University of California, Berkeley
U.C. Berkeley Division of Biostatistics Working Paper Series

Year 2011 *Paper* 290

Targeted Minimum Loss Based Estimation of an Intervention Specific Mean Outcome

Mark J. van der Laan* Susan Gruber†

*Division of Biostatistics, School of Public Health, University of California, Berkeley, laan@berkeley.edu

†Division of Biostatistics, School of Public Health, University of California, Berkeley, sgruber65@yahoo.com

http://biostats.bepress.com/ucbbiostat/paper290

Targeted Minimum Loss Based Estimation of an Intervention Specific Mean Outcome

Mark J. van der Laan and Susan Gruber

Abstract

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)). TMLE requires 1) writing the target parameter as a particular mapping from a typically infinite dimensional parameter of the probability distribution of the unit data structure into the parameter space, 2) computing the canonical gradient/efficient influence curve of the pathwise derivative of the target parameter mapping, 3) specifying a loss function for this parameter that is possibly indexed by unknown "nuisance" parameters, 4) a least favorable parametric submodel/path through an initial/current estimator of the parameter chosen so that the linear span of the generalized loss-based score at zero fluctuation includes the efficient influence curve, and 5) an updating algorithm involving the iterative minimization of the loss-specific empirical risk over the fluctuation parameters of the least favorable parametric submodel/path. By the generalized loss-based score condition 4) on the submodel and loss function, it follows that the resulting estimator of the infinite dimensional parameter solves the efficient influence curve (i.e., efficient score) equation, providing the basis for the double robustness and asymptotic efficiency of the corresponding substitution estimator of the target parameter obtained by plugging in the updated estimator of the infinite dimensional parameter in the target parameter mapping.

To enhance the finite sample performance of the TMLE of the target parameter, it is of interest to choose the parameter and the nuisance parameter of the

loss function as low dimensional as possible. Inspired by this goal, we present a particular closed form TMLE of an intervention specific mean outcome based on general longitudinal data structures. %We also present its generalization of this type of TMLE to other causal parameters. This TMLE provides an alternative to the closed form TMLE presented in van der Laan and Gruber (2010) and Stitelman and vanderLaan (2011) based on the log-likelihood loss function. 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.

1 Introduction.

Many studies generate data sets that can be represented as n independent and identically distributed observations on a specified longitudinal data structure. By specifying a causal graph (Pearl (1995), Pearl (2000)), or equivalently, a system of structural equations specifying the observed variables as a function of a set of observed parent variables and an unmeasured exogenous error term, one codes the assumptions needed to be able to define a post-intervention distribution of this longitudinal structure that represents the distribution the data would have had under a specified intervention on a subset of the nodes defining the observed longitudinal data structure. Causal effects are defined as parameters of a collection of post intervention distributions.

A current and important topic is the estimation of causal effects of setting the value of multiple time point intervention-nodes on some final outcome of interest based on observing n independent and identically distributed copies of a longitudinal data structure. In particular, one might be concerned with estimation of the mean of the outcome under the post-intervention distribution for a specified multiple time point intervention. Under a causal graph and a so called sequential randomization and positivity assumption, one can identify the latter by the so called G-computation formula which maps the distribution of the observed longitudinal data structure on the experimental unit into the post-intervention distribution of the outcome. In this article we consider estimation of this intervention specific mean outcome in a semiparametric model that only makes statistical assumptions about the intervention mechanism, where the latter is defined by the conditional distribution of the intervention node, given the parent nodes of the intervention node, across the intervention nodes.

Different type of estimators of the intervention specific mean outcome in such a semiparametric model 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 such as the augmented IPTW estimating equation, maximum likelihood based G-computation estimators based on parametric models or data adaptive loss-based learning algorithms, and targeted maximum likelihood (or more general, minimum loss-based) estimators defined in terms of an initial estimator, loss function and least favorable fluctuation submodel through an initial or current estimator that is used to iteratively update the initial estimator till convergence. 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 are 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 proposed in Bang and Robins (2005) into the framework of targeted minimum loss based estimation. The resulting estimator 1) incorporates data adaptive estimation in place of parametric models, 2) can be applied to parameters for which there exists no mapping of the efficient influence curve into an estimating equation, thus also avoiding the potential problem of estimating equations having no or multiple solutions, and 3) has flexibility to incorporate 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). This results in a new TMLE based on a loss function that 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): see our discussion for more details on this. We generalize this new TMLE to causal parameters defined by projections on working marginal structural models.

This article is organized as follows. In Section 2 we define 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. We show that the target parameter can be defined as a function of an iteratively defined sequence of conditional means of the outcome under the distribution specified by the G-computation formula, one for each intervention node. In Section 2 we also derive 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. 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 5. A generalization of the TMLE for causal parameters defined by working marginal structural models is presented in the Appendix. The Appendix also provides R-code implementing the newly proposed TMLE.

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

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, $L(j)$ is an intermediate covariate realized after $A(j-1)$ and before $A(j)$, $j = 0, \ldots, K$, and Y is a final outcome of interest.

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

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

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. Here we used the notation $\bar{L}(k) = (L(0), \ldots, L(k))$. Note also that $Q_{0,L(k)}$ denotes the conditional distribution of $L(k)$, given $Pa(L(k))$, and, $g_{0,A(k)}$ denotes the 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 the random variable with probability distribution P^a, and let Y^a be its final component.

Our statistical target parameter is the mean of Y^a: $\Psi(P) = E_{P^a}Y^a$, where $\Psi : \mathcal{M} \to \mathbb{R}$. 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 NPSEM $L(k) = f_{L(k)}(Pa(L(k)), U_{L(k)}), A(k) = f_{A(k)}(Pa(A(k)), U_{A(k)})$ in terms of a set of functions $(f_{L(k)} : k = 0, \ldots, K+1), (f_{A(k)} : k = 0, \ldots, K)$, and an exogenous vector of 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. The probability distribution of this counterfactual 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_0^a defined by the true distribution P_0 of O under this system. 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 $EY_{\bar{a}}$ implies a statistical model $\mathcal{M}$ defined as the set of possible probability distribution P of O, and a statistical target parameter $\Psi : \mathcal{M} \to \mathbb{R}$. For the sake of estimation 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.

2.1 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 into the real line defined as follows. Compute $\bar{Q}_Y^a = E_{Q_Y^a}Y \equiv E(Y \mid \bar{A}(K) = \bar{a}(K), \bar{L}(K))$ by computing the integral of Y with respect to (w.r.t.) conditional distribution Q_Y^a of Y, given $\bar{L}(K), \bar{A}(K) = \bar{a}(K)$. Given $\bar{Q}_Y^a$, we compute $\bar{Q}_{L(K)}^a = E_{Q_{L(K)}^a}\bar{Q}_Y^a$, obtained by integrating out $L(K)$ in $\bar{Q}_Y^a$ w.r.t. the conditional distribution $Q_{L(K)}^a$ of $L(K)$, given $\bar{L}(K-1), \bar{A}(K-$

1) = $\bar{a}(K-1)$. This process is iterated: Given $\bar{Q}^a_{L(k)}$, we 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 till 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: $\bar{Q}^a_{0,L(0)} = E_0 Y^a = \Psi(Q_0)$.

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.

2.2 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 $\hat{\Psi}$ (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., $\hat{\Psi}(P_n) - \Psi(P_0) = 1/n \sum_{i=1}^n D^*(P_0)(O_i) + o_P(1/\sqrt{n})$. 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 (**?**). 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. This form of the efficient influence curve was established in Bang and Robins (2005).

Theorem 1 *Let $D(Q,g)(O) = Y\frac{I(\bar{A}(K)=\bar{a}(K))}{g_{0:K}} - \Psi(Q)$. 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 functions on Q, g. 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 mean $\bar{Q}^a_{L(k)}$ under $Q^a_{L(k)}$, and $g_{0:k-1}$.

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

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

Note also that

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

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

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

Note that

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

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

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

Finally,

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

□

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 $g_K(\bar{A}(K) = \bar{a}(K), \bar{L}(K)) > 0$ 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.

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 $\bar{Q}^a_0 = \arg\min_{\bar{Q}^a} P_0\mathcal{L}_{\eta_0}(\bar{Q}^a)$, 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 $P = Q * g$. 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 the 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 $\epsilon_{K,n} = \arg\min_\epsilon P_n \mathcal{L}(\bar{Q}^a_{L(K+1),n}(\epsilon, g_n))$. It iterates this updating process going backwards till 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 $\epsilon_{0,n} = \arg\min_\epsilon P_n \mathcal{L}_{\bar{Q}^{a,*}_{L(1),n}}(\bar{Q}^a_{L(0),n}(\epsilon, g_n))$ 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 describing the template for construction of the TMLE, we first present the summary of the practical implementation of the proposed TMLE.

3.1 Summary of practical implementation of TMLE.

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 can assume that $Y \in [0, 1]$. A special case would be that Y is binary valued with values in $\{0, 1\}$. Firstly, we carry out a logistic regression regressing 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. Let g_n be an estimator of g_0. Subsequently, we use this initial estimator of $\bar{Q}^a_{Y,0} = E_0(Y \mid \bar{A}(K) = \bar{a}(K), \bar{L}(K))$ as an off-set 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 $\bar{A}(K) = \bar{a}(K)$. 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 $\bar{A}(K-1) = \bar{a}(K-1), \bar{L}(K-1)$. This initial estimator of $\bar{Q}^a_{L(K)} = E(Y^a \mid \bar{L}^a(K-1))$ is used as an off-set 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, 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 $\bar{A}(k-1) = \bar{a}(k-1), \bar{L}(k-1)$, and use this fit as an off-set 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 $\bar{Q}^a_{L(0)} = \Psi(\bar{Q}^a_0)$.

3.2 Loss function for $\bar{Q}^a_0$.

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 can assume that $Y \in [0,1]$. A special case would be that Y is binary valued with values in $\{0,1\}$. As a consequence, 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

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

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 not respecting 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 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 $k = K+2, \ldots, 1$, $\bar{Q}^a_{0,L(k-1)} = \arg\min_{\bar{Q}^a_{L(k-1)}} E_0 \mathcal{L}_{\bar{Q}^a_{0,L(k)}}(\bar{Q}^a_{L(k-1)})$. 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 a cross-validation selector below for the sake of construction of an initial estimator of $\bar{Q}^a_0$.

3.3 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.

3.4 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 logistic regression model for this conditional mean using one of the standard software implementations of logistic regression, ignoring that the outcome Y might not be binary. 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 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 $\bar{L}(K-1)$, ignoring that the outcome is not binary. 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.

For each of these loss-based learning algorithms 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

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

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.

3.5 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

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

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 $\bar{Q}^a_0 = (\bar{Q}^a_{L(0)}, \ldots, \bar{Q}^a_{L(K+1)})$.

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 $\epsilon^j_n = \arg\min_\epsilon P_n \mathcal{L}_{\bar{Q}^{a,j}_n}(\bar{Q}^{a,j}_n(\epsilon, g_n))$, 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 $\bar{Q}^a_{L(k)}$:

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

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 $\bar{Q}_n^{a,*}$ or g_n is consistent. In addition, under regularity conditions, if $g_n = g_0$, $\Psi(\bar{Q}_n^{a,*})$ 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 $\bar{Q}_n^{a,*}$. As shown in van der Laan and Robins (2003), 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 $\Psi(\bar{Q}_n^{a,*})$ 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 $T_g(P_0) \subset L_0^2(P_0)$ within the Hilbert space $L_0^2(P_0)$ with inner product $\langle h_1, h_2 \rangle_{P_0} = P_0 h_1 h_2$. Note that if $\bar{Q}^{a,*} = \bar{Q}_0^a$, then the latter influence curve is the efficient influence curve $D^*(\bar{Q}_0^a, g_0, \psi_0)$, 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 $\sqrt{n}(\Psi(\bar{Q}_n^{a,*}) - \Psi(\bar{Q}_0^a))$ with

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

If one is only willing to assume that either $\bar{Q}_n^{a,*}$ or g_n is consistent, then the influence curve is more complex (see van der Laan and Robins (2003), 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 an 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 the inverse-probability-of-treatment-weighting (IPTW) estimator, and a parametric maximum likelihood estimator (MLE_p) obtained by plugging untargeted estimates of $\bar{Q}^a_{L(k)}$ into the G-computation formula given in Eq. 1. Influence curve-based estimates of the variance of the TMLE are reported, and compared with the empirical variance of the Monte Carlo estimates.

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 randomized controlled trial (RCT) to assess drug effectiveness on a continuous-valued outcome, for example, the effect of an asthma medication on airway constriction after one year of adherence to treatment. Suppose a subset of subjects in the treatment group discontinue the treatment drug in response to results of an intermediate biomarker assay or clinical test midway through the trial (e.g. forced expiratory volume). The diagram in Figure 1 shows the time ordering of intervention nodes (A) and covariate/event nodes (L). $A(0)$ and $A(2)$ represent treatment nodes, and $A(1), A(3)$ represent censoring nodes.

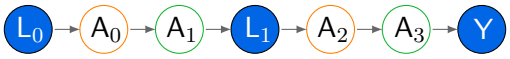

Figure 1: Simulation 1: Time ordering of intervention and non-intervention nodes, baseline covariates $\mathsf{L}_0 = (W_1, W_2, W_3)$, treatment nodes ($\mathsf{A}_0$, A_2), censoring nodes (A_1, A_3), time-dependent covariate L_1, outcome Y.

Our target parameter is the mean outcome under treatment $A(0) = A(2) = 1$ and no censoring $A(1) = A(3) = 1$ minus the mean outcome under control $A(0) = A(2) = 0$ and no censoring: $\psi_0 = E_0\{Y(1,1,1,1) - Y(0,1,0,1)\}$. With this scenario in mind, data were generated as follows:

$$\begin{aligned}
W_1, W_2 &\sim Bernouli(0.5) \\
W_3 &\sim N(4,1) \\
g_0(1 \mid Pa(A_0)) &= P(A_0 = 1 \mid L_0) = expit(0) = 0.5 \\
g_1(1 \mid Pa(A_1)) &= P(A_1 = 1 \mid A_0, L_0) \\
&= expit(0.1 + 0.5W_1 + W_2 - 0.1W_3 + A_0) \\
L_1 &= 3 + A_0 - 2W_1W_2 - 0.5W_3 + \epsilon_1
\end{aligned}$$

$$
\begin{aligned}
g_2(1 \mid Pa(A_2), A_1 = 1) &= P(A_2 = 1 \mid A_0, A_1 = 1, \bar{L}(1)) \\
&= expit(-1.2 - 0.2W_1 - 0.2W_2 + 0.1W_3 + 0.4L_1) \\
g_3(1 \mid Pa(A_3), A_1 = 1) &= P(A_3 = 1 \mid A_0, A_1 = 1, A_2, \bar{L}(1)) \\
&= expit(1.8 + 0.1W_2 - 0.05W_3 - 0.4L_1 - 1.5A_2) \\
Y &= expit(3 - 0.3A_0 + 0.1W_2 - 0.4L_1 - 0.5A_2 + \epsilon_2)
\end{aligned}
$$

with $\epsilon_1, \epsilon_2 \sim_{i.i.d.} N(0,1)$. Results were obtained for 500 samples of size $n_1 = 100$, and $n_2 = 1000$.

Because the study mimics an RCT, the initial treatment assignment probabilities are known by design, however censoring and intermediate treatment assignment probabilities are unknown and must be estimated from the data. In one simulation setting, correctly specified regression models were used to estimate each of the four factors of g: initial treatment assignment probabilities (estimated as the empirical proportion assigned to treatment and control), censoring (loss to follow-up) at baseline, intermediate switching from treatment to control, and subsequent loss to follow-up before measuring the outcome at one year. Two approaches were used to estimate the g-factors. The first 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 covariates measured prior to A_k in the time ordering shown in Figure 1. This alternate formulation contains the truth, but in finite samples can lead to violations of the positivity assumption, and poor fits of the true regression coefficients: $g_{n,k}$ was not bounded away from $(0,1)$ in this simulation. For convenience, in Table 1 these are referred to as correct and misspecified models for g, respectively.

Three separate sets of logistic regression models were used to estimate conditional means $\bar{Q}^a_{L(k)}$: 1) including all terms used to generate the data at each node that gives practically unbiased estimation of ψ_0 (Q_c), 2) including main term baseline covariates only (Q_{m_1}), and 3) an intercept-only model (Q_{m_2}).

The IPTW estimator is consistent when g_n is a consistent estimator of g_0. Thus, IPTW results are expected to be unbiased only when g is correctly specified. Consistency of MLE_p relies on consistent estimation of $\bar{Q}^a_{L(k)}$. TMLE estimates were also obtained based on each of these three initial parametric model specifications, in conjunction with the correct and mis-specified models for g. The TMLE of the targeted causal effect was defined as the difference of the two TMLEs for the two treatment specific means.

Results: Results in Table 1 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 g_0 is consistently estimated, misspecification bias in MLE_p estimates that rely on specification Q_{m1} or Q_{m2} is greatly reduced by the TMLE procedure. However, relative to the correctly specified MLE and IPTW estimator, some bias remains at the larger sample size. When g_0 is misspecified the bias of the IPTW estimator is extreme relative to the true parameter value ($\psi_0 = -0.1779$), and variance is three to four times that of MLE_p and TML estimators, even when $n = 1000$. TMLE's ability to reduce the bias is impaired by misspecification of g_0, but because the submodel and quasi-log-likelihood loss function used in the estimation procedure respect the bounds on the statistical model $\mathcal{M}$, the variance does not suffer(Gruber and van der Laan, 2010a).

Table 1: Simulation 1 results, $\psi_0 = -0.1779$.

		$n = 100$				$n = 1000$			
		Rel bias	Bias	Var	MSE	Rel bias	Bias	Var	MSE
					g correctly specified				
	IPTW	5.93	−0.011	0.010	0.010	0.62	−0.0011	0.0003	0.0003
Q_c	MLE_p	1.25	−0.002	0.003	0.003	0.66	−0.0012	0.0002	0.0002
	TMLE	0.49	−0.001	0.004	0.004	0.66	−0.0012	0.0002	0.0002
Q_{m_1}	MLE_p	12.15	−0.022	0.004	0.005	12.13	−0.0216	0.0003	0.0008
	TMLE	5.21	−0.009	0.004	0.005	4.96	−0.0088	0.0003	0.0003
Q_{m_2}	MLE_p	20.27	−0.036	0.003	0.005	20.20	−0.0359	0.0003	0.0016
	TMLE	6.11	−0.011	0.004	0.004	6.49	−0.0115	0.0003	0.0004
					g misspecified				
	IPTW	−84.50	0.150	0.013	0.036	−87.56	0.1557	0.0010	0.0252
Q_c	MLE_p	1.25	−0.002	0.003	0.003	0.66	−0.0012	0.0002	0.0002
	TMLE	0.92	−0.002	0.003	0.003	0.62	−0.0011	0.0002	0.0002
Q_{m_1}	MLE_p	12.15	−0.022	0.004	0.005	12.13	−0.0216	0.0003	0.0008
	TMLE	8.96	−0.016	0.004	0.004	7.25	−0.0129	0.0003	0.0004
Q_{m_2}	MLE_p	20.27	−0.036	0.003	0.005	20.20	−0.0359	0.0003	0.0016
	TMLE	12.07	−0.021	0.003	0.004	10.01	−0.0178	0.0003	0.0006

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), and time-dependent covariate/event nodes (L) for one such study design is shown in Figure 2.

Figure 2: Simulation 2: Time ordering of intervention and non-intervention nodes, baseline covariates $\mathsf{L}_0 = (W_1, W_2, W_3, W_4, W_5)$), treatment node A_0, censoring nodes (A_1, A_2, A_3), time-dependent covariates (L_2, L_3, L_5, L_6), intermediate and final outcome (L_1, L_4, Y).

IPTW, 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:

$$
\begin{aligned}
W_1, W_2, W_3, W_4, W_5 &\sim N(0,1) \\
P(A_0 = 1 \mid Pa(A_0)) &= expit(0.1W_1 + 0.2W_2 + 0.1W_3 + 0.2W_4 + 0.1W_5) \\
P(A_1 = 1 \mid Pa(A_1)) &= expit(0.1 + 0.2W_1 + 0.4W_4 + 0.2W_5 + 0.1A_0) \\
L_1 &= expit(-2 + 0.4W_1W_2 + 0.3W_3 + 0.4W_4 - 0.3W_5 - A_0) \\
L_2 &= 1 + 0.2W_2 + 0.7W_4 + 0.1W_5 + 0.5A_0 + \epsilon_1 \\
L_3 &= 1 + 0.1W_1 + 0.2W_3 + 0.5W_5 + 0.2A_0 + 0.2L_2 + \epsilon_2 \\
P(A_2 = 1 \mid Pa(A_2)) &= expit(-0.6 + 0.3W_2 + A_0 + 0.1L_2 + 0.5L_3) \\
L_4 &= expit(-0.5 + 0.1W_2 - A_0 + 0.3L_2 - 0.7L_3) \\
L_5 &= 1 - 1.5W_2 + 0.4A0 + 0.5L_2 + 0.1L_3 + \epsilon_3 \\
L_6 &= 1 + 0.1W_3 + 2A_0 + 0.6L_3 + 0.2L_5 + \epsilon_4 \\
P(A_3 = 1 \mid Pa(A_3)) &= expit(1.2 - 0.4A_0 - 0.2L_2 + 0.3L_50.2L_6) \\
P(Y = 1 \mid Pa(Y))) &= expit(-3 - 0.2W_1W_2 - 0.2A_0 + 0.3L_3 + 0.6L_5 + 0.2L_6),
\end{aligned}
$$

with $\epsilon_1, \epsilon_2, \epsilon_3, \epsilon_4 \sim_{i.i.d.} N(0,1)$. The values of censoring nodes A(t) for subjects for whom an event was observed at previous time $t' < t$ were deterministically

set to 1 to reflect the fact that the outcome at t_k is known even in the absence of additional follow-up time. The values at all nodes following censoring or an observed outcome event at time t were set to 0. As in Simulation 1, results were obtained for correct and misspecified regression models for $\bar{Q}^a_{L(k)}$ and g_k. The conditional means $\bar{Q}^q_{L(k)}$ 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 all 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, $\psi_0 = 0.386$.

		$n = 100$				$n = 1000$			
		Rel bias	Bias	Var	MSE	Rel bias	Bias	Var	MSE
					g correctly specified				
	IPTW	−4.58	−0.018	0.044	0.045	0.07	$3e-4$	0.003	0.003
Q_c	MLE_p	12.25	0.047	0.024	0.026	−2.70	−0.010	0.002	0.002
	TMLE	7.64	0.030	0.030	0.030	−0.29	−0.001	0.002	0.002
Q_{m_1}	MLE_p	13.17	0.051	0.025	0.028	−0.30	−0.001	0.002	0.002
	TMLE	7.63	0.029	0.031	0.032	−0.18	−0.001	0.002	0.002
Q_{m_2}	MLE_p	2.79	0.011	0.020	0.020	1.98	0.008	0.002	0.002
	TMLE	0.63	0.002	0.030	0.030	0.19	0.001	0.002	0.002
					g misspecified				
	IPTW	−27.15	−0.105	0.021	0.032	−18.80	−0.073	0.002	0.007
Q_c	MLE_p	12.25	0.047	0.024	0.026	−2.70	−0.010	0.002	0.002
	TMLE	9.98	0.039	0.030	0.032	−0.22	−0.001	0.002	0.002
Q_{m_1}	MLE_p	13.17	0.051	0.025	0.028	−0.30	−0.001	0.002	0.002
	TMLE	10.52	0.041	0.033	0.035	1.19	0.005	0.002	0.002
Q_{m_2}	MLE_p	2.79	0.011	0.020	0.020	1.98	0.008	0.002	0.002
	TMLE	4.64	0.018	0.030	0.031	5.94	0.023	0.002	0.003

Results: Sparsity in the data at small sample size increases bias in all estimators, with the exception of TMLE under dual misspecification Q_{m_2}, g_{mis}, in comparison with performance at the larger sample size, where the positivity assumption is met (Table 2). The G-computation estimators using the

specification Q_{m_2} (intercept-only model) are least impacted by this violation. Sparsity again inflates IPTW variance relative to the other estimators. When g_0 is correctly specified all estimators have comparable MSE at $n = 1000$. At that sample size, if g_0 is misspecified the variance dominates the MSE for all estimators, except for the IPTW.

5.3 Inference.

Table 3 allows us to compare the empirical variance of the Monte Carlo TMLE estimates obtained above, with influence curve-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

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

	$n = 100$			$n = 1000$		
	Emp var	$\hat{\sigma}^2_{IC}/n$	Coverage	Emp var	$\hat{\sigma}^2_{IC}/n$	Coverage
			g correct			
Simulation 1						
Q_c	0.0037	0.0024	0.87	0.00024	0.00028	0.96
Q_{m_1}	0.0044	0.0025	0.85	0.00026	0.00032	0.95
Q_{m_2}	0.0037	0.0037	0.93	0.00027	0.00037	0.93
Simulation 2						
Q_c	0.0296	0.0115	0.77	0.00189	0.00193	0.96
Q_{m_1}	0.0312	0.0065	0.77	0.00188	0.00189	0.96
Q_{m_2}	0.0300	0.0180	0.75	0.00190	0.00216	0.96
			g misspecified			
Simulation 1						
Q_c	0.0034	0.0018	0.83	0.00024	0.00019	0.93
Q_{m_1}	0.0040	0.0017	0.78	0.00026	0.00021	0.83
Q_{m_2}	0.0034	0.0024	0.87	0.00027	0.00024	0.79
Simulation 2						
Q_c	0.0304	0.0106	0.66	0.00190	0.00122	0.87
Q_{m_1}	0.0331	0.0055	0.67	0.00192	0.00108	0.88
Q_{m_2}	0.0304	0.0098	0.67	0.00212	0.00116	0.80

are consistently estimated, and it results in 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 quite 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, if g_0 is misspecified but Q_0 is correctly estimated, coverage is close to the nominal rate for Simulation 1.

6 Concluding remarks.

TMLE is a general template for construction of semiparametric efficient substitution estimators requiring writing the target parameter as a function of an infinite dimensional parameter (e.g., $\Psi(\bar{Q}^a)$) for an infinite dimensional parameter (e.g., $\bar{Q}^a$), 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 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 di-

mensional 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 targeted minimum loss-based estimation (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.

It will be of interest to further 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 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, 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.

In this paper we made use of sequential loss-based learning defined as follows: Let $Q_{k0} = \arg\min_{Q_k} P_0 \mathcal{L}_{k,Q_{k+1,0}}(Q_k)$ be defined as the minimizer of the risk of a loss function that is indexed by $Q_{k+1,0}$, $k = K+1, \ldots, 0$. Let $Q = (Q_0, \ldots, Q_{K+1})$. The parameter of interest is a parameter $\Psi(Q)$. For $k = K+1, \ldots, 0$, given an estimator of $Q_{k+1,0}$, one applies loss-based learning of $Q_{k,0}$ based on the loss function $\mathcal{L}_{k,Q_{k+1,0}}(Q_k)$. The statistical properties of such estimators of $Q(P_0)$ based on sequential cross-validation estimator selection remain to be studied.

If $\Psi(Q)$ is a pathwise differentiable parameter with efficient influence curve $D^*(Q, g)$, we demonstrated how to augment this sequential learning proce-

dure into a targeted minimum loss-based (sequential) learning algorithm: start with initial estimator $Q_n = (Q_{k,n} : k = 0, \ldots, K+1)$ and g_n, construct submodels $\{Q_k(\epsilon, g) : \epsilon\}$ through Q_k at $\epsilon = 0$ so that $\frac{d}{d\epsilon}\mathcal{L}_{k,Q_{k+1}}(Q_k(\epsilon, g))\big|_{\epsilon=0} = D_k^*(Q, g)$ and $D^*(Q, g) = \sum_k D_k^*(Q, g)$, and for $k = K+1, \ldots, 0$, $\epsilon_{k,n} = \arg\min_{\epsilon_k} P_n \mathcal{L}_{k,Q^*_{k+1,n}}(Q_{k,n}(\epsilon, g_n))$, where $Q^*_{k,n} = Q_{k,n}(\epsilon_{k,n})$. The final TMLE is the plug-in estimator $\Psi(Q^*_n)$, and we have $P_n D^*(Q^*_n, g_n) = 0$. In particular, 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: TMLE for causal parameters defined by working MSM without baseline covariates.

Consider the same longitudinal data structure $O = (L(0), A(0), \ldots, L(K), A(K), Y = L(K+1))$, and statistical model $\mathcal{M}$. Let L^a be the random variable with distribution equal to the G-computation formula $P^a = \prod_{k=0}^{K+1} Q^a_{L(k)}$, where $Q^a_{L(k)}$ is the conditional distribution of $L(k)$, given $\bar{L}(k-1), \bar{A}(k-1) = \bar{a}(k-1)$. In this article we presented a TMLE for EY^a. Suppose now that our target parameter is $\Psi(Q) = f(EY^a : a \in \mathcal{A})$ for some multivariate Euclidean valued function f, and a collection of static regimens $\mathcal{A}$. For example, given a working model $\{m_\beta : \beta\}$ for EY^a, and function h, we could define

$$\Psi(Q_0) = \arg\min_\beta E_0 \sum_{a \in \mathcal{A}} h(a)(Y^a - m_\beta(a))^2,$$

or if $Y \in [0, 1]$,

$$\Psi(Q_0) = \arg\min_\beta -E_0 \sum_{a \in \mathcal{A}} h(a) \left\{ Y^a \log m_\beta(a) + (1 - Y^a) \log(1 - m_\beta(a)) \right\}.$$

In this article we presented the efficient influence curve $D^{a,*} = \sum_k D_k^{a,*}$ of the target parameter EY^a with an orthogonal decomposition given in Theorem 1, where $D_k^{a,*}$ is a score of the conditional distribution $Q_{L(k)}$ of $L(k)$, given $\bar{L}(k-1), \bar{A}(k-1)$. The efficient influence curve of the target parameter Ψ is thus given by $D^* = \sum_{a \in \mathcal{A}} f'_a D_a^*$, where $f'_{a'} = \frac{d}{dEY^{a'}} f(EY^a : a \in \mathcal{A})$. We note that this can be decomposed as

$$D^* = \sum_a f'_a \left\{ \sum_k D_k^{a,*} \right\} \equiv \sum_k D_k^*,$$

where

$$D_k^* = \sum_a f'_a D_k^{a,*}.$$

This represents an orthogonal decomposition of the efficient influence curve D^* of Ψ in terms of scores of $Q_{L(k)}$. Specifically,

$$D_{K+1} = \sum_{a \in \mathcal{A}} f'_a \frac{I(\bar{A}(K) = a)}{g_{0:K}} (Y - \bar{Q}_{K+1}^a),$$

and

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

We note that $\Psi(Q)$ depends on Q through $\bar{Q} \equiv (\bar{Q}^a : a \in \mathcal{A})$, where $\bar{Q}^a = (\bar{Q}_k^a = E(Y^a \mid \bar{L}^a(k-1)) : k = 0, \ldots, K+1)$. In this article we proposed a sum loss function $\mathcal{L}_{a,\bar{Q}^a}(\bar{Q}^a) = \sum_{k=0}^{K+1} \mathcal{L}_{a,k,\bar{Q}_{k+1}^a}(\bar{Q}_k^a)$ for $\bar{Q}_0^a$, where

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

As a consequence, $\mathcal{L}_{\bar{Q}}(\bar{Q}) \equiv \sum_{a \in \mathcal{A}} f'_a \mathcal{L}_{a,\bar{Q}^a}(\bar{Q}^a)$ is a valid loss function for $\bar{Q}_0 = (\bar{Q}_0^a : a \in \mathcal{A})$. Note

$$\begin{aligned} \mathcal{L}_{\bar{Q}}(\bar{Q}) &= \sum_{a \in \mathcal{A}} f'_a \sum_{k=0}^{K+1} \mathcal{L}_{a,k,\bar{Q}_{k+1}^a}(\bar{Q}_k^a) \\ &= \sum_{k=0}^{K+1} \left\{ \sum_{a \in \mathcal{A}} f'_a \mathcal{L}_{a,k,\bar{Q}_{k+1}^a}(\bar{Q}_k^a) \right\} \\ &\equiv \sum_{k=0}^{K+1} \mathcal{L}_{k,\bar{Q}_{k+1}}(\bar{Q}_k), \end{aligned}$$

where $\bar{Q}_k = (\bar{Q}_k^a : a \in \mathcal{A})$ and $\mathcal{L}_{k,\bar{Q}_{k+1}}(\bar{Q}_k) = \sum_{a\in\mathcal{A}} f_a' \mathcal{L}_{a,k,\bar{Q}_{k+1}^a}(\bar{Q}_k^a)$ is a loss function for $\bar{Q}_k$.

Consider the submodel $\bar{Q}_k(\epsilon_k, g) = (\bar{Q}_k^a(\epsilon_k, g) : a \in \mathcal{A})$ defined by

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

where we define $g_{0:-1} = 1$. This submodel does indeed satisfy the generalized score-condition

$$\frac{d}{d\epsilon_k}\mathcal{L}_{k,\bar{Q}_{k+1}}(\bar{Q}_k(\epsilon_k, g))\Big|_{\epsilon_k=0} = D_k^*,\ k = 0, \ldots, K+1.$$

In particular, the submodel $\bar{Q}(\epsilon_0, \ldots, \epsilon_{K+1}, g)$ defined by these k-specific submodels through $\bar{Q}_k$, $k = 0, \ldots, K+1$, and the above sum loss function $\mathcal{L}_{\bar{Q}}(\bar{Q})$ satisfies the condition that the generalized score spans the efficient influence curve:

$$D^*(Q, g) \in \left\langle \frac{d}{d\epsilon}\mathcal{L}_{\bar{Q}}(\bar{Q}(\epsilon, g))\Big|_{\epsilon=0} \right\rangle. \tag{4}$$

Finally, we present the TMLE-algorithm. Suppose we already obtained an initial estimator $\bar{Q}_{k,n}$, $k = 0, \ldots, K+1$ and g_n. Let $\bar{Q}_{K+2,n}^{a,*} \equiv Y$ for each $a \in \mathcal{A}$.

For $k = K+1$ to $k = 1$, we compute

$$\epsilon_{k,n} \equiv \arg\min_{\epsilon_k} P_n \mathcal{L}_{k,\bar{Q}_{k+1,n}^*}(\bar{Q}_{k,n}(\epsilon_k, g_n)),$$

and the corresponding update $\bar{Q}_{k,n}^* = \bar{Q}_{k,n}(\epsilon_{k,n}, g_n)$. Finally, for each $a \in \mathcal{A}$, define $\bar{Q}_{L(0),n}^{a,*} = 1/n \sum_{i=1}^n \bar{Q}_{1,n}^{a,*}(L_i(0))$, providing the TMLE $\bar{Q}_{L(0),n}^*$ of $\bar{Q}_{L(0),0} = (\bar{Q}_{L(0),0}^a : a)$. This defines the TMLE $\bar{Q}_n^* = (\bar{Q}_{k,n}^*, k = 0, \ldots, K+1)$ of $\bar{Q}_0$.

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

$$\Psi(\bar{Q}_n^*) = f(\bar{Q}_{L(0),n}^*) = f\left(\frac{1}{n}\sum_{i=1}^n \bar{Q}_{1,n}^{a,*}(L_i(0)) : a \in \mathcal{A}\right).$$

The TMLE solves the efficient influence curve equation $0 = P_n D^*(\bar{Q}_n^*, g_n, \Psi(\bar{Q}_n^*))$, thereby making it a double robust locally efficient substitution estimator, under regularity conditions.

APPENDIX: TMLE for causal parameters defined by working MSM with baseline covariates.

Consider the same longitudinal data structure $O = (L(0), A(0), \ldots, L(K), A(K), Y = L(K+1))$, and statistical model $\mathcal{M}$. Let L^a be the random variable with distribution equal to the G-computation formula $P^a = \prod_{k=0}^{K+1} Q^a_{L(k)}$, where $Q^a_{L(k)}$ is the conditional distribution of $L(k)$, given $\bar{L}(k-1), \bar{A}(k-1) = \bar{a}(k-1)$. Let $V \subset L(0)$ be a user supplied vector of baseline covariates. Suppose now that our target parameter is $\Psi(Q) = f(Q_V, (E(Y^a|V = v) : a \in \mathcal{A}, v \in \mathcal{V}))$ for some multivariate real valued function f, a collection of static regimens $\mathcal{A}$, and a collection $\mathcal{V}$ of possible values for V. Here Q_V denotes the distribution of V. For example, given a working model $\{m_\beta : \beta\}$ for $E(Y^a \mid V)$, we could define

$$\begin{aligned}&\Psi(Q_0) = \arg\min_\beta E_0 \textstyle\sum_{a\in\mathcal{A}} h(a,V)(Y^a - m_\beta(a,V))^2 \\ &= \arg\min_\beta \textstyle\sum_{a\in\mathcal{A}} \sum_v h(a,v)(E(Y^a \mid V = v) - m_\beta(a,v))^2 Q_V(v).\end{aligned}$$

Such a $\psi_0 = \beta_0$ solves the equation $0 = E_0 \sum_{a\in\mathcal{A}} h(a,V) \frac{d}{d\beta_0} m_{\beta_0}(a,V)(E_0(Y^a \mid V) - m_{\beta_0}(a,V))$. Alternatively, if $Y \in [0,1]$,

$$\Psi(Q_0) = \arg\min_\beta -E_0 \sum_{a\in\mathcal{A}} h(a,V) \left\{Y^a \log m_\beta(a,V) + (1-Y^a)\log(1 - m_\beta(a,V))\right\}.$$

This $\psi_0 = \beta_0$ solves the equation $0 = E_0 \sum_{a\in\mathcal{A}} h(a,V) \frac{\frac{d}{d\beta_0} m_{\beta_0}(a,V)}{m_{\beta_0}(1-m_{\beta_0})}(E_0(Y^a \mid V) - m_{\beta_0}(a,V))$.

Recall the definitions of $\bar{Q}^a_{L(k)} = E(Y^a \mid \bar{L}^a(k-1)) = E_{L(k)}(\bar{Q}^a_{L(k+1)} \mid \bar{L}(k-1), \bar{A}(k-1) = \bar{a}(k-1))$ as an iteratively defined conditional mean. Note that $E(Y^a \mid V) = E(\bar{Q}^a_{L(1)}(L(0)) \mid V)$ is obtained by integrating $\bar{Q}^a_{L(1)}$ over $L(0)$ w.r.t the distribution of $L(0)$, given V. Define $\bar{Q}^a_{L(0)|V} = E_{L(0)|V}(\bar{Q}^a_{L(1)} \mid V)$ as this conditional mean. Thus, $E(Y^a \mid V)$ is a function of the following vector of iteratively defined conditional means $\bar{Q}^a \equiv (\bar{Q}^a_{L(K+1)}, \ldots, \bar{Q}^a_{L(1)}, \bar{Q}^a_{L(0)|V})$. In particular, $\bar{Q}^a_{L(0)|V,0}(v) = E_{Q_0}(Y^a \mid V = v)$. In the two examples above $\psi_0 = \beta_0$ is a function of $\bar{Q}_{L(0)|V,0} \equiv (\bar{Q}^a_{L(0)|V,0} : a \in \mathcal{A})$ and the marginal distribution $Q_{V,0}$ of V. To conclude, $\Psi(Q_0)$ can also be represented as $\Psi(\bar{Q}_{L(0)|V}, Q_V)$ or $\Psi(\bar{Q}, Q_V)$, where $\bar{Q} = (\bar{Q}^a : a \in \mathcal{A})$.

The next theorem is the analogue of Theorem 1 and provides the desired orthogonal composition of the efficient influence curve of $\Psi : \mathcal{M} \to \mathbb{R}^d$.

Theorem 3 *Let* $D(Q,g)(O) = \sum_{a\in\mathcal{A}} I(A = a) \frac{h_1(a,V)}{g_{0:K}}(Y - m_{\Psi(Q)}(a,V))$, *where* h_1 *is such that* $P_0 D(Q_0, g_0) = 0$. *This is a gradient of* $\Psi : \mathcal{M} \to \mathbb{R}^d$ *in the*

model with g_0 known. The efficient influence curve of Ψ is given by (up till a constant normalizing matrix) $D^* = \sum_{k=1}^{K+1} D_k^* + D_{L(0)|V}^* + D_V^*$, where for $k = 1, \ldots, K+1$, $D_k^* = \Pi(D \mid T_k)$ 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 $Q_{L(k)}$, $D_{L(0)|V}^*$ is the projection of D onto the tangent space $T_{L(0)|V} = \{h(L(0), V) : E(h \mid V) = 0\}$ of the conditional distribution of $L(0)$, given V, and D_V^* is the projection of D onto $T_V = \{h(V) : Eh(V) = 0\}$. These projections are defined in the Hilbert space $L_0^2(P)$ with inner-product $\langle h_1, h_2 \rangle_P = Ph_1h_2$.

Recall the definition $\bar{Q}_{L(k)}^a = E(Y^a \mid \bar{L}^a(k-1))$, and the recursive relation $\bar{Q}_{L(k)}^a = E_{Q_{L(k)}^a} \bar{Q}_{L(k+1)}^a$ defined by integrating w.r.t. the distribution of $L(k)$ given $\bar{L}(k-1), \bar{A}(k-1) = \bar{a}(k-1)$, $k = K+1, \ldots, 1$. For notational convenience, we also use the notation $\bar{Q}_{L(0)}^a = E(Y^a \mid V) = E_{Q_{L(0)|V}}(\bar{Q}_{L(1)}^a \mid V)$ instead of $\bar{Q}_{L(0)|V}^a$. Let $\bar{Q} = (\bar{Q}_{L(k)}^a : k = K+1, \ldots, 0 : a \in \mathcal{A})$, and let Q_V be the marginal distribution of V. Recall that $\Psi(Q) = \Psi(\bar{Q}, Q_V)$ or $\Psi(Q) = \Psi(\bar{Q}_{L(0)}, Q_V)$.

We have

$$D_{K+1}^* = \sum_{a \in \mathcal{A}} I(A = a) \frac{h_1(a, V)}{g_{0:K}} (Y - \bar{Q}_{L(K+1)}^a(\bar{L}(K)),$$

and for $k = K, \ldots, 1$,

$$\begin{aligned} D_k^* &= \sum_{a \in \mathcal{A}} I(\bar{A}(k-1) = \bar{a}(k-1)) \frac{h_1(a, V)}{g_{0:k-1}} \left\{ \bar{Q}_{L(k+1)}^a - E_{Q_{L(k)}^a} \bar{Q}_{L(k+1)}^a \right\}, \\ &= \sum_{a \in \mathcal{A}} I(\bar{A}(k-1) = \bar{a}(k-1)) \frac{h_1(a, V)}{g_{0:k-1}} \left\{ \bar{Q}_{L(k+1)}^a - \bar{Q}_{L(k)}^a \right\}. \end{aligned}$$

In addition, we have (recall $\bar{Q}_{L(0)}^a(V) = E_Q(Y^a \mid V)$)

$$\begin{aligned} D_{L(0)|V}^* &= \sum_{a \in \mathcal{A}} h_1(a, V) \left\{ \bar{Q}_{L(1)}^a - E_{Q_{L(0)|V}}(\bar{Q}_{L(1)}^a \mid V) \right\} \\ &= \sum_{a \in \mathcal{A}} h_1(a, V) \left\{ \bar{Q}_{L(1)}^a - \bar{Q}_{L(0)}^a \right\}, \end{aligned}$$

and

$$D_V^* = \sum_{a \in \mathcal{A}} h_1(a, V) \left\{ \bar{Q}_{L(0)}^a - m_{\Psi(Q)}(a, V) \right\}.$$

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

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

depends on Q, g only through $\bar{Q}^a_{L(k+1)}$, its mean $\bar{Q}^a_{L(k)}$ under $Q^a_{L(k)}$, across all $a \in \mathcal{A}$, and $g_{0:k-1}$. Similarly, $D^*_{L(0)|V}(Q,g) = D^*_{L(0)|V}(\bar{Q}_{L(0)}, \bar{Q}_{L(1)})$, and $D^*_V(Q,g) = D^*_V(\bar{Q}_{L(0)}, \Psi(\bar{Q}))$.

Loss function for $(\bar{Q}, Q_V)$: We will alternate notation $\bar{Q}^a_k$ and $\bar{Q}^a_{L(k)}$. Recall that $\Psi(Q)$ depends on Q through Q_V, and $\bar{Q} \equiv (\bar{Q}^a : a \in \mathcal{A})$, where $\bar{Q}^a = (\bar{Q}^a_{L(0)|V)}, (\bar{Q}^a_k = E(Y^a \mid \bar{L}^a(k-1)) : k = 1, \ldots, K+1))$.

Note $\bar{Q}^a_k$ is a function of $\bar{l}(k-1)$, $k = 1, \ldots, K+1$. We use

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

as loss function for $\bar{Q}^a_k$, indexed by "nuisance" parameter $\bar{Q}^a_{k+1}$. Note that as a function of O we have

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

Indeed, this is a valid loss function since

$$E_0(\bar{Q}^a_{k+1} \mid \bar{A}(k-1) = \bar{a}(k-1), \bar{L}(k-1)) = \arg\min_{\bar{Q}^a_k} E_0 \mathcal{L}_{a,k,\bar{Q}^a_{k+1}}(\bar{Q}^a_k)(O).$$

In particular,

$$\mathcal{L}_{a,0,\bar{Q}^a_1}(\bar{Q}^a_0) = -\left\{ \bar{Q}^a_1(L(0)) \log \bar{Q}^a_0(V) + (1 - \bar{Q}^a_1(L(0))) \log(1 - \bar{Q}^a_0(V)) \right\}$$

is the loss function for $\bar{Q}^a_0(V) = E_{L(0)|V}(\bar{Q}^a_1(L(0)) \mid V)$. We use

$$\mathcal{L}_{k,\bar{Q}_{k+1}}(\bar{Q}_k) \equiv \sum_{a \in \mathcal{A}} h_1(a,V) \mathcal{L}_{a,k,\bar{Q}^a_{k+1}}(\bar{Q}^a_k),\ k = 0, \ldots, K+1$$

as loss function for $\bar{Q}_k = (\bar{Q}^a_k : a \in \mathcal{A})$. One can view $\sum_{k=0}^{K+1} \mathcal{L}_{k,\bar{Q}_{k+1}}(\bar{Q}_k)$ as loss function for $\bar{Q} = (\bar{Q}_k : k = 0, \ldots, K+1)$, which is indeed a valid loss function for $\bar{Q}_0$: i..e, the expectation of this loss function as a function of O at the correct "nuisance parameters" is minimized by $\bar{Q}_0$. We will use the log-likelihood loss $\mathcal{L}(Q_V) = -\log Q_V$ as loss function for the distribution $Q_{V,0}$ of V, but this loss will play no role since we will estimate $Q_{V,0}$ with the empirical distribution function $Q_{V,n}$. This finalizes the loss function for all components of $(\bar{Q}, Q_V)$, and the sum loss function $\mathcal{L}_{\bar{Q}}(Q_V, \bar{Q}) \equiv \mathcal{L}_{\bar{Q}}(\bar{Q}) - \log Q_V$ is a valid loss function for $(\bar{Q}, Q_V)$ as a whole.

Least favorable submodel: Consider the submodel $\bar{Q}_k(\epsilon_k, g) = (\bar{Q}_k^a(\epsilon_k, g) : a \in \mathcal{A})$ with parameter ϵ_k defined by

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

where we define $g_{0:-1} = 1$. This defines a submodel $\bar{Q}(\epsilon, g)$ with parameter $\epsilon = (\epsilon_k : k = 0, \ldots, K+1)$. Note that

$$\frac{d}{d\epsilon_k} h_1(a, V) \mathcal{L}_{a,k,\bar{Q}_{k+1}^a}(\bar{Q}_k^a(g, \epsilon_k))\Big|_{\epsilon_k=0} = h_1(a, V) \frac{I(\bar{A}(k-1) = \bar{a}(k-1))}{g_{0:k-1}} (\bar{Q}_{k+1}^a - \bar{Q}_k^a).$$

This shows that

$$\frac{d}{d\epsilon_k} \mathcal{L}_{k,\bar{Q}_{k+1}}(\bar{Q}_k(\epsilon_k, g))\Big|_{\epsilon_k=0} = D_k^*,\ k = 1, \ldots, K+1,$$

and similarly this holds for $k = 0$ with $D_{L(0)|V}^*$. Consider also a submodel $Q_V(\epsilon_V)$ of Q_V with score D_V^*, but this one will play no role in the TMLE-algorithm. This defines our submodel $(Q_V(\epsilon_V), \bar{Q}(\epsilon, g) : \epsilon_V, \epsilon)$. The sum loss function $\mathcal{L}_{\bar{Q}}(Q_V, \bar{Q})$ and this submodel satisfy the condition that the generalized score spans the efficient influence curve:

$$D^*(Q, g) \in \left\langle \frac{d}{d\epsilon} \mathcal{L}_{\bar{Q}}(Q_V(\epsilon_V), \bar{Q}(\epsilon, g)\Big|_{\epsilon=0} \right\rangle. \tag{5}$$

TMLE algorithm: Finally, we present the TMLE algorithm. Suppose we already obtained an initial estimator $\bar{Q}_{k,n},\ k = 0, \ldots, K+1$ and g_n, and we estimate the marginal distribution of V with the empirical distribution function. Let $\bar{Q}_{K+2,n}^{a,*} \equiv Y$ for each $a \in \mathcal{A}$, and let $Q_{V,n}$ be the empirical distribution of V_i, $i = 1, \ldots, n$.

For $k = K+1$ to $k = 0$, we compute

$$\epsilon_{k,n} \equiv \arg\min_{\epsilon_k} P_n \mathcal{L}_{k,\bar{Q}_{k+1,n}^*}(\bar{Q}_{k,n}(\epsilon_k, g_n)),$$

and the corresponding update $\bar{Q}_{k,n}^* = \bar{Q}_{k,n}(\epsilon_{k,n}, g_n)$. This defines the TMLE $\bar{Q}_n^* = (\bar{Q}_{k,n}^* : k = 0, \ldots, K+1)$. In particular, $\bar{Q}_{0,n}^* = \bar{Q}_{L(0),n}^*$ is the TMLE of $\bar{Q}_{L(0)} = (E(Y^a \mid V) : a \in \mathcal{A})$.

Finally, we compute the TMLE of ψ_0 as the plug-in estimator corresponding with $\bar{Q}_{L(0),n}^*$ and $Q_{V,n}$:

$$\psi_n^* = \Psi(Q_{V,n}, \bar{Q}_{L(0),n}^*).$$

The TMLE solves the efficient influence curve equation $0 = P_n D^*(\bar{Q}_n^*, g_n, \Psi(Q_{V,n}, \bar{Q}_n^*))$, thereby making it a double robust locally efficient substitution estimator, under regularity conditions.

This demonstrates that our TMLE as presented in this article for a single intervention specific mean outcome can be generalized to general causal parameters. In particular, the TMLE presented in this article is easily generalized to the TMLE for the causal effect $EY^{a(0)=1,0} - EY^{a(0)=0,0}$ of a point treatment $A(0)$ on a survival outcome $Y = I(T > t_0)$ with T subject to right-censoring, by defining $A(k)$, $k = 1, \ldots, K+1$, as a right-censoring indicator at the k-th time point, which is intervened upon by setting it uncensored (i.e, $a(k) = 0$)). The above TMLE for working marginal structural models can also applied to working marginal structural models for dynamic treatment regimens.

Appendix: 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). These functions plus additional software to run the simulations presented above are available online at `www.stat.berkeley.edu/~laan/ Software`.

```
#-----------------------------------------------------------------------------
# function: getEstimates
# purpose: IPTW, Parametric MLE, TMLE estimates of tmt-specific mean outcome
# arguments:
#   d: dataset, wide format, following the time-ordering of the nodes
#   Anodes: tmt and censoring node  columns in d
#   Lnodes: time-dependent covariate and outcome columns in d
#   Ynodes: intermediate and final event node columns (subset of Lnodes)
#   Qform: regression formulas for Q_1 through Q_K+1
#   gform: regression formulas for each treatment and censoring event
#   gbds:  lower and upper bounds on estimated probabilities for g-factors
#-----------------------------------------------------------------------------
getEstimates <- function(d, Anodes, Lnodes, Ynodes, Qform, gform, gbds=c(0,1))
{
     n <- nrow(d)
     n.Q <- length(Lnodes)
     n.g <- length(Anodes)

     g1W <- estg(d, gform, Anodes, Ynodes)
     cum.g1W <- bound(t(apply(g1W, 1, cumprod)), gbds)
     cum.g1W[is.na(cum.g1W)] <- Inf
```

```
     iptw <- mean(d[,Lnodes[n.Q]] * d[,Anodes[n.g]]/cum.g1W[,n.g])

     # Gcomp and TMLE
     Qstar <- Qinit <- d[, Lnodes[n.Q]]
     IC <- rep(0, n)
     for (i in n.Q:1){
          Anode.cur  <- which.max(Anodes[Anodes < Lnodes[i]])
          uncensored <- d[,Anodes[Anode.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)

          Qstar.kplus1 <- Qstar
          Qstar <- estQ(Qstar.kplus1, d, Qform[i], uncensored, deterministic,
                    h = 1/cum.g1W[,Anode.cur])
          IC[uncensored] <- (IC + (Qstar.kplus1 - Qstar)/
                                cum.g1W[,Anode.cur])[uncensored]
     }
     return(c(iptw=iptw, Gcomp=mean(Qinit), tmle=mean(Qstar),
              var.tmle=var(IC)/n))
}

# Utility functions
#---------------------------------------------------------------------------
# function: estQ
# purpose: parametric estimation of  Q_k, targeted when h is supplied
#---------------------------------------------------------------------------
estQ <- function(Q.kplus1, d, Qform, uncensored, deterministic, h=NULL){
     Qform <- update.formula(Q.kplus1 ~ .)
     m <- glm(as.formula(Qform),
          data=data.frame(Q.kplus1, d)[uncensored & !deterministic,],
          family="quasibinomial")
     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, Anodes, 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 < Anodes[i])){
               Ynode.prev <- max(Ynodes[Ynodes < Anodes[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[,Anodes[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 like-

lihood estimation: application to the treatment of antiretroviral resistant HIV infection. *Stat Med*, 28:152–72, 2009.

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 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.