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

Year 2011 *Paper* 281

A General Implementation of TMLE for Longitudinal Data Applied to Causal Inference in Survival Analysis

Ori M. Stitelman* Victor De Gruttola†
Mark J. van der Laan‡

*University of California, Berkeley, DIvision of Biostatistics, ostitelman@berkeley.edu

†Harvard School of Public Health, degrut@hsph.harvard.edu

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

http://biostats.bepress.com/ucbbiostat/paper281

A General Implementation of TMLE for Longitudinal Data Applied to Causal Inference in Survival Analysis

Ori M. Stitelman, Victor De Gruttola, and Mark J. van der Laan

Abstract

In many randomized controlled trials the outcome of interest is a time to event, and one measures on each subject baseline covariates and time-dependent covariates until the subject either drops-out, the time to event is observed, or the end of study is reached. The goal of such a study is to assess the causal effect of the treatment on the survival curve. Standard methods (e.g., Kaplan-Meier estimator, Cox-proportional hazards) ignore the available baseline and time-dependent covariates, and are therefore biased if the drop-out is affected by these covariates, and are always inefficient. We present a targeted maximum likelihood estimator of the causal effect of treatment on survival fully utilizing all the available covariate information, resulting in a double robust locally efficient substitution estimator that will be consistent and asymptotically linear if either the censoring mechanism is consistently estimated, or if the maximum likelihood based estimator is already consistent. In particular, under the independent censoring assumption assumed by current methods, this TMLE is always consistent and asymptotically linear so that it provides valid confidence intervals and tests. Furthermore, we show that when both the censoring mechanism and the initial maximum likelihood based estimator are mis-specified, and thus inconsistent, the TMLE exhibits stability when inverse probability weighted estimators and double robust estimating equation based methods break down The TMLE is used to analyze the Tshepo study, a study designed to evaluate the efficacy, tolerability, and development of drug resistance of six different first-line antiretroviral therapies. Most importantly this paper presents a general algorithm that may be used to create targeted maximum likelihood estimators of a large class of parameters of interest for general longitudinal data structures.

1 Introduction

Many clinical trials are designed to assess the causal effect of different treatments on the time it takes for a particular outcome to occur, such as death, viral progression, or symptom relief. Such trials collect on each experimental unit a longitudinal data structure involving baseline covariates, treatment assignment, and time-dependent covariate processes up till the minimum of drop-out and time to event of interest. The typical approach for assessing these causal effects in the literature and as mandated by the FDA for pharmaceutical drug development is to employ a Cox-proportional hazards model only including treatment, and testing for the coefficient of treatment to be equal to zero. These methods ignore the available covariate information. It is well known that this test is biased if censoring depends on baseline covariates or even time-dependent covariates that are also predictive of survival.

Time dependent confounding, in the form of informative censoring, is a major obstacle that stands in the way for getting an unbiased estimator of causal effects, even in randomized controlled trials. If there are time dependent covariates that both predict censoring and the time to event, then the causal effect on the time to event may not be unbiasedly estimated by only accounting for baseline covariates. This is a common issue in many clinical trials where treatment is initially randomized but subjects are differentially lost to follow up among the treatment arms. Adjusting for time-dependent post-treatment covariates in a multiplicative intensity model results in non-interpretable coefficients in front of treatment, even if the multiplicative intensity model would be correctly specified. That is, standard regression methods cannot be employed.

Moreover, even if the Cox-proportional hazards model is correctly specified, the Cox-proportional hazard model does typically not represent the causal effect of treatment of interest, such as the additive causal effect of treatment on survival, or the causal relative risk. Ideally, the parameter being estimated should be easily interpreted by both non-statisticians and statisticians alike. In other words, the parameter being estimated should be a quantity that a subject matter expert and not a statistician could make informed treatment decisions on.

The causal inference literature allows one to define the actual causal quantities of interest, and establish identifiability results of these causal quantities so that they can be identified as a target parameter of the data generating distribution of the experimental unit under clearly stated causal (non-testable) assumptions. Specifically, under a causal model such as the Neyman-Rubin model or the nonparametric structural equation model (Pearl (2008)), the assumption that the treatment and censoring nodes are sequentially randomized given the observed history, and a positivity assumption, one can identify the post intervention distribution, under setting treatment and enforcing no censoring, from the so called G-computation formula

(Robins (1987)). The statistical estimation problem is now defined as the estimation (based on observing n i.i.d. copies of the experimental unit) of the target parameter of the data generating distribution under a semiparametric statistical model that represents realistic statistical assumptions.

Particular classes of estimators that may be used to estimate such target parameters of interest are a MLE of the G-computation formula parameter based on parametric models, the Inverse Probability of Censoring Weighted (IPCW) Estimator, the Augmented-IPCW (A-IPCW) estimator (Robins and Rotnitzky (1992), van der Laan and Robins (2003)), and the Targeted Maximum Likelihood Estimator (TMLE).

The MLE is a substitution estimator of the target parameter of the data generating distribution. The MLE relies on a correctly specified parametric model for the relevant factor of the data generating distribution, which can be factored in terms of an intensity of the time to event process, and the conditional distributions of the time-dependent covariate processes. If one utilizes likelihood based adaptive estimation to estimate the data generating distribution, then there is no theory that supports the construction of valid 95-percent confidence intervals based on this approach: in fact, it is easily shown that such a data adaptive MLE of the target parameter will be overly biased so that the bias will not converge to zero at a root-n rate and thus cannot be ignored in statistical inference (see e.g., van der Laan and Rubin (2006)).

The IPCW estimator re-weights the observed data by the inverse of the product of the propensity score and censoring probability in order make the treatment arms among the uncensored subjects comparable w.r.t. confounders, and then applies standard estimation as if treatment was randomized and censoring was non-informative. The consistency of these estimators rely on consistent estimation of the treatment and censoring mechanism. These estimators are highly unstable in situations when the parameter of interest is weakly identifiable [1], such as when there is a level of covariates that is predictive of treatment or censoring. The instability of IPCW estimators becomes even more extreme as the dimensionality of the observed data structure increases, as is the case when there are time dependent

[1]The parameter of interest is weakly identifiable when there are levels of covariates that are almost completely predictive of treatment or censoring. In situations where there are levels of covariates that are completely predictive of censoring certain parameters of interest are not identifiable. Experimental designs are created in order to make causal parameters as identifiable as possible through either randomization or the ability to assign covariates and treatment. However, in certain situations it is impossible to randomize treatment. Even in randomized trials it is impossible to randomize censoring. As a result there may be levels of covariates that are almost completely predictive of censoring. This informative censoring often makes parameters in the time-to-event setting weakly identifiable

covariates that must be accounted for.

Even though locally efficient double robust A-IPCW estimators, based on estimating equation theory, of the causal effect of treatment on survival incorporating time-dependent covariates have been proposed (e.g., Robins and Rotnitzky (1992), Hubbard, van der Laan, and Robins (1999), van der Laan and Robins (2003)), these estimators have not gained traction in the literature, due to their complexity as well as the above mentioned instability w.r.t. the choice of estimator of the censoring and treatment mechanism. The IPCW and A-IPCW estimators are based on solving an estimating equation and have shown to be unstable in situations where the parameter of interest is weakly identifiable due to not respecting global constraints implied by the statistical model and the fact that the target parameter is a particular function of the true data generating distribution.

For this reason Targeted Maximum Likelihood Estimation (TMLE), has been proposed (van der Laan and Rubin (2006)), which provides estimators that are double robust locally efficient and also respect the global constraints on the target parameter by being a substitution estimator. The advantages of applying TMLE for estimating causal effects, in general, has been addressed in many articles (see e.g. the seminal paper on the topic van der Laan and Rubin (2006), van der Laan (2010a), van der Laan (2010b), and a forthcoming book van der Laan and Rose (2011)). For articles that demonstrate the advantages of TMLE compared to IPCW, A-IPCW and other estimators in simulation studies for estimation of the additive causal effect of a point treatment on a completely observed outcome, estimation of the mean of an outcome under missingness, and estimation of a causal effect in case-control studies, we refer to van der Laan and Rubin (2006), Gruber and van der Laan (2010a), Gruber and van der Laan (2010b), Porter, Gruber, van der Laan, and Sekhon (forthcoming 2011), Rose and van der Laan (2010) among others. In a series of earlier papers we developed the TMLE for estimating causal effects of treatment on time to event subject to right-censoring incorporating baseline covariates (Moore and van der Laan (2009), Stitelman and van der Laan (2011), Stitelman and van der Laan (2010)), and the advantages of this TMLE relative to the other classes of estimators was demonstrated through simulations and data analysis. The advantages of TMLE relative to MLE and estimating equation based estimators, as observed for these relatively simple data structures, can be expected to be more strongly expressed for more complex longitudinal data structures.

In this article we propose a TMLE for estimating the treatment specific survival curve, and other closely related parameters that are functions of treatment specific survival, that accounts for informative censoring due to *time-dependent covariates*. In addition, the TMLE presented here may also be used to gain efficiency when censoring is independent. In a two part article entitled Targeted Maximum Likelihood Based Causal Inference Mark van der Laan proposes a general template

for constructing targeted maximum likelihood estimators (TMLEs) of parameters of the G-computation formula. In this article we use that template to construct a TMLE for the treatment specific survival curve that adjusts for possible confounding due to intermediate time-dependent variables. The TMLE is locally efficient, double robust, and may be implemented using standard statistical software. Moreover, the resulting estimator, like all TMLEs, benefits from the advantages of being a substitution estimator as opposed to being defined as a solution of an estimating equation. In addition, we propose solutions that address the computational difficulties of constructing a TMLE for this longitudinal data structure and illustrate how those solutions result in an algorithm that performs extremely well in terms of computation time. A simulation study is presented that compares the characteristics of this TMLE, a double robust estimating equation estimator, IPCW estimator, and versions of these three estimators that only account for baseline confounders. The stability of the TMLE that incorporates time dependent covariates is displayed. We even demonstrate the stability of this TMLE compared to other methods when the initial censoring mechanism, and outcome and intermediate variable processes are mis-specified. Finally, we present an analysis assessing the causal effect modification of cART therapies by gender using the Tshepo study, a study designed to evaluate the efficacy, tolerability, and development of drug resistance of six different first-line cART regimens. In analyzing the Tshepo study we contrast how the TMLE presented here compares to the other common methods for estimating time to event parameters.

The algorithm and approach that we develop here with minor adjustments will be able to address even more complicated questions of interest for longitudinal data structures such as the effect of time dependent treatments strategies or dynamic treatment rules on time to event outcomes. The value of this algorithm is that it is a general approach that may be used for estimating many different parameters of interest for a wide range of longitudinal data structures. Current approaches for estimating parameters of interest in longitudinal data structures rely on either IPCW based estimates or MLE based methods whose drawbacks we addressed above (see e.g. Samore et al. (2005), Hernan et al. (2005, 2006, 2009), Lok (2009)). In fact, simple parameters such as the effect of a point treatment on a single outcome with no time component can be expressed as a specific instance of the proposed approach. Moreover, increasingly complex parameters of interest, such as estimating the treatment specific survival curves accounting for baseline covariates (Moore and van der Laan (2009)) can be evaluated using this approach. Finally, parameters of interest, in longitudinal data structures, for which there have been no computationally feasible and efficient approach may be estimated using the algorithm presented here.

1.1 Organization

In the next section we review the general targeted minimum loss based estimator. In Section 3 we apply this template for TMLE to our setting, and present the details required for a fast and effective implementation of this TMLE. Section 4 presents a simulation study, and Section 5 a data analysis of the Tshepo study. We conclude with a discussion. In the Appendix we provide a generalization of our fast implementation of TMLE to general longitudinal data structures, and parameters defined by marginal structural working models for static or dynamic interventions.

2 Targeted Minimum Loss Based Estimation (TMLE)

Let $O_1, \ldots, O_n$ be n i.i.d. copies of a random variable O with probability distribution P_0 that is known to be an element of a statistical model $\mathscr{M}$. Suppose that our target parameter of interest is $\Psi : \mathscr{M} \to \mathbb{R}$, so that $\psi_0 = \Psi(P_0)$ represents the true target parameter value we desire to estimate from the data. Assume that this target parameter mapping is pathwise differentiable, and let $D^*(P)$ be the efficient influence curve (also called the canonical gradient) of the pathwise derivative at P. Let $P \to Q(P)$ be a parameter such that $\Psi(P)$ only depends on P through $Q(P)$. For notational convenience, we will use the notation $\Psi(Q)$ and $\Psi(P)$ interchangeably. Let $D^*(P)$ only depend on P through $Q(P)$ and another nuisance parameter $g(P)$. Again, we will use the notation $D^*(P)$ and $D^*(Q(P), g(P))$ interchangeably.

Let $L_g(Q)$ be a loss function for $Q_0 = Q(P_0)$, possibly indexed by nuisance parameter $g = g(P)$, so that

$$Q_0 = \arg\min_{Q \in \mathscr{Q}} E_{P_0} L_{g_0}(Q)(O),$$

where $\mathscr{Q} = \{Q(P) : P \in \mathscr{M}\}$ is the parameter space for Q. Note that $L_{g_0}(Q)$ is a function of O. In addition, given Q, g, let $\{Q_g(\varepsilon) : \varepsilon\} \subset \mathscr{Q}$ be a parametric submodel through Q, with ε representing the finite dimensional parameter, so that the linear span of the components of $\frac{d}{d\varepsilon} L_g(Q_g(\varepsilon))\big|_{\varepsilon=0}$ includes the components of $D^*(Q,g)$:

$$D^*(Q,g) \in \langle \frac{d}{d\varepsilon} L_g(Q_g(\varepsilon))\Big|_{\varepsilon=0} \rangle.$$

Here we used the notation $\langle f = (f_1, \ldots, f_k) \rangle$ for the linear span of the k components of f. We refer to this condition on the choice of loss function and submodel as the generalized score condition.

A TMLE can now be defined by starting with an initial estimator Q_n^0, g_n, and iteratively updating $Q_n^k = Q_n^{k-1}(\varepsilon_n^k)$, where $\varepsilon_n^k = \arg\min_\varepsilon P_n L_{g_n}(Q_{n,g_n}^{k-1}(\varepsilon))$, till $\varepsilon_n^k \approx$

0. The final update is the TMLE Q_n^* and solves $P_n D^*(Q_n^*, g_n) = 0$. This particular way of iteratively updating is just one approach, but we will employ the backwards single step updating approach instead explained below and proposed in van der Laan (2010a).

In our implementation of the TMLE, we have $Q = \prod_{l=1}^N Q_l$ represents the relevant factors of the density of O, $L_g(Q) = -\sum_l \{\log Q_l\} w_l(g)$ is a weighted log-likelihood loss function, and $Q(\varepsilon) = \prod_l Q_l(\varepsilon_l)$ with $\varepsilon = (\varepsilon_l : l)$ is multivariate. We use the notation $L_g(Q_l) = -\{\log Q_l\} w_l(g)$ so that $L_g(Q) = \sum_l L_g(Q_l)$. This choice happens to have some computational advantages relative to selecting the log-likelihood loss $L(Q) = -\sum_l \{\log Q_l\}$ and submodel $Q_g(\varepsilon) = \prod_l Q_{l,g}(\varepsilon_l)$ in which the fluctuation is indexed by g, even though both couples of loss function and submodel satisfies the required generalized score condition. Both types of TMLE will be presented below, and the computational advantage of using the weighted-log-likelihood loss will be explained.

We start out with an initial estimator Q_n^0, g_n of Q_0, g_0. Since Q is defined by N orthogonal factors, the updating algorithm can be modified by updating one factor at the time and using the most recent update of all factors in the calculation of the next update. One could compare a single update of all the N-factors in this manner with the first step of the TMLE above that updates all factors simultaneously. In addition, we employ a TMLE using a *backwards* updating algorithm that first updates the last factor, and then proceeds backwards: Let

$$\varepsilon_n^N = \arg\min_{\varepsilon^N} P_n L_{g_n}(Q_{n,N}^0(\varepsilon^N)).$$

This yields an update Q_n^1 obtained by updating the (last) N-th factor $Q_{n,N}^0$ with $Q_{n,N}^0(\varepsilon_n^N)$. We now proceed with updating the $N-1$-th (next to last) factor:

$$\varepsilon_n^{N-1} = \arg\min_{\varepsilon^{N-1}} P_n L_{g_n}(Q_{n,N-1}^1(\varepsilon^{N-1})).$$

This yields a second update Q_n^2 obtained by updating the $N-1$-th factor $Q_{n,N-1}^1$ with $Q_{n,N-1}^1(\varepsilon_n^{N-1})$. This updating process is iterated resulting in sequence $Q_n^1, Q_n^2, \ldots, Q_n^N$ of N subsequent updates of the initial estimator Q_n^0. Note that the first $J-1$ factors in the update Q_n^J are still equal to the corresponding $J-1$ factors in Q_n^0, $J = 1, \ldots, N$. The last N-th update involves the update of the first factor. This process can start over by updating the final factor again till the fluctuation parameter is estimated as zero. In our applications, due to the backwards updating approach, in a next round all ε-fits will be equal to zero. As a consequence, the TMLE of Q_0 is given by $Q_n^* = Q_n^N$. This general closed form TMLE was presented in van der Laan (2010a). The TMLE of $\psi_0 = \Psi(Q_0)$ is the corresponding substitution estimator $\psi_n^* = \Psi(Q_n^*)$.

Because this backwards iterative updating approach allows a closed form TMLE we decided to implement this type of TMLE.

3 Following the road map for constructing a TMLE for longitudinal data (van der Laan (2010b))

In this section we follow the road map for constructing the TMLE in van der Laan (2010b). We use that template to construct a TMLE for the treatment specific survival curve adjusting for time-dependent confounders of the right-censoring process. In particular, we focus on the Tshepo analysis in which the time-dependent confounders are CD4+ and Viral Load over time. The notation used here is the same as used in van der Laan (2010a) and van der Laan (2010b).

3.1 Code Data in terms of binary indicators

The 650 subjects in the Tshepo study are viewed as i.i.d. observations of a random variable O that represents a longitudinal time-ordered data structure:

$$O = (L(0), A(0), L(1), A(1) \ldots, L(K), A(K), L(K+1)). \tag{1}$$

It is assumed that $L(t)$ occurs before $A(t)$, and we are interested in the effects of interventions on the A-nodes. $L(0)$ corresponds with the variables collected at baseline. In the Tshepo study this variable includes gender, body-mass-index, baseline viral load, baseline cd4-count, among others. $A(0)$ is the baseline treatment, a binary variable that corresponds with an individual's baseline cART therapy, EFV or NVP. $A(t)$ is a binary variable that equals 1 when an individual is censored at time t and 0 when the individual is not censored at time t. $L(t)$ includes the failure time event process, as well the CD4+ and viral load process as time-dependent covariate processes. Let $\bar{A}(t)$ and $\bar{L}(t)$ denote the history of $A()$ and $L()$ up until time t. Specifically, $L(t)$ is decomposed as follows:

$$L(t) = (L(t, j) : j = 1, ..., n(t)), \tag{2}$$

where at each time point, t, there are $n(t)$ different components indexed by j. In the Tshepo analysis $L(t,1)$ corresponds with the event process, $L(t,2)$ corresponds with the CD4+ process, and $L(t,3)$ corresponds with the viral load process. Thus,

in the Tshepo application $n(t)$ is equal to 3. $L(t,j)$ are further decomposed into binary variables $L(t,j,l)$ in the following way:

$$L(t,j) = (L(t,j,l) : l = 1,...,n(t,j)), \tag{3}$$

where at each time point, t, and for each component, j, there are $n(t,j)$ different ordered categorical levels of the process $L(t,j)$ indexed by l. So in the case of a survival process, like $L(t,1)$, $n(t,j)$ is equal to 2 since there are two levels of this process at each time point. $L(t,2)$ and $L(t,3)$ can have multiple levels. The number of levels chosen and there cut points should be based on subject matter knowledge of the observed process, or one could select many levels thereby approximating the continuous process. For example, let us consider $L(t,2)$, the CD4+ process. We could assume 3 ordered categorical levels of $L(t,2)$ so that $n(t,2) = 3$. Those levels are $CD4(t) < 200$, $200 < CD4(t) < 400$, $400 < CD4(t)$. These levels correspond with the following $L(t,2,l)$ which are indicator functions of $CD4(t)$ at each time point: $L(t,2,1) = 1(CD4(t) \leq 200)$, $L(t,2,2) = 1(200 < CD4(t) < 400)$, and $L(t,2,3) = 1(400 \leq CD4(t))$. However, as always when coding a categorical variable, the categorical variable, $L(t,j)$ can be represented by $n(t,j) - 1$ binary variables. For example, the values of the indicators $L(t,2,1)$ and $L(t,2,2)$ define the value of the third indicator $L(t,2,3)$. The binary indicators $L(t,3,l)$ for the viral load process are constructed in the same way. For the sake of factorizing the likelihood, we also need to agree on an ordering of the components that code $L(t)$. Within each $L(t)$, we order the processes as follows: event, then CD4+, and then viral load.

The data structure has now been ordered in terms of baseline covariates $L(0)$, baseline treatment $A(0)$, censoring indicators $A(t)$, and binary covariates $L(t,j,l)$ indexed by (t,j,l). For each variable X in this ordered sequence, we define the parents $Pa(X)$ of that variable X as the variables that precede X in the ordered sequence.

3.2 Define factorization of likelihood in terms of binary conditional distributions

The ordering of all variables as proposed in the previous section, which respects the time ordering, implies the following factorization of the observed data likelihood:

$$
\begin{aligned}
P(O) &= \overbrace{P(L(0))}^{Q_{L(0)}} && (4)\\
&\quad \overbrace{P(A(0) \mid L(0))}^{g_{A(0)}} && (5)\\
&\quad \prod_{t=1}^{K+1} \prod_{j=1}^{n(t)} \prod_{l=1}^{n(t,j)-1} \overbrace{P(L(t,j,l) \mid Pa(L(t,j,l)))}^{Q_{L(t,j,l)}} && (6)\\
&\quad \prod_{t=1}^{K} \overbrace{P(A(t) \mid Pa(A(t)))}^{g_{A(t)}} && (7)
\end{aligned}
$$

3.3 Define statistical model and target parameter

The post-intervention distribution of L, obtained by intervening on all of the $A(t)$ nodes by setting $A(0)$ to $a \in \{0,1\}$ and $A(t)$ to 0, or no censoring, is defined by the so called G-computation formula:

$$
\begin{aligned}
P_{a,0}(L) &= \overbrace{P(L(0))}^{Q_{L_0}}\\
&\quad \prod_{t=1}^{t_k} \prod_{j=1}^{n(t)} \prod_{l=1}^{n(t,j)-1} \overbrace{P(L(t,j,l) \mid Pa(L(t,j,l)), A=a, \bar{A}(t-1)=0)}^{Q_{L(t,j,l),a,0}}
\end{aligned}
$$

For the purpose of causal inference, we can assume a nonparametric structural equation model (NPSEM) stating $L(t) = f_{L(t)}(Pa(L(t)), U_{L(t)})$, and $A(t) = f_{A(t)}(Pa(A(t)), U_{A(t)})$ for a collection of deterministic functions $f_{L(t)}$ indexed by $t = 0, \ldots, K+1$, and $f_{A(t)}$ indexed by $t = 0, \ldots, K$. Here $U = (U_{L(t)}, U_{A(t)} : t)$ represents unobserved error terms, $Pa(L(t)) = (\bar{A}(t-1), \bar{L}(t-1))$, and $Pa(A(t)) = (\bar{A}(t-1), \bar{L}(t))$. This NPSEM can now be used to define a post intervention distribution of a counterfactual $L_{a,0}$ by setting $A = a$ and $A(t) = 0$ in the NPSEM. Under the assumption that the treatment and censoring nodes are sequentially randomized, and a positivity assumption so that all conditioning events in this G-computation formula

have a non-zero probability, we have that the G-computation formula $P_{a,0}$ identifies the probability distribution of the counterfactual $L_{a,0}$ defined by intervening on the NPSEM.

Now the treatment specific survival curve at time $K+1$ may be cast as a mapping from $P_{a,0}$ to the real line as follows:

$$\Psi(Q_0) =$$

$$E_{L(0)} \sum_{L(K+1,1,1)=0} \prod_{t=1}^{K+1} \prod_{j=1}^{n(t)} \prod_{l=1}^{n(t,j)-1} \overbrace{P(L(t,j,l) \mid Pa(L(t,j,l)), A=a, A(t-1)=0)}^{Q_{L(t,j,l)},a,0},$$

where the sum over $L(K+1,1,1)=0$ represents all possible paths of $L()$ that end in $L(K+1,1,1)=0$, which are thus all possible combinations of $L(t,j,l)$ that result in a subject surviving past time $K+1$. This treatment specific survival curve at time $K+1$ is represented as a conditional treatment specific survival curve, given $L(0)$, averaged w.r.t. the distribution of $L(0)$. This probability of survival for a given $L(0)$ is evaluated by summing over all possible paths defined by CD4+, viral load, and time to events larger than K of the post-intervention probabilities assigned to these paths conditional on $L(0)$.

Note that $P = Qg$, where $Q = Q_{L(0)} \prod_{tjl} Q_{L(t,j,l)}$, and $g = g_{A(0)} \prod_t g_{A(t)}$. Let $P_0 = Q_0 g_0$ denote the true data generating distribution. Let $\mathscr{M}$ be a statistical model that leaves Q_0 unspecified, but might impose assumptions on $g_{A(0),0}$ as well as the censoring mechanism $g_{A(t),0}$. For example, in the Tshepo study the treatment mechanism $g_{A,0}$ is known, and a model might be assumed on the censoring mechanism based on subject matter knowledge. The statistical model $\mathscr{M}$, and the target parameter $\Psi : \mathscr{M} \to \mathbb{R}$ is now defined.

3.4 Determine efficient influence curve

The third step in the template is to determine the efficient influence curve for the target parameter mapping $\Psi : \mathscr{M} \to \mathbb{R}$. The efficient influence curve can be represented as a projection of an IPCW-estimating function D_{IPCW} onto the tangent space of the parameter Q:

$$D^*(Q,g) = \Pi(D_{IPCW} | T_Q);$$

where T_Q is the tangent space of the Q-factor of the density $P = Qg$ of O. The D_{IPCW} is given by

$$D_{IPCW}(O) = \frac{1(A=a)1(C>K)1(T>K)}{g_{A(0)}(a) \prod_{t=1}^{K} g_{A(t)}(0 \mid Pa(A(t)))},$$

where C is the censoring time and T is the event time. Thus the efficient influence curve can be decomposed as

$$D^*(Q,g) = \prod(D_{IPCW}(Q,g) \mid T_Q) = D_0 + \sum_{tjl} D_{tjl},$$

where, D_0 and D_{tjl} are given by:

$$\begin{aligned} D_0 &= P[L_{a,0}(K+1,1,1) = 0 \mid A = a, L(0)] \\ D_{tjl} &= C_{tjl}(Q,g)[L(t,j,l) - Q_{L(t,j,l)}(1 \mid Pa(L(t,j,l)))]. \end{aligned}$$

The function $C_{tjl}(Q,g)$ is only a function of O through $Pa(L(t,j,l))$ and can be factorized into a part that is a function of g and a part that is a function of Q:

$$C_{tjl}(Q,g) = C_{tjl}(Q)C_{tjl}(g),$$

where

$$\begin{aligned} C_{tjl}(Q) = \{&P[L_{a,0}(K+1,1,1) = 0 \mid L(t,j,l) = 1, Pa(L(t,j,l))] - \\ &P[L_{a,0}(K+1,1,1) = 0 \mid L(t,j,l) = 0, Pa(L(t,j,l))]\}, \end{aligned} \tag{8}$$

and

$$C_{tjl}(g) = \frac{I(A=a)I(C > t_-)}{g_{A(0)}(1 \mid L(0)) \prod_{s=1}^{t-1} g_{A(s)}(0 \mid Pa(A(s)))}. \tag{9}$$

The benefit of this factorization will be described below and plays a major part in making the TMLE algorithm for the observed data structure computationally feasible for large t and many binary covariates.

3.5 The loss function and submodel for the TMLE

For notational convenience, we will now and then use the notation Q_{tjl} instead of $Q_{L(t,j,l)}$. Formulas (9) and (9) are used to define the loss function and submodel used in the definition of the TMLE of the treatment specific survival curve. If one uses the standard log-likelihood loss function $L(Q) = -\log Q$ in the TMLE, then one fluctuates the initial estimator of the conditional distribution $Q_{L(t,j,l),n}$ by adding the clever covariate extension $\varepsilon C_{tjl}(Q_n, g_n)$ on the logit scale. If on the other hand, one uses a weighted-log-likelihood loss $L_g(Q) = -\log Q_{L(0)} - \sum_{tjl}\{logQ_{tjl}\}C_{tjl}(g)$, then one fluctuates the initial estimator of the conditional distribution $Q_{L(t,j,l),n}$ by adding the clever covariate extension $\varepsilon C_{tjl}(Q_n)$ on the logit scale.

3.6 Initial estimator of Q_0 and treatment mechanism

We need to construct an initial estimator of $Q_{L(0),0}$ and the conditional distributions $Q_{L(t,j,l),0}$, and we also need an estimator of the conditional distribution $g_{A(0),0}$ of the treatment $A(0)$, as well as the conditional distribution $g_{A(t),0}$ of the censoring indicators $A(t)$. The marginal distribution of $L(0)$ is estimated with the empirical probability distribution.

Define clusters of Q-probabilities that need to be considered for pooling

For the purpose of constructing the initial estimator Q_n it is necessary to determine what $L(t,j,l)$ should be pooled into an appropriate repeated measures data set. Each variable indexed by j is treated separately, since it makes no sense to smooth across different variables such as CD4-count, Viral load and death. We pool $L(t,1,1)$ across time t which results in a repeated measures data set for $(L(t,1,1), Pa(L(t,1,1)) : t)$, in which $L(t,1,1)$ represents the binary outcome and the parents of $L(t,1,l)$ represents the covariates. The indicators $L(t,2,l)$ for CD4-count are pooled over time t and levels l resulting in a repeated measures data set for $(L(t,2,l), Pa(L(t,2,l)) : t,l)$, in which $L(t,2,l)$ represents the binary outcome, and the parents of $L(t,2,l)$ represents the covariates. Finally, we also create such a repeated measures data set for $(L(t,3,l), Pa(L(t,3,l)) : t,l)$, for the viral load process. Each of these three data sets are used to estimate the clusters $Q_1 = (Q_{t,1,1} : t)$, $Q_2 = (Q_{t,2,l} : t,l)$, and $Q_3 = (Q_{t,3,l} : t,l)$ of Q-probabilities, respectively.

Apply loss-based super learner to repeated measures data set to estimate each cluster of Q-probabilities

We can estimate these conditional binary probability distributions with logistic regression applied to the pooled data sets. Instead we recommend using loss-based super learning in which parametric logistic regression models might be included in the library of the super learner (van der Laan, Polley, and Hubbard (2007)). We use the following three log-likelihood loss functions for these intensities of the time to

death, CD4, and viral load:

$$\begin{aligned} L_1(Q_1)(O) &= \sum_t R_1(t) \log \bar{Q}_{t,1,1}^{L(t,1,1)} (1-\bar{Q}_{t,1,1})^{1-L(t,1,1)} \\ L_2(Q_2)(O) &= \sum_{t,l} R_2(t,l) \log \bar{Q}_{t,2,l}^{L(t,2,l)} (1-\bar{Q}_{t,2,l})^{1-L(t,2,l)} \\ L_3(Q_3)(O) &= \sum_{t,l} R_3(t,l) \log \bar{Q}_{t,3,l}^{L(t,3,l)} (1-\bar{Q}_{t,3,l})^{1-L(t,3,l)}, \end{aligned}$$

where $R_1(t) = I(L(t-1,1,1)=0)$, $R_2(t,l)$ is the indicator of $L(t,2,l)$ still being at risk (i.e., it is not a deterministic function of $Pa(L(t,2,l))$), and similarly $R_3(t,l)$ is the indicator of $L(t,3,l)$ being at risk. Here $\bar{Q}_{t,1,1}$ is short-hand notation for $Q_{t,1,1}(1 \mid Pa(L(t,1,1)))$, and we used the same short-hand notation for the other binary conditional distributions.

Estimate treatment/censoring mechanism:

Similarly, we can estimate the conditional binary distribution of $A(0)$, given W, with logistic regression, and we can estimate the conditional binary probability distributions of $A(t)$, given $Pa(A(t))$, based on logistic regression based on a pooled data set for $(A(t), Pa(A(t))) : t)$, pooling over time. Again, we recommend loss-based super learning with the following two log-likelihood loss functions:

$$\begin{aligned} L(g_{A(0)})(O) &= -\log g_{A(0)}(A(0) \mid W) \\ L(g)(O) &= -\sum_t R_g(t) \log g_{A(t)}(A(t) \mid Pa(A(t))), \end{aligned}$$

where $R_g(t)$ is the indicator of $A(t)$ being at risk (i.e., $A(t)$ is not yet a deterministic function of $Pa(A(t))$).

The resulting estimates $g_{A(0),n}$ and $g_{A(t),n}$ are then used for the construction of $C_{tjl}(g_n)$ below.

3.7 Targeted MLE algorithm based on log-likelihood loss function

Now that we have obtained the initial estimates, $Q_{1,n}$, $Q_{2,n}$, $Q_{3,n}$, $g_{A(0),n}$, and $g_{A(t),n}$ we are able to implement the TMLE algorithm. For notational convenience, let $\bar{Q}_{tjl}$

denote the conditional probability of $L(t,j,l)=1$, given $Pa(L(t,j,l))$. Consider the following parametric fluctuations of the initial estimator Q_n:

$$\begin{aligned} logit\bar{Q}_{(t,1,1),n}(\varepsilon) &= logit\bar{Q}_{(t,1,1),n}+\varepsilon_1 C_{t11}(Q_n,g_n) \\ logit\bar{Q}_{(t,2,l),n}(\varepsilon) &= logit\bar{Q}_{(t,2,l),n}+\varepsilon_2 C_{t2l}(Q_n,g_n) \\ logit\bar{Q}_{(t,3,l),n}(\varepsilon) &= logit\bar{Q}_{(t,3,l),n}+\varepsilon_3 C_{t3l}(Q_n,g_n). \end{aligned}$$

In addition, we fluctuate $Q_{L(0),n}$ with a parametric submodel $Q_{L(0),n}(\varepsilon_0)$ that has score $D^*_{L(0)}(Q_n)$ at $\varepsilon_0=0$, but this submodel will play no role in the TMLE since the MLE of ε_0 will be equal to zero. This defines a submodel $\{Q_n(\varepsilon):\varepsilon=(\varepsilon_0,\varepsilon_1,\varepsilon_2,\varepsilon_3)\}$ through the estimator Q_n. We can combine this submodel with the log-likelihood loss $L(Q)=-logQ$, so that the score condition indeed holds: the linear span of the components of the score at zero fluctuation includes the efficient influence curve at (Q_n,g_n).

A variety of iteratively updating algorithms are proposed in van der Laan (2010a). In a TMLE defined by this log-likelihood loss and logistic regression parametric submodels, one can fit ε_j with logistic regression by regressing the binary indicator $L(t,j,l)$ onto the clever covariate $C_(t,j,l)$ with the logit of the initial $Q_{(t,j,l),n}$ as an offset. This can be done with the repeated measure data set for the corresponding j. This approach corresponds with the approach in van der Laan (2010a) that uses the same pooling to fit ε as was used for the initial estimator. Alternatively, we enforce $\varepsilon_1=\varepsilon_2=\varepsilon_3=\varepsilon$, and the single common ε is fitted with single pooled repeated measures data set resulting in a common fluctuation ε_n for all t, j, and k, as also presented in van der Laan (2010a). This updating process is iterated till convergence (i.e., $\varepsilon_n=0$), and the resulting $Q^*_{(t,j,l),n}$ represent the TMLE Q^*_n of Q_0, and thereby the TMLE $\Psi(Q^*_n)$ of ψ_0.

Finally, one can also use the closed form single step TMLE in van der Laan (2010a) that updates one factor at the time, starting at the last factor, and ending at the update of the first binary conditional distribution in the ordering, always using the most recent update of Q_n in the calculation of the clever covariates. This algorithm was shown to converge in a single round, and therefore represents a closed form implementation of TMLE. In this case, each parametric submodel through $\bar{Q}_{(t,j,l),n}$ has its own fluctuation parameter $\varepsilon(t,j,l)$. We will use this single step TMLE algorithm approach, which is computationally the most attractive TMLE among these different types of TMLE.

3.8 Markov property for initial estimate to speed up algorithm

The implementation of the one-step TMLE above is still time consuming and non-trivial to implement for general longitudinal data structures. An examination of $C_{tjl}(Q_n)$ (Equation (9)) reveals the computational complexity. Each clever covariate requires the evaluation of both $P[L_{a,0}(K+1,1,1)=0 \mid L(t,j,l)=1, Pa(L(t,j,l))]$ and $P[L_{a,0}(K+1,1,1)=0 \mid L(t,j,l)=0, Pa(L(t,j,l))]$ under the most recent update of Q_n. These evaluations may be calculated using either a Monte Carlo approach or by integration over all paths in the G-computation formula for the conditional probability that $L_{a,0}(K+1,1,1)=0$, given $L(t,j,l), Pa(L(t,j,l))$. In either case, the fact that each $Q_{(t,j,l),n}$ is a function of the entire history through $Pa(L(t,j,l))$ makes these approaches computationally costly since the number of paths resulting in no event up till time $K+1$ is exponential in the number nodes (in particular, in the number of time-points). This computational cost is further exaggerated by the fact that each clever covariate has to be evaluated for each subject in the data set for each $L(t,j,k)$. One way to simplify the amount of necessary computations is to enforce that the conditional distribution of $L(t,j,l)$ (under Q_n) is not a function of the entire history, but rather, a function of the most recent history or some subset of the subject's history. By enforcing this Markov type property on the estimate Q_n, each conditional probability $Q_{(t,j,l),n}(1 \mid Pa(L(t,j,l))$ has only few possible realizations as a function of $Pa(L(t,j,,l))$, so that the number of values of Q_n over which to integrate in the expression for $C_{tjl}(Q_n)$ is linear in the number of binary variables.

For the purpose of our analysis of the Tshepo study we enforce the following Markov property on the initial estimate Q_n. Firstly, it is assumed that $\bar{Q}_{(t,1,1),n}$ is only a function of $Pa(L(t,1,1))$ with $\bar{A}(t-1)=0, \bar{L}(t-1,1)=0$ (i.e., no censoring, no event, yet) through $L(0), A(0)$ and CD4 $L(t-1,2)$ and viral load $L(t-1,3)$ at previous time-point $t-1$. Secondly, it is assumed that $\bar{Q}_{(t,2,l),n}$ is only a function of $Pa(L(t,2,l))$ with $\bar{A}(t-1)=0, \bar{L}(t-1,1)=0$, and $L(t,2,1)=\ldots,=L(t,2,l-1)=0$ (i.e., no censoring, no event, yet, and CD4 has at least level l) through $L(0), A(0)$ and CD4 $L(t-1,2)$ and viral load $L(t-1,3)$ at previous time-point $t-1$. Finally, it is assumed that $\bar{Q}_{(t,3,l),n}$ is only a function of $Pa(L(t,3,l))$ with $\bar{A}(t-1)=0, \bar{L}(t-1,1)=0$, and $L(t,3,1)=\ldots,=L(t,3,l-1)=0$ (i.e., no censoring, no event, yet, and viral load has at least level l) through $L(0), A(0)$ and CD4 $L(t,2)$ at time t, and viral load $L(t-1,3)$ at previous time-point $t-1$.

Thus, for the purpose of our analysis we assumed that each conditional distribution of $L(t,j,l)$ was a function of the time dependent covariates through the most recently observed levels. In other analyses these assumptions may be relaxed to include multiple time points, summary metrics such as functions of the most recent history of the time dependent covariates (e.g., a slope of past CD4-count

process). In addition, cross-validation can be used to adaptively select the degree of dimension reduction applied to the histories, such as the degree of the Markov property, so that, if it is necessary to incorporate more time points of the past, then the algorithm will select accordingly. In this way, the number of calculations are controlled, but still adaptive to what is needed to fit Q_0 well.

3.9 TMLE at weighted-log-likelihood loss function

Even if we enforce this Markov assumption on the initial estimate Q_n, note that $C_{tjl}(g_n)$ is still a function of the full history $Pa(A(t))$ through $g_{A(t)}$, so that the updates of Q_n during the single step TMLE algorithm would still map into k-step updates Q_n^k that will *not* satisfy the Markov property. This issue will be addressed by moving $C_{tjl}(g_n)$ from being a factor of the clever covariate to being a weight in the log-likelihood loss function. That is, we use a weighted logistic regression for each update with weights equal to $C_{tjl}(g_n)$, and a new clever covariate $C_{tjl}(Q_n)$ instead of $C_{tjl}(Q_n, g_n)$. This corresponds with using a weighted-log-likelihood loss $L_g(Q) = -\log Q_{L(0)} - \sum_{tjl}\{logQ_{tjl}\}C_{tjl}(g)$ and fluctuating the initial estimator of the conditional distribution $Q_{L(t,j,l),n}$ by adding the clever covariate extension $\varepsilon C_{tjl}(Q_n)$ on the logit scale. Thus, we now use the following parametric fluctuations of the initial estimator Q_n:

$$\begin{aligned} logit\bar{Q}_{(t,1,1),n}(\varepsilon) &= logit\bar{Q}_{(t,1,1),n} + \varepsilon_{t,1,1}C_{t11}(Q_n) \\ logit\bar{Q}_{(t,2,l),n}(\varepsilon) &= logit\bar{Q}_{(t,2,l),n} + \varepsilon_{t,2,l}C_{t2l}(Q_n) \\ logit\bar{Q}_{(t,3,l),n}(\varepsilon) &= logit\bar{Q}_{(t,3,l),n} + \varepsilon_{t,3,l}C_{t3l}(Q_n). \end{aligned}$$

We still fluctuate $Q_{L(0),n}$ with a parametric submodel $Q_{L(0),n}(\varepsilon_0)$ that has score $D^*_{L(0)}(Q_n)$ at $\varepsilon_0 = 0$, but this submodel will play no role in the TMLE since the MLE of ε_0 will be equal to zero. This defines now a submodel $\{Q_n(\varepsilon) : \varepsilon\}$, and this submodel is combined with the weighted-log-likelihood loss $L_g(Q)$.

This weighted log-likelihood $L_g(Q)$ loss and parametric submodel $Q(\varepsilon)$ map into the same desired score $\frac{d}{d\varepsilon}L_g(Q(\varepsilon))$ at zero fluctuation $\varepsilon = 0$ as the unweighted log-likelihood and the parametric submodel using the $C_{tjl}(Q_n, g_n)$ as clever covariates. Thus, also this weighted-log-likelihood and submodel satisfies that its generalized score at zero fluctuation spans the components of the efficient influence curve at (Q, g). The major advantage of moving $C_{tjl}(g_n)$ into the weight of the loss-function is that it only requires that $C_{tjl}(g_n)$ be evaluated for each observed history and not at all possible histories as required for evaluation of the clever covariates. Thus, with the Markov property on Q_n, and changing the clever covariate

to $C_{tjl}(Q_n)$, the dependence of the clever covariate in the logistic regression fluctuations on the entire past has been removed. The TMLE algorithm now increases in time linearly with each additional $L(t,j,l)$ added to the graph as opposed to exponentially. As a result, the algorithm is now computationally feasible without making major restrictive assumptions. In fact, the resulting algorithm is faster than the iterative TMLE algorithm used in Stitelman and van der Laan (2011) and Stitelman and van der Laan (2010) that only adjusted for baseline covariates.

3.10 Using the iterative conditional expectation formula to speed up computations and minimize memory

Another important way to simplify the algorithm is to take advantage of iterative conditional expectations so that when one works back through factors of the likelihood all necessary evaluations of $Q^*_{(t,j,l),n}$ that are needed for subsequent steps are evaluated directly after the update. In this way each $P[L_{a,0}(K+1,1,1)=0 \mid L(t,j,l)=\delta, Pa(L(t,j,l))]$ may be written as a simple iterative conditional expectation of the already evaluated conditional expectations and the newly updated $Q^*_{(t,j,l),n}$ for the binary variable $L(t,j,l)$. To understand this, represent the longitudinal data structure O as the ordered sequence $O(l)$, $l=0,\ldots,L$, where $O(0)=W$, $O(1)=A$, and several of subsequent $O(l)$ correspond with $A(t)$, and all other $O(l)$ are indicators coding the death-process, viral load process and CD4-count process. Suppose that $O(k)$ is an L-indicator and we already evaluated the clever covariate for this L-indicator and also computed the TMLE update for the conditional distribution of this L-indicator. We now wish to determine the clever covariate and TMLE-update for the next L-indicator in the sequence, going backwards. Now, we note that

$$P(L_{a,0}(K+1,1,1)=0 \mid O(k-1), Pa(O(k-1))) =$$
$$\sum_{o(k)} P(L_{a,0}(K+1,1,1)=0 \mid O(k)=o(k), Pa(O(k)))P(O(k)=o(k) \mid Pa(O(k))).$$

If $O(k-1)$ is also an L-indicator, then the above relation allows us to map the previous clever covariate and the last updated conditional probability of $O(k)$ into the clever covariate for the conditional distribution of $O(k-1)$. If $O(k-1)$ is a censoring $A(t)$-node, then it follows that, at the only relevant value zero for this censoring node (thus equal to the intervention used in the G-computation formula), the left-hand side equals $P(L_{a,0}(K+1,1,1)=0 \mid O(k-2), Pa(O(k-2)))$, so that we have

$$P(L_{a,0}(K+1,1,1)=0 \mid O(k-2), Pa(O(k-2))) =$$
$$\sum_{o(k)} P(L_{a,0}(K+1,1,1)=0 \mid O(k)=o(k), Pa(O(k)))P(O(k)=o(k) \mid Pa(O(k))).$$

So, again, this allows us to map the previous clever covariate and the last updated conditional probability of $O(k)$ into the clever covariate for the next Q-conditional distribution of $O(k-2)$.

3.11 Evaluation of target parameter of targeted MLE

Note also that at the final step of the iterative backwards one-step TMLE algorithm, we have to evaluate the clever covariate expression $P(L_{a,0}(K+1,1,1)=0 \mid W,A=a)$ for all n observed values of W. The empirical mean of the latter is now the TMLE $\Psi(Q_n^*)$! Thus the final evaluation of the target parameter of interest is a natural by-product of completing the iterative backwards single step TMLE algorithm.

3.12 Statistical Inference

Influence curve based inference is obtained in exactly the same manner as presented in van der Laan (2010b). In particular, if one assumes that g_n is a consistent estimator of g_0, then one can use the conservative influence curve $D^*(Q,g_0)$, where Q represents the limit of the TMLE Q_n^*. Thus, the asymptotic variance of $\sqrt{n}(\Psi(Q_n^*)-\Psi(Q_0))$ can be conservatively estimated with

$$\sigma_n^2 = \frac{1}{n}\sum_{i=1}^{n}\{D^*(Q_n^*,g_n)(O_i)\}^2.$$

Note that during the backward solving TMLE algorithm one has to evaluate each of the contributions $D^*_{tjl}(Q_n^*,g_n)$ and $D^*_0(Q_n^*,g_n)$ that make up the efficient influence curve $D^*(Q_n^*,g_n)$, as represented in equation (8). Thus, both the final evaluation of the parameter of interest and the statistical inference for that parameter are provided by implementing the backward solving algorithm proposed here.

4 Simulation Study

In this section we present the results of simulation studies that compare the bias and efficiency of six different estimators of the treatment specific survival curve $S_1(t_0)$: Baseline TMLE, Baseline IPCW, Baseline A-IPCW, Time-Dependent TMLE, Time-Dependent IPCW, Time-Dependent EE. Baseline refers to the data structure that excludes the time-dependent covariates, and EE is an an abbreviation for an estimating equation based estimator we developed for the complete longitudinal data structure (it can be viewed as an A-IPCW of the type presented in van der Laan and Robins (2003), but it is based on the representation of the efficient influence curve

as used in the TMLE). The EE involves representing the efficient influence curve for the longitudinal data structure as an estimating function in the target parameter ψ_0, and defining the estimator as the solution of the corresponding estimating equation, estimating the nuisance parameters with the initial estimators as used in the TMLE. No similar estimating equation based estimators have gained traction in the literature due to the computational difficulties of constructing such an estimate when there are many time points and intermediate variables. The representation of the efficient influence curve (8) and the corresponding estimating equation based estimator of ψ_0 as we implemented here make this estimating equation based estimator computationally feasible. The EE is just like the TMLE a double robust locally efficient estimator, but the TMLE is also a substitution estimator, while the EE is not. The Time-Dependent IPCW is defined as the empirical mean of

$$D_{IPCW}(O) = \frac{I(T > t_0, A = 1, C > t_0)}{\bar{G}_n(t_0- \mid X, A = 1) g_{A(0),n}(A(0) \mid W)},$$

where $g_{A(0),n}$ is an estimator of the treatment mechanism g_0, conditional on baseline covariates, $\bar{G}_n(t- \mid X, A = 1) = \prod_{s<t}\{1 - g_{A(s),n}(1 \mid Pa(A(s)))\}$ is the estimator of the survivor function of censoring, conditional on baseline treatment, baseline covariates, and time-dependent covariates.

The goal of the first set of simulations presented here is to illustrate the bias reduction that occurs when one adjusts for time-dependent covariates that affect drop-out beyond the effect of the baseline covariates on time to drop-out. The second set of simulations show that if censoring is non-informative, a TMLE and EE incorporating the available time-dependent covariates improve efficiency relative to an estimator that ignores the time-dependent covariates, even though in this independent censoring scenario the latter is still a valid asymptotically linear estimator. Furthermore, our simulations also demonstrate that a locally efficient double-robust substitution estimator (Time Dependent TMLE) performs better in finite samples than both a locally efficient double-robust non-substitution estimator (Time Dependent EE) and the current standard for accounting for time-dependent covariates (Time Dependent IPCW). In fact, the simulations suggest that the benefit of targeted learning increases quickly, and dramatically, when the complexity (e.g., dimension of data structure) of the estimation problems increases.

In our simulations we simulate a longitudinal data structure

$$O = (W(0), A(0), N(1), W_4(1), W_5(1), A(1) ..., N(K), W_4(K), W_5(K), A(K), N(K+1)),$$

for $t = 1, ..., K+1$. Here $W(0) = (W_1(0), W_2(0), W_3(0), W_4(0), W_5(0))$ are the baseline covariates, $A(0)$ is the binary baseline treatment randomized with probability 0.5, $N(t)$ is the indicator of observing a failure time event at time t, $A(t)$ is the

indicator of observing a censoring event at time t, and $W_4(t)$ and $W_5(t)$ are the continuous time-dependent covariates. In each simulation, 500 simulated data sets with sample size $n = 500$ were generated, the treatment specific survival curve $S_1(t_0)$ at time point $t_0 = 3$ was estimated using each of the six different estimators, and estimates of bias and MSE were reported. In each simulation the true treatment specific survival $S_1(t_0)$ equals .469. All estimators were supplied consistent estimators of the conditional intensity of the censoring process, and failure-time process, while the conditional distributions of the time-dependent covariates were estimated inconsistently by discretizing the continuous covariates $W_4(t)$, $W_5(t)$, coding these discretized covariates with binary indicators, and estimating the conditional distribution of the binary indicators with logistic parametric regression.

4.1 Simulations with Informative Censoring

The precise data generating mechanism is described as follows.
1. The drawing of the baseline covariates $W(0)$ involved first generating from a mean zero multivariate normal and truncating any component from above by 2 and from below by -2. The covariance matrix was defined as 1 on the diagonal and 0.2 off-diagonal. The truncation was enforced to ensure that the censoring mechanisms was not suffering too much from practical violations of the positivity assumption, as required for identifiability of $S_1(t_0)$.
2. The two time-dependent covariates $W_4(t)$ and $W_5(t)$ are generated as follows:

$$W_4(t) = .2A(0) + .5W_1(0) - .4W_2(0) - .4W_3(0) + 2W_4(t-1) + 2W_5(t-1) + U_4$$
$$W_5(t) = .1A(0) + .1W_1(0) + .1W_2(0) - .4W_3(0) + 2W_4(t) + 2W_5(t-1) + U_5,$$

where U_4 and U_5 are i.i.d. $N(0, \sigma = 0.4)$.
3. The event indicators, $N(t)$, were generated as Bernoulli-indicators with the probability defined by the following conditional intensity of time to failure T:

$$\lambda_T(t) = \text{expit}(-3 + .3A(0) + .3W_1(0) - .3W_2(0) - .3W_3(0) + 2W_4(t-1) + 2W_5(t-1)).$$

4. The censoring indicators, $A(t)$, were generated as Bernoulli-indicators with the probability defined by the following conditional intensity for censoring for the low and high informative censoring case, respectively:

$$\lambda_C(t) = \text{expit}(-4 + .8A(0) + .3W_1(0) - .3W_2(0) - .3W_3(0) + .1W_4(t) + .1W_5(t-1))$$
$$\lambda_C(t) = \text{expit}(-4 + .8A(0) + .3W_1(0) - .3W_2(0) - .3W_3(0) + 1W_4(t) + 1W_5(t-1)).$$

The results are presented in Table 1. Each table below presents the mean of the estimates, mean of the influence curve based standard errors, mean square error,

and the coverage probabilities for 95 percent wald-type influence curve based confidence intervals for each of the estimators investigated. The low-informative censoring results show 1) that the TMLE and EE estimators that only use the baseline-covariates are very similar to the estimators that incorporate the time-dependent covariates, and 2) the Time-Dependent IPCW is highly inefficient relative to the other estimators. The simulation for the high-informative censoring shows some interesting results. Firstly, the estimators that only incorporate the baseline-covariates are highly biased: the MSE of the Baseline estimators are over 13 times larger than the MSE of the Time-Dependent TMLE. Secondly, the Time-Dependent TMLE has an MSE that is almost 75% smaller than the MSE of the Time-Dependent EE, demonstrating the crucial benefit of being a substitution estimator beyond being a double robust efficient estimator.

Interestingly, in this particular case, the Time-Dependent IPCW estimator performs remarkably well. However, it can be explained as a lucky scenario where a biased estimator happens to nail the right answer. This has to do with the fact that the covariates that strongly effect the event are also very predictive of censoring causing the IPCW estimator to do artificially well in this scenario. This is because the High Censoring scenario is a simulation where the informative censoring is so extreme that there are levels of covariates that are so predictive of censoring that in finite samples it is extremely rare to find an uncensored individual at those levels of the covariates. Moreover, those same levels of the covariates are extremely predictive of the event, so much so that by the time point of interest the event will have occurred with almost probability of 1 for those individuals. As a result the contributions to the IPCW estimator for individuals that have a high probability of being censored is always zero, which is exactly the right contribution, since the probability of the event for those individuals happening before the time of interest is also essentially 1. We show below that if the direction of the effect of the baseline variables on the censoring is switched the IPCW does very poorly. Apparently, a change in the censoring mechanism dramatically affects the MSE of the IPCW-estimator, demonstrating that this initial finding represents a-typical behavior of the IPCW-estimator. This is because those individuals that are at levels of W that are almost completely predictive of an event before the time of interest are no longer at levels of W that are almost completely predictive of censoring.

In our modified simulation, we generated the censoring events for the low and high informative censoring case as follows:

$$\lambda_C(t) = \text{expit}(-4 + .8A(0) + .3W_1(0) - .3W_2(0) - .3W_3(0) - .01W_4(t) - .01W_5(t-1)),$$
$$\lambda_C(t) = \text{expit}(-4 + .8A(0) + .3W_1(0) - .3W_2(0) - .3W_3(0) - .1W_4(t) - .1W_5(t-1)).$$

Table 2 presents the results for this simulation. Again, the incorporation of the time-dependent covariates results in an important bias reduction (and MSE) for the

Low Informative Censoring Scenario

	Time Dependent			Baseline		
	TMLE	EE	IPCW	TMLE	A-IPCW	IPCW
Mean of Estimates	0.469	0.469	0.486	0.475	0.475	0.475
Mean SE	0.027	0.027	0.041	0.027	0.027	0.040
Mean Square Error	0.00070	0.00070	0.00113	0.00076	0.00076	0.00077
Coverage	0.942	0.942	0.986	0.940	0.938	0.996

High Informative Censoring Scenario

	Time Dependent			Baseline		
	TMLE	EE	IPCW	TMLE	A-IPCW	IPCW
Mean of Estimates	0.479	0.470	0.475	0.587	0.585	0.595
Mean SE	0.029	0.035	0.039	0.034	0.034	0.059
Mean Square Error	0.00112	0.00440	0.00073	0.01485	0.01453	0.01740
Coverage	0.898	0.898	0.996	0.066	0.074	0.352

Table 1: Simulation Results For Informative Censoring: Mean of Estimates and Mean Square Error for All Six Estimators

TMLE and EE estimators. In the low informative censoring simulation, the Time-Dependent IPCW estimator has an MSE that is 1.6 times as large as the MSE of the Time-Dependent TMLE and EE estimator. In the high informative censoring scenario, the MSE of the Time-Dependent IPCW estimator is 128 (!) times as large as the MSE of the Time-Dependent TMLE and EE estimator. The latter demonstrates a complete break down of the IPCW-estimator, reflecting that it is simply a very unreliable estimator, even though it represents current practice.

4.2 Simulations with Independent Censoring

The data generating distribution is as above, except that the censoring mechanism was modified again. The hazard of censoring was now only a function of time, so that censoring is independent of the evolving processes, but three different hazards were considered representing different levels of independent censoring: no censoring, medium censoring, and high censoring. In the first scenario every individual was left uncensored. In the second and third scenario each subject was censored with 20 percent probability (Medium Censoring Scenario) and 60 percent probability (High Censoring Scenario), respectively.

The results are presented in Table 3. We know that under independent censoring all 6 estimators are consistent. Indeed, the results demonstrate that all estimators are unbiased across the three simulations, so that the estimators only differ in

Low Informative Censoring Scenario

	Time Dependent			Baseline		
	TMLE	EE	IPCW	TMLE	A-IPCW	IPCW
Mean of Estimates	0.470	0.470	0.452	0.469	0.469	0.470
Mean SE	0.027	0.027	0.040	0.027	0.027	0.042
Mean Square Error	0.00065	0.00066	0.00105	0.00068	0.00067	0.00077
Coverage	0.960	0.960	0.974	0.956	0.958	1.000

High Informative Censoring Scenario

	Time Dependent			Baseline		
	TMLE	EE	IPCW	TMLE	A-IPCW	IPCW
Mean of Estimates	0.468	0.468	0.174	0.432	0.433	0.396
Mean SE	0.027	0.027	0.026	0.033	0.033	0.067
Mean Square Error	0.00068	0.00067	0.08732	0.00251	0.00241	0.00731
Coverage	0.960	0.960	0.000	0.798	0.810	0.836

Table 2: Simulation Results For Informative Censoring: Mean of Estimates and Mean Square Error for All Six Estimators

their efficiency (i.e., variance). In the no-censoring scenario, all estimators behave similarly, with the exception of the IPCW-estimators that are somewhat inefficient. Gains in efficiency due to utilizing the time-dependent covariates can only be expected if a significant proportion of the subjects are right-censored, since an efficient estimator treats a censored subject that is very sick at the censoring time differently than a censored subject that was relatively healthy at the censoring time. Indeed, the table shows that as the amount of independent censoring increases, the IPCW-estimators become more and more inefficient relative to the efficient TMLE and EE estimators. It is also of interest to note that, for the high censoring scenario, the Time Dependent TMLE is almost 1.8 times as efficient as the Baseline TMLE. This demonstrates the substantial gain in efficiency one can obtain by utilizing time-dependent covariates. Furthermore, we note that in the high censoring scenario the locally efficient double-robust non-substitution estimator (Time Dependent EE) has a mean square error of almost 2.25 (!) times the MSE of the locally efficient double-robust substitution estimator (Time Dependent TMLE). This demonstrates, once again, the enormous importance of being a substitution estimator. This gain is most likely due to estimated censoring probabilities that are empirically imbalanced across strata of the covariates, so that the estimators behave similarly as in a high-informative censoring simulation. Finally, it is noteworthy that the Time Dependent IPCW estimator has a mean square error over six times as large as the MSE of the Time Dependent TMLE.

No Censoring Scenario

	Time Dependent			Baseline		
	TMLE	EE	IPCW	TMLE	A-IPCW	IPCW
Mean of Estimates	0.468	0.468	0.468	0.468	0.468	0.468
Mean SE	0.027	0.027	0.038	0.027	0.027	0.038
Mean Square Error	0.00067	0.00068	0.00073	0.00069	0.00069	0.00073
Coverage	0.952	0.952	0.990	0.950	0.950	0.990

Medium Censoring Scenario

	Time Dependent			Baseline		
	TMLE	EE	IPCW	TMLE	A-IPCW	IPCW
Mean of Estimates	0.469	0.470	0.471	0.469	0.469	0.470
Mean SE	0.028	0.028	0.051	0.029	0.029	0.051
Mean Square Error	0.00070	0.00072	0.00120	0.00081	0.00081	0.00106
Coverage	0.960	0.960	0.996	0.952	0.952	1.000

High Censoring Scenario

	Time Dependent			Baseline		
	TMLE	EE	IPCW	TMLE	A-IPCW	IPCW
Mean of Estimates	0.474	0.481	0.474	0.467	0.467	0.466
Mean SE	0.044	0.047	0.114	0.043	0.042	0.112
Mean Square Error	0.00110	0.00248	0.00712	0.00196	0.00197	0.00496
Coverage	0.988	0.988	0.978	0.940	0.940	0.984

Table 3: Simulation Results For Independent Censoring: Mean of Estimates and Mean Square Error for All Six Estimators

4.3 Simulations - Confidence Interval Coverage and Width

In each of the tables above we show the 95 percent confidence interval coverage probabilities. These confidence intervals were constructed by relying on the fact that the TMLE solves the efficient influence curve estimating equation. In three of the four informative censoring scenarios the TMLE produces valid 95 percent confidence intervals. In the fourth, the high informative censoring scenario for the first simulation (Table 1), the TMLE has a coverage probability of 89.8 percent, which is less than ideal. In all cases the confidence intervals constructed for the estimators using only baseline variables are far less than the desired 95 percent coverage for the high informative censoring scenarios.

For each scenario we explored if the coverage was different if the true g_0 was used instead of the g_n-fit obtained by using logistic regression with the known model for treatment and censoring. This exercise was carried out because asymptotically the coverage probabilities should be 95 percent when using the true g_0 (van der

Laan and Robins (2003)). The coverage probabilities remained unchanged when using the true g_0. Thus, in the high informative censoring scenario for the first simulation the coverage probability was still significantly less than the desired 95 percent coverage at 88.8 percent (First line of Table 4). This suggests that the departure from 95 percent in the coverage probabilities is due to the fact that the sample size is too small to produce the asymptotically valid coverage, and not that there is an issue in the way g_0 was estimated. This is confirmed by Table 4 which shows the results for simulations with different sample sizes at the true g_0. As the sample size increases the coverage probabilities begin to approach the desired 95 percent coverage. Even with a sample size of 5,000 the coverage is not 95 percent and this has to do with the large violation in the positivity assumption for certain levels of baseline covariates W in this simulation.

In Stitelman and van der Laan (2010) we show that quantile based bootstrap estimates of the confidence intervals are the preferred method for inference, especially in situations where positivity violations are an issue (as is the case in the high informative censoring scenario for the first simulation above). In that article it was shown that in finite samples the departure from normality was manifested in a skewed distribution of the bootstrap estimates and thus quantile based bootstrap estimates of the confidence intervals performed better than Wald type confidence intervals. Moreover, such an approach takes into account the finite sample variance associated with data adaptively estimating Q_0. This suggests that quantile based bootstrap estimates of the 95 percent confidence intervals are the preferred approach for producing confidence intervals. Bootstrap based confidence intervals are computationally feasible for a single data analysis; however, they are time consuming in simulation and their benefit was already displayed in Stitelman and van der Laan (2010), so they were not explored here.

Sample Size	Mean of Estimates	Mean SE	Mean Square Error	Coverage
500	0.480	0.028	0.0012	0.888
1,000	0.475	0.020	0.0005	0.906
5,000	0.471	0.010	0.0003	0.933

Table 4: Simulation Results For High Informative Censoring Using True g_0: By Sample Size

4.4 Simulations - Mis-specifying both of the initial estimates, Q_n and g_n

The simulations presented in the previous sections have been based on initial estimates of Q_n and g_n that incorporated all potential confounders and used a correct model for g_0, and an approximately correct model for Q_0. For the three estimators that only used baseline information the known model was used excluding the time-dependent components. The intention for the simulation study presented in the current subsection is to illustrate the effect of mis-specifying both of the initial estimates, Q_n and g_n, on the behavior of the different estimators of the target parameter.

The data used for this simulation study were simulated in the same way as the data simulated for the modified high informative censoring scenario of section 4.1. For the study here we evaluate what happens to the simulation results when the time dependent covariates, W_4, W_5, and then both W_4 and W_5 are removed from the models for the initial estimates of Q_n and g_n. This allows us to observe how the different estimators behave when the initial estimates for Q_n and g_n are both initially mis-specified.

Table 5 displays the results of this simulation study. As for the original simulation all of the estimators that only incorporate baseline information continue to perform poorly. These methods initially used mis-specified models for their initial Q_n and g_n since they only incorporate baseline covariates, so it should be no surprise that further mis-specifying the initial models causes the estimators to behave even more poorly in terms of both bias and mean square error. The time dependent IPCW estimator, which was very unstable even when g_n was correctly specified, behaves as poorly as before in terms of both bias, variance and coverage of its confidence intervals with mis-specification. A direct comparison of the time dependent TMLE and EE reveals the stability of the TMLE even when both the initial Q_n and initial g_n are estimated based off of a mis-specified model. Both methods produce slightly biased estimates and in one case the TMLE does slightly better and in the other the EE does slightly better. However, these two methods of the six are the only ones that produce estimates that are on average anywhere close to the truth, .469. In the case where either W_4 or W_5 are removed from the model specification the TMLE is 9 to 12 times more efficient than the EE. Furthermore the TMLE and EE produce similar confidence interval coverage, using Wald type influence curve based variance estimates, but the confidence interval lengths for the TMLE are about half the size of those for the EE (Mean SE differences of 0.034 vs 0.063 and 0.034 vs. 0.066). When both W_4 and W_5 are removed from the specification of the initial model the relative stability of the TMLE is displayed and the fact that the

EE can produce estimates that don't obey the proper model is made obvious. The EE completely breaks down in this situation and the method is very biased, with a mean estimate of 1.243 (outside the proper range!). However, the TMLE remains stable and produces a mean estimate of 0.462. This is because the TMLE is able to adjust the initial Q_n by updating it at each node in the causal graph through the backward passing algorithm. So even though the initial Q_n is mis-specified in terms of the relationship of W_4 and W_5 it is able to readjust in the updating steps, while the EE does not posses this quality. However, the Mean SE which is based on the misspecified estimate of the influence curve does blow up in this situation for both the TMLE and the EE. Thus, the coverage probabilities are 1 but they are so large that they are useless in practice. However, the stability of the TMLE estimates suggests that quantile based confidence intervals constructed with the nonparametric bootstrap would still produce reasonable confidence intervals. This demonstrates an important advantage (i.e., robustness property) of the nonparametric bootstrap relative to influence curve based inference that relies on consistent estimation of g_0.

Correctly Specifying Initial Models

	Time Dependent			Baseline		
	TMLE	EE	IPCW	TMLE	A-IPCW	IPCW
Mean of Estimates	0.468	0.468	0.174	0.432	0.433	0.396
Mean SE	0.027	0.027	0.026	0.033	0.033	0.067
Mean Square Error	0.00068	0.00067	0.08732	0.00251	0.00241	0.00731
Coverage	0.960	0.960	0.000	0.798	0.810	0.836

Removing $W_4(t)$ From Initial Model Specification

	Time Dependent			Baseline		
	TMLE	EE	IPCW	TMLE	A-IPCW	IPCW
Mean of Estimates	0.457	0.455	0.172	0.420	0.421	0.411
Mean SE	0.034	0.063	0.026	0.035	0.036	0.067
Mean Square Error	0.00133	0.01211	0.08893	0.00360	0.00359	0.00512
Coverage	0.900	0.900	0.000	0.740	0.740	0.910

Removing $W_5(t)$ From Initial Model Specification

	Time Dependent			Baseline		
	TMLE	EE	IPCW	TMLE	A-IPCW	IPCW
Mean of Estimates	0.459	0.461	0.173	0.411	0.411	0.396
Mean SE	0.034	0.066	0.026	0.038	0.038	0.065
Mean Square Error	0.00133	0.01649	0.08840	0.00467	0.00465	0.00725
Coverage	0.920	0.920	0.000	0.640	0.650	0.810

Removing $W_4(t)$ and $W_5(t)$ From Initial Model Specification

	Time Dependent			Baseline		
	TMLE	EE	IPCW	TMLE	A-IPCW	IPCW
Mean of Estimates	0.462	1.243	0.357	0.405	0.405	0.403
Mean SE	0.616	0.619	0.056	0.038	0.038	0.063
Mean Square Error	0.00472	1.02729	0.01415	0.00549	0.00549	0.00604
Coverage	1.000	1.000	0.440	0.590	0.600	0.870

Table 5: Simulation Results For Independent Censoring: Mean of Estimates and Mean Square Error for All Six Estimators

5 Tshepo Analysis Revisited

In an earlier paper we used a targeted maximum likelihood estimator (TMLE) to assess the causal effects of different cART treatments on the time until HIV viral progression. That analysis was based on the Tshepo study, an open-label, randomized, 3x2x2 factorial design HIV study conducted at Princess Marina Hospital in Gaborone, Botswana to evaluate the efficacy, tolerability, and development of drug resistance of six different first-line cART regimens. In particular we focused on the effect of two NNRTI-based cART therapies to which subjects were randomized. The two therapies of interest were efavirenz (EFV) and nevirapine (NVP) and we assessed the causal effect of treatment as well as whether gender modified the effect of the therapy. The initial paper illustrated the advantages of using TMLE to estimate causal effects on time to event outcomes as opposed to the Cox proportional hazards model, the typical approach in this setting.

Our initial analysis of the Tshepo study was based on TMLEs of the causal effect of the treatment on survival, and corresponding effect-modification parameters, *only adjusting for the baseline covariates* (Stitelman and van der Laan (2011)). We extend here this TMLE to account for potential bias due to informative censoring by time-dependent covariates CD4 and viral load that have an effect on both time to drop-out and the time to event of interest. We will directly compare results using this TMLE that only incorporates the baseline covariates to the TMLE that accounts for time dependent confounding in the form of informative censoring due to the time-dependent covariates. Moreover, we will compare these results to results based on an IPCW estimator and a locally efficient double robust estimating equation based estimator.

For the analysis performed here we evaluate the effect modification of gender on the two cART treatments for two outcomes of interest:

1. Time to death censored by treatment modification or end of study (DEATH).
2. Time to minimum of virologic failure, death, or treatment modification censored by end of study (TLOVR).

For each of the two time to event outcomes we will estimate the mean (over the 36 months) of the difference in additive risk by gender over the first 36 months after randomization to cART therapy. For each of the two time to event outcomes we will estimate this parameter using the six estimators examined in the simulation analysis in the previous section. We will also report the difference in additive risk by gender at 36 months. Prior to doing this analysis we expected that utilizing the time-dependent covariates should have a small effect on the estimates for the TLOVR outcome since censoring is independent for this time to event outcome. On the other

hand, the time to death is subject to censoring by time to treatment modification which is expected to be informed by CD4 and viral load, so that one might expect a bias reduction for the new TMLE relative to the previously implemented TMLE that only incorporated the baseline covariates.

Table 6 shows the results for the TMLE of treatment effect modification by gender for the TLOVR outcome. As expected there is little difference in the TMLE with only baseline covariates and the TMLE which also incorporates the time-dependent covariates, in the sense that the point estimate, standard error (SE), and p-value are similar. Furthermore, all six of the methods return similar estimates for the mean risk difference over the 36 months. However, the two double robust locally efficient estimators have much lower estimates of the standard error. The method used to estimate the SE for the IPCW estimator is known to be conservative, so a direct comparison in this situation is not appropriate. However, if one looks at the risk difference at 36 months the point estimates do change slightly for the two double robust locally efficient estimators that take into account time dependent covariates(TD TMLE and TD EE) compared to their baseline counterparts(BASE TMLE and BASE EE). Given our simulation results and the supporting theory, this change in the point estimate may be attributed to an efficiency gain due to an adjustment for empirical confounding, chance imbalance between the confounders for different levels of censoring. The fact that the IPCW estimator does not change is just further evidence of this estimator's inability to efficiently extract information from the data. Overall, these changes do not make an appreciable difference in the conclusions drawn from the results. However, in alternative analyses these differences could be larger especially in situations with more censoring as seen in the simulations above. The results as a whole indicate that gender does in fact modify the effect of drug treatment on the TLOVR outcome. The same conclusion that was determined based on an analysis that just accounts for baseline confounding.

Mean Risk Difference						
	TD TMLE	TD DR-EE	TD IPCW	BASE TMLE	BASE DR-EE	BASE IPCW
Est	0.132	0.133	0.129	0.126	0.126	0.130
SE	0.039	0.038	0.101	0.038	0.038	0.100
p	0.001	0.001	0.199	0.001	0.001	0.196
Risk Difference @ 36 Months						
	TD TMLE	TD DR-EE	TD IPCW	BASE TMLE	BASE DR-EE	BASE IPCW
Est	0.200	0.201	0.183	0.189	0.189	0.183
SE	0.050	0.049	0.103	0.049	0.049	0.103
p	0.000	0.000	0.074	0.000	0.000	0.074

Table 6: Gender Effect Modification on TLOVR

Mean Risk Difference						
	TD TMLE	TD DR-EE	TD IPCW	BASE TMLE	BASE DR-EE	BASE IPCW
Est	0.039	0.039	0.043	0.033	0.032	0.037
SE	0.017	0.017	0.117	0.017	0.017	0.117
p	0.021	0.019	0.717	0.055	0.058	0.753
Risk Difference @ 36 Months						
Est	0.063	0.065	0.052	0.051	0.051	0.052
SE	0.023	0.023	0.125	0.024	0.024	0.125
p	0.005	0.004	0.680	0.029	0.030	0.680

Table 7: Gender Effect Modification on Death

Table 7 shows the results for treatment effect modification by gender for the death outcome. In this Table 7 we see an appreciable difference in TD TMLE versus BASE TMLE. In fact, this difference changes the way in which the results may be interpreted. In this case we know that there is a large amount of informative censoring since treatment modification is one of the censoring events and individuals modify treatment for many reasons including that there are side effects or the treatment is not working. The point estimate of the mean risk difference using TMLE moves from 3.3% to 3.9% and the significance level changes from .055 to .021, due to the incorporation of the time-dependent covariates. Thus, the significance changes from close to significant at the 95 percent level to significant at the 95 percent level. If one looks at the risk difference at the last time point, 36 months, the difference between TD TMLE and BASE TMLE is even more striking and the change in significance moves from significant at the 95 percent level to significant at the 99.5 percent level. The TD TMLE results indicate that gender does in fact modify the effect of the drug treatment EFV/NVP and the difference in the effect between males and females is on average 3.9 percent and at time 36 months is 6.3 percent. Figure 1 shows the survival curves upon which these parameter estimates are based. The IPCW estimator, due to its instability and inefficient use of the data, is unable to produce any bias reduction by accounting for time dependent covariates in this situation. Figure 2 more clearly depicts the instability of IPCW in situations with sparse outcomes like this one. The figure compares the TD TMLE to the TD IPCW treatment specific survival curve for men treated with EFV. It is clear from these plots that the IPCW estimator is unable to stay stable and produce a monotonic survival curve, while the TMLE remains stable and produces sensible results. In other situations we have observed the IPCW estimator to produce estimates of survival probabilities that exceed 1. These characteristics of the IPCW estimator in estimating the treatment specific survival curve are very detrimental to it being used in practice.

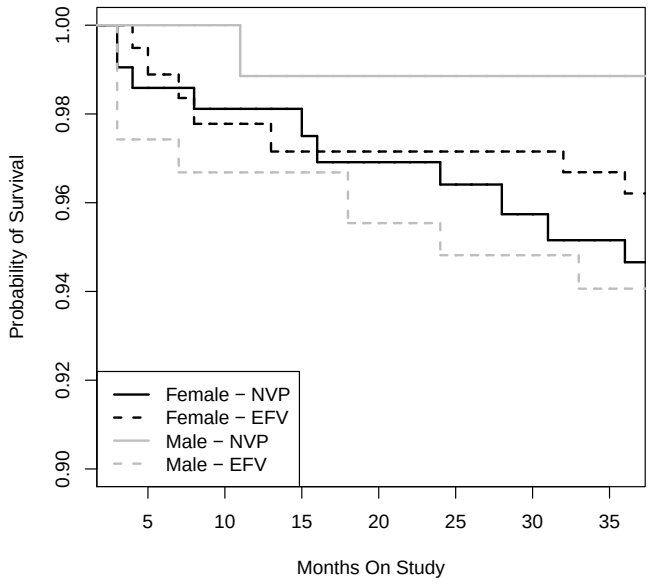

Figure 1: Gender Specific Treatment Specific Survival Curves: Death Outcome

6 Discussion

This article represents the first implementation of TMLE of a causal effect of a multiple time point intervention that is subject to time-dependent confounding by intermediate variables. In this particular case, the multiple time point intervention is represented by a point treatment at baseline, and a time-dependent process that can only jumps once from zero to one, where the latter represents the censoring process. The TMLE presented here generalizes to TMLE of causal effects of any other multiple time point intervention that is subject to time-dependent confounding. This generalization includes the TMLE of the causal effect of a time-dependent treatment or exposure on a time to event outcome that might also be subject to right-censoring, incorporating time-dependent covariate processes to improve efficiency and remove bias.

The enormous challenge in semiparametric estimation of causal effects of multiple time-point intervention has been that incorporating an estimate of the treatment and censoring mechanism can easily do more harm than good. Even estimating equation based estimators, known to be double robust and asymptotically

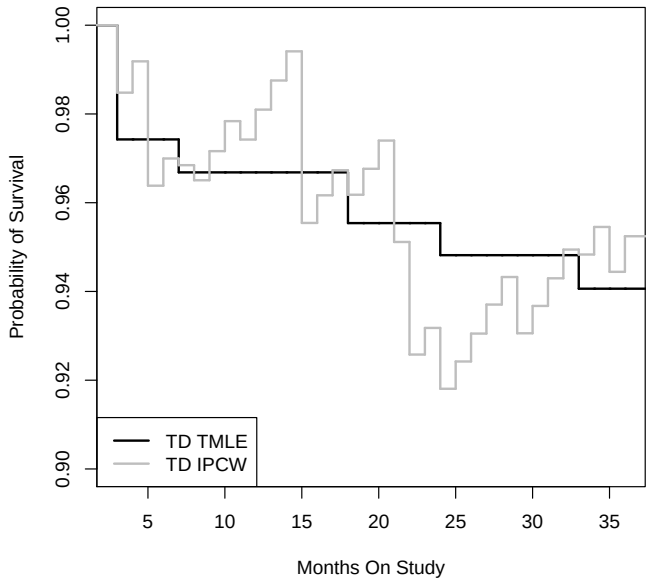

Figure 2: Treatment Specific Survival Curves For Men Treated With EFV: TD TMLE vs TD IPCW

locally efficient, suffer from this instability due to not respecting known global constraints implied by the statistical model. On the other hand, by being a substitution estimator, TMLE fully respects all global constraints implied by the statistical model and the target parameter mapping, while being double robust and locally efficient. For example, consider the TMLE implemented in this article. If at any point in time t for a particular subject the censoring probability approaches 1, then that subject will contribute at that time point t large weights for the TMLE-updates of $Q_{sjl,n}$ for $s \geq t$. That means, such subjects can cause large values of the fluctuation parameters ε_{sjl}. However, these potentially large values of the fluctuation parameters enter on the logistic scale, and can at most cause predicted probabilities for some of the binary variables to approach 1 or 0.

Our simulations and data analyses results demonstrate the remarkable stability of the TMLE that incorporates all measured covariates, reproducing results obtained with robust methods that ignore time-dependent covariates when it is known that censoring is exogenous, while it properly adjusts for time-dependent confounding in the case that the outcome is subject to informative censoring. It is shown that

this stands in sharp contrast to the currently popular IPCW-estimator that is typically not able to properly utilize the measured time-dependent covariates. Moreover, we have shown that even when both the initial estimates of Q_n and g_n are mis-specified the TMLE remains very stable relative to other methods.

We suggest that this TMLE should replace the current analysis of randomized controlled trials with time to event outcomes based on Cox-proportional hazards analysis ignoring both baseline as well as time-dependent covariates, thereby known to be biased whenever there is informative censoring, and also known to be very inefficient. TMLE improves on IPCW, augmented IPCW-estimation as well as on maximum likelihood based methods such as multiple imputation methods, but provides an important marriage between the camps that pursue double robust semiparametric efficient estimators, and the camp that prefers the practically robust maximum likelihood based substitution estimators based on parametric models.

In future work we plan to extend this TMLE to other causal inference problems, and incorporate the C-TMLE extension of TMLE that allows the selection of covariates into the fits of the censoring and treatment mechanism based on the log-likelihood of the resulting TMLE (van der Laan and Gruber (2010)).

References

Gruber, S. and M. van der Laan (2010a): "An application of collaborative targeted maximum likelihood estimation in causal inference and genomics," *The International Journal of Biostatistics*, 6.

Gruber, S. and M. van der Laan (2010b): "A targeted maximum likelihood estimator of a causal effect on a bounded continuous outcome," *The International Journal of Biostatistics*, 6.

Hernan, M. et al. (2005): "Structural accelerated failure time models for survival analysis in studies with time-varying treatments," *Pharmacoepidemiology and Drug Safety*, 14, 477–491.

Hernan, M. et al. (2006): "Comparison of dynamic treatment regimes via inverse probability weighting," *Basic and Clinical Pharmacology and Epidemiology*, 98, 237–242.

Hernan, M. et al. (2009): "Observation plans in longitudinal studies with time-varying treatments," *Statistical Methods in Medical Research*, 18, 27–52.

Hubbard, A., M. van der Laan, and J. Robins (1999): "Nonparametric locally efficient estimation of the treatment specific survival distribution with right censored data and covariates in observational studies," in M. Halloran and D. Berry, eds., *Statistical Models in Epidemiology, The Environment and Clinical Trials*, *IMA*

Volumes in Mathematics and its Applications, Ed. M.E. Halloran and D. Berry, volume 116, Springer Verlag, 135–178.

Lok, J. (2009): "Statistical modeling of causal effects in continuous time," *Annals of statistics*, 36, 1464–1507.

Moore, K. and M. van der Laan (2009): "Application of time-to-event methods in the assessment of safety in clinical trials," in K. E. Peace, ed., *in Design, Summarization, Analysis & Interpretation of Clinical Trials with Time-to-Event Endpoints*, Chapman and Hall.

Pearl, J. (2008): *Causality: Models, Reasoning, and Inference*, Cambridge University Press, Cambridge.

Porter, K., S. Gruber, M. van der Laan, and J. Sekhon (forthcoming 2011): "The relative performance of targeted maximum likelihood estimators," *The International Journal of Biostatistics*.

Robins, J. (1987): "Addendum to: "A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect" [Math. Modelling **7** (1986), no. 9-12, 1393–1512; MR 87m:92078]," *Comput. Math. Appl.*, 14, 923–945.

Robins, J. and A. Rotnitzky (1992): "Recovery of information and adjustment for dependent censoring using surrogate markers," in *AIDS Epidemiology*, Methodological issues, Bikhäuser.

Rose, S. and M. van der Laan (2010): "Why match? investigating matched case-control study designs with causal effect estimation," *The International Journal of Biostatistics*, 5.

Rosenblum, M. and M. van der Laan (2010): "Targeted maximum likelihood estimation of the parameter of a marginal structural model," *The International Journal of Biostatistics*, 6.

Samore, M. et al. (2005): "A simulation-based evaluation of methods to estimate the impact of an adverse event on hospital length of stay." *Medical Care*, 45, 108–115.

Stitelman, O. and M. van der Laan (2010): "Collaborative targeted maximum likelihood for time to event data," *The International Journal of Biostatistics*, 6.

Stitelman, O. and M. van der Laan (2011): "Targeted maximum likelihood estimation of time-to-event parameters with time-dependent covariates," *Forthcoming*.

van der Laan, M. (2010a): "Targeted maximum likelihood based causal inference: Part 1," *The International Journal of Biostatistics*, 6.

van der Laan, M. (2010b): "Targeted maximum likelihood based causal inference: Part 2," *The International Journal of Biostatistics*, 6.

van der Laan, M. and S. Gruber (2010): "Collaborative double robust penalized targeted maximum likelihood estimation," *The International Journal of Biostatistics*.

van der Laan, M., E. Polley, and A. Hubbard (2007): "Super learner," *Statistical Applications in Genetics and Molecular Biology*, 6.

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

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

van der Laan, M. and D. Rubin (2006): "Targeted maximum likelihood learning," *The International Journal of Biostatistics*, 2.

A Appendix A: The general single-step recursive TMLE algorithm

In the paper above we presented an algorithm for producing a TMLE of the treatment specific survival curve in a particular longitudinal data structure. A broad number of questions of interest, in different longitudinal data structures, may be answered using a generalization of the algorithm proposed above. We present here the general single-step recursive algorithm for TMLE in van der Laan (2010a,b) with the modifications introduced in the current article, so that it will be clear how this TMLE can be generalized to a large class of estimation problems based on general longitudinal data structures. In van der Laan (2010b) we demonstrated the TMLE for fitting marginal structural working models for dynamic treatment regimens, using the when to start treatment for HIV-infected patients as the principle scientific question of interest. Therefore, we also use here this general class of estimation problems to demonstrate the general algorithm.

Data and Statistical Model: Let $O = (L(0), A(0), \ldots, L(K), A(K), Y = L(K+1)) \sim P_0$. Assume a nonparametric model or a model that only makes assumptions on the conditional distributions of $g_{0,A(t)}(A(t) \mid Pa(A(t))) = P_0(A(t) \mid Pa(A(t)))$, $t = 0, \ldots, K$. We denote this statistical model for P_0 with $\mathscr{M}$. Let $g_0(A \mid X) \equiv \prod_t g_{0,A(t)}(A(t) \mid Pa(A(t)))$. Assume $Y \in [0,1]$, and Y can be either continuous, discrete, or binary. The variable Y represents the outcome of interest one is able to calculate after having observed the complete longitudinal data structure. For example, Y could be defined as an outcome defined at a fixed time such as $Y = I(T > t)$ for some timepoint t and survival time T coded by the L-process. In the latter case, Y may already be measured before the complete longitudinal data structure O is collected. Alternatively, Y may also be defined as the time to event T itself, or an outcome collected at the end of the study. We will also use the notation $W = L(0)$, and let V be a subset of W. The A-process represents the intervention nodes, and it can have multiple components. For example, $A(t) = (A_1(t), A_2(t))$ has two components, one treatment component and one indicator of drop-out.

Target parameter: Let $L \to d_\theta(L)$ be a rule for assigning values to $A = (A(t) : t = 0, \ldots, K)$. Let L_θ be the random variable whose probability distribution equals the G-computation formula:

$$P(L_\theta = l) = \prod_{t=0}^{K+1} P(L(t) = l(t) \mid \bar{L}(t-1) = \bar{l}(t-1), \bar{A}(t-1) = d_\theta(\bar{L}(t-1))).$$

Let Y_θ be the Y-component of L_θ. Given a working model $\{m_\beta : \beta\}$, we define the target parameter of interest as

$$\Psi(P) = \arg\min_\beta E_P \sum_\theta \left\{E_P(Y_\theta \mid V) - m_\beta(\theta, V)\right\}^2 h(\theta, V),$$

or

$$\Psi(P) = \arg\min_\beta -E_P \sum_\theta \left\{Y(\theta)\log m_\beta(\theta, V) + (1 - Y(\theta))\log\{1 - m_\beta(\theta, V)\}\right\} h(\theta, V).$$

We will consider the second target parameter, but the first is treated in the same manner. Let $\psi_0 = \Psi(P_0)$, and sometimes we will also denote this true value with β_0. We note that, if $\{m_\beta : \beta\}$ would be a correctly specified model for $E_0(Y(\theta) \mid V)$, then $m_{\beta_0}(\theta, V) = E_0(Y(\theta) \mid V)$. In general, m_{β_0} represents the above defined projection of $E_0(Y(\theta) \mid V)$ on the working model. Let $m_\beta(\theta, V) = 1/(1 + \exp(\beta\phi(\theta, V)))$ for a vector of real valued functions $\phi = (\phi_1, \ldots, \phi_d)$.

Gradient for Target Parameter: Let $\mathscr{M}(g_0)$ be the statistical model that assumes g_0 is known. The canonical gradient of Ψ at P_0 is the same in model $\mathscr{M}(g_0)$ as it is in the actual model $\mathscr{M}$. Thus, the canonical gradient of $\Psi : \mathscr{M} \to \mathbb{R}^d$ can be defined as a projection of a gradient of $\Psi : \mathscr{M}(g_0) \to \mathbb{R}^d$ onto the tangent space of model $\mathscr{M}(g_0)$. A gradient of $\Psi : \mathscr{M}(g_0) \to \mathbb{R}^d$ is given by

$$D(P_0)(O) = \sum_\theta \frac{I(A = d_\theta(L))}{g_0(A \mid X)} \phi(\theta, V) h(\theta, V)(Y - m_{\beta_0}(\theta, V)).$$

Factorization of likelihood: Analogue as in van der Laan (2010a,b) and in this article, consider the factorization of the likelihood in terms of binary variables $\{L(t, j, l) : (t, j, l)\}$, $\{A(t, j, l) : (t, j, l)\}$, according to a specified ordering, and we do not discretize the final node Y:

$$P = Q_{L(0)} \left\{ \prod_{t \le K, j, l} Q_{L(t,j,l)} \right\} Q_Y \prod_{t,j,l} g_{A(t,j,l)}.$$

Given the ordering of these variables, we define $Pa(L(t, j, l))$ and $Pa(A(t, j, l))$ accordingly as the predecessors in the ordered sequence. Let $\bar{Q}_{t,j,l}(Pa(L(t, j, l))) \equiv Q_{L(t,j,l)}(1 \mid Pa(L(t, j, l)))$.

Initial estimator: We estimate $Q_{L(0)}$ with the empirical distribution $Q_{L(0),n}$. Let

g_n be an estimator of g_0. Consider an initial estimator Q_n^0, where each conditional distribution of a binary variable $L(t,j,l)$ can be represented as a logistic fit:

$$\bar{Q}^0_{n,(t,j,l)} = \text{Expit}\{\text{Logit}\bar{Q}^0_{n,(t,j,l)}\}.$$

The conditional mean $\bar{Q}_{n,Y}$ of Y, given $Pa(Y)$ is also represented as a logistic fit:

$$E_{Q_{n,Y}}(Y \mid Pa(Y)) = \text{Expit}\{\text{Logit}\bar{Q}_{n,Y}(Pa(Y))\}.$$

Canonical gradient (vanderLaan (2010a)): We will use some short-hand notation. The projection $D^*_{L(0)}$ of $D(P)$ on the tangent space $T_{L(0)}$ of $Q_{L(0)}$ is given by

$$D^*_{L(0)} = \Pi(D \mid T_{L(0)}) = E_Q\left\{\sum_\theta \phi(\theta,V)h(\theta,V)(Y_\theta - m_\beta(\theta,V)) \mid L(0)\right\}.$$

The projection D^*_{tjl} of $D(P)$ on the tangent space $T_{L(t,j,l)}$ of $Q_{L(t,j,l)}$ is given by $\Pi(D \mid T_{L(t,j,l)}) = C_{tjl}(P)(L(t,j,l) - \bar{Q}_{L(t,j,l)})$, where

$$C_{tjl}(P) = E_P(D(P) \mid L(t,j,l) = 1, Pa(L(t,j,l))) - E_P(D(P) \mid L(t,j,l) = 0, Pa(L(t,j,l))).$$

Let $C_{tjl}(P)(\delta)$, $\delta \in \{0,1\}$, represent the two terms corresponding with $L(t,j,l) = \delta$. As shown in (van der Laan (2010a)), it follows that

$$C_{tjl}(Q,g)(\delta) = \sum_\theta E_Q\left\{\frac{I(\bar{A}(t-1)=d_\theta(\bar{L}(t-1)))}{g(\bar{A}(t-1)|X)}h(\theta,V)\phi(\theta,V)Y_\theta \mid L(t,j,l) = \delta, Pa(L(t,j,l))\right\}.$$

We have $C_{tjl}(Q,g)(\delta) = \frac{1}{g(\bar{A}(t-1)|X)}C_{tjl}(Q)(\delta)$, where

$$C_{tjl}(Q)(\delta) = \sum_\theta I(\bar{A}(t-1) = d_\theta(\bar{L}(t-1)))E_Q(h(\theta,V)\phi(\theta,V)Y_\theta \mid L_\theta(t,j,l) = \delta, \bar{L}_\theta(t,j,l)).$$

We have $D^*_Y = \Pi(D \mid T_Y)$ is given by $C_{Y,g}C_Y(Y - E_Q(Y \mid Pa(Y)))$, where $C_{Y,g} = 1/g_0(A \mid X)$, and

$$C_Y = \sum_\theta I(A = d_\theta(L))\phi(\theta,V)h(\theta,V).$$

The canonical gradient is given by $D^* = D^*_{L(0)} + \sum_{t,j,l} D^*_{t,j,l} + D^*_Y$.

Parametric working model: The TMLE requires a choice of working model and loss function whose score generates the canonical gradient. Given an initial Q, we use the logistic working models $\text{Logit}\bar{Q}_Y(\varepsilon_Y) = \text{Logit}\bar{Q}_Y + \varepsilon_Y C_Y$, and

$\text{Logit}\bar{Q}_{t,j,l}(\varepsilon) = \text{Logit}\bar{Q}_{t,j,l} + \varepsilon_{t,j,l} C_{t,j,l}(Q)$. This defines a working submodel $\{Q(\varepsilon) : \varepsilon\}$ through Q at $\varepsilon = 0$.

Loss function: As loss functions we use the weighted log-likelihood loss functions $L_{g_0}(Q_{t,j,l}) = -\{\log Q_{t,j,l}\} 1/g_0(\bar{A}(t-1) \mid X)$, regular log-likelihood loss $L(Q_{L(0)}) = -\log Q_{L(0)}$, and the weighted quasi-log-likelihood loss
$L_{g_0}(\bar{Q}_Y) = -1/g_0(\bar{A} \mid X)\{Y \log \bar{Q}_Y + (1-Y)\log(1-\bar{Q}_Y)\}$. The loss function for Q is now the sum-loss function $L_{g_0}(Q) = L(Q_{L(0)}) + \sum_{t,j,l} L_{g_0}(Q_{t,j,l}) + L_{g_0}(\bar{Q}_Y)$. We note that this is a valid loss-function in the sense that $Q_0 = \arg\min_Q E_0 L_{g_0}(Q)$. The loss function and working model satisfy that the linear span of $\frac{d}{d\varepsilon} L_{g_0}(Q(\varepsilon))$ at $\varepsilon = 0$ contains the components of the canonical gradient $D^*(Q, g_0)$.

Ordering of binary components that make up longitudinal data structure: We ordered the binary variables $\{L(t,j,l)\}$, and Y and factorized the Q-factor of the likelihood accordingly. Let $B(m)$, $m = 1, \ldots, M$, denote this ordered sequence of variables, where $B(M) = Y$. Thus, $L = (L(0), B(m), m = 1, \ldots, M)$. This representation of L in terms of a sequence of ordered binaries $B(m)$, $m = 1, \ldots, M$, will allow us to clarify the algorithm.

Markov property is preserved by parametric working model: Suppose that the conditional distribution $Q_m = Q_{B(m)}$ of $B(m)$ is a function of $Pa(B(m))$ only through $(B(m-1), \ldots, B(m-k))$ and the A-nodes in $Pa(B(m))$, for a fixed small integer k. If $Q = (Q_m : m = 0, \ldots, M)$ satisfies this Markov property, then the fluctuated $Q(\varepsilon)$ also satisfies this Markov property. The purpose of the Markov property is that it controls the amount of computations and storage in the algorithm, since it identifies Q_m by 2^k possible history values for the L-process, instead of 2^m. If m gets large, this results in enormous savings.

Recursive relation for clever covariates: The clever covariate for node $B(m)$ requires us to calculate $I(\bar{A}(t_m - 1) = d_\theta(\bar{L}(t_m - 1)))E(Y_\theta \mid B(m), Pa(B(m)))$, where we defined t_m as the time-point for node $B(m)$. We will now prove the recursive relation for this conditional expectation which allows us to obtain the m-th clever covariate from the $m+1$-th clever covariate, by simply integrating over $B(m+1)$, given $Pa(B(m+1))$. We will ignore the indicator in front of conditional expectation in our proof, but we will use that $(B(m), Pa(B(m)))$ is only evaluated at values for which $\bar{A}(t_m - 1) = d_\theta(\bar{L}(t_m - 1))$. We also assume the strong SRA which states that for each $t = 0, \ldots, K$, $A(t)$ is independent of L_θ, given $Pa(A(t))$, although this assumption is not a statistical assumption, and thereby will not affect the validity of the whole estimation procedure w.r..t to the estimand.

Lemma 1 *If* $\bar{A}(t_m - 1) = d_\theta(\bar{L}(t_m - 1)$*, then*

$$E(Y_\theta \mid B(m), Pa(B(m))) =$$
$$\sum_b E(Y_\theta \mid B(m+1) = b, Pa(B(m+1)))P(B(m+1) = b \mid Pa(B(m+1))).$$

If there is a node $A(t_m)$ between $B(m)$ and $B(m+1)$, then the left-hand side equals $E(Y_\theta \mid A(t_m), B(m), Pa(B(m)))$.

Proof: Let $B_\theta(m)$ denote the counterfactual counterpart of $B(m)$, defined by L_θ, and let $Pa(B_\theta(m)) = (B_\theta(1), \ldots, B_\theta(m-1))$. Let $pa(B_\theta(m))$ be the values for the B-nodes implied by $pa(B(m))$. We have

$$\begin{aligned}
&E(Y_\theta \mid B(m) = b(m), Pa(B(m)) = pa(B(m))) = \\
&E(Y_\theta \mid B_\theta(m) = b(m), Pa(B_\theta(m)) = pa(B_\theta(m))) \\
&= \textstyle\sum_b E(Y_\theta \mid B_\theta(m+1) = b, Pa(B_\theta(m+1)) = pa(B_\theta(m+1))) \\
&\qquad P(B_\theta(m+1) = b \mid Pa(B_\theta(m+1) = pa(B_\theta(m+1))). (*)
\end{aligned}$$

The first equality used $\bar{A}(t_m - 1) = d_\theta(\bar{L}(t_m - 1))$ and SRA, and the second equality uses the standard iterative conditional expectation rule. Suppose there is no A-node between $B(m)$ and $B(m+1)$. Then, the latter expression (*) equals (by same arguments as above for first equality)

$$\begin{aligned}
&\textstyle\sum_b E(Y_\theta \mid B(m+1) = b, Pa(B(m+1)) = pa(B(m+1))) \\
&\qquad P(B(m+1) = b \mid Pa(B(m+1)) = pa(B(m+1))).
\end{aligned}$$

Suppose now that there is an A-node $A(t_m)$ between $B(m)$ and $B(m+1)$. In this case, by SRA, the expression (*) equals

$$\begin{aligned}
&\textstyle\sum_b E(Y_\theta \mid A(t_m), B_\theta(m+1) = b, Pa(B_\theta(m+1)) = pa(B_\theta(m+1))) \\
&\qquad P(B_\theta(m+1) = b \mid A(t_m), Pa(B_\theta(m+1)) = pa(B(m+1))) \\
&= \textstyle\sum_b E(Y_\theta \mid A(t_m), B(m+1) = b, Pa(B(m+1) = pa(B(m+1))) \\
&\qquad P(B(m+1) = b \mid A(t_m), Pa(B(m+1)) = pa(B(m+1))) \\
&= \textstyle\sum_b E(Y_\theta \mid B(m+1) = b, Pa(B(m+1)) = pa(B(m+1))) \\
&\qquad P(B(m+1) = b \mid Pa(B(m+1)) = pa(B(m+1))),
\end{aligned}$$

since $Pa(B(m+1))$ includes $A(t_m)$. This completes the proof of the lemma. □

Evaluation of target parameter at final step of TMLE-algorithm: In the final step of the recursive algorithm below, after having updated $Q_{B(1)}$, and integrated out over $B(1)$, we will have evaluated $E_{Q_n^*}(Y_\theta \mid L(0))$ for each θ at the current updated $(Q^*_{B(1),n}, \ldots, Q^*_{B(M),n})$. $Q_{L(0),n}$ is not updated and thus remains equal to the empirical distribution function of $L_i(0)$, $i = 1, \ldots, n$. We now evaluate $\Psi(Q_n^*)$ as the solution of $P_n D^*_{L(0)}(Q_n^*, \psi) = 0$ in ψ:

$$0 = \frac{1}{n} \sum_{i=1}^{n} E_Q \left(\sum_\theta \phi(\theta, V_i) h(\theta, V_i)(Y_\theta - m_\beta(\theta, V_i)) \mid L_i(0) \right).$$

This can be rewritten as

$$0 = \frac{1}{n}\sum_{i=1}^{n}\sum_{\theta}\phi(\theta,V_i)h(\theta,V_i)(E_Q(Y_\theta \mid L_i(0)) - m_\beta(\theta,V_i)).$$

This solution corresponds with (see Rosenblum and van der Laan (2010)) the minimizer β_n of

$$\sum_{i=1}^{n}\sum_{\theta}E_Q(Y(\theta)\mid L_i(0))\log m_\beta(\theta,V_i) + (1-E_Q(Y(\theta)\mid L_i(0)))\log\{1-m_\beta(\theta,V_i)\}h(\theta,V_i).$$

As a consequence, β_n can be computed with standard logistic regression of the outcome $E_Q(Y(\theta) \mid L_i(0)) \in (0,1)$ on the logistic regression model m_β in (θ,V_i) based on the pooled data set for the repeated measures data structure $(E_Q(Y(\theta) \mid L_i(0)),\theta,V_i:\theta)$, using weights $h(\theta,V_i)$, $i=1,\ldots,n$. This solution β_n represents the TMLE $\Psi(Q_n^*)$.

TMLE-single-step recursive algorithm: Firstly, we compute an initial fit Q_n^0 of Q_0. Recall that Q_n^0 is represented by the conditional distributions $Q_{m,n}$ of B_m, given $Pa(B_m)$, $m=1,\ldots,M$, and the empirical distribution $Q_{L(0),n}$ of $L_i(0)$, $i=1,\ldots,n$. For $m=M$ to $m=1$ (i.e., from final node to first node)

- Let $k=M-m$. This step will involve updating the estimator of m-th factor $Q_{0,m}$, and it corresponds with the $k+1$-th iterative update of Q_n^0.
- If $m=M$ (and thus $k=0$), then compute clever covariate $C_m = C_Y = \sum_\theta I(A = d_\theta(L))\phi(\theta,V)h(\theta,V)$ for updating $Q_{n,m}^k$.
- Suppose $m<M$. Compute clever covariate $C_m(Q_n^k)$ for updating $Q_{n,m}^k$ using the recursive relation presented in Lemma 1: In previous $k-1$-th step (i.e., at $m+1$-th factor) we calculated $E_{Q_n^k}(Y_\theta \mid B(m+1),Pa(B(m+1)))$, and we also obtained the update $Q_{n,m+1}^k$ of $m+1$-th factor. The recursive mapping of Lemma 1 allows us to map these two ingredients into $E_{Q_n^k}(Y_\theta \mid B(m),Pa(B(m)))$, and thereby into $C_m(Q_n^k)$.
- Update m-th factor $Q_{n,m}^k$ with $\text{Logit}\bar{Q}_{n,m}^k(\varepsilon_{m,n}) = \text{Logit}\bar{Q}_{n,m}^k + \varepsilon_{m,n}C_m(Q_n^k)$, where
$$\varepsilon_{m,n} = \arg\min_{\varepsilon} P_n L_{g_n^0}(Q_{n,m}^k(\varepsilon)).$$
- Define the $k+1$-th iterative update Q_n^{k+1} as Q_n^k but with the m-th factor $Q_{n,m}^k$ replaced by $Q_{n,m}^k(\varepsilon_m)$. For the next step we store Q_n^{k+1}, and $E_{Q_n^{k+1}}(Y_\theta \mid B(m),Pa(B(m)))$ for all θ.

The final fit Q_n^{M-1} at the final step $m=1$ (with $k=M-1$), defines the TMLE of Q_0: $Q_n^* = Q_n^{M-1}$. Finally, we evaluate the target parameter in the manner presented above, giving $\Psi(Q_n^*)$.

This appendix illustrates how the general algorithm we propose and demonstrated in the main article may be used to estimate a wide range of parameters of interest for general longitudinal data structures.