The International Journal of Biostatistics

Volume 8, *Issue* 1 2012 *Article* 9

Targeted Minimum Loss Based Estimation of Causal Effects of Multiple Time Point Interventions

Mark J. van der Laan, *University of California - Berkeley*
Susan Gruber, *Harvard University*

Recommended Citation:
van der Laan, Mark J. and Gruber, Susan (2012) "Targeted Minimum Loss Based Estimation of Causal Effects of Multiple Time Point Interventions," *The International Journal of Biostatistics*: Vol. 8: Iss. 1, Article 9.
DOI: 10.1515/1557-4679.1370

Targeted Minimum Loss Based Estimation of Causal Effects of Multiple Time Point Interventions

Mark J. van der Laan and Susan Gruber

Abstract

We consider estimation of the effect of a multiple time point intervention on an outcome of interest, where the intervention nodes are subject to time-dependent confounding by intermediate covariates.

In previous work van der Laan (2010) and Stitelman and van der Laan (2011a) developed and implemented a closed form targeted maximum likelihood estimator (TMLE) relying on the log-likelihood loss function, and demonstrated important gains relative to inverse probability of treatment weighted estimators and estimating equation based estimators. This TMLE relies on an initial estimator of the entire probability distribution of the longitudinal data structure. To enhance the finite sample performance of the TMLE of the target parameter it is of interest to select the smallest possible relevant part of the data generating distribution, which is estimated and updated by TMLE. Inspired by this goal, we develop a new closed form TMLE of an intervention specific mean outcome based on general longitudinal data structures. The target parameter is represented as an iterative sequence of conditional expectations of the outcome of interest. This collection of conditional means represents the relevant part, which is estimated and updated using the general TMLE algorithm. We also develop this new TMLE for other causal parameters, such as parameters defined by working marginal structural models. The theoretical properties of the TMLE are also practically demonstrated with a small scale simulation study.The proposed TMLE is building upon a previously proposed estimator Bang and Robins (2005) by integrating some of its key and innovative ideas into the TMLE framework.

KEYWORDS: Asymptotic linearity of an estimator, causal effect, efficient influence curve, confounding, G-computation formula, influence curve, longitudinal data, loss function, marginal structural working model, nonparametric structural equation model, positivity assumption, randomization assumption, semiparametric statistical model, treatment regimen, targeted maximum likelihood estimation, targeted minimum loss based estimation, TMLE

Author Notes: This work was supported by NIH grants R01 AI074345-04 and R01 AI073127.

1 Introduction

Many studies generate data sets that can be represented as n independent and identically distributed observations on a specified longitudinal data structure. For example, the longitudinal data structure might represent measurements collected on a randomly sampled subject over a certain time period. Some of these measurements might represent an exposure or treatment, and one might be concerned with assessing the effect of this exposure or treatment on a final outcome of interest. For that purpose, one may code the observed data structure as a time-ordered $O = (L(0), A(0), \ldots, L(K), A(K), Y = L(K+1))$, where $L(0)$ are baseline covariates, $A(t)$ denotes an exposure or treatment at "time" t, and $L(t)$ denotes covariates measured between two subsequent treatments $A(t-1)$ and $A(t)$, while Y is the final outcome measured after the final treatment.

We will use the notation $\bar{L}(k) = (L(0), \ldots, L(k))$ and similarly $\bar{A}(k) = (A(0), \ldots, A(k))$ to denote histories of a time-dependent process. By specifying a causal graph (Pearl (1995), Pearl (2000)), or equivalently, a system of nonparametric structural equations, it is assumed that each component of the observed longitudinal data structure (e.g, $L(k)$) is a function of a set of observed parent variables (often the whole history, e.g. $\bar{A}(k-1), \bar{L}(k-1)$) and an unmeasured exogenous error term. Such a causal model provides a parameterization of the distribution of the observed data structure O, and allows one to define a post-intervention distribution that represents the distribution O would have had under a specified intervention on the nodes $A = (A(0), \ldots, A(K))$. These nodes $(A(0), \ldots, A(K))$ are called the intervention nodes. Causal effects of interventions on these intervention nodes are defined as parameters of a collection of post-intervention distributions defined by a specified set of interventions.

For the sake of presentation, in the main part of this article we focus on estimation of the intervention specific mean outcome for a single static intervention that sets $\bar{A} = \bar{a}$. The post-intervention distribution will be denoted with P_a and let L_a be a random variable with this distribution, which is also called a counterfactual. Under the assumption that the intervention nodes are sequentially randomized, conditional on the observed parent nodes, and a positivity assumption, one can identify the post-intervention distribution P_a of the counterfactual L_a by the so called G-computation formula P^a. This G-computation formula P^a is defined by the product over all $L(k)$-nodes of the conditional distribution of the $L(k)$-node, given its parents, and $\bar{A}(k-1) = \bar{a}(k-1)$. Thus, the latter probability distribution P^a is defined by the distribution P of the data O, and only equals the desired distribution P_a of the counterfactual L_a under the above mentioned nonparametric structural equation model, sequential randomization assumption and positivity assumption. This G-computation formula P^a for the distribution of L_a also pro-

vides a G-computation formula $E_{P^a}Y^a$ for the mean of Y_a. This establishes that, under the causal model and required above mentioned identifiability conditions, the causal quantity of interest can be written as a function Ψ of the data generating distribution, namely $\Psi(P) = E_{P^a}Y^a$, where the latter represents the statistical target parameter or estimand of interest.

Assuming these identifiability conditions, a current and important topic is the statistical estimation of the G-computation estimand $E_{P^a}Y^a$ based on observing n independent and identically distributed copies of the longitudinal data structure O. The statistical estimation problem is defined by specifying the statistical model $\mathcal{M}$ (the set of possible probability distributions of O) and the definition of the statistical target parameter $\Psi : \mathcal{M} \to \mathbb{R}$. We consider a statistical model that only makes statistical assumptions about the intervention mechanism, where the latter is defined by the conditional distribution of the intervention node $A(k)$, given the parent nodes $(\bar{A}(k-1), \bar{L}(k))$ of the intervention node, across the intervention nodes, while we put no restrictions on any of the conditional distributions of $L(k)$, given $(\bar{A}(k-1), \bar{L}(k-1))$.

The G-computation formula for the causal quantity of interest can now be represented as a target parameter mapping the statistical model to the real line, and this target parameter is a pathwise differentiable functional of the data generating distribution on this statistical model. The canonical gradient $D^*(P)$ at a $P \in \mathcal{M}$ of the pathwise derivative of $\Psi : \mathcal{M} \to \mathbb{R}$ at P, which defines the pathwise derivative along paths through P, is also called the efficient influence curve at P (formally defined in Section 2.2 below). An estimator is asymptotically efficient at P if and only if it is asymptotically linear at P with influence curve equal to this canonical gradient $D^*(P)$. Thus, it is no surprise that the construction of an efficient estimator of $\Psi(P)$ under i.i.d sampling from P will need to involve the utilization of the canonical gradient/efficient influence curve $D^*(P)$. In fact, an efficient estimator will at least need to approximately solve the so called efficient influence curve equation (also called the efficient score equation) defined by setting the empirical mean of the efficient influence curve at the estimator equal to zero. However, this requirement, by no means, defines an efficient estimator: for example, there are infinite possible estimators $\mathbf{P}_n$ of P_0 that will solve the (one-dimensional) efficient influence curve equation $\sum_i D^*(P)(O_i) = 0$ in the infinite dimensional P. In particular, a nonparametric maximum likelihood estimator (NPMLE), assuming it would be well defined, would solve all score equations, including this efficient influence curve equation. However, due to the curse of dimensionality, an NPMLE is often ill defined, and has poor practical performance. Another class of estimators are so called estimating equation based estimators that aim to represent $D^*(P)$ as an estimating function $D^*(\Psi(P), \eta(P))$, and given an estimator η_n of the nuisance parameter η, it defines the estimator of ψ_0 as the solution of $\sum_i D^*(\psi, \eta_n)(O_i) = 0$. This approach

for constructing an efficient estimator has various previously outlined disadvantages (see e.g. van der Laan and Rubin (2006); van der Laan and Rose (2011)) such as 1) the efficient influence curve $D^*(P)$ might not allow a representation as an estimating function, 2) the estimating equation might have no or multiple solutions, 3) it might be the case that even in the limit $E_0 D^*(\psi, \eta_0) = 0$ has multiple solutions so that the efficient influence curve estimating function simply does not identify the true ψ_0, 4) the resulting estimator ψ_n and η_n are not compatible with a single probability distribution, and, more generally, the estimator ψ_n does not necessarily respect the global constraints implied by the statistical model $\mathcal{M}$ and the target parameter mapping $\Psi : \mathcal{M} \to \mathbb{R}$. The latter explains the often erratic behavior of these estimators in the context of practical violations of the positivity assumptions (for example, resulting in an estimated probability that is negative or larger than 1).

Targeted minimum loss based estimation (TMLE) provides a template for the construction of semiparametric locally efficient double robust substitution estimators of the target parameter of the data generating distribution in a semiparametric censored data or causal inference model based on a sample of independent and identically distributed copies from this data generating distribution (van der Laan and Rubin (2006); van der Laan (2008); van der Laan and Rose (2011)). It relies on an initial estimator of a relevant part of the data generating distribution (defined as a minimizer of the risk of a loss function), and updates this estimator in a targeted manner using a least-favorable parametric fluctuation model whose loss-based score at zero fluctuation spans the efficient influence curve. The estimator is defined by iteratively maximizing an empirical risk over this least favorable parametric submodel through the current estimator, and often exists in closed form by only requiring a finite number of iterations. Since TMLE is a substitution estimator it respects the global constraints implied by the statistical model and target parameter mapping. By construction, the update of the initial estimator solves the efficient influence curve equation. As a consequence, the resulting plug-in estimator of the target parameter is double robust and asymptotically efficient, under appropriate regularity conditions. The choices of the so called relevant part, its loss function, the least-favorable submodel through this relevant part, and the choice of iterative updating algorithm, define the actual TMLE.

1.1 Existing TMLE approach for estimating an intervention specific mean outcome

In previous work van der Laan (2010) and Stitelman and van der Laan (2011a) developed and implemented a closed form targeted maximum likelihood estimator of the intervention specific mean outcome, defining the relevant part as the relevant

factor of the density of O, using the log-likelihood loss function. Stitelman and van der Laan (2011a) demonstrated dramatic gains of this TMLE relative to inverse probability of treatment weighted estimators and estimating equation based estimators, based on simulated and an analysis of a randomized controlled trial with a time until event outcome that is subject to drop-out informed by time-dependent biomarkers. This TMLE relies on an initial estimator of the G-computation formula P^a: it writes $\Psi(P) = \Psi_1(P^a) = E_{P^a}Y^a$, obtains an initial estimator of P^a, updates this initial estimator based on a least favorable fluctuation indexed by an estimator of the intervention mechanism, and plugs it in Ψ_1. Thus, this estimator requires estimating the entire density of the longitudinal data structure.

To enhance the finite sample performance of TMLE of the target parameter it is of interest to select the smallest possible relevant part of the data generating distribution in the definition of TMLE, which is estimated and updated by TMLE. That is, one might have two representations of the target parameter as a function of a relevant part of P: for each $P \in \mathcal{M}$, $\Psi(P) = \Psi_1(Q_1(P))$ and $\Psi(P) = \Psi_2(Q_2(P))$ for different parameters/relevant parts $Q_1(P)$ and $Q_2(P)$ defined as minimizers of the risk of different loss functions. Each of these representations will imply a TMLE involving an initial estimator of the relevant part $Q_j(P)$, an update Q_{jn}^* based on iterative minimization of an empirical risk along a least favorable fluctuation model through the current update of the estimator of $Q_j(P)$, and corresponding plug-in TMLE $\Psi_j(Q_{jn}^*)$ of the target parameter $\psi_0 = \Psi(P_0)$, $j = 1, 2$. In addition, suppose that $Q_1(P)$ is a smaller parameter than $Q_2(P)$ in the sense that $Q_1(P)$ is a (many to one) function of $Q_2(P)$. The initial estimator and updated estimator (TMLE) of Q_2 also yields an initial and updated estimator of Q_1 (and both TMLE's solve the efficient influence curve equation), so from that point of view the main difference between the two resulting TMLEs is the behavior of the estimator of the (more) relevant part Q_1. Even though both TMLE are double robust and asymptotically efficient under regularity conditions, this behavior can easily affect the finite sample performance of the TMLE of ψ_0, and, in fact, it can also affect its asymptotic behavior (by converging to a different limit or by violating the regularity conditions allowing for the asymptotic linearity). If the initial estimator of Q_2 is based on a bias-variance trade-off with respect to (w.r.t.) the larger parameter $Q_2(P)$ (e.g., by using cross-validation w.r.t. the loss function for Q_2), then this plug-in estimator of Q_1 based on this estimator of Q_2 will generally be worse than an estimator of Q_1 that directly addresses the bias-variance trade-off w.r.t. Q_1. For example, if the target parameter $\Psi(P)$ only depends on the conditional mean of the outcome Y, given its parents, then an estimator of this conditional mean based on a loss function for this conditional mean might achieve a better rate of convergence than an estimator based on an estimator of the conditional density of Y, given its parents. It is *not* a mistake to use a plug-in estimator of the conditional mean based on an estimate of

the conditional density, but one wants to fit this conditional density based on a criterion for a candidate estimator that reflects the performance of the resulting plug-in estimator of the conditional mean. In addition, by focusing on what really needs to be estimated, the resulting TMLE can also become much simpler to implement.

Inspired by this goal of selecting a small relevant part Q (and obtaining a simple to implement TMLE), we develop a new closed form TMLE of an intervention specific mean outcome based on general longitudinal data structures. The target parameter is represented as an iterative sequence of conditional expectations of the outcome of interest, and this collection of conditional means represents the relevant part, which is estimated and updated using the general TMLE algorithm. We also develop this new TMLE for other causal parameters, such as parameters defined by working marginal structural models (MSM), parametric models for the marginal mean of counterfactual outcomes. The theoretical properties of the TMLE are also practically demonstrated with a small scale simulation study. The proposed TMLE builds upon a previously proposed estimator by Bang and Robins (2005) by integrating some of its key and innovative ideas into the TMLE framework.

It will be of interest to further study and evaluate the practical performance of this TMLE in future studies, in particular, in comparison with other TMLEs such as the one proposed in van der Laan (2010) and Stitelman and van der Laan (2011a) based on the log-likelihood loss function. A practical advantage of the TMLE presented in this article is that it is easier to implement since it only involves fitting K (iteratively defined) regressions, while the TMLE in van der Laan (2010) based on the log-likelihood involves fitting K conditional densities of $L(K)$. It should be noted again that by using a more targeted loss function for the initial estimator such as the one in this article, the TMLE based on fitting conditional densities can still be as good as a TMLE based on only fitting the required conditional means (see also the Appendix in van der Laan and Rose (2011) and van der Laan and Gruber (2010) for efficient influence curve based targeted loss functions that can be used to build the initial estimator). In other words, as remarked above, it is not a mistake to use a plug-in estimator based on an estimate of the whole density of the data, but one wants to fit this density based on a criterion for a candidate estimator that reflects the performance of the resulting plug-in estimator of the target parameter. Nonetheless, a gain in simplicity for the implementation is of great interest, even when both types of TMLE would be similar w.r.t. their statistical behavior.

1.2 Immediately relevant literature overview

Different type of estimators of the causal effect of a multiple time point intervention have been proposed. These estimators can be categorized as

- inverse probability of treatment/censoring weighted (IPTW) estimators,
- estimating equation based estimators based on solving an estimating equation in the parameter of interest, for a given estimator of nuisance parameters of the estimating equation, such as the augmented IPTW estimating equation,
- nonparametric maximum likelihood estimators, maximum likelihood based estimators based on parametric models (also called G-computation estimators), or data adaptive loss-based learning algorithms,
- targeted maximum likelihood (or more general, minimum loss-based) estimators (TMLE) defined in terms of an initial estimator of relevant part of data generating distribution, loss function for this relevant part, least favorable fluctuation submodel through an initial or current estimator that is used to iteratively update the initial estimator until convergence, and plugging this updated estimator into the parameter mapping.

We will now briefly discuss these different types of estimators and then highlight the new contribution of this article to the current literature.

The IPTW estimator relies on an estimator of the intervention mechanism, the maximum likelihood estimator relies on an estimator of the relevant factor of the likelihood, while the augmented IPTW estimator and TMLE utilize both estimators. The augmented IPTW and the TMLE both solve the efficient influence curve equation, and are thereby so called double robust, and locally asymptotically efficient. The TMLE is also a substitution estimator and is therefore guaranteed to respect the global constraints of the statistical model and target parameter mapping.

IPTW estimation is presented and discussed in detail in (Robins, 1999; Hernan et al., 2000). Augmented IPTW is originally developed in Robins and Rotnitzky (1992). Further development on estimating equation methodology and double robustness is presented in (Robins et al., 2000; Robins, 2000; Robins and Rotnitzky, 2001) and van der Laan and Robins (2003). For a detailed bibliography on locally efficient estimating equation methodology we refer to Chap. 1 in van der Laan and Robins (2003).

For the original paper on TMLE we refer to van der Laan and Rubin (2006). Subsequent papers on TMLE in observational and experimental studies include Bembom and van der Laan (2007), van der Laan (2008), Rose and van der Laan (2008, 2009, 2011), Moore and van der Laan (2009a,b,c), Bembom et al. (2009),

Polley and van der Laan (2009), Rosenblum et al. (2009), van der Laan and Gruber (2010), Stitelman and van der Laan (2010), Gruber and van der Laan (2010b), Rosenblum and van der Laan (2010), Wang et al. (2010), and Stitelman and van der Laan (2011b). For a general comprehensive book on this topic, which includes most of these applications on TMLE and many more, we refer to van der Laan and Rose (2011). An original example of a particular type of TMLE (based on a double robust parametric regression model) for estimation of a causal effect of a point-treatment intervention was presented in Scharfstein et al. (1999) and we refer to Rosenblum and van der Laan (2010) for a detailed review of this earlier literature and its relation to TMLE.

van der Laan (2010) and Stitelman and van der Laan (2011a) (see also van der Laan and Rose (2011)) present a closed form TMLE, based on the log-likelihood loss function, for estimation of a causal effect of a multiple time point intervention on an outcome of interest (including survival outcomes that are subject to right-censoring) based on general longitudinal data structures. In this article we integrate some key ideas from the double robust estimating equation method for longitudinal data proposed in Bang and Robins (2005) into the framework of targeted minimum loss based estimation. The resulting TMLE enhances the Bang and Robins (2005) estimator by 1) incorporating data adaptive estimation in place of parametric models, 2) generalizing it to parameters for which there exists no mapping of the efficient influence curve into an estimating equation, 3) avoiding the potential problem of estimating equations having no or multiple solutions, and 4) incorporating robust choices of loss functions and hardest parametric submodels so that the resulting TMLE is a robust substitution estimator (e.g., the squared error loss and linear fluctuation for conditional means is replaced by a robust loss and logistic fluctuation function). The new TMLE may have advantages relative to the TMLE based on the log-likelihood loss function as developed in van der Laan (2010) and Stitelman and van der Laan (2011a), as explained above. We also generalize this new TMLE to causal parameters defined by projections on working marginal structural models.

1.3 Organization

This article is organized as follows. In Section 2 we start out with defining the estimation problem in terms of the longitudinal unit data structure, the statistical model for the probability distribution of this unit data structure, the G-computation formula for the distribution of the data under a multiple time point intervention, and the corresponding target parameter being the intervention specific mean outcome. Subsequently, we define the target parameter as a function of an iteratively defined sequence of conditional means of the outcome under the distribution spec-

ified by the G-computation formula, one for each intervention node. For the sake of developing the TMLE based on this representation of the target parameter, we present a particular orthogonal decomposition of the canonical gradient/efficient influence curve of the target parameter mapping, where each component corresponds with a "score" of these conditional means (due to Bang and Robins (2005)). In Section 3 we present the TMLE of this target parameter in terms of an iteratively defined sequence of loss functions for the iteratively defined sequence of conditional means, an initial estimator using iterative loss-based learning to estimate each of the subsequently defined conditional means, an iteratively defined sequence of least favorable parametric submodels that are used for fluctuating each conditional mean subsequently, and finally the TMLE algorithm that updates the initial estimator by iteratively minimizing the loss-based empirical risk along the least favorable parametric submodel through the current estimator. The TMLE solves the efficient influence curve estimating equation, which provides a basis for establishing the double robustness of TMLE and statistical inference. In Section 4 we review the statistical properties of this TMLE and statistical inference. In Section 5 we carry out a small scale simulation study comparing this TMLE with an IPTW and a parametric MLE-based estimator. We conclude with some remarks in Section 6. The more technical results and generalization are deferred to the Appendix. A formal presentation of sequential loss-based cross-validation is presented in the Appendix. In the Appendix we also demonstrate that this new TMLE is compatible with a probability distribution and is thus a true substitution estimator. The Appendix also present a generalization of the TMLE for causal parameters defined by working marginal structural models. The Appendix concludes with R-code implementing the newly proposed TMLE.

2 Longitudinal data structure, model, target parameter, efficient influence curve

We start out with formally defining the estimation problem in terms of data, statistical model, and target parameter mapping. For the sake of TMLE, we need to decide on representing $\Psi(P)$ as a $\Psi_1(Q(P))$ for a mapping Ψ_1 of some relevant part $Q(P)$ of P, present a loss function for $Q(P)$, and a least favorable fluctuation through Q. Instead of defining the relevant part of the distribution of O as its relevant factor Q of its density, we define it as an iteratively defined sequence $\bar{Q}^a$ of conditional means indexed by the intervention a. The latter is a function of the relevant factor of the likelihood and thus represents a strictly smaller relevant part of P than its factor: i.e., for each $P \in \mathcal{M}$, $\bar{Q}^a(P)$ is a function of $Q(P)$. Finally, we utilize

a representation of the efficient influence curve in terms of scores of these iteratively defined conditional means, due to Bang and Robins (2005), fully preparing us to select the loss function and least favorable fluctuation models that defines the TMLE.

2.1 The statistical estimation problem in terms of data, model, and target parameter

We observe n i.i.d. copies of a longitudinal data structure

$$O = (L(0), A(0), \ldots, L(K), A(K), Y = L(K+1)),$$

where $A(j)$ denotes a discrete valued intervention node, $L(0)$ baseline covariates, $L(j)$ is a time-dependent confounder realized after $A(j-1)$ and before $A(j)$, $j = 1, \ldots, K$, and Y is a final univariate outcome of interest. There are no restrictions on the dimension and support of $L(j)$, $j = 0, \ldots, K$.

The probability distribution P_0 of O can be factorized according to the time-ordering as

$$\begin{aligned} P_0(O) &= \prod_{k=0}^{K+1} P_0(L(k) \mid Pa(L(k))) \prod_{k=0}^{K} P_0(A(k) \mid Pa(A(k))) \\ &\equiv \prod_{k=0}^{K+1} Q_{0,L(k)}(O) \prod_{k=0}^{K} g_{0,A(k)}(O) \\ &\equiv Q_0(O) g_0(O), \end{aligned}$$

where $Pa(L(k)) \equiv (\bar{L}(k-1), \bar{A}(k-1))$ and $Pa(A(k)) \equiv (\bar{L}(k), \bar{A}(k-1))$ denote the parents of $L(k)$ and $A(k)$ in the time-ordered sequence, respectively. Recall that $\bar{L}(k) = (L(0), \ldots, L(k))$, and $\bar{A}(k) = (A(0), \ldots, A(k))$. Note also that $Q_{0,L(k)}$ denotes the true conditional distribution of $L(k)$, given $Pa(L(k))$, and, $g_{0,A(k)}$ denotes the true conditional distribution of $A(k)$, given $Pa(A(k))$. We will also use the notation $g_{0:k} \equiv \prod_{j=0}^{k} g_{A(j)}$. We consider a statistical model $\mathcal{M}$ for P_0 that possibly assumes knowledge on g_0. If $\mathcal{Q}$ is the set of all values for Q_0 and $\mathcal{G}$ the set of possible values of g_0, then this statistical model can be represented as $\mathcal{M} = \{P = Qg : Q \in \mathcal{Q}, g \in \mathcal{G}\}$. In this statistical model $\mathcal{Q}$ puts no restrictions on the conditional distributions $Q_{0,L(k)}$, $k = 0, \ldots, K+1$.

Let

$$P^a(l) = \prod_{k=0}^{K+1} Q^a_{L(k)}(\bar{l}(k)), \tag{1}$$

where $Q^a_{L(k)}(\bar{l}(k)) = Q_{L(k)}(l(k) \mid \bar{l}(k-1), \bar{A}(k-1) = \bar{a}(k-1))$. This is the so called G-computation formula for the post-intervention distribution corresponding with the intervention that set all intervention nodes $\bar{A}(K)$ equal to $\bar{a}(K)$. Let $L^a = (L(0), L^a(1), \ldots, Y^a = L^a(K+1))$ denote a random variable with probability distribution P^a with final component Y^a. Our statistical target parameter is the mean of Y^a, $\Psi(P) = E_{P^a}Y^a$, thus defining $\Psi : \mathcal{M} \to \mathbb{R}$. This target parameter only depends on P through $Q = Q(P)$. Therefore, we will also denote the target parameter mapping with $\Psi : \mathcal{Q} = \{Q(P) : P \in \mathcal{M}\} \to \mathbb{R}$, acknowledging the abuse of notation.

Consider the nonparametric structural equation model (NPSEM) defined by

$$L(k) = f_{L(k)}(Pa(L(k)), U_{L(k)}),\ k = 0, \ldots, K+1,$$

and

$$A(k) = f_{A(k)}(Pa(A(k)), U_{A(k)})\ k = 0, \ldots, K,$$

in terms of a set of deterministic functions $(f_{L(k)} : k = 0, \ldots, K+1), (f_{A(k)} : k = 0, \ldots, K)$, and an exogenous vector of random errors $U = (U_{L(0)}, \ldots, U_{L(K+1)}, U_{A(0)}, \ldots, U_{A(K)})$ (Pearl (1995), Pearl (2000)). This allows one to define the counterfactual $L_{\bar{a}}$ by deterministically setting all the $A(k)$ equal to $a(k)$ in this system of structural equations. Here we used the notation $L_{\bar{a}}$ for the counterfactual, while above we used the notation $L^{\bar{a}}$ for the random variable with probability distribution defined by the G-computation formula P^a. The probability distribution of this counterfactual $L_{\bar{a}}$ is called the post-intervention distribution of L. Under the sequential randomization assumption stating that $A(k)$ is independent of $L_{\bar{a}}$, given $Pa(A(k))$, and the positivity assumption, $P(A(k) = a(k) \mid \bar{L}(k), \bar{A}(k-1) = \bar{a}(k-1)) > 0$ a.e., the probability distribution of $L_{\bar{a}}$ is identified and given by the G-computation formula P^a_0 defined by the true distribution P_0 of O under this system. Thus under these causal assumptions $L_{\bar{a}}$ and $L^{\bar{a}}$ have the same probability distribution. In particular, for any underlying distribution defined by the distribution of the exogenous errors U and the collection of functions (i.e., $f_{L(k)}$ and $f_{A(k)}$), we have that $EY_{\bar{a}} = E_{P^a}Y^a = \Psi(P)$ for the distribution P of O implied by this underlying distribution. Thus the causal model and causal parameter $EY_{\bar{a}}$ (and its identifiability) implies a statistical model $\mathcal{M}$ defined as the set of possible probability distributions P of O, and a statistical target parameter $\Psi : \mathcal{M} \to \mathbb{R}$. For the sake of defining the estimation problem of $EY_{\bar{a}}$ in this causal model, only the statistical model $\mathcal{M}$ and the statistical target parameter are relevant. As a consequence, the estimation of $\Psi(P_0)$ based on the statistical knowledge $P_0 \in \mathcal{M}$ as developed in this article also applies to estimation of the intervention specific mean $EY_{\bar{a}}$ in this causal model.

2.2 Representation of target parameter as function of an iteratively defined sequence of conditional means

By the iterative conditional expectation rule (tower rule), we can represent $E_{P^a}Y^a$ as an iterative conditional expectation, first conditioning on $\bar{L}^a(K)$, then conditioning on $\bar{L}^a(K-1)$, and so on, until the conditional expectation given $L(0)$, and finally taking the mean over $L(0)$. Formally, this defines a mapping from Q onto the real line defined as follows:

- Compute $\bar{Q}^a_Y = E_{Q^a_Y}Y \equiv E(Y \mid \bar{A}(K) = \bar{a}(K), \bar{L}(K))$ by computing the integral of Y w.r.t. conditional distribution Q^a_Y of Y, given $\bar{L}(K), \bar{A}(K) = \bar{a}(K)$.
- Given $\bar{Q}^a_Y$, next compute $\bar{Q}^a_{L(K)} = E_{Q^a_{L(K)}}\bar{Q}^a_Y$, obtained by integrating out $L(K)$ in $\bar{Q}^a_Y$ w.r.t. the conditional distribution $Q^a_{L(K)}$ of $L(K)$, given $\bar{L}(K-1), \bar{A}(K-1) = \bar{a}(K-1)$.
- This process is iterated: Given $\bar{Q}^a_{L(k)}$, compute $\bar{Q}^a_{L(k-1)} = E_{Q^a_{L(k-1)}}\bar{Q}^a_{L(k)}$, starting at $k = K+2$ and moving backwards until the final step,
- $\bar{Q}^a_{L(0)} = E_{Q_{L(0)}}\bar{Q}^a_{L(1)}$ at $k = 1$.

For notational convenience, here we define $\bar{Q}^a_{L(K+2)} \equiv Y$. Note that $\bar{Q}^a_{L(k)} = \bar{Q}^a_{L(k)}(\bar{L}(k-1))$ is a function of O through $\bar{L}(k-1)$, and, in particular, $\bar{Q}^a_{L(0)}$ is a constant. We also note that in terms of counterfactuals or the distribution of P^a we have $\bar{Q}^a_{L(k)} = E_Q(Y^a \mid \bar{L}^a(k-1))$. Of course, if this process is applied to the true distribution Q_0, then we indeed obtain the desired intervention specific mean: $\bar{Q}^a_{0,L(0)} = E_0Y^a = \Psi(Q_0)$.

Instead of representing our target parameter as a function of $Q = (Q_Y, Q_{L(K)}, \ldots, Q_{L(0)})$, we will view it as a function of an iteratively defined sequence of conditional means $\bar{Q}^a \equiv (\bar{Q}^a_Y, \bar{Q}^a_{L(K)}, \ldots, \bar{Q}^a_{L(0)})$, where $\bar{Q}^a_{L(k)}$ is viewed as a parameter (i.e., $E_{Q^a_{L(k)}}\bar{Q}^a_{L(k+1)}$) of $Q^a_{L(k)}$, given the previous $\bar{Q}^a_{L(k+1)}$. We will write $\Psi(\bar{Q}^a)$ if we want to stress that our target parameter only depends on Q through this iteratively defined $\bar{Q}^a$. Note that indeed $\bar{Q}^a$ is a function of Q.

2.3 Representation of efficient influence curve of target parameter as sum of iteratively defined scores of iteratively defined conditional means

Given the statistical model $\mathcal{M}$, and target parameter $\Psi : \mathcal{M} \to \mathbb{R}$, efficiency theory teaches us that an estimator $\widehat{\Psi}$ (viewed as mapping from empirical distribution into

$\mathbb{R}$) is asymptotically efficient at P_0 among the class of regular estimators of $\Psi(P_0)$ if and only if the estimator is asymptotically linear at P_0 with influence curve equal to the canonical gradient $D^*(P_0)$ of the pathwise derivative of $\Psi : \mathcal{M} \to \mathbb{R}$ at P_0: i.e., $\widehat{\Psi}(P_n) - \Psi(P_0) = 1/n \sum_{i=1}^n D^*(P_0)(O_i) + o_P(1/\sqrt{n})$. We remind the reader that a pathwise derivative for a path $\{P(\epsilon) : \epsilon\} \subset \mathcal{M}$ through P at $\epsilon = 0$ is defined by $\frac{d}{d\epsilon}\Psi(P(\epsilon))\Big|_{\epsilon=0}$. If for all paths through P this derivative can be represented as $PD^*(P)S \equiv \int D^*(P)(o)S(o)dP(o)$, where S is the score of the path at $\epsilon = 0$, and $D^*(P)$ is an element of the tangent space at P, then the target parameter mapping is pathwise differentiable at P and its canonical gradient is $D^*(P)$. The canonical gradient forms a crucial ingredient for the construction of double robust semiparametric efficient estimators, and, in particular, for the construction of a TMLE. We note that, due to the factorization of $P = Qg$ and that the target parameter only depends on P through Q, the canonical gradient does not depend on the model choice for g. In particular, the canonical gradient in the model in which g_0 is known equals the canonical gradient in our model $\mathcal{M}$, which assumes some model $\mathcal{G}$, possibly a nonparametric model (Bickel et al. (1997)). The following theorem provides the canonical gradient and presents a particular representation of the canonical gradient that will be utilized in the definition of our TMLE presented in the next section. The proof is presented in the appendix. This form of the efficient influence curve was previously established in Bang and Robins (2005).

Theorem 1 *Let $D(Q,g)(O) = Y\frac{I(\bar{A}(K)=\bar{a}(K))}{g_{0:K}(O)} - \Psi(Q)$, where we use the notation $I(B)$ for the indicator of an event B. This is a gradient of the pathwise derivative of Ψ in the model in which g is known. For notational convenience, in this theorem we often use a notation that suppresses the dependence of function-evaluations on Q, g and O. The efficient influence curve is given by $D^* = \sum_{k=0}^{K+1} D_k^*$, where $D_k^* = \Pi(D \mid T_k)$ is the projection of D onto the tangent space $T_k = \{h(L(k), Pa(L(k)) : E_Q(h \mid Pa(L(k))) = 0\}$ of $Q_{L(k)}$ in the Hilbert space $L_0^2(P)$ with inner-product $\langle h_1, h_2\rangle_P = Ph_1h_2$. Recall the definition $\bar{Q}^a_{L(k)} = E(Y^a \mid \bar{L}^a(k-1))$, and the recursive relation $\bar{Q}^a_{L(k)} = E_{Q^a_{L(k)}}\bar{Q}^a_{L(k+1)}$.*

We have

$$D^*_{K+1} = \frac{I(\bar{A}(K) = \bar{a}(K))}{g_{0:K}}(Y - \bar{Q}^a_{K+1}),$$

and

$$\begin{aligned} D_k^* &= \frac{I(\bar{A}(k-1) = \bar{a}(k-1))}{g_{0:k-1}}\left\{\bar{Q}^a_{L(k+1)} - E_{Q^a_{L(k)}}\bar{Q}^a_{L(k+1)}\right\}, \\ &= \frac{I(\bar{A}(k-1) = \bar{a}(k-1))}{g_{0:k-1}}\left\{\bar{Q}^a_{L(k+1)} - \bar{Q}^a_{L(k)}\right\},\ k = K, \ldots, 0. \end{aligned}$$

In particular,

$$D_0^* = \bar{Q}^a_{L(1)} - E_{L(0)}\bar{Q}^a_{L(1)} = \bar{Q}^a_{L(1)} - \Psi(\bar{Q}^a).$$

We note that for each $k = K+1, \ldots, 0$*,*

$$D_k^*(Q, g) = D_k^*(\bar{Q}^a_{L(k)}, \bar{Q}^a_{L(k+1)}, g_{0:k-1})$$

depends on Q, g *only through* $\bar{Q}^a_{L(k+1)}$*, its conditional mean* $\bar{Q}^a_{L(k)}$ *under* $Q^a_{L(k)}$*, and* $g_{0:k-1}$*.*

The following theorem states the double robustness of the efficient influence curve as established previously (e.g, van der Laan and Robins (2003)).

Theorem 2 *Consider the representation* $D^*(\bar{Q}^a, g, \Psi(\bar{Q}^a))$ *of the efficient influence curve as provided in Theorem 1 above. We have for any* g *for which* $g(\bar{A}(K) = \bar{a}(K), \bar{L}(K)) > 0$ *a.e.,*

$$P_0 D^*(\bar{Q}^a, g, \Psi(\bar{Q}^a_0)) = 0 \text{ if } \bar{Q}^a = \bar{Q}^a_0 \text{ or } g = g_0.$$

3 TMLE of intervention specific mean

In this section we develop the TMLE. Firstly, we present an overview of the general procedure in terms of initial estimator of the relevant part $\bar{Q}^a$, loss function, least favorable fluctuation model, and the iterative updating algorithm. In subsection 3.2 we present the practical implementation of the TMLE, while the remaining subsections present the main ingredients of the TMLE in more detail.

3.1 Overview of TMLE

The first step of the TMLE involves writing our target parameter as $\Psi(\bar{Q}^a)$, as done above. Secondly, we construct an initial estimator $\bar{Q}^a_n$ of $\bar{Q}^a_0$, and g_n of g_0. In addition, we need to present a loss function $\mathcal{L}_\eta(\bar{Q}^a)$ for $\bar{Q}^a_0$, possibly indexed by a nuisance parameter η, satisfying $\bar{Q}^a_0 = \underset{\bar{Q}^a}{\operatorname{argmin}}\ P_0\mathcal{L}_{\eta_0}(\bar{Q}^a)$, and a parametric submodel $\{\bar{Q}^a(\epsilon, g) : \epsilon\}$ in the parameter space of $\bar{Q}^a$, so that the linear span of the loss-based score $\frac{d}{d\epsilon}\mathcal{L}_{\eta_0}(\bar{Q}^a(\epsilon, g))$ at $\epsilon = 0$ includes the efficient influence curve $D^*(Q, g)$ of the target parameter mapping at $P = Qg$. Specifically, for each component $\bar{Q}^a_{0,L(k)}$ of $\bar{Q}^a = (\bar{Q}^a_{L(0)}, \ldots, \bar{Q}^a_{L(K+1)})$ we propose a loss function $\mathcal{L}_{k,\bar{Q}^a_{L(k+1)}}(\bar{Q}^a_{L(k)})$ indexed by "nuisance" parameter $\bar{Q}^a_{L(k+1)}$, and a corresponding submodel $\bar{Q}^a_{L(k)}(\epsilon, g)$ through $\bar{Q}^a_{L(k)}$ at $\epsilon = 0$ so that $\frac{d}{d\epsilon}\mathcal{L}_{\bar{Q}^a_{L(k+1)}}(\bar{Q}^a_{L(k)}(\epsilon, g))$ at $\epsilon = 0$ equals the k-th

component $D_k^*(\bar{Q}^a_{L(k)}, \bar{Q}^a_{L(k+1)}, g)$ of the efficient influence curve D^* as defined in Theorem 1, $k = 0, \ldots, K+1$. The sum loss function $\sum_{k=0}^{K+1} \mathcal{L}_{k,\bar{Q}^a_{L(k+1)}}(\bar{Q}^a_{L(k)})$ is now a loss function for $(\bar{Q}^a_{L(0)}, \ldots, \bar{Q}^a_{L(K+1)})$ and the corresponding "score" of the submodel through $\bar{Q}^a$ defined by all these k-specific submodels spans the complete efficient influence curve.

Finally, we will present a particular closed form targeted minimum loss-based estimation algorithm that iteratively minimizes the empirical mean of the loss function over this parametric submodel through the current estimator of $\bar{Q}^a_0$ (starting with initial estimator), updating one component at a time. This algorithm starts with updating the initial estimator $\bar{Q}^a_{L(K+1),n}$ of $\bar{Q}^a_{L(K+1)}$ based on the $(K+1)$-th loss function $\mathcal{L}_{K+1}(\bar{Q}^a_{L(K+1)})$ resulting in update $\bar{Q}^{a,*}_{L(K+1),n} = \bar{Q}^a_{L(K+1),n}(\epsilon_{K,n}, g_n)$ with $\epsilon_{K,n} = \underset{\epsilon}{\operatorname{argmin}}\ P_n \mathcal{L}(\bar{Q}^a_{L(K+1),n}(\epsilon, g_n))$. It iterates this updating process going backwards until obtaining the update $\bar{Q}^{a,*}_{L(0),n} = \bar{Q}^a_{L(0),n}(\epsilon_{0,n}, g_n)$ of the initial estimator $\bar{Q}^a_{L(0),n}$ of $\bar{Q}^a_{L(0)}$, where $\epsilon_{0,n} = \underset{\epsilon}{\operatorname{argmin}}\ P_n \mathcal{L}_{\bar{Q}^{a,*}_{L(1),n}}(\bar{Q}^a_{L(0),n}(\epsilon, g_n))$ using the most recent updated estimator $\bar{Q}^{a,*}_{L(1),n}$ of $\bar{Q}^a_{0,L(1)}$. This yields the TMLE $\bar{Q}^{a,*}_n$ of the vector of conditional means $\bar{Q}^a_0$. In particular, its first component $\bar{Q}^{a,*}_{L(0),n}$ is the TMLE of $\Psi(\bar{Q}^a_0) = \bar{Q}^a_{0,L(0)}$.

By the fact that the MLE of ϵ_k solves the score equation, it follows that the TMLE solves $P_n D_k^*(\bar{Q}^{a,*}_{L(k),n}, \bar{Q}^{a,*}_{L(k+1),n}, g_{0:k-1,n})$ for each $k = K+1, \ldots, 0$. In particular, this implies that $(\bar{Q}^{a,*}_n, g_n)$ solves the efficient influence curve equation: $P_n D^*(\bar{Q}^{a,*}_n, g_n, \Psi(\bar{Q}^{a,*}_n)) = 0$. Before we proceed with specifying the TMLE in detail, we first present the summary of the practical implementation of the proposed TMLE.

In the following, we will assume that Y is bounded (i.e, $P_0(Y \in (a,b)) = 1$ for some $a < b < \infty$), and thereby, without loss of generality, we will assume that $Y \in [0,1]$. A special case would be that Y is binary valued with values in $\{0,1\}$.

3.2 Summary of practical implementation of TMLE

Let g_n be an estimator of g_0. Firstly, we carry out a regression of Y onto $\bar{A}(K) = \bar{a}(K), \bar{L}(K)$. For example, we might fit a multivariate linear logistic regression of Y_i onto a set of main terms that are univariate summary measures Z_i extracted from $\bar{L}_i(K)$ among the observations with $\bar{A}_i(K) = \bar{a}(K)$. Alternatively, we use data adaptive machine learning algorithms to fit this underlying regression.

Subsequently, we fluctuate this initial estimate to better target the fit towards the parameter of interest by using the initial estimator of $\bar{Q}^a_{Y,0} = E_0(Y \mid \bar{A}(K) = \bar{a}(K), \bar{L}(K))$ as an offset in a univariate logistic regression with clever covariate

$I(\bar{A}(K) = \bar{a}(K))/g_{0:K,n}$, and fit the corresponding univariate logistic regression of Y among the observations with $\bar{A}(K) = \bar{a}(K)$, minimizing the negative log likelihood loss function. This yields the TMLE $\bar{Q}^{a,*}_{Y,n}$ of the last component $\bar{Q}^a_{Y,0}$ of $\bar{Q}^a_0$.

We now run a logistic regression of $\bar{Q}^{a,*}_{Y,n}$ onto $\bar{A}(K-1) = \bar{a}(K-1), \bar{L}(K-1)$, among the observations with $\bar{A}_i(K-1) = \bar{a}(K-1)$. This initial estimator of $\bar{Q}^a_{L(K)} = E(Y^a \mid \bar{L}^a(K-1))$ is used as an offset in a univariate logistic regression of $\bar{Q}^{a,*}_{Y,n}$ with clever covariate $I(\bar{A}(K-1) = \bar{a}(K-1))/g_{0:K-1,n}$. Let $\bar{Q}^{a,*}_{L(K),n}$ be the resulting fit of $\bar{Q}^a_{L(K)}$. This is the TMLE of $\bar{Q}^a_{L(K),0}$ (second from last component of $\bar{Q}^a_0$).

This process of subsequent estimation of the next conditional mean, given the TMLE fit of the previous conditional mean, followed by a targeting flucutation, is iterated. Thus, for any $k \in \{K+1, \ldots, 1\}$, run a logistic regression of the previous TMLE fit $\bar{Q}^{a,*}_{L(k+1),n}$ onto $\bar{A}(k-1) = \bar{a}(k-1), \bar{L}(k-1)$, among the observations $\bar{A}_k(k-1) = \bar{a}(k-1)$, and use this fit as an offset in a univariate logistic regression of $\bar{Q}^{a,*}_{L(k+1),n}$ with clever covariate $I(\bar{A}(k-1) = \bar{a}(k-1))/g_{0:k-1,n}$. Let $\bar{Q}^{a,*}_{L(k),n}$ be the resulting logistic regression fit of $\bar{Q}^a_{L(k)}$. This is the TMLE of $\bar{Q}^a_{L(k),0}$, $k = K+1, \ldots, 1$.

Consider now the fit $\bar{Q}^{a,*}_{L(1),n}$ at the $k = 1$ step. This is a function of $L(0)$. We estimate $\bar{Q}^a_{L(0)}$ with the empirical mean $\frac{1}{n}\sum_{i=1}^n \bar{Q}^{a,*}_{L(1),n}(L_i(0))$. Let $\bar{Q}^{a,*}_n = (\bar{Q}^{a,*}_{L(k),n}, k = 0, \ldots, K+1)$ be the TMLE of $\bar{Q}^a_0$. The last estimate $\frac{1}{n}\sum_{i=1}^n \bar{Q}^{a,*}_{L(1),n}(L_i(0))$ is the TMLE $\bar{Q}^*_{L(0),n} = \Psi(\bar{Q}^{a,*}_n)$ of our target parameter $\bar{Q}^a_{L(0)} = \Psi(\bar{Q}^a_0)$.

3.3 Loss function for $\bar{Q}^a_0$

For each k, $\bar{Q}^a_{L(k)}$ is a function that maps $\bar{L}(k-1)$ into $(0,1)$. For each $k = K+1, \ldots, 0$, we define the following loss function for $\bar{Q}^a_{L(k)}$, indexed by "nuisance" parameter $\bar{Q}^a_{L(k+1)}$:

$$\begin{aligned}\mathcal{L}_{k,\bar{Q}^a_{L(k+1)}}(\bar{Q}^a_{L(k)}) &= -I(\bar{A}(k-1) = \bar{a}(k-1)) \times \\ &\quad \left\{\bar{Q}^a_{L(k+1)} \log \bar{Q}^a_{L(k)} + (1 - \bar{Q}^a_{L(k+1)}) \log(1 - \bar{Q}^a_{L(k)})\right\}.\end{aligned}$$

For notational convenience, here we define $\bar{Q}^a_{L(K+2)} \equiv Y$, so that the loss function for $\bar{Q}^a_Y$ is given by

$$\mathcal{L}_{K+1}(\bar{Q}^a_Y) = -I(\bar{A}(K) = \bar{a}(K)) \left\{Y \log \bar{Q}^a_Y + (1-Y) \log(1 - \bar{Q}^a_Y)\right\}.$$

Indeed, we have that

$$E_0(\bar{Q}^a_{L(k+1)}(L(k),\bar{L}(k-1)) \mid \bar{A}(k-1)=\bar{a}(k-1),\bar{L}(k-1)) \\ = \underset{\bar{Q}^a_{L(k)}}{\operatorname{argmin}}\ E_{P_0}\mathcal{L}_{k,\bar{Q}^a_{L(k+1)}}(\bar{Q}^a_{L(k)}).$$

In other words, given any function $\bar{Q}^a_{L(k+1)}$ of $L(k),\bar{L}(k-1)$, the minimizer of the expectation of the loss function $\mathcal{L}_{k,\bar{Q}^a_{L(k+1)}}$ over all candidates $\bar{Q}^a_{L(k)}$, is the actual conditional mean under $Q^a_{0,L(k)}$ of $\bar{Q}^a_{L(k+1)}$ (see e.g., Gruber and van der Laan (2010a)). In particular, if the "nuisance" parameter $\bar{Q}^a_{L(k+1)}$ of this loss function is correctly specified, then this minimizer equals the desired $\bar{Q}^a_{0,L(k)}$.

An alternative choice of loss function is a (possibly weighted) squared error loss function:

$$\mathcal{L}_{k,\bar{Q}^a_{L(k+1)}}(\bar{Q}^a_{L(k)}) = I(\bar{A}(k-1)=\bar{a}(k-1))\left(\bar{Q}^a_{L(k+1)}-\bar{Q}^a_{L(k)}\right)^2.$$

However, this choice combined with linear fluctuation submodels (as in Bang and Robins (2005)) will yield a non-robust TMLE that does not respect the global constraints of the model and target parameter, for the same reason as presented in Gruber and van der Laan (2010a).

These loss functions for $\bar{Q}^a_{L(k)}$ across k can be combined into a single loss function $\mathcal{L}_\eta(\bar{Q}^a) = \sum_{k=0}^{K+1}\mathcal{L}_{k,\eta_k}(\bar{Q}^a_{L(k)})\Big|_{\eta_k=\bar{Q}^a_{L(k+1)}}$ indexed by a nuisance parameter $\eta=(\eta_k : k=0,\ldots,K+1)$. This can be viewed as a sum loss function indexed by nuisance parameters η_k, and, at correctly specified nuisance parameters, it is indeed minimized by $\bar{Q}^a_0$. However, the nuisance parameters are themselves minimizers of the risk of these loss functions, so that it is sensible to define $\bar{Q}^a_0$ as the solution of the iterative minimization of the risks of the loss functions: $Y=\bar{Q}^a_{0,L(K+2)}$, for $k=K+2,\ldots,1$, $\bar{Q}^a_{0,L(k-1)} = \underset{\bar{Q}^a_{L(k-1)}}{\operatorname{argmin}}\ E_0\mathcal{L}_{\bar{Q}^a_{0,L(k)}}(\bar{Q}^a_{L(k-1)})$. This is indeed the way we utilize this loss function for $\bar{Q}^a$ in both the definition of the TMLE, as well as in the definition of the initial estimator of $\bar{Q}^a_0$.

3.4 Least favorable parametric submodel

In order to compute a TMLE we wish to determine a submodel $\{\bar{Q}^a_{L(k)}(\epsilon_k,g):\epsilon_k\}$ through $\bar{Q}^a_{L(k)}$ at $\epsilon_k=0$ so that

$$\frac{d}{d\epsilon_k}\mathcal{L}_{k,\bar{Q}^a_{L(k+1)}}(\bar{Q}^a_{L(k)}(\epsilon_k,g))\Big|_{\epsilon_k=0} = D^*_k(Q,g). \tag{2}$$

Recall the definition of $D_k^*(Q,g)$ in Theorem 1. We can select the following submodel

$$\text{logit}\bar{Q}^a_{L(k)}(g,\epsilon_k) = \text{logit}\bar{Q}^a_{L(k)} + \epsilon_k \frac{1}{g_{0:k-1}},\ k = K+1,\ldots,0,$$

where we define $g_{0:-1} = 1$. This submodel does indeed satisfy the generalized score-condition (2). In particular, the submodel $\bar{Q}^a(\epsilon_0,\ldots,\epsilon_{K+1},g)$ defined by these k-specific submodels through $\bar{Q}^a_{L(k)}$, $k = 0,\ldots,K+1$, and the above sum loss function $\mathcal{L}_{\bar{Q}^a}(\bar{Q}^a) = \sum_{k=0}^{K+1} \mathcal{L}_{k,\bar{Q}^a_{L(k+1)}}(\bar{Q}^a_{L(k)})$ satisfies the condition that the generalized score spans the efficient influence curve:

$$D^*(Q,g) \in \left\langle \frac{d}{d\epsilon}\mathcal{L}_{\bar{Q}^a}(\bar{Q}^a(\epsilon,g))\Big|_{\epsilon=0} \right\rangle. \tag{3}$$

Here we used the notation $\langle(h_0,\ldots,h_{K+1})\rangle = \{\sum_k c_k h_k : c_k\}$ for all linear combinations spanned by the components of h.

3.5 Initial estimator

For notational convenience, in the remainder of the paper we will interchangeably use the notation $\bar{Q}^a_{L(k)}$ and $\bar{Q}^a_k$. Firstly, we fit $\bar{Q}^a_{K+1}$ based on a loss-based learning algorithm with loss function $\mathcal{L}_{K+1}(\bar{Q}^a_{K+1})$, or the squared error loss function. Note that this loss function is not indexed by an unknown nuisance parameter. For example, one could fit $\bar{Q}^a_{K+1}$ by fitting a parametric regression model for this conditional mean using standard software. However, in general, we recommend the utilization of machine learning algorithms based on this same loss function. Given an estimator $\bar{Q}^a_{K+1,n}$ of $\bar{Q}^a_{K+1}$, we can fit $\bar{Q}^a_K$ based on a loss-based learning algorithm with loss function $\mathcal{L}_{K,\bar{Q}^a_{K+1,n}}(\bar{Q}^a_K)$. For example, a fit could be obtained by fitting a linear or logistic regression model for the conditional mean of $\bar{Q}^a_{K+1,n}$ as a linear function of a set of main terms extracted from $\bar{L}(K-1)$. This process can be iterated. So for $k = K+1$ to $k = 1$, we fit $\bar{Q}^a_k$ with a loss-based learning algorithm based on loss function $\mathcal{L}_{k,\bar{Q}^a_{k+1}}(\bar{Q}^a_k)$, given the previously selected estimator of the nuisance parameter $\bar{Q}^a_{k+1}$ in this loss function. Finally, $\bar{Q}^a_{L(0),n} = 1/n\sum_{i=1}^n \bar{Q}^a_{a,n}(L_i(0))$. In this manner, we obtain a fit $\bar{Q}^a_n$ of $\bar{Q}^a_0 = (\bar{Q}^a_{L(0)},\ldots,\bar{Q}^a_{L(K+1)})$. We can estimate g_0 with a log-likelihood based learning algorithm, which results in an estimator g_n of g_0. We refer to the Appendix for a formal presentation of sequential loss-based cross-validation as an ingredient for sequential loss-based learning.

3.6 TMLE algorithm

We already obtained an initial estimator $\bar{Q}^a_{k,n}$, $k = 0, \ldots, K+1$ and g_n. Let $\bar{Q}^{a,*}_{K+2,n} \equiv Y$. For $k = K+1$ to $k = 1$, we compute

$$\epsilon_{k,n} \equiv \underset{\epsilon_k}{\operatorname{argmin}} \, P_n \mathcal{L}_{k,\bar{Q}^{a,*}_{k+1,n}}(\bar{Q}^a_{k,n}(\epsilon_k, g_n)),$$

and the corresponding update $\bar{Q}^{a,*}_{k,n} = \bar{Q}^a_{k,n}(\epsilon_{k,n}, g_n)$. Finally, $\bar{Q}^{a,*}_{L(0),n} = 1/n \sum_{i=1}^n \bar{Q}^{a,*}_{1,n}(L_i(0))$. This defines the TMLE $\bar{Q}^{a,*}_n = (\bar{Q}^{a,*}_{k,n}, k = 0, \ldots, K+1)$ of $\bar{Q}^a_0 = (\bar{Q}^a_{0,L(0)}, \ldots, \bar{Q}^a_{0,L(K+1)})$.

Finally, we compute the TMLE of ψ_0 as the plug-in estimator corresponding with $\bar{Q}^{a,*}_n$:

$$\Psi(\bar{Q}^{a,*}_n) = \bar{Q}^{a,*}_{L(0),n} = \frac{1}{n} \sum_{i=1}^n \bar{Q}^{a,*}_{1,n}(L_i(0)).$$

We note that this single step recursive TMLE is an analogue to the recursive algorithm in Bang and Robins (2005) (operating on estimating functions), and the single step recursive TMLE in van der Laan (2010) and Stitelman and van der Laan (2011a).

Remark: Iterative TMLE based on common fluctuation parameter. One could have used a hardest parametric submodel $\bar{Q}^a(\epsilon, g) = (\bar{Q}^a_k(\epsilon, g) : k = 0, \ldots, K+1)$ with a common $\epsilon_k = \epsilon$ for all $k = 0, \ldots, K+1$, and use the sum-loss function $\mathcal{L}_{\bar{Q}^a}(\bar{Q}^a)$ so that the generalized score $\frac{d}{d\epsilon}\mathcal{L}_{\bar{Q}^a}(\bar{Q}^a(\epsilon, g))$ at zero fluctuation equals the efficient influence curve. An iterative TMLE is now defined as follows: Set $j = 0$, compute $\epsilon^j_n = \underset{\epsilon}{\operatorname{argmin}} \, P_n \mathcal{L}_{\bar{Q}^{a,j}_n}(\bar{Q}^{a,j}_n(\epsilon, g_n))$, compute the update $\bar{Q}^{a,j+1}_n = \bar{Q}^{a,j}_n(\epsilon^j_n, g_n)$, and iterate this updating step till convergence (i.e., $\epsilon^j_n \approx 0$). Notice that the common ϵ^j_n now provides an update of all $K+1$ components of $\bar{Q}^{a,j}_n$, and that the nuisance parameter in the loss function is also updated at each step. The final $\bar{Q}^{a,*}_n$ solves the efficient influence curve equation $P_n D^*(\bar{Q}^*_n, g_n)$ again. However, the above TMLE algorithm with the multivariate ϵ-fluctuation parameter using the backwards (recursive) updating algorithm, converges in one single step and thus exists in closed form. Therefore, we prefer this single step TMLE (analogue to the expressed preference of the single step (backwards updating) TMLE above the common-ϵ iterative TMLE in van der Laan (2010)).

Remark: TMLE using Inverse probability of treatment weighted loss function. Alternatively, we can select the submodels

$$\text{logit}\bar{Q}^a_{L(k)}(\epsilon_k) = \text{logit}\bar{Q}^a_{L(k)} + \epsilon_k 1, \; k = K+1, \ldots, 0,$$

and, for each $k = K+1, \ldots, 0$, given $\bar{Q}^a_{L(k+1)}$ and g, the following loss function for $\bar{Q}^a_{L(k)}$:

$$\mathcal{L}_{k,\bar{Q}^a_{L(k+1)},g}(\bar{Q}^a_{L(k)}) = -\frac{I(\bar{A}(k-1)=\bar{a}(k))}{g_{0:k-1}} \left\{ \bar{Q}^a_{L(k+1)} \log \bar{Q}^a_{L(k)} + (1 - \bar{Q}^a_{L(k+1)}) \log\{1 - \bar{Q}^a_{L(k)}\} \right\}.$$

This choice of loss function and submodel also satisfies the generalized score condition (2). The same single step recursive (backwards) TMLE applies.

4 Statistical properties and inference for TMLE

The TMLE $\bar{Q}^{a,*}_n$ solves $P_n D^*(\bar{Q}^{a,*}_n, g_n, \Psi(\bar{Q}^{a,*}_n)) = 0$, where the efficient influence curve $D^*(\bar{Q}^a, g, \Psi(\bar{Q}^a))$ is presented in Theorem 1. Due to the double robustness stated in Theorem 2, the estimator $\Psi(\bar{Q}^{a,*}_n)$ will be consistent for ψ_0 if either $\bar{Q}^{a,*}_n$ or g_n is consistent. In addition, under regularity conditions, if $g_n = g_0$, $\Psi(\bar{Q}^{a,*}_n)$ will also be asymptotically linear with influence curve $D^*(\bar{Q}^{a,*}, g_0, \psi_0)$, where $\bar{Q}^{a,*}$ is the possibly misspecified limit of $\bar{Q}^{a,*}_n$ (e.g., Theorem A5 in Appendix A18, van der Laan and Rose (2011)). As shown in van der Laan and Robins (2003) (Theorem 2.3 section 2.3.7), if g_n is a maximum likelihood based consistent estimator of g_0 according to a model $\mathcal{G}$ with tangent space $T_g(P_0)$, then under similar regularity conditions, the TMLE $\Psi(\bar{Q}^{a,*}_n)$ is asymptotically linear with influence curve $D^*(\bar{Q}^{a,*}, g_0, \psi_0) - \Pi(D^*(\bar{Q}^{a,*}, g_0, \psi_0) \mid T_g(P_0))$, where $\Pi(\cdot \mid T_g(P_0))$ is the projection operator onto $T_g(P_0) \subset L^2_0(P_0)$ within the Hilbert space $L^2_0(P_0)$ with inner product $\langle h_1, h_2 \rangle_{P_0} = P_0 h_1 h_2$. Note that if $\bar{Q}^{a,*} = \bar{Q}^a_0$, then the latter influence curve is the efficient influence curve $D^*(\bar{Q}^a_0, g_0, \psi_0)$, so that, in this case, the TMLE is asymptotically efficient. Therefore, under the assumption that $\mathcal{G}$ contains the true g_0, we can conservatively estimate the asymptotic covariance matrix of $\sqrt{n}(\Psi(\bar{Q}^{a,*}_n) - \Psi(\bar{Q}^a_0))$ with

$$\Sigma_n = P_n D^*(\bar{Q}^{a,*}_n, g_n, \psi^*_n) D^*(\bar{Q}^{a,*}_n, g_n, \psi^*_n)^\top.$$

If one is only willing to assume that either $\bar{Q}^{a,*}_n$ or g_n is consistent, then the influence curve is more complex (see Theorem 2.5 in van der Laan and Robins (2003), and Theorem A5 in Appendix A18 in van der Laan and Rose (2011)), and we recommend the bootstrap, although one can still use Σ_n as a first approximation, and confirm findings of interest with the bootstrap.

Formal asymptotic linearity theorems with precise conditions can be established by imitating the proof in Zheng and van der Laan (2011) for the natural direct

effect parameter, and Zheng and van der Laan (2010) and van der Laan and Rose (2011) for the additive causal effect parameter. In fact, the asymptotic linearity theorem for the TMLE presented in this article will have very similar structure and conditions to the asymptotic linearity theorem stated in the above referenced articles. General templates for establishing asymptotic linearity are provided in van der Laan and Robins (2003) and van der Laan and Rose (2011) as well.

5 Simulation studies

The TMLE presented in this paper provides a streamlined approach to the analysis of longitudinal data that reduces bias introduced by informative censoring and/or time-dependent confounders. Simulation studies presented in this section illustrate its application in two important areas, the estimation of the effect of treatment in a randomized controlled trial (RCT) with informative drop-out and time-dependent treatment modification, and estimation of the effect of treatment on survival in an observational study setting. TMLE performance is compared with two IPTW estimators. The first is an unstabilized IPTW estimator defined as $\psi_u^{IPTW} = E[I(\bar{A}_i(K) = \bar{a}(K))/g_{0:K,i}Y_i]$. The second IPTW estimator uses normalized weights, and is defined as $\psi_n^{IPTW} = E[I(\bar{A}_i(K) = \bar{a}(K))m/g_{0:k,i}Y_i]$, where $m = 1/n\sum_{i=1}^{n} I(A_i(K) = \bar{a}(K))/g_{0:K,i}$ is the empirical mean of the unstabilized weights, taken over subjects who are uncensored at the final time point. Estimates are also obtained for the parametric G-formula maximum likelihood estimator (MLE$_p$) obtained by plugging untargeted estimates of $\bar{Q}^a_{L(k)}$ into the G-computation formula (1). Influence curve (IC) based estimates of the variance of the TMLE are compared with the empirical variance of the Monte Carlo estimates, and coverage of the IC-based confidence intervals is reported.

5.1 Simulation 1: Additive effect of treatment in RCT with non-compliance and informative drop-out

Treatment decisions made over time can make it difficult to assess the effect of a particular drug regimen on a subsequent outcome. Consider a RCT to assess drug effectiveness on a continuous-valued outcome. Our target parameter is the mean outcome under the treatment regime $(A(0) = 1, A(1) = 1)$ minus the mean outcome under control, $(A(0) = 0, A(1) = 0)$, and no censoring $(C(0) = 0, C(1) = 0)$: $\psi_0 = E_0\{Y(1,0,1,0) - Y(0,0,0,0)\}$.

The diagram in Figure 1 shows the time ordering of intervention nodes (A, C) and covariate/event nodes (L, Y) for simulation 1. This time ordering corresponds to a study designed to estimate the effect of an asthma medication on

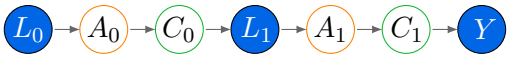

Figure 1: Simulation 1: Time ordering of intervention and non-intervention nodes, baseline covariates $L_0 = (W_1, W_2, W_3)$, treatment nodes (A_0, A_1), censoring nodes (C_0, C_1), time-dependent covariate L_1, outcome $Y = L_2$.

airway constriction after one year of adherence to treatment. Suppose that in response to results of an intermediate biomarker assay or clinical test (e.g. $L_1 =$ forced expiratory volume measured at six months) a subset of subjects in the treatment group discontinue treatment midway through the trial. Thus $A_0 = 1, A_1 = 0$ for these subjects, who are following neither the treatment protocol of interest, nor the control protocol. MLE$_p$ can consistently estimate ψ_0 as long as the Q factors of the likelihood (regression models for the conditional expectation of Y given the parents of Y, at all L and Y nodes), are correctly specified. IPTW also provides consistent estimates of ψ_0 providing the g factors of the likelihood (conditional distributions for all intervention nodes, conditional on the past), are correctly specified. TMLE is consistent when either one of these requirements is met. These estimators were applied to estimate the additive treatment effect in 500 samples ($n_1 = 100$, $n_2 = 1000$), drawn from the following data generating distribution:

$$\begin{aligned}
W_1, W_2 &\sim \text{Bernouli}(0.5) \\
W_3 &\sim N(4,1) \\
g_{0,1}(1 \mid Pa(A_0)) &= P_0(A_0 = 1 \mid L_0) = 0.5 \\
g_{0,2}(0 \mid Pa(C_0)) &= P_0(C_0 = 0 \mid A_0, L_0) \\
&= \text{expit}(0.1 + 0.5W_1 + W_2 - 0.1W_3) \\
L_1 &= 3 + A_0 - 0.5W_1W_3 - 0.5W_3 + \epsilon_1 \\
g_{0,3}(1 \mid Pa(A_1), A_0 = 1, C_0 = 0) &= P_0(A_1 = 1 \mid A_0 = 1, C_0 = 0, \bar{L}(1)) \\
&= \text{expit}(-1.2 - 0.2W_2 + 0.1W_3 + 0.4L_1) \\
g_{0,4}(0 \mid Pa(C_1), C_0 = 0) &= P_0(C_1 = 0 \mid A_0, C_0 = 0, A_1, \bar{L}(1)) \\
&= \text{expit}(2 - 0.05W_3 - 0.4L_1 - 1.5A_1) \\
Y &= \text{expit}(3 - 0.3A_0 + 0.1W_2 - 0.5L_1 - 0.5A_1 + \epsilon_2)
\end{aligned}$$

with $\epsilon_1, \epsilon_2 \sim_{i.i.d.} N(0,1)$. Under this data generating distribution the true value of the additive treatment effect is $\psi_0 = -0.160$, and the semi-parametric efficiency bound on the variance of the estimator is $\sigma^2 = 0.39/n$.

Performance is illustrated under correct and misspecified models for the Q and g factors of the likelihood. The label 'Q_c' denotes a set of logistic regression models that includes all terms used to generate the data at each covariate and event node. Using these models to estimate conditional means $\bar{Q}^a_{L(k)}$ gives practically unbiased estimation of ψ_0. 'Q_{m_1}' is a set of logistic regression models that includes

main term baseline covariates only. 'Q_{m_2}' is a set of more severely misspecified intercept-only logistic regression models for the outcome.

Two approaches were used to estimate $g_{0,1}, g_{0,2}, g_{0,3}, g_{0,4}$, initial treatment assignment probabilities, censoring (loss to follow-up) at baseline, intermediate switching from treatment to control, and subsequent loss to follow-up before the outcome is ascertained. The first approach relied on correctly specified logistic regression models to regress A_k on the parents of A_k. The second used main terms logistic regression models that included all baseline covariates. For convenience, in Table 1 these are referred to as correct and misspecified models for g, respectively. Estimated values for g_n, were not bounded away from $(0, 1)$. For each estimator, the additive causal effect estimate was defined as the difference in the estimated treatment specific means.

Results: Table 1 lists the empirical bias, percent bias relative to the true parameter value, variance, and mean squared error (MSE) of the Monte Carlo esti-

Table 1: Simulation 1 results, $\psi_0 = -0.160$.

		$n = 100$				$n = 1000$			
		%Rel bias	Bias	Var	MSE	%Rel bias	Bias	Var	MSE
					g correctly specified				
	IPTW$_u$	13.40	-0.021	0.032	0.033	-0.40	0.001	0.0007	0.0007
	IPTW$_n$	-13.27	0.021	0.004	0.004	-15.45	0.025	0.0002	0.0008
Q_c	MLE$_p$	0.01	0.000	0.006	0.006	-0.82	0.001	0.0003	0.0004
	TMLE	-0.72	0.001	0.006	0.006	-0.60	0.001	0.0003	0.0003
Q_{m_1}	MLE$_p$	8.93	-0.014	0.008	0.008	6.21	-0.010	0.0004	0.0005
	TMLE	6.34	-0.010	0.008	0.008	1.81	-0.003	0.0004	0.0004
Q_{m_2}	MLE$_p$	11.23	-0.018	0.006	0.006	12.61	-0.020	0.0005	0.0009
	TMLE	1.96	-0.003	0.006	0.006	3.13	-0.005	0.0004	0.0005
					g misspecified				
	IPTW$_u$	26.08	-0.042	0.049	0.051	6.73	-0.011	0.0005	0.0006
	IPTW$_n$	-8.28	0.013	0.004	0.004	-12.93	0.021	0.0002	0.0007
Q_c	MLE$_p$	0.01	0.000	0.006	0.006	-0.82	0.001	0.0003	0.0004
	TMLE	0.40	-0.001	0.006	0.006	-0.58	0.001	0.0004	0.0004
Q_{m_1}	MLE$_p$	8.93	-0.014	0.008	0.008	6.21	-0.010	0.0004	0.0005
	TMLE	9.51	-0.015	0.008	0.008	6.15	-0.010	0.0004	0.0005
Q_{m_2}	MLE$_p$	11.23	-0.018	0.006	0.006	12.61	-0.020	0.0005	0.0009
	TMLE	5.90	-0.009	0.007	0.007	7.26	-0.012	0.0005	0.0006

mates. The results confirm that all estimators are unbiased under correct parametric model specification, although sparsity in the data inflates IPTW variance at the smaller sample size when unstabilized weights are used (IPTW$_u$). Normalizing the weights decreases variance significantly, however at the cost of increased bias when $n = 1000$. As anticipated, MLE$_p$ estimates are biased under misspecified models Q_{m_1} and Q_{m_2}. When these misspecified models are targeted by TMLE using the correctly specified models for g_0 bias is greatly reduced.The large relative bias and variance of IPTW$_u$ at sample size $n = 100$ illustrates one problem inherent to longitudinal studies: sparsity in the data, i.e., a lack of information for learning the target parameter. This sparsity impairs TMLE performance as well, but to a lesser degree, since the submodel and quasi-log-likelihood loss function used in the estimation procedure respect the bounds on the statistical model $\mathcal{M}$, and thus the variance does not greatly suffer (Gruber and van der Laan, 2010a).

5.2 Simulation 2: Causal effect of treatment on survival with right-censoring and time-dependent covariates

Consider an observational study in which we wish to estimate the treatment-specific survival probability at time t_k, $\psi_0 = P(T_{\bar{a}} > t_k)$, where treatment is assigned at baseline, time-dependent covariates and mortality are assessed periodically during follow-up and at the end of study. During the trial some subjects experience the event, and others drop out due to reasons related to treatment or covariate information, thereby confounding a naive effect estimate. The time-ordering of the intervention nodes (A, C), and time-dependent covariate/event nodes (L, Y) for one such study design is shown in Figure 2.

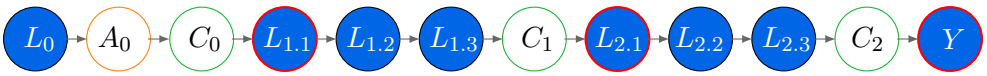

Figure 2: Simulation 2: Time ordering of intervention and non-intervention nodes, baseline covariates ($L_0 = W_1, W_2, W_3, W_4, W_5$), treatment node ($A_0$), censoring nodes ($C_0, C_1, C_2$), time-dependent covariates ($L_{1.2}, L_{1.3}, L_{2.2}, L_{2.3}$), intermediate and final outcomes ($L_{1.1}, L_{2.1}, Y = L_{3.1}$).

As a concrete example, consider a study to assess the effect of prostate surgery on mortality at the end of follow-up, where covariates measured periodically include PSA level (prostate-specific antigen) and a summary health measure ($L_{t.2}$ and $L_{t.3}$). IPTW$_u$, IPTW$_n$, MLE$_p$, and TMLE were applied to 500 samples of size $n_1 = 100$, $n_2 = 1000$, to estimate mean survival under treatment at time

$t_k = 3$. Data were generated as follows:

$$\begin{aligned}
W_1 &\sim N(67, 16) \\
\log(W_2) &\sim N(-1, 2) \\
\log(W_3) &\sim N(-2, 1) \\
g_{0,1}(1 \mid Pa(A_0)) &= P_0(A_0 = 1 \mid Pa(A_0)) \\
&= \text{expit}(1 - 0.3W_1 + 0.1W_2 + 0.2W_3)
\end{aligned}$$

$$\begin{aligned}
g_{0,2}(1 \mid Pa(C_0)) &= P_0(C_0 = 0 \mid Pa(C_0)) \\
&= \text{expit}(0.1 + 0.2W_1 + 0.02W_1 + 0.01W_3 + 0.4A_0) \\
L_{1.1} &= \text{expit}(-2.5 + 0.1W_1 + 0.1W_3 + 0.4A_0) \\
L_{1.2} &= W_2 - W_2A_0 + 0.05\log(W_1) + 0.02W_2W_3 + \epsilon_1 \\
L_{1.3} &= 0.02W_1W_2 + \epsilon_2 \\
g_{0,3} &= P_0(C_1 = 0 \mid Pa(C_1)) \\
&= \text{expit}(-0.6 + 0.3W_2 + A_0 + 0.1L_{1.2} + 0.5L_{1.3}) \\
L_{2.1} &= \text{expit}(-1 + 0.01W_1 - A_0 + 0.1L_{1.2} - 0.2L_{1.3}) \\
L_{2.2} &= 0.01\log(W_1) + L_{1.2} + 0.02L_{1.2}L_{1.3} + \epsilon_3 \\
L_{2.3} &= 0.1\log(W_1) + 0.2L_{1.3}\epsilon_4 \\
g_{0,4} &= P_0(C_2 = 0 \mid Pa(C_2)) \\
&= \text{expit}(1.3 - 0.4A_0 - 0.2L_{1.3} + 0.3L_{2.2} - 0.2L_{2.3}) \\
P_0(Y = 1 \mid Pa(Y)) &= \text{expit}(-2 - A_0 - 0.1L_{1.3} + 0.7L_{2.2} + 0.8L_{2.3})
\end{aligned}$$

with $\epsilon_1, \epsilon_3 \sim_{i.i.d.} N(0, 0.5), \epsilon_2, \epsilon_4 \sim_{i.i.d.} N(0, 1)$. The outcome at time t_k is known for a subject who is observed to experience the outcome at time $t' < t_k$. This knowledge is encoded in the dataset by deterministically setting the values of subsequent censoring nodes to 0, and assigning the value 1 to all subsequent event nodes. When a subject is censored before an event is observed, all subsequent event nodes are set to 0 and subsequent censoring nodes are set to 1. The true parameter value is $\psi_0 = 0.348$, with an efficiency bound $\sigma_2^2 = 1.19/n$.

Results were obtained for correct and misspecified regression models for $\bar{Q}^a_{L(k)}$ and $g_{0,k}$. The conditional means $\bar{Q}^a_{L(k)}$ were estimated with logistic regression models including all terms used to generate the actual data (Q_c), including main term baseline covariates only (Q_{m_1}), and an intercept-only model (Q_{m_2}). The g factors were estimated by using correctly specified logistic regression models to regress A_k on the parents of A_k, and a second time, using main terms logistic regression models that included baseline covariates measured prior to A_k in the time ordering shown in Figure 2. Again, the censoring and treatment probabilities were not truncated from below.

Table 2: Simulation 2 results, $\psi_0 = 0.348$.

		$n = 100$				$n = 1000$			
		%Rel bias	Bias	Var	MSE	%Rel bias	Bias	Var	MSE
					g correctly specified				
	IPTW$_u$	3.46	0.012	0.020	0.020	3.61	0.013	0.001	0.001
	IPTW$_n$	−85.95	−0.299	0.001	0.090	−84.13	−0.293	0.000	0.086
Q_c	MLE$_p$	4.35	0.015	0.012	0.013	−1.86	−0.006	0.001	0.001
	TMLE	1.72	0.006	0.013	0.013	0.10	0.000	0.001	0.001
Q_{m_1}	MLE$_p$	5.20	0.018	0.013	0.013	−1.67	−0.006	0.001	0.001
	TMLE	3.48	0.012	0.014	0.014	2.21	0.008	0.001	0.001
Q_{m_2}	MLE$_p$	−5.65	−0.020	0.009	0.010	−5.40	−0.019	0.001	0.001
	TMLE	2.15	0.007	0.014	0.014	1.80	0.006	0.001	0.001
					g misspecified				
	IPTW$_u$	−4.02	−0.014	0.013	0.014	0.78	0.003	0.001	0.001
	IPTW$_n$	−95.89	−0.334	0.001	0.112	−94.33	−0.328	0.001	0.108
Q_c	MLE$_p$	4.35	0.015	0.012	0.013	−1.86	−0.006	0.001	0.001
	TMLE	−0.67	−0.002	0.014	0.014	0.35	0.001	0.001	0.001
Q_{m_1}	MLE$_p$	5.20	0.018	0.013	0.013	−1.67	−0.006	0.001	0.001
	TMLE	0.00	0.000	0.016	0.016	0.72	0.002	0.001	0.001
Q_{m_2}	MLE$_p$	−5.65	−0.020	0.009	0.010	−5.40	−0.019	0.001	0.001
	TMLE	8.69	0.030	0.022	0.023	0.90	0.003	0.001	0.001

Results: Table 2 lists the empirical bias, percent bias relative to the true parameter value, variance, and mean squared error (MSE) of the Monte Carlo estimates. IPTW$_u$ has larger MSE under correct specification of $g_{0,A(k)}$ than under misspecification when $n = 100$ due to sparsity in the data that inflates the variance. Normalizing the weights decreases this variance, but the corresponding increase in bias overshadows this gain (even at the larger sample size). TMLE remains stable, and the targeting step reduces the bias observed in the MLE$_p$ estimator when $Q_{0,L(k)}$ is misspecified and $g_{0,A(k)}$ is correctly specified. When sparsity is not an issue ($n = 1000$) even the misspecified g contains information that is helpful in reducing the bias.

5.3 Inference

Table 3 allows us to compare the empirical variance of the Monte Carlo TMLE estimates obtained above, with IC-based variance estimates, $\widehat{var}(\psi_n) = \hat{\sigma}^2_{IC}/n$, and

lists coverage of 95% IC-based confidence intervals. As an estimate of the influence curve we use the estimated efficient influence curve, which is known to be asymptotically correct if both Q_0 and g_0 are consistently estimated, and provides asymptotically conservative variance estimates if g_0 is consistently estimated. When g_0 is correctly specified sparsity in the data leads to anti-conservative confidence intervals. However when sample size is increased to 1000, observed coverage is close to the nominal rate. As predicted by theory, when both Q_0 and g_0 are misspecified, (efficient-)influence curve-based inference is not reliable. However, at the larger sample size coverage is close to the nominal rate, and the analytical variance estimate closely approximates the empirical variance. The somewhat counter-intuitive failure to achieve the desired coverage in simulation 2 with $n = 100$ and g misspecified, given that the mean estimated variance is enormous, is due to the fact that sparsity caused a single outlying estimates of the variance to blow up, while all others were better behaved. When this outlier is removed from the calculation, the mean estimated variances are 0.0092 (Q_c), 0.0087 (Q_{m_1}), and 0.011 (Q_{m_2}). Failure to attain the nominal coverage rate now makes sense, since these estimated variances are smaller under sparsity than the corresponding empirical variances.

6 Concluding remarks

TMLE is a general template for construction of semiparametric efficient substitution estimators that requires writing the target parameter as a function of an infinite dimensional parameter (e.g., $\Psi(\bar{Q}^a)$), a loss function for this parameter possibly indexed by a nuisance parameter (e.g., $\mathcal{L}_\eta(\bar{Q}^a)$), a parametric submodel with loss function-specific score spanning the efficient influence curve (and/or any other desired estimating function), and a specification of a resulting iterative targeted minimum loss-based estimation algorithm that minimizes the loss function-specific empirical risk along the parametric submodel until no further update improves the empirical risk. *Since the nuisance parameters in the loss function are a function of $\bar{Q}^a$ itself, the estimator of the nuisance parameters in the loss function are also updated at each step to reflect their last fitted value*. The TMLE is a two-stage procedure, where the first stage involves loss-based learning of the infinite dimensional parameter, and the subsequent stage is a targeted iterative update of this initial estimator that is only concerned with fitting the target parameter, and which guarantees that the TMLE of the infinite dimensional parameter solves the efficient influence curve equation. The influence curve of the TMLE is defined by the fact that it solves this estimating equation.

As apparent from a formal analysis of the TMLE, whether the conditions for asymptotic linearity are met depends on how well (e.g., at what rate) the TMLE

Table 3: Empirical variance of Monte Carlo estimates, mean IC-based variance estimates, and coverage of nominal 95% confidence intervals.

	$n = 100$			$n = 1000$		
	Emp var	$\bar{\hat{\sigma}}^2_{IC}/n$	Coverage	Emp var	$\bar{\hat{\sigma}}^2_{IC}/n$	Coverage
g correctly specified						
Simulation 1						
Q_c	0.0058	0.0030	0.82	0.0004	0.0004	0.96
Q_{m_1}	0.0078	0.0028	0.73	0.0004	0.0004	0.96
Q_{m_2}	0.0059	0.0051	0.88	0.0004	0.0006	0.98
Simulation 2						
Q_c	0.0130	0.0099	0.89	0.0011	0.0013	0.97
Q_{m_1}	0.0137	0.0100	0.87	0.0011	0.0014	0.97
Q_{m_2}	0.0143	0.0132	0.91	0.0011	0.0014	0.96
g misspecified						
Simulation 1						
Q_c	0.0060	0.0031	0.81	0.0004	0.0004	0.96
Q_{m_1}	0.0081	0.0027	0.73	0.0004	0.0004	0.94
Q_{m_2}	0.0065	0.0048	0.88	0.0005	0.0006	0.95
Simulation 2						
Q_c	0.0138	$4.2e+5$	0.87	0.0011	0.0012	0.95
Q_{m_1}	0.0164	$4.2e+5$	0.84	0.0011	0.0012	0.95
Q_{m_2}	0.0219	$4.2e+5$	0.80	0.0011	0.0012	0.95

estimates these nuisance parameters of the efficient influence curve. The latter also affects the finite sample performance of the TMLE. Therefore, if the initial estimator of the infinite dimensional parameter in the TMLE involves trading-off bias and variance w.r.t. an infinite dimensional parameter that is much richer than needed for evaluation of the target parameter, then finite sample performance is degraded relative to a TMLE that uses an initial estimator that involves trading-off bias and variance for a smaller infinite dimensional parameter that is more relevant for the target parameter. From this perspective, the TMLE proposed in this article, inspired by the double robust estimator of Bang and Robins (2005), appears to be based on an excellent choice of loss function and parametric submodel.

By the same token, a substitution estimator obtained by plugging in a log-likelihood based super learner will be less targeted than a substitution estimator obtained by plugging in a loss-based super learner based on a more targeted loss function. Therefore, loss-based learning provides fundamental improvements on

log-likelihood based learning by allowing the selection of a targeted loss function, and TMLE provides the additional bias reduction so that the resulting estimators allow for statistical inference in terms of a central limit theorem, under appropriate regularity conditions.

The TMLE we presented in this article can be generalized to any pathwise differentiable parameter of the distribution of Y^a, possibly conditional on $L(0)$ or $L^a(k)$ at a particular k (as in history adjusted marginal structural models), by applying the conditional iterative expectation rule to $(P(Y^a = y) : a, y)$ as in this article for EY^a, and applying the above TMLE framework with the decomposition of the efficient influence curve, the loss functions and submodels. Precise demonstrations for causal parameter defined by marginal structural models are presented in the Appendices below.

For future research it will also be of interest to develop a collaborative TMLE based on the TMLE presented here, thereby also allowing the targeted estimation of the intervention mechanism based on the collaborative double robustness of the efficient influence curve as presented in van der Laan and Gruber (2010) and van der Laan and Rose (2011).

Appendix

Introduction to Appendix

This appendix begins by presenting the proof of the representation of the efficient influence curve of Theorem 1 in Appendix A. In Appendix B we make some general remarks about sequential loss-based (machine) learning, which allows the incorporation of data adaptive estimators to obtain the initial estimator of the sequence of iteratively defined conditional means. In Appendix C we show that the TMLE of the sequence of iteratively defined conditional means is compatible with a probability distribution in the statistical model, showing that the proposed TMLE fully respects the global constraints of the statistical model. In Appendix D we provide the actual R code that was used to implement this TMLE in the presented simulations.

Additional appendices are available in a longer technical report version of this paper (van der Laan and Gruber, 2011). Appendix E of that paper demonstrates how the TMLE is quickly generalized for target parameters defined as a multivariate real valued function of a collection of intervention specific means EY^a indexed by interventions $a \in \mathcal{A}$. In particular, this demonstrates the TMLE of the unknown coefficients of a working marginal structural model. In Appendix F of the technical report we further generalize this TMLE to also apply to working marginal structural models for the causal curve $E(Y^a \mid V)$ in a, conditional on some user

supplied baseline covariate V. The general roadmap of TMLE is apparent from these applications, which also demonstrates how this TMLE can be developed for all other causal parameters of interest, such as history adjusted marginal structural working models, dynamic treatments, and so on.

A Proof of Theorem 1

The formula for D^*_{K+1} is obvious. Note,

$$\begin{aligned}
&D^*_K = E(D \mid L(K), \bar{A}(K-1), \bar{L}(K-1)) - E(D \mid \bar{A}(K-1), \bar{L}(K-1)) \\
&= \frac{I(\bar{A}(K-1)=\bar{a}(K-1)}{g_{0:K-1}} \left\{ E\Big(\frac{YI(A(K)=a(K))}{g_K} \,|\, L(K), \bar{A}(K-1) = \bar{a}(K-1), \bar{L}(K-1) \Big) \right. \\
&\qquad \left. - E\left(\frac{YI(A(K)=a(K))}{g_K} \mid \bar{A}(K-1) = \bar{a}(K-1), \bar{L}(K-1) \right) \right\}.
\end{aligned}$$

Note also that

$$\begin{aligned}
&E(YI(A(K)=a(K))/g_K \mid L(K), \bar{A}(K-1) = \bar{a}(K-1), \bar{L}(K-1)) \\
&= E(E(Y \mid \bar{L}(K), A(K), \bar{A}(K-1)) \frac{I(A(K)=a(K))}{g_K} \mid \bar{L}(K), \bar{A}(K-1) = \bar{a}(K-1)) \\
&= E(\bar{Q}^a_Y(\bar{L}(K)) I(A(K)=a(K))/g_K \mid L(K), \bar{A}(K-1) = \bar{a}(K-1), \bar{L}(K-1)) \\
&= E(\bar{Q}^a_Y(\bar{L}(K)) \mid L(K), \bar{A}(K-1) = \bar{a}(K-1), \bar{L}(K-1)) \\
&= \bar{Q}^a_Y(\bar{L}(K)).
\end{aligned}$$

Thus,

$$E(YI(A(K)=a(K))/g_K \mid \bar{A}(K-1) = \bar{a}(K-1), \bar{L}(K-1)) = E_{Q^a_{L(K)}} \bar{Q}^a_Y.$$

Thus, we found the following representation of D_K:

$$D_K = \frac{I(\bar{A}(K-1) = \bar{a}(K-1))}{g_{0:K-1}} \left\{ \bar{Q}^a_Y - E_{Q^a_{L(K)}} \bar{Q}^a_Y \right\}.$$

Consider now

$$\begin{aligned}
&D_{K-1} = E(D \mid L(K-1), \bar{A}(K-2), \bar{L}(K-2)) - E(D \mid \bar{A}(K-2), \bar{L}(K-2)) \\
&= \frac{I(\bar{A}(K-2)=\bar{a}(K-2))}{g_{0:K-2}} \Big\{ E(Y \frac{I(A(K)=a(K), A(K-1)=a(K-1))}{g_{K-1:K}} \mid L(K-1), \bar{A}(K-2), \bar{L}(K-2)) \\
&\qquad - E(YI(A(K)=a(K), A(K-1)=a(K-1))/g_{K-1:K} \mid \bar{A}(K-2), \bar{L}(K-2)) \Big\}.
\end{aligned}$$

Note that

$$\begin{aligned}
&E(Y \frac{I(A(K)=a(K), A(K-1)=a(K-1))}{g_{K-1:K}} \mid L(K-1), \bar{A}(K-2) = \bar{a}(K-2), \bar{L}(K-2)) \\
&= E(Y^a \mid L(K-1), \bar{A}(K-1) = \bar{a}(K-1), \bar{L}(K-2)) \\
&= E(Y^a \mid \bar{L}^a(K-1)) \\
&= \bar{Q}^a_{L(K)}.
\end{aligned}$$

This shows

$$D_{K-1} = \frac{I(\bar{A}(K-2) = \bar{a}(K-2))}{g_{0:K-2}} \left\{ \bar{Q}^a_{L(K)} - E_{Q^a_{L(K-1)}} \bar{Q}^a_{L(K)} \right\}.$$

In general, for $k = 1, \ldots, K+1$, we have

$$\begin{aligned}
&D_k = E(D \mid L(k), \bar{A}(k-1), \bar{L}(k-1)) - E(D \mid \bar{A}(k-1), \bar{L}(k-1)) \\
&= \frac{I(\bar{A}(k-1)=\bar{a}(k-1))}{g_{0:k-1}} \left\{ E(Y^a \mid L(k), \bar{A}(k-1), \bar{L}(k-1)) \right. \\
&\left. -E(Y^a \mid \bar{A}(k-1), \bar{L}(k-1)) \right\} \\
&= \frac{I(\bar{A}(k-1)=\bar{a}(k-1))}{g_{0:k-1}} \left\{ \bar{Q}^a_{L(k+1)} - E_{Q^a_{L(k)}} \bar{Q}^a_{L(k+1)} \right\} \\
&= \frac{I(\bar{A}(k-1)=\bar{a}(k-1))}{g_{0:k-1}} \left\{ \bar{Q}^a_{L(k+1)} - \bar{Q}^a_{L(k)} \right\}.
\end{aligned}$$

Finally,

$$\begin{aligned}
&D_0 = E(D \mid L(0)) = E(Y^a \mid L(0)) - \Psi(\bar{Q}^a) = \bar{Q}^a_{L(1)} - E_{Q_{L(0)}} \bar{Q}^a_{L(1)} \\
&= \bar{Q}^a_{L(1)} - \bar{Q}^a_{L(0)}.
\end{aligned}$$

⊠

B Sequential Super Learning

For each of these sequential regressions we could employ a super learning algorithm (van der Laan et al. (2007) and Chapter 3 in van der Laan and Rose (2011) based on Polley and van der Laan (2010)), which is defined in terms of a library of candidate estimators and it uses cross-validation to select among these candidate estimators. For that purpose it is appropriate to review the cross-validation selector among candidate estimators based on a loss function with a nuisance parameter, as originally presented and studied in van der Laan and Dudoit (2003). Consider the loss function $\mathcal{L}_{\bar{Q}^a_{k+1}}(\bar{Q}^a_k)$ for $\bar{Q}^a_{k,0}$ with nuisance parameter $\bar{Q}^a_{k+1}$. Given an estimator $\hat{\bar{Q}}^a_{k+1}$ of the nuisance parameter, given a candidate estimator $\hat{\bar{Q}}^a_k$ of $\bar{Q}^a_{k,0}$ (or, more precisely, $E_{Q_{L(k),0}} \bar{Q}^a_{k+1,n}$), the cross-validated risk of this candidate estimator is evaluated as

$$E_{B_n} P^1_{n,B_n} \mathcal{L}_{k,\hat{\bar{Q}}^a_{k+1}(P^0_{n,B_n})}(\hat{\bar{Q}}^a_k(P^0_{n,B_n})).$$

Here $B_n \in \{0,1\}^n$ is a cross-validation scheme splitting the sample of n observations in a training sample $\{i : B_n(i) = 0\}$ and validation sample $\{i : B_n(i) = 1\}$, and P^1_{n,B_n}, P^0_{n,B_n} are the corresponding empirical distributions. Typically, we select B_n to correspond with V-fold cross-validation by giving it a uniform distribution

on V vectors with np 1's and $n(1-p)$ 0's. Thus, in this cross-validated risk the nuisance parameter is estimated with the previously selected estimator, but applied to the training sample within each sample split. In particular, given a set of candidate estimators $\hat{\bar{Q}}^a_{k,j}$ of $\bar{Q}^a_{k,0}$ indexed by $j = 1, \ldots, J$, the cross-validation selector is given by

$$J_n \equiv \underset{j}{\operatorname{argmin}}\ E_{B_n} P^1_{n,B_n} \mathcal{L}_{k,\hat{\bar{Q}}^a_{k+1}(P^0_{n,B_n})}(\hat{\bar{Q}}^a_{k,j}(P^0_{n,B_n})).$$

Given the cross-validation selector J_n, one would estimate $\bar{Q}^a_{k,0}$ with $\hat{\bar{Q}}^a_{k,J_n}(P_n)$. (Note that the latter represents now an estimator $\hat{\bar{Q}}^a_k$ of the nuisance parameter $\bar{Q}^a_k$ in the loss function of the next parameter $\bar{Q}^a_{k-1}$, and the same cross-validation selector can now be employed.) The oracle inequality for the cross-validation selector in van der Laan and Dudoit (2003) applies to this cross-validation selector J_n. However, specific theoretical study of the resulting estimator of (e.g.) $\bar{Q}^a_{L(0)}$ based on the sequential cross-validation procedure (given collections of candidate estimators $\hat{\bar{Q}}^a_{k,j}$, $j = 1, \ldots, J_k$, $k = K+1, \ldots, 1$) described above is warranted and is an area for future research.

To save computer time, one could decide to estimate the nuisance parameters in these loss functions with the selected estimator based on the whole sample. We suggest that this may not harm the practical performance of the cross-validation selector, but this remains to be investigated.

C The TMLE is a substitution estimator

In this section we demonstrate that the TMLE $(\bar{Q}^{a,*}_{k,n} : k = 0, \ldots, K+1)$ is compatible with a probability distribution $\mathbf{P}_n \in \mathcal{M}$ of P_0, which then shows that the TMLE presented in this article is indeed a substitution estimator, and thereby respects all global constraints of the statistical model $\mathcal{M}$. This is shown under a "range" condition specified below.

Specifically, we will show that there exists a sequence $(Q^a_{k,n} : k = 0, \ldots, K+1)$, where $Q^a_{k,n}$ is a conditional distribution of $L(k)$, given $\bar{L}(k-1)$, $\bar{A}(k-1) = \bar{a}(k-1)$, so that the corresponding iteratively defined sequence of conditional means equals $(\bar{Q}^{a,*}_{k,n} : k = 0, \ldots, K+1)$. This then shows that all sequential regressions are compatible with any probability distribution whose conditional distributions of $L(k)$, given $\bar{A}(k-1) = \bar{a}(k-1), \bar{L}(k-1)$ coincide with these conditional distributions $(Q^a_{k,n} : k)$. A condition is that, for each k, given $\bar{L}(k-1)$, (the next regression) $\bar{Q}^{a,*}_{k,n}(\bar{L}(k-1))$ is in the range of (previous regression) $L(k) \rightarrow \bar{Q}^{a,*}_{k+1,n}(L(k), \bar{L}(k-1))$ over a support of $L(k)$.

Our fit $\bar{Q}^{a,*}_{K+1,n}$ of the conditional mean of Y, given $\bar{A}(K) = \bar{a}(K), \bar{L}(K)$, is obviously compatible with a conditional distribution $Q^a_{K+1,n}$ of Y, given $\bar{A}(K) = \bar{a}(K), \bar{L}(K)$, where we use that our TMLE fit of the conditional mean is a bounded function with values in $(0, 1)$. Let $h(L^a(K) \mid \bar{L}^a(K-1)) = E(Y^a \mid L^a(K), \bar{L}^a(K-1))$ be this first conditional regression function, where the expectation is under $Q^a_{K+1,n}$. Let $\mu(\bar{L}^a(K-1)) = \bar{Q}^{a,*}_{K,n}$ be the next regression of our TMLE.

Given $\bar{L}^a(K-1)$, the function h of $L^a(K)$ maps into $(0, 1)$, and μ is a scalar in $(0, 1)$. Suppose that we have shown that for any given function h of $L^a(K)$ mapping into $(0, 1)$, and any scalar $\mu \in (0, 1)$, there exists a distribution of $L^a(K)$ so that the expectation of $h(L^a(K))$ equals μ: this result can then be applied, conditional on each value $\bar{L}^a(K-1)$. Here it will be assumed that μ is in the range of h.

We have then shown that, given the previous regression $h(L^a(K) \mid \bar{L}^a(K-1))$ and the current regression $\mu(\bar{L}^a(K-1))$, there exists a conditional distribution $Q^a_{K,n}$ of $L^a(K)$, given $\bar{L}^a(K-1)$ (i.e., a conditional distribution of $L(K)$, given $\bar{L}(K-1), \bar{A}(K-1) = \bar{a}(K-1)$) for which the mean of $h(L^a(K) \mid \bar{L}^a(K-1))$ w.r.t. this $Q^a_{k,n}$ equals $\mu(\bar{L}^a(K-1))$. This now proves that the first two conditional means $\bar{Q}^{a,*}_{K+1,n}$ and $\bar{Q}^{a,*}_{K,n}$ are compatible with two conditional distributions $Q^a_{K+1,n}$ and $Q^a_{K,n}$ of $L^a(K+1)$, given $\bar{L}^a(K)$, and $L^a(K)$, given $\bar{L}^a(K-1)$, respectively. By induction (i.e, we apply this same proof to the next pair of conditional means, and so forth), we have now shown the desired result that there exists conditional distributions $(Q^a_{k,n} : k)$ compatible with $(\bar{Q}^{a,*}_{k,n} : k)$, which completes the proof.

It remains to show the desired result. We show the result for the case that all random variables are discrete. That is, we need to show that for a $h(x)$ mapping in $(0, 1)$ and $\mu \in (0, 1)$, we can define a distribution of X such that the mean of $h(X)$ equals μ. First define a distribution of $Z = h(X)$ that has mean μ. Now, we note that $P(Z = z) = P(X \in h^{-1}(z))$. Thus this specified probability mass on $Z = z$ implies that X puts the same probability mass on the set $h^{-1}(z)$, across z (note all the sets $h^{-1}(z)$ across z are disjoint). This defines a non-empty set of possible distributions of X that are compatible with this distribution of Z.

As is apparent from this proof, there are in fact plenty of conditional distributions compatible with our sequential regressions, which fits our intuition since a mean is far from sufficient for identifying a whole distribution. The only condition in this proof was that the next regression fit of the previous conditional regression has to be in the range of the previous conditional regression, which would be achieved by any reasonable estimation procedure: i.e., one likes the estimate of the mean of a random variable to be in the set of possible values of that random variable. This condition is obviously satisfied by estimating the conditional mean under $L(k)$, given $\bar{L}(k-1), \bar{A}(k-1) = \bar{a}(k-1)$, by plugging in an estimator of the

conditional density. In Gruber and van der Laan (2010a) it is shown how one can construct a regression that is fully respecting the range of the outcome, by using the logistic link function. In particular, this method can be used to enforce conditional bounds on the outcome, given the conditioning variables. In this manner, we can enforce each of the sequential regressions to satisfy any known bounds on the relevant outcome as implied by the statistical model for the conditional distributions that make up the likelihood. We plan to present these methods in more detail in a future article.

D R code

The functions below implement TMLE, IPTW, and MLE_p estimators of the treatment-specific mean outcome for the R statistical programming environment (R Development Core Team, 2010). Modifications to the functions *estg* and *estQ* would allow for pooling data across time points to estimate conditional means, data adaptive estimation, and other customizations.

```
#-----------------------------------------------------------------
# function: getEstimates
#   IPTW, MLE, TMLE for intervention-specific mean outcome
# arguments:
#   d: dataset, wide format, using time-ordering of the nodes
#   Inodes: treatment and censoring node columns in d
#   Lnodes: time-dependent covariate and outcome columns in d
#   Ynodes: intermediate and final event node columns
#                   (a subset of Lnodes)
#   Qform: regression formulas for Q_1 through Q_K+1
#   gform: regression formulas for each tmt and censoring event
#   gbd:   lower bound on estimated probabilities for g-factors
#   family: regression family for initial estimates of Q nodes
# encoding binary intervention nodes:
#   treatment: A(t) = 1, subject  following tmt regime at time t
#   censoring: A.C(t)  = 1, subject observed at time t
#-----------------------------------------------------------------
getEstimates <- function(d, Inodes, Lnodes, Ynodes, Qform,
     gform, gbd = 0,  family = "quasibinomial")
{
     n <- nrow(d)
     n.Q <- length(Lnodes)
     n.g <- length(Inodes)
     g1W <- estg(d, gform, Inodes, Ynodes)
     cum.g1W <- bound(t(apply(g1W, 1, cumprod)), c(gbd,1))
     empirical.meanwt <- mean(1/cum.g1W[,n.g], na.rm=TRUE)
     cum.g1W[is.na(cum.g1W)] <- Inf
```

```
     iptw <- mean(d[,Lnodes[n.Q]] * d[,Inodes[n.g]]/cum.g1W[,n.g])
     wt.n <- 1 / cum.g1W[,n.g] / empirical.meanwt
     iptw.wt.n <- mean(d[,Lnodes[n.Q]] * d[,Inodes[n.g]] * wt.n)
     # Gcomp and TMLE
     Qstar <- Qinit <- d[, Lnodes[n.Q]]
     IC <- rep(0, n)
     for (i in n.Q:1){
          Inode.cur  <- which.max(Inodes[Inodes < Lnodes[i]])
          uncensored <- d[,Inodes[Inode.cur]] == 1

          if(any(Ynodes < Lnodes[i])){
               Ynode.prev <- max(Ynodes[Ynodes < Lnodes[i]])
               deterministic <- d[,Ynode.prev]==1
          } else {
               deterministic <- rep(FALSE, n)
          }
          Qinit <- estQ(Qinit, d, Qform[i], uncensored,
                    deterministic, family = family)
          Qstar.kplus1 <- Qstar
          Qstar <- estQ(Qstar.kplus1, d, Qform[i], uncensored,
                    deterministic, h = 1/cum.g1W[,Inode.cur],
                    family = family)
          IC[uncensored] <- (IC + (Qstar.kplus1 - Qstar)/
                    cum.g1W[,Inode.cur])[uncensored]
     }
     IC <- IC + Qstar - mean(Qstar)
     return(c(iptw = iptw, iptw.wt.n=iptw.wt.n, Gcomp=mean(Qinit),
              tmle = mean(Qstar), var.tmle = var(IC)/n))
}

# Utility functions
#---------------------------------------------------------------------
# function: estQ
# purpose: parametric estimation of Q_k, targeted if h is supplied
#---------------------------------------------------------------------
estQ <- function(Q.kplus1, d, Qform, uncensored, deterministic,
     h=NULL, family){
     Qform <- update.formula(Q.kplus1 ~ .)
     m <- glm(as.formula(Qform),
        data=data.frame(Q.kplus1,d)[uncensored & !deterministic,],
        family = family)
     Q1W <- predict(m, newdata = d, type = "response")
     if(!is.null(h)){
          off <- qlogis(bound(Q1W, c(.0001, .9999)))
          m <- glm(Q.kplus1 ~ -1 + h + offset(off),
               data = data.frame(Q.kplus1, h, off),
               subset = uncensored&!deterministic,
               family ="quasibinomial")
```

```
          Q1W  <- plogis(off + coef(m)*h)
     }
     Q1W[deterministic] <- 1
     return(Q1W)
}

#----------------------------------------------------------------
# function: estg
# purpose: parametric estimation of each g-factor
#----------------------------------------------------------------
estg <- function(d, form, Inodes, Ynodes){
     n <- nrow(d)
     n.g <- length(form)
     gmat <- matrix(NA, nrow = n, ncol = n.g)
     uncensored <- rep(TRUE, n)
     deterministic <- rep(FALSE, n)
     for (i in 1:n.g) {
          if(any(Ynodes < Inodes[i])){
               Ynode.prev <- max(Ynodes[Ynodes < Inodes[i]])
               deterministic <- d[,Ynode.prev] == 1
          }
          m <- glm(as.formula(form[i]), data = d,
               subset = uncensored & !deterministic,
                family = "binomial")
          gmat[uncensored,i] <- predict(m, newdata=d[uncensored,],
               type = "response")
          gmat[deterministic,i] <- 1
          uncensored <- d[,Inodes[i]] == 1
     }
     return(gmat)
}

#----------------------------------------------------------------
# function: bound
# purpose: truncate values within supplied bounds
#----------------------------------------------------------------
bound <- function(x, bounds){
     x[x < min(bounds)] <- min(bounds)
     x[x > max(bounds)] <- max(bounds)
     return(x)
}
```

References

H. Bang and J.M. Robins. Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61:962–972, 2005.

O. Bembom and M.J. van der Laan. A practical illustration of the importance of realistic individualized treatment rules in causal inference. *Electron J Stat*, 1: 574–596, 2007.

O. Bembom, M.L. Petersen, S.-Y. Rhee, W.J. Fessel, S.E. Sinisi, R.W. Shafer, and M.J. van der Laan. Biomarker discovery using targeted maximum likelihood estimation: application to the treatment of antiretroviral resistant HIV infection. *Stat Med*, 28:152–72, 2009.

P.J. Bickel, C.A.J. Klaassen, Y. Ritov, and J. Wellner. *Efficient and Adaptive Estimation for Semiparametric Models*. Springer-Verlag, 1997.

S. Gruber and M.J. van der Laan. A targeted maximum likelihood estimator of a causal effect on a bounded continuous outcome. *International Journal of Biostatistics*, 6:article 26, www.bepress.com/ijb/vol6/iss1/26, 2010a.

S. Gruber and M.J. van der Laan. An application of collaborative targeted maximum likelihood estimation in causal inference and genomics. *Int J Biostat*, 6(1), 2010b.

M.A. Hernan, B. Brumback, and J.M. Robins. Marginal structural models to estimate the causal effect of zidovudine on the survival of HIV-positive men. *Epidemiol*, 11(5):561–570, 2000.

K.L. Moore and M.J. van der Laan. Application of time-to-event methods in the assessment of safety in clinical trials. In Karl E. Peace, editor, *Design, summarization, analysis & interpretation of clinical trials with time-to-event endpoints*, Boca Raton, 2009a. Chapman & Hall.

K.L. Moore and M.J. van der Laan. Covariate adjustment in randomized trials with binary outcomes: targeted maximum likelihood estimation. *Stat Med*, 28(1): 39–64, 2009b.

K.L. Moore and M.J. van der Laan. Increasing power in randomized trials with right censored outcomes through covariate adjustment. *J Biopharm Stat*, 19(6): 1099–1131, 2009c.

J. Pearl. Causal diagrams for empirical research. *Biometrika*, 82:669–710, 1995.

J. Pearl. *Causality: Models, reasoning, and inference*. New York: Cambridge University Press, 2000.

E.C. Polley and M.J. van der Laan. Predicting optimal treatment assignment based on prognostic factors in cancer patients. In K.E. Peace, editor, *Design, summarization, analysis & interpretation of clinical trials with time-to-event endpoints*, Boca Raton, 2009. Chapman & Hall.

E.C. Polley and M.J. van der Laan. Super learner in prediction. Technical Report 266, Division of Biostatistics, University of California, Berkeley, 2010.

R Development Core Team. *R: a language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, 2010. URL `http://www.R-project.org`.

J.M. Robins. Marginal structural models versus structural nested models as tools for causal inference. In *Statistical models in epidemiology: the environment and clinical trials*. Springer, Berlin Heidelberg New York, 1999.

J.M. Robins. Robust estimation in sequentially ignorable missing data and causal inference models. In *Proceedings of the American Statistical Association*, 2000.

J.M. Robins and A. Rotnitzky. Recovery of information and adjustment for dependent censoring using surrogate markers. In *AIDS epidemiology*. Birkhäuser, Basel, 1992.

J.M. Robins and A. Rotnitzky. Comment on the Bickel and Kwon article, "Inference for semiparametric models: some questions and an answer". *Stat Sinica*, 11(4): 920–936, 2001.

J.M. Robins, A. Rotnitzky, and M.J. van der Laan. Comment on "On profile likelihood". *J Am Stat Assoc*, 450:431–435, 2000.

S. Rose and M.J. van der Laan. Simple optimal weighting of cases and controls in case-control studies. *Int J Biostat*, 4(1):Article 19, 2008.

S. Rose and M.J. van der Laan. Why match? Investigating matched case-control study designs with causal effect estimation. *Int J Biostat*, 5(1):Article 1, 2009.

S. Rose and M.J. van der Laan. A targeted maximum likelihood estimator for two-stage designs. *Int J Biostat*, 7(17), 2011.

M. Rosenblum and M.J. van der Laan. Targeted maximum likelihood estimation of the parameter of a marginal structural model. *Int J Biostat*, 6(2):Article 19, 2010.

M. Rosenblum, S.G. Deeks, M.J. van der Laan, and D.R. Bangsberg. The risk of virologic failure decreases with duration of HIV suppression, at greater than 50% adherence to antiretroviral therapy. *PLoS ONE*, 4(9): e7196.doi:10.1371/journal.pone.0007196, 2009.

D.O. Scharfstein, A. Rotnitzky, and J.M. Robins. Adjusting for nonignorable drop-out using semiparametric nonresponse models, (with discussion and rejoinder). *J Am Stat Assoc*, 94:1096–1120 (1121–1146), 1999.

O.M. Stitelman and M.J. van der Laan. Collaborative targeted maximum likelihood for time-to-event data. *Int J Biostat*, 6(1):Article 21, 2010.

O.M. Stitelman and M.J. van der Laan. Targeted maximum likelihood estimation of time-to-event parameters with time-dependent covariates. Technical Report, Division of Biostatistics, University of California, Berkeley, 2011a.

O.M. Stitelman and M.J. van der Laan. Targeted maximum likelihood estimation of effect modification parameters in survival analysis. *Int J Biostat*, 7(1), 2011b.

M.J. van der Laan. Estimation based on case-control designs with known prevalance probability. *Int J Biostat*, 4(1):Article 17, 2008.

M.J. van der Laan. Targeted maximum likelihood based causal inference: Part I. *Int J Biostat*, 6(2):Article 2, 2010.

M.J. van der Laan and S. Dudoit. Unified cross-validation methodology for selection among estimators and a general cross-validated adaptive epsilon-net estimator: finite sample oracle inequalities and examples. Technical Report 130, Division of Biostatistics, University of California, Berkeley, 2003.

M.J. van der Laan and S. Gruber. Collaborative double robust penalized targeted maximum likelihood estimation. *Int J Biostat*, 6(1), 2010.

M.J. van der Laan and S. Gruber. Targeted maximum loss based estimation of an intervention specific mean. Technical report 290, Division of Biostatistics, University of California, Berkeley, 2011.

M.J. van der Laan and J.M. Robins. *Unified methods for censored longitudinal data and causality*. Springer, Berlin Heidelberg New York, 2003.

M.J. van der Laan and S. Rose. *Targeted Learning: Causal Inference for Observational and Experimental Data*. Springer, Berlin Heidelberg New York, 2011.

M.J. van der Laan and Daniel B. Rubin. Targeted maximum likelihood learning. *Int J Biostat*, 2(1):Article 11, 2006.

M.J. van der Laan, E.C. Polley, and A.E. Hubbard. Super learner. *Stat Appl Genet Mol*, 6(1):Article 25, 2007.

H. Wang, S. Rose, and M.J. van der Laan. Finding quantitative trait loci genes with collaborative targeted maximum likelihood learning. *Stat Prob Lett*, published online 11 Nov (doi: 10.1016/j.spl.2010.11.001), 2010.

W. Zheng and M.J. van der Laan. Asymptotic theory for cross-validated targeted maximum likelihood estimation. Technical Report 273, Division of Biostatistics, University of California, Berkeley, 2010.

W. Zheng and M.J. van der Laan. Targeted maximum likelihood estimation of natural direct effect. Technical Report 288, Division of Biostatistics, University of California, Berkeley, 2011.