\} \left\{ m_{s}'(A_{s}, H_{s}) - m_{s}(A_{s}, H_{s}) \right\} | A_{t} = a_{t}, H_{t} = H_{t}, \underline{A}_{s+1} \in \underline{\mathcal{B}}_{s+1}]$$

$$(6)$$

For t = 0 the remainder term is conditioned on $\overline{A} \in \overline{\mathcal{B}}$. When $t = \tau$, let $\operatorname{Rem}_t(a_t, h_t; \eta') = 0$. **Theorem 5.** For all $t \in \{1, \ldots, \tau\}$,

$$m_t(a_t, h_t) = \mathbb{E} \left[\phi_{t+1}(Z, \eta') \mid A_t = a_t, H_t = h_t, \underline{A}_{t+1} \in \underline{\mathcal{B}}_{t+1} \right] \\ + \operatorname{Rem}_t(a_t, h_t; \eta').$$

Proof The proof is similar to that of (Díaz et al., 2023, Lemma 1) with adjustments to handle the conditional structure of the Local LMTP parameter. First, note the identity

$$\mathbb{E}\left[m_s(A_s, H_s) \mid A_{s-1}, H_{s-1}, \underline{A}_s \in \underline{\mathcal{B}}_s\right] = \mathbb{E}\left[r_s(A_s, H_s)m_s(A_s^d, H_s) \mid A_{s-1}, H_{s-1}, \underline{A}_s \in \underline{\mathcal{B}}_s\right].$$

By this identity and the tower rule, the following recursive relationship holds:

$$\begin{split} & \mathbb{E}[m'_{s}(A^{d}_{s}, H_{s}) - m_{s}(A^{d}_{s}, H_{s}) \mid A_{s-1}, H_{s-1}, \underline{A}_{s} \in \underline{\mathcal{B}}_{s}] \\ &= -\mathbb{E}\left\{\mathbb{1}[\underline{A}_{s+1} \in \underline{\mathcal{B}}_{s+1}][m'_{s+1}(A^{d}_{s}, H_{s}) - m_{s}(A^{d}_{s}, H_{s})] \mid A_{s-1}, H_{s-1}, \underline{A}_{s} \in \underline{\mathcal{B}}_{s}\right\} \\ &+ \mathbb{E}\left\{\mathbb{1}[\underline{A}_{s+1} \in \underline{\mathcal{B}}_{s+1}]\left[r_{s}(A_{s}, H_{s}) - r'_{s}(A_{s}, H_{s})\right]\left[m'(A_{s}, H_{s}) - m_{s}(A_{s}, H_{s})\right] \mid A_{s-1}, H_{s-1}, \underline{A}_{s} \in \underline{\mathcal{B}}_{s}\right\} \\ &+ \mathbb{E}\left\{\mathbb{1}[\underline{A}_{s+1} \in \underline{\mathcal{B}}_{s+1}]r'_{s}(A_{s}, H_{s})\mathbb{E}\left[m'_{s+1}(A^{d}_{s+1}, H_{s+1}) - m_{s+1}(A^{d}_{s+1}, H_{s+1}) \mid A_{s-1}, H_{s-1}, \underline{A}_{s} \in \underline{\mathcal{B}}_{s}\right]\right\} \end{split}$$

Recursively applying the above relation from $s = t + 1, \ldots, \tau$ yields the stated result.

A.4 **Proof of Theorem 3**

Proof The proof follows that of (Díaz et al., 2023, Theorem 2), therefore we focus on the parts where the proof differs from theirs. The strategy of the proof is to first derive a putative EIF assuming the data are discrete, and then show that this is indeed the correct EIF in the general setting.

Let $\{P_{\epsilon} : \epsilon \in \mathbb{R}\}$ be a parametric submodel satisfying $P_0 = P$. The functional $\theta(P)$ is pathwise differentiable with EIF D(Z; P) if

$$\left. \frac{d}{d\epsilon} \Theta(P_{\epsilon}) \right|_{\epsilon=0} = \mathbb{E}_{P}[D(Z;P)s(Z)],\tag{7}$$

where the score s(Z) is given by

$$s(Z) = \frac{dP_{\epsilon}}{d\epsilon}\Big|_{\epsilon=0}.$$

We derive the EIF for discrete data using the functional Delta method (van der Laan and Rose, 2011), which states that for a substitution estimator $\hat{\Theta}$ that can be written as $\hat{\Theta}^*(P_n f : f \in \mathcal{F}$ for a set of functions \mathcal{F} , then the influence function of $\hat{\Theta}(P_n)$ can be written as

$$\sum_{f \in \mathcal{F}} \frac{d\hat{\Theta}^*(P)}{dPf} \left\{ f(O) - Pf \right\}.$$

Begin by writing the non-parametric MLE of the target parameter assuming all data are discrete:

$$\Theta(P_n) = \sum_{\bar{a}_{\tau}, \bar{l}_{\tau+1}} l_{\tau+1} \frac{\mathbb{1}[\bar{a}_{\tau} \in \bar{\mathcal{B}}_{\tau}]}{P_n f_{\bar{a}_{\tau+1}}} P_n f_{l_{\tau+1}, a_{\tau}, h_{\tau}} \prod_{k=1}^{\tau} \frac{P_n f_{d_k, h_k}}{P_n f_{a_k, h_k}},$$

where

$$f_{\bar{a}_{\tau}} = \mathbb{1}(\bar{A}_{\tau} \in \bar{\mathcal{B}}_{\tau}),$$

$$f_{l_{\tau+1},a_{\tau},h_{\tau}}(Z) = \mathbb{1}(L_{\tau+1} = l_{\tau+1}, A_{\tau} = a_{\tau}, H_{\tau} = h_{\tau}, \bar{A}_{\tau} \in \bar{\mathcal{B}}_{\tau}),$$

$$f_{d_k,h_k} = \mathbb{1}[d(A) = a, W = w],$$

$$f_{a_k,h_k} = \mathbb{1}[A = a, W = w].$$

Let $\mathcal{F} = \{f_{\bar{a}_{\tau}}, f_{l_{\tau+1}, a_{\tau}, h_{\tau}}, f_{d_k, h_k}, f_{a_k, h_k}\}$. The derivatives of $\Theta(P)$ with respect to each $f \in \mathcal{F}$ are given by

$$\begin{split} \frac{d\Theta(P)}{dP\bar{a}_{\tau}} &= -\sum_{\bar{a}_{\tau},\bar{l}_{\tau+1}} l_{\tau+1} \frac{\mathbb{1}[\bar{a}_{\tau} \in \bar{\mathcal{B}}_{\tau}]}{(Pf_{\bar{a}_{\tau+1}})^2} Pf_{l_{\tau+1},a_{\tau},h_{\tau}} \prod_{k=1}^{\tau} r(a_k,h_k), \\ \frac{d\Theta(P)}{dPf_{l_{\tau+1},a_{\tau},h_{\tau}}} &= l_{\tau+1} \frac{\mathbb{1}[\bar{a}_{\tau} \in \bar{\mathcal{B}}_{\tau}]}{Pf_{\bar{a}_{\tau+1}}} \prod_{k=1}^{\tau} r(a_k,h_k), \\ \frac{d\Theta(P)}{dPf_{d_s,h_s}} &= \frac{\mathbb{1}[\underline{a}_{s+1} \in \underline{\mathcal{B}}_{s+1}]}{Pf_{\bar{a}_{\tau+1}}} m_s(a_s,h_s) \prod_{k=1}^{s-1} r(a_k,h_k) \\ \frac{d\Theta(P)}{dPf_{a_s,h_s}} &= -\frac{\mathbb{1}[\underline{a}_{s+1} \in \underline{\mathcal{B}}_{s+1}]}{Pf_{\bar{a}_{\tau+1}}} m_s(a_s,h_s) \prod_{k=1}^{s} r(a_k,h_k). \end{split}$$

Note that

$$\sum_{f \in \mathcal{F}} \frac{d\Theta(P)}{dPf} Pf = 0.$$

The conjectured EIF is therefore given by

$$\sum_{f \in \mathcal{F}} \frac{d\Theta(P)}{dPf} f(O).$$

Using the derivatives above and summing over s yields the stated EIF.

The functional $\Theta(P_{\epsilon})$ can be expanded as

$$\Theta(P_{\epsilon}) = \Theta(P) + \int D(Z, P) dP_{\epsilon} - \operatorname{Rem}_{0}(\eta_{\epsilon}, \eta),$$

where $\text{Rem}_0(\eta_{\epsilon}, \eta)$ is defined as in (6). Next, we evaluate the derivative in (7)

$$\frac{d}{d\epsilon}\Theta(P_{\epsilon})\Big|_{\epsilon=0} = \int D(Z;P)\left(\frac{dP_{\epsilon}}{d\epsilon}\right)\Big|_{\epsilon=0} - \frac{d}{d\epsilon}\operatorname{Rem}_{0}(\eta_{\epsilon},\eta)\Big|_{\epsilon=0},$$
$$= \int D(Z;P)s(Z)\Big|_{\epsilon=0} - \frac{d}{d\epsilon}\operatorname{Rem}_{0}(\eta_{\epsilon},\eta)\Big|_{\epsilon=0}.$$

Note that due to the product structure of $\operatorname{Rem}_0(\eta_{\epsilon}, \eta)$ and the fact that $\eta_{\epsilon}|_{\epsilon=0} = \eta$, it follows that

$$\left. \frac{d}{d\epsilon} \operatorname{Rem}_0(\eta_{\epsilon}, \eta) \right|_{\epsilon=0} = 0.$$

This completes the proof.

A.5 Additional Simulation Results

ſ				95% Coverage	$MAE \times 100$			$ME \times 100$		
	\mathcal{B}^{a}	Scenario	N	TMLE	IPW	Sub	TMLE	IPW	Sub	TMLE
	{1}	1	250	97.0%	9.03	4.58	5.10	-6.46	-3.50	-0.50
			500	96.0%	4.80	2.72	4.31	-1.70	-1.58	0.63
			1000	93.0%	3.25	2.12	2.84	-1.62	-1.27	-0.17
		2	250	76.0%	15.91	4.49	4.72	-15.91	-3.48	-3.71
			500	84.0%	15.25	2.56	2.62	-15.25	-1.54	-1.57
			1000	83.0%	15.56	2.10	2.08	-15.56	-1.28	-1.29
		3	250	97.0%	8.63	15.92	5.34	-6.56	-15.92	-0.56
			500	96.0%	4.98	15.28	4.47	-2.18	-15.28	0.68
			1000	93.0%	3.16	15.61	2.84	-1.59	-15.61	-0.26
		4	250	0.0%	15.94	15.92	16.03	-15.94	-15.92	-16.03
			500	0.0%	15.23	15.28	15.35	-15.23	-15.28	-15.35
			1000	0.0%	15.56	15.61	15.66	-15.56	-15.61	-15.66
	{ 2 }	1	250	98.0%	17.92	4.55	6.72	-15.61	-2.60	0.98
			500	97.0%	7.47	2.56	5.17	-5.35	-0.73	1.89
			1000	96.0%	4.15	1.88	3.52	-2.16	-0.05	1.28
		2	250	74.0%	16.91	4.88	4.91	-16.91	-2.47	-2.60
			500	85.0%	16.25	2.58	2.48	-16.25	-0.78	-0.82
			1000	87.0%	16.58	1.86	1.81	-16.58	-0.08	-0.08
		3	250	96.0%	18.96	16.94	7.11	-17.38	-16.94	0.98
			500	99.0%	7.46	16.30	4.93	-4.71	-16.30	2.12
			1000	95.0%	3.74	16.63	3.42	-1.96	-16.63	1.28
		4	250	0.0%	16.98	16.94	16.95	-16.98	-16.94	-16.95
			500	0.0%	16.23	16.30	16.33	-16.23	-16.30	-16.33
			1000	0.0%	16.58	16.63	16.64	-16.58	-16.63	-16.64
	{ 3 }	1	250	99.0%	21.12	3.39	8.84	-19.23	-1.05	1.86
			500	96.0%	9.04	2.25	5.66	-6.56	0.49	3.32
			1000	89.0%	4.15	1.70	4.52	-1.14	0.71	2.82
		2	250	83.0%	8.06	3.85	3.87	-8.06	-0.89	-1.02
			500	92.0%	7.41	2.21	2.16	-7.41	0.49	0.50
		2	1000	90.0%	7.74	1.65	1.73	-7.74	0.54	0.54
		3	250	96.0%	21.89	8.12	8.25	-20.32	-8.12	2.99
			500	94.0%	8.26	7.48	5.81	-5.64	-7.48	3.59
			1000	93.0%	4.42	7.81	4.30	-1.29	-7.81	2.96
		4	250	25.0%	8.03	8.12	8.12	-8.02	-8.12	-8.12
			500	7.0%	7.43	7.48	7.49	-7.43	-7.48	-7.49
	()		1000	0.0%	7.75	7.81	7.81	-7.75	-7.81	-7.81
	{4}	1	250	100.0%	15.15	3.16	6.56	-12.62	0.49	1.48
			500	97.0%	6.51	2.68	5.35	-2.76	1.13	2.72
			1000	91.0%	4.17	1.77	4.25	1.23	0.83	2.68
		2	250	87.0%	2.52	3.29	3.35	1.18	0.74	0.50

		500	80.0%	2.41	2.63	2.60	1.80	1.33	1.18
		1000	83.0%	1.71	1.59	1.70	1.46	0.72	0.73
	3	250	96.0%	14.89	2.52	6.80	-13.47	1.07	3.17
		500	92.0%	6.67	2.35	5.65	-2.49	1.71	3.86
		1000	89.0%	4.25	1.65	4.34	1.12	1.38	3.71
	4	250	94.0%	2.54	2.52	2.59	1.20	1.07	1.18
		500	88.0%	2.44	2.35	2.36	1.84	1.71	1.75
		1000	86.0%	1.70	1.65	1.66	1.46	1.38	1.40
{ 5 }	1	250	95.0%	6.73	4.36	6.05	-3.87	3.22	2.12
		500	91.0%	4.80	3.98	4.16	-0.88	3.15	2.12
		1000	91.0%	3.18	2.86	2.98	1.15	2.37	1.89
	2	250	75.0%	11.69	4.66	4.58	11.69	3.56	3.20
		500	53.0%	12.25	4.19	4.12	12.25	3.38	3.36
		1000	58.0%	11.90	2.70	2.83	11.90	2.27	2.22
	3	250	94.0%	7.12	11.50	6.28	-3.85	11.50	2.96
		500	92.0%	4.30	12.14	3.97	-0.62	12.14	2.36
		1000	91.0%	3.02	11.81	3.23	0.94	11.81	2.51
	4	250	2.0%	11.67	11.50	11.71	11.67	11.50	11.71
		500	0.0%	12.27	12.14	12.22	12.27	12.14	12.22
		1000	0.0%	11.91	11.81	11.86	11.91	11.81	11.86

Table 4: Results of Simulation Study 1 showing empirical coverage of the 95% confidence intervals, Mean Absolute Error (MAE), and Mean Error (ME) for the inverse probability weighted estimator (IPW), substitution estimator (Sub), and Targeted minimum loss-based estimator (TMLE).

References

- James Heckman, Justin L Tobias, and Edward Vytlacil. Four parameters of interest in the evaluation of social programs. *Southern Economic Journal*, 68(2):210–223, 2001.
- Iván Díaz and Mark J van der Laan. Population intervention causal effects based on stochastic interventions. *Biometrics*, 68(2):541–549, 2012.
- Sebastian Haneuse and Andrea Rotnitzky. Estimation of the effect of interventions that modify the received treatment. *Statistics in Medicine*, 2013.
- Iván Díaz, Nicholas Williams, Katherine L. Hoffman, and Edward J. Schenck. Nonparametric causal effects based on longitudinal modified treatment policies. *Journal of the American Statistical Association*, 118(542):846–857, 2023. doi: 10.1080/01621459.2021.1955691. URL https://doi.org/10.1080/01621459.2021.1955691.
- Katherine L. Hoffman, Diego Salazar-Barreto, Kara E. Rudolph, and Iván Díaz. Introducing longitudinal modified treatment policies: a unified framework for studying complex exposures, 2023.
- Iván Díaz, Katherine L Hoffman, and Nima S Hejazi. Causal survival analysis under competing

risks using longitudinal modified treatment policies. *Lifetime Data Analysis*, 30(1):213–236, 2024.

- Jessica G Young, Miguel A Hernán, and James M Robins. Identification, estimation and approximation of risk under interventions that depend on the natural value of treatment using observational data. *Epidemiologic methods*, 3(1):1–19, 2014.
- Alexander R Luedtke, Oleg Sofrygin, Mark J van der Laan, and Marco Carone. Sequential double robustness in right-censored longitudinal models. *arXiv preprint arXiv:1705.02459*, 2017.
- Heejung Bang and James M Robins. Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61(4):962–973, 2005.
- Victor Chernozhukov, Whitney Newey, Víctor M Quintas-Martínez, and Vasilis Syrgkanis. RieszNet and ForestRiesz: Automatic debiased machine learning with neural nets and random forests. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 3901–3914. PMLR, 17–23 Jul 2022a. URL https://proceedings.mlr.press/v162/chernozhukov22a.html.
- James J Heckman. Instrumental variables: A cautionary tale. Working Paper 185, National Bureau of Economic Research, September 1995. URL http://www.nber.org/papers/t0185.
- Jinyong Hahn. On the role of the propensity score in efficient semiparametric estimation of average treatment effects. *Econometrica*, pages 315–331, 1998.
- Ilya Shpitser and Judea Pearl. Effects of treatment on the treated: identification and generalization. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, UAI '09, page 514–521, Arlington, Virginia, USA, 2009. AUAI Press. ISBN 9780974903958.
- Finbarr P. Leacy and Elizabeth A. Stuart. On the joint use of propensity and prognostic scores in estimation of the average treatment effect on the treated: a simulation study. *Statistics in Medicine*, 33(20):3488–3508, 2014. doi: https://doi.org/10.1002/sim.6030. URL https:// onlinelibrary.wiley.com/doi/abs/10.1002/sim.6030.
- Aolin Wang, Roch A. Nianogo, and Onyebuchi A. Arah. G-computation of average treatment effects on the treated and the untreated. *BMC Medical Research Methodology*, 17(1):3, Jan 2017. ISSN 1471-2288. doi: 10.1186/s12874-016-0282-4. URL https://doi.org/10.1186/s12874-016-0282-4.
- Roland A Matsouaka, Yi Liu, and Yunji Zhou. Variance estimation for the average treatment effects on the treated and on the controls. *Statistical Methods in Medical Research*, 32(2): 389–403, 2023. doi: 10.1177/09622802221142532. URL https://doi.org/10.1177/09622802221142532. PMID: 36476035.
- Alan E. Hubbard, Nicholas P. Jewell, and Mark J. van der Laan. *Direct Effects and Effect Among the Treated*, pages 133–143. Springer New York, New York, NY, 2011. ISBN 978-1-4419-9782-1. doi: 10.1007/978-1-4419-9782-1_8. URL https://doi.org/10.1007/978-1-4419-9782-1_8.

- Erica E. M. Moodie, Olli Saarela, and David A. Stephens. A doubly robust weighting estimator of the average treatment effect on the treated. *Stat*, 7(1):e205, 2018. doi: https://doi.org/10.1002/sta4.205. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/sta4.205. e205 sta4.205.
- Tyler J. VanderWeele and Miguel A. Hernan. Causal inference under multiple versions of treatment. *Journal of Causal Inference*, 1(1):1–20, 2013. doi: doi:10.1515/jci-2012-0002. URL https://doi.org/10.1515/jci-2012-0002.
- Eric J Tchetgen Tchetgen and Stijn Vansteelandt. Alternative identification and inference for the effect of treatment on the treated with an instrumental variable. *Harvard University Biostatistics Working Paper Series*, November 2013. URL https://biostats.bepress.com/harvardbiostat/paper166. Working Paper 166.
- Lan Liu, Wang Miao, Baoluo Sun, James M Robins, and Eric J Tchetgen Tchetgen. Doubly robust estimation of a marginal average effect of treatment on the treated with an instrumental variable. *Harvard University Biostatistics Working Paper Series*, June 2015. URL https://biostats.bepress.com/ucbbiostat/paper268. Working Paper 191.
- Michael Schomaker and Philipp FM Baumann. Doubly robust estimation of average treatment effects on the treated through marginal structural models. *Observational Studies*, 9(3):43–57, 2023.
- James M Robins, Miguel A Hernán, and Uwe Siebert. Effects of multiple interventions. *Comparative quantification of health risks: global and regional burden of disease attributable to selected major risk factors*, 1:2191–2230, 2004.
- Thomas S Richardson and James M Robins. 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(30):2013, 2013.
- Judea Pearl. Myth, Confusion, and Science in Causal Analysis. Technical Report R-348, Cognitive Systems Laboratory, Computer Science Department University of California, Los Angeles, Los Angeles, CA, May 2009.
- Peter J Bickel. On adaptive estimation. The Annals of Statistics, pages 647-671, 1982.
- Aad W van der Vaart and Jon A Wellner. *Weak Convergence and Emprical Processes*. Springer-Verlag New York, 1996.
- Andrea Rotnitzky, James Robins, and Lucia Babino. On the multiply robust estimation of the mean of the g-functional. *arXiv preprint arXiv:1705.08582*, 2017.
- Mark J van der Laan and Daniel Rubin. Targeted maximum likelihood learning. *The International Journal of Biostatistics*, 2(1), 2006.
- Mark J van der Laan and Sherri Rose. Targeted Learning: Causal Inference for Observational and Experimental Data. Springer, New York, 2011.

- Victor Chernozhukov, Whitney K. Newey, Victor Quintas-Martinez, and Vasilis Syrgkanis. Automatic debiased machine learning via neural nets for generalized linear regression, 2021.
- Victor Chernozhukov, Whitney K Newey, and Rahul Singh. Debiased machine learning of global and local parameters using regularized Riesz representers. *The Econometrics Journal*, 25(3): 576–601, 04 2022b. ISSN 1368-4221. doi: 10.1093/ectj/utac002. URL https://doi.org/ 10.1093/ectj/utac002.
- Victor Chernozhukov, Whitney Newey, Rahul Singh, and Vasilis Syrgkanis. Automatic debiased machine learning for dynamic treatment effects and general nested functionals, 2023.
- Iván Díaz and Mark J van der Laan. Stochastic treatment regimes. In *Targeted Learning in Data Science*, pages 219–232. Springer, 2018.
- Environmental Protection Agency. National ambient air quality standards (naaqs) for pm, 2023. URL https://www.epa.gov/pm-pollution/ national-ambient-air-quality-standards-naaqs-pm.
- Maya L Petersen, Kristin E Porter, Susan Gruber, Yue Wang, and Mark J Van Der Laan. Diagnosing and responding to violations in the positivity assumption. *Statistical methods in medical research*, 21(1):31–54, 2012.
- Janet M Begun, WJ Hall, Wei-Min Huang, Jon A Wellner, et al. Information and asymptotic efficiency in parametric-nonparametric models. *The Annals of Statistics*, 11(2):432–452, 1983.
- Peter J Bickel, Chris AJ Klaassen, YA'Acov Ritov, and Jon A Wellner. *Efficient and Adaptive Estimation for Semiparametric Models*. Springer-Verlag, 1997.
- A. W. van der Vaart. Asymptotic Statistics. Cambridge University Press, 1998.
- Edward H. Kennedy. Semiparametric Theory and Empirical Processes in Causal Inference, pages 141–167. Springer International Publishing, Cham, 2016. ISBN 978-3-319-41259-7. doi: 10.1007/978-3-319-41259-7_8. URL https://doi.org/10.1007/978-3-319-41259-7_8.
- James Robins, Lingling Li, Eric Tchetgen Tchetgen, and Aad W van der Vaart. Quadratic semiparametric Von Mises calculus. *Metrika*, 69(2-3):227–247, 2009.
- R von Mises. On the asymptotic distribution of differentiable statistical functions. *The annals of mathematical statistics*, 18(3):309–348, 1947.
- J Pfanzagl and W Wefelmeyer. Contributions to a general asymptotic statistical theory. *Statistics* & *Risk Modeling*, 3(3-4):379–388, 1985.
- Michel Emery, Dan Voiculescu, and Arkadi Nemirovski. Lectures on probability theory and statistics : Ecole d'Eté de Probabilités de Saint-Flour XXIII - 1998. Lecture Notes in Mathematics, Lectures from the 28th summer school on probability theory held in Saint-Flour, vol. 1738. Springer, Berlin, 2000.
- Martin J. Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2019.

- Wenjing Zheng and Mark J van der Laan. Cross-validated targeted minimum-loss-based estimation. In *Targeted Learning*, pages 459–474. Springer, 2011.
- Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 01 2018. ISSN 1368-4221. doi: 10.1111/ectj.12097. URL https://doi.org/10.1111/ectj.12097.
- Chris AJ Klaassen. Consistent estimation of the influence function of locally asymptotically linear estimators. *The Annals of Statistics*, 15(4):1548–1562, 1987.