

The International Journal of Biostatistics

Volume 8, Issue 1

2012

Article 26

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

Ori M. Stitelman, *University of California - Berkeley*

Victor De Gruttola, *Harvard School of Public Health*

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

Recommended Citation:

Stitelman, Ori M.; De Gruttola, Victor; and van der Laan, Mark J. (2012) "A General Implementation of TMLE for Longitudinal Data Applied to Causal Inference in Survival Analysis," *The International Journal of Biostatistics*: Vol. 8: Iss. 1, Article 26.
DOI: 10.1515/1557-4679.1334

©2012 De Gruyter. All rights reserved.

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. 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.

KEYWORDS: survival analysis, causal inference, double robust, targeted maximum likelihood estimation, time-dependent covariates, informative censoring

Author Notes: Portions of the work in this article were supported in part by the Russell M. Grossman Endowment, NIEHS grant ES015493, NIAID grant R01 AI74345-5, and NIH grants R01 AI074345-04 and R01 AI51164.

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 or symptom relief. Such trials collect on each subject a longitudinal data structure which includes baseline covariates, the randomized treatment assignment, and other time-dependent covariates for as long as the subject is observed prior to experiencing the outcome of interest. The typical approach for assessing the causal effect of the randomized treatment in the literature and as mandated by the FDA for pharmaceutical drug development, is to employ a Cox-proportional hazards model only including treatment in the model. Once the parameters of the Cox-proportional hazards model is estimated a test is conducted to determine if the coefficient on treatment is not equal to zero. This approach ignores 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 of 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 outcome 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 for treatment, even if the multiplicative intensity model is correctly specified. That is, standard regression methods cannot be employed.

Moreover, even if the Cox-proportional hazards model is correctly specified, the estimated parameter does not typically 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 establishes ways to define the actual causal quantities of interest. The literature also establishes identifiability results for these causal quantities so that they can be identified as a target parameter of the data generating distribution of the data 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)), 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 these 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 of the G-computation formula parameter is a substitution estimator of the target parameter of the data generating distribution. This estimator 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 to make the treatment arms among the uncensored subjects comparable with respect to confounders, and then applies standard estimation as if treatment was randomized and censoring was non-informative. The consistency of these estimators relies on consistent estimation of the treatment and censoring mechanism. These estimators are highly unstable in situations when the parameter of interest is weakly identifiable¹, such as when there is a level of covariates that is predictive of treatment or censoring. The in-

¹The 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

stability 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 covariates that must be accounted for.

Even though locally efficient double robust A-IPCW estimators 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 with respect to 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 (see e.g. Kang and Schafer (2007), Bang and Robins (2005), and Robins, Sued, Lei-Gomez, and Rotnitzky (2007)). This is due to the fact that these estimators do not respect the global constraints implied by the statistical model and thus do not acknowledge that the target parameter is a particular function of the true data generating distribution.

Targeted Maximum Likelihood Estimation (TMLE), proposed by (van der Laan and Rubin (2006)), 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 van der Laan and Rose (2011)). For articles that demonstrate the advantages of TMLE compared to IPCW, A-IPCW and other estimators in 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), Rose and van der Laan (2010), Porter, Gruber, van der Laan, and Sekhon (2011) among others. In a series of earlier papers we presented 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)), where we presented the advantages of this TMLE relative to the other classes of estimators through simulation study 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 co-*

variates. 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. 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 has been no computationally feasible and efficient approach may be estimated using the algorithm presented here.

1.1 Organization

In the next section we describe the targeted maximum likelihood estimator for the treatment specific survival curve incorporating time dependent covariates. In Section 3 we describe adjustments to this TMLE that make it computationally feasible. Section 4 presents a simulation study, and Section 5 a data analysis of an HIV study, the Tshepo study. We conclude with a discussion.

2 Targeted Maximum Likelihood Estimation Of Treatment Specific Survival Incorporating Time Dependent Covariates

In this section we will describe the process of constructing the TMLE of the treatment specific survival curve that incorporates time dependent covariates. For the purpose of clarity we will present the approach for a specific example with a clearly defined data structure. In particular, we will focus on an analysis of the Tshepo study, a study designed to evaluate the efficacy, tolerability, and development of drug resistance of six different first-line cART regimens.

The primary scientific question of interest for the presented analysis is "what is the causal effect of cART treatment on the time until death?" The observed data consists of n i.i.d observations, O_i , from the data generating distribution, P_0 . Let t be a discrete time index, ($t=0,1,2,\dots$), where $t=0$ indexes the baseline time before treatment. Each observation O_i consists of W_i , the baseline covariates, $A_i(t=0)$, a binary indicator of cART treatment, $CD4_i(t)$, a time dependent process of CD4 level, $VL_i(t)$, a time dependent viral load process, $A_i^c(t)$, the censoring process, and $Y_i(t)$ the death process. Specifically, these variables are as follows:

- W_i is a vector of random variables that capture the state of subject i at baseline. These include gender, baseline CD4 count, baseline VL, age, body mass index, past health history, etc.
- $A_i(0)$ is binary random variable that indicates the cART treatment assigned to subject i . When $A_i(0) = 1$ the subject was treated with efavirenz (EFV) and when $A_i(0) = 0$ the subject was treated with nevirapine(NVP). For the remainder of the paper $A(0)$ will be referred to as A .
- $CD4_i(t)$ is a continuous random process that captures the CD4 count of each subject i at time t . In the observed data this variable ranges from 3 to 1,105.

- $VL_i(t)$ is a continuous random process that captures the viral load of each subject i at time t . In the observed data this variable ranges from 400 to 750,000. All individuals with a viral load less than 400 were recorded as 400.
- $A_i^c(t)$ is the censoring process and remains zero at each time point t at which the subject i is observed and once the subject is no longer observed, or lost to follow up it jumps to 1. After the subject is no longer observed the process remains at 1.
- $Y_i(t)$ is the death process and remains zero at each time point t at which the subject i is alive and once the subject is no longer alive it jumps to 1. After death the process remains at 1.

The entire random process is assumed to be generated by the following time ordering: $W \rightarrow \mathbf{A} \rightarrow Y(1) \rightarrow CD_4(1) \rightarrow VL(1) \rightarrow \mathbf{A}^c(\mathbf{1}) \dots \rightarrow Y(k-1) \rightarrow CD_4(k-1) \rightarrow VL(k-1) \rightarrow \mathbf{A}^c(\mathbf{k}-\mathbf{1}) \rightarrow Y(k)$. $t = k$ is the time point at which one is interested in understanding the effect of the cART therapy on death. The time ordering corresponds with a causal graph with each node having an arrow coming into it from all of its ancestors, that is all nodes preceding it in the time ordering.² The random variables displayed in bold in the time ordering correspond with those variables for which we would like to have the ability to intervene in the causal graph.

In the time ordering above variables preceding a node we will refer to as the node's parents and will denote that set of nodes as $Pa(\cdot)$. For example $Pa(CD_4(1))$ is equal to the set of nodes $\{W, A, Y(1)\}$. Also, let a bar over a variable denote the history of a variable up until and including time t . So $\bar{A}^c(t)$ is the entire history of censoring until t . We will use $\bar{A}^c(t) = 0$ to express the fact that an individual was not censored up until and including time t and similarly $\bar{Y}(t) = 0$ will express that an individual did not experience the event through time t . The causal graph posits a set of causal assumptions in the form of a non-parametric structural equation model (NPSEM) in which our data structure corresponds with the displayed nodes and the exogenous nodes are suppressed. This is common practice for displaying that each displayed node has an error/exogenous node with an arrow going into it and no arrows going into any other nodes.

Necessary conditions typically imposed to make causal parameters identifiable may be made through the use of the causal graph, i.e. the assumptions of consistency and coarsening at random (CAR). The former states that the observed data consist of the counterfactual outcome corresponding with the intervention actually observed. This assumption is a direct consequence of defining the observed

²The full graph is not presented because the amount of arrows necessary to show all edges between nodes is visually unappealing and does not add to the understanding of the graph.

data structure as particular nodes in the causal graph/NPSEM. The CAR assumption is imposed by assuming the strong sequential randomization assumption on the intervention nodes of the causal graph, i.e. for each intervention node, the conditional density of an intervention node, given the collection of all counterfactuals and the intervention nodes that are parents of the intervention node, is a function only of the observed parents of the intervention node. This is directly implied by the causal graph, as the only arrows into the treatment and censoring process variables are from their observed ancestors and no other nodes. That is there are no exogenous error nodes(unobserved variables) that affect both the censoring and/or treatment nodes and other nodes in the causal graph. For the treatment variable this assumption is sometimes referred to as the randomization of treatment assumption or no unmeasured confounders assumption.

The causal graph acts as a tool that allows us to consider potential outcomes under some desired set of interventions on those variables. Namely, we are interested in observing the value $Y(k)$ for each individual in the observed data had we set treatment A to specified level, a , and all censoring $A^c(t)$ to 0, or uncensored.

The ordering of all variables as proposed implies the following factorization of the observed data likelihood:

$$P(O) = P(W)\mathbf{P}(\mathbf{A}|\mathbf{W}) \quad (1)$$

$$\prod_{t=1}^{k-1} [P(Y(t)|Pa(Y(t)))P(CD_4(t)|Pa(CD_4(t)))P(VL(t)|Pa(VL(t)))\mathbf{P}(\mathbf{A}^c(\mathbf{t})|\mathbf{Pa}(\mathbf{A}^c(\mathbf{t})))]$$

$$P(Y(k)|Pa(Y(k)))$$

Thus, P is factorized into two distinct components Q , corresponding to the non-bolded conditional distributions, and g , the bolded conditional distributions. Q is the factors of the likelihood associated with the full data process. These are the factors separate from our intervention on the system. The g factors of P are the contribution to the likelihood of the variables on which we wish to intervene. These are the variables that act to coarsen, or hide future paths of the full data structure. For example, if a particular subject O_i is observed to have taken treatment EFV, or $A = 1$, the entire future process for that subject under NVP, or $A = 0$, is hidden. Similarly, if a subject is censored at a particular time point their entire future process is coarsened or hidden. For this reason treatment, $A(0)$, and the censoring process, $\{A^c(t) : 1, \dots, k\}$, make up the coarsening factor of P_0, g_0 .

Now we can consider the distribution of the data under interventions on the coarsening variables, treatment and censoring. This distribution is known as the g -computation formula and for our data structure is of the following form:

$$\begin{aligned}
 P_{A=a, A^c(t)=0} &= P(W) \\
 &\prod_{t=1}^{k-1} [P(Y(t)|Pa(Y(t)), A=a, A^c(t-1)=0) \\
 &P(CD_4(t)|Pa(CD_4(t)), A=a, A^c(t-1)=0) \\
 &P(VL(t)|Pa(VL(t)), A=a, A^c(t-1)=0)] \\
 &P(Y(k)|Pa(Y(k)), A=a, A^c(k-1)=0).
 \end{aligned} \tag{2}$$

Notice that all of the g factors have been removed from the distribution since they are now set to the intervened levels with probability 1. While all of the other distributions are conditioned as though A was set to the desired level a and no censoring $A^c(t) = 0$ for all t . This corresponds with Pearl's do calculus for causal graphs (Pearl (2008)).

Since the distribution of the data under the desired interventions is defined, we can now propose interesting parameters of interest under that distribution. In particular, these parameters should directly answer the scientific question of interest. By contrasting the probability of survival past the chosen time, k , under different interventions we can define a parameter of interest that quantifies the causal effect of the different cART therapies. An example of such a parameter is the probability of surviving past k when the cART therapy is EFV and there is no censoring minus the probability of surviving past k had the cART therapy been NVP and there is no censoring. This parameter of interest, which we will call ψ_0^{RD} , can be written as a function of the intervention distributions $P_{A=a, A^c(t)=0}$, of which we have two (one for EFV, when $A = 1$, and one for NVP, when $A = 0$). In particular ψ_0^{RD} can be written as a function of two treatment specific survival curves, one at each of the levels of treatment. The fact that each treatment specific survival curve may be written as a function of the intervention distributions $P_{A=1, A^c(t)=0}$ and $P_{A=0, A^c(t)=0}$, that is the distributions when setting $A = 1$ and $A = 0$, will be expressed by writing the treatment specific survival curves as as a function of the intervention distributions in the following way: $\Psi_{A=1}(P_{A=1, A^c(t)=0})$ and $\Psi_{A=0}(P_{A=0, A^c(t)=0})$, where $\Psi_{A=a}(P_{A=a, A^c(t)=0})$ is the mean probability of survival past time k when treated at level a .

Since the death process, $Y(k)$ is binary, it equals one if an individual has died and zero otherwise, the probability that $Y(k)$ equals zero under a specified intervention is the treatment specific survival curve. Thus, the treatment specific survival curve at time k may be cast as a mapping from the intervention distributions, $P_{A=a, A^c(t)=0}$, to the real line as $P_{A=a, A^c(t)=0}(Y_k = 0)$. The causal effect of the

cART treatment may than be expressed as a risk difference of these treatment specific survival curves³:

$$\Psi^{RD}(P_{A=1,A^c(t)=0}, P_{A=0,A^c(t)=0}) = P_{0,A=1,A^c(t)=0}(Y(k) = 0) - P_{0,A=0,A^c(t)=0}(Y(k) = 0).$$

Alternatively, the parameter of interest can be defined as an expected value under the intervention distributions, $P_{A=a,A^c(t)=0}$:

$$\psi_0 = [1 - E_{P_{0,A=1,A^c(t)=0}}[Y(k)]] - [1 - E_{P_{0,A=0,A^c(t)=0}}[Y(k)]].$$

Up until this point we have not specified a model for the underlying data generating distribution. Had the true conditional distributions $P_{0,A=a,A^c(t)=0}$ been known, $\Psi_{A=a}(P_0)$ could simply be evaluated. However since $P_{0,A=1,A^c(t)=0}$ and $P_{0,A=0,A^c(t)=0}$ are not known, they must be estimated. One possible model for the conditional distributions presented in the likelihood above ((1)) is a simple linear regression model for the continuous distributions, VL and CD4, and logistic regression model for the other distributions. Instead of using a simple regression model, we advocate the use of a less restrictive non-parametric model. Estimates of $P_{0,A=a,A^c(t)=0}$ according to this non-parametric model, may be generated using a nonparametric cross-validated learning algorithm. Smoothing over multiple time points may be implemented in the same way as for simple regression. In practice we like to use Super Learner for constructing Q_n and g_n .⁴ We will refer to the estimates of Q_0 as Q_n , the estimates of g_0 as g_n and the estimates of the entire distribution P_0 as P_n . Q_n is composed of the following estimates:

$$\begin{aligned} &P_n(W) \\ &P_n(Y(t)|Pa(Y(t))) \\ &P_n(CD_4(t)|Pa(CD_4(t))) \\ &P_n(VL(t)|Pa(VL(t))), \end{aligned}$$

³Other contrasts of the treatment specific survival curves may be used to quantify the causal effect of the cART treatment such as the ratio of the two curves; however, for the purpose of clarity in this paper we will focus on the risk difference. Additionally, parameters that average over a set of t may be of interest.

⁴Super Learner is a system that constructs estimates in a non-parametric model by combining many candidate learners in a way that creates a final estimate of the conditional distribution with better properties than any of the single candidate learners (van der Laan, Polley, and Hubbard (2007)).

and g_n includes:

$$\begin{aligned} P_n(A|W) \\ P_n(A^c(t)|Pa(A^c(t))). \end{aligned}$$

Whether one employs simple regression techniques or a non-parametric estimation procedure such as super learner, the estimates produced have been chosen in order to optimize prediction and not to optimize the estimate of the parameter of interest. This is because these procedures use a loss function to minimize the distance between the estimate and the true distribution globally. This is true whether the loss function is mean squared error, log-likelihood, or any other sensible loss function for estimating conditional distributions. Updating this distribution to a distribution targeted toward estimating the parameter of interest as well as possible is the primary goal of TMLE.

The next step of the TMLE procedure is to update the estimates of $P_{0,A=a,A^c(t)=0}$ so they are optimized for estimating the parameter of interest. This is the step for which TMLE got its name and is the reason TMLEs exhibit their robustness and efficiency characteristics. This update step ensures that the TMLE targets the parameter of interest. This is accomplished by enforcing that the estimates of $P_{0,A=a,A^c(t)=0}$ solve the efficient influence curve equation for the target parameter. The theory that establishes this result is presented in van der Laan and Rubin (2006) and van der Laan and Rose (2011). The efficient influence curve can be represented as a projection of an IPCW-estimating function D_{IPCW} onto the tangent space of the Q factor of the density:

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

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

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

where C is the censoring time and T is the event time.

In general the update process relies on the use of a loss function for Q_0 , $\mathcal{L}(Q)$, and a parametric submodel $\{Q(\varepsilon, g) : \varepsilon\}$. The fluctuation, ε , from the initial estimate Q_n is chosen in the direction that has the greatest benefit for estimating the parameter of interest. This is done by choosing ε_n such that the linear span of the loss-based score $\frac{d}{d\varepsilon} \mathcal{L}(Q(\varepsilon, g))$ at $\varepsilon = 0$ includes the efficient influence curve

$D^*(Q, g)$. In many situations it also makes sense for the loss function to be indexed by a particular nuisance parameter as will be done below since the loss function is also dependent on the estimate of g . We will express the loss function as indexed by g_n as $\mathcal{L}_{g_n}(Q)$. For more details concerning this updating process we refer the reader to van der Laan and Rose (2011) as they are outside the immediate scope of this paper.

In van der Laan (2010a) an approach for constructing targeted maximum likelihood estimates of parameters in general longitudinal data structures is presented. The proposed approach is based on a binary factorization of the Q_0 components of the likelihood. It then proposes a backward passing algorithm that updates the initial estimate Q_n at each binary factor through the use of a logistic parametric submodel.

There are three distinct advantages of this binary factorization: (1) it allows for a closed form TMLE of the parameter of interest and clever covariates (2) it allows for the use of the multitude of data adaptive machine learning algorithms that are available for estimation of binary conditional distributions in order to get a flexible estimate of the initial distributions and (3) it allows for a single formulation of the clever covariate across nodes.

We will now recast this general algorithm in terms of the HIV example presented here. The first step in creating a binary factorization of the Q_0 components of the likelihood is to decide on cut points for the intermediate continuous random variables. The continuous random variables in W_i may remain continuous. In the case of the HIV study we are using as an example, this involves breaking the variables $CD_4(t)$ and $VL(t)$ into ordered levels. For the purpose of our analysis, each continuous variable was divided into three ordered categorical levels. The number of levels chosen and their cut points should be based on subject matter knowledge of the observed process, or one could adaptively select the number of levels using cross-validation.

Let us consider the CD_4 process. We could assumed ordered categorical levels. The levels are: $CD_4(t) < 200$, $200 < CD_4(t) < 400$, $400 < CD_4(t)$. These levels correspond with the following indicator functions of $CD_4(t)$ at each time point: $CD_4(t, 1) = 1(CD_4(t) \leq 200)$, $CD_4(t, 2) = 1(200 < CD_4(t) < 400)$, and $CD_4(t, 3) = 1(400 \leq CD_4(t))$. However, as always, when coding a categorical variable, the categorical variable may be expressed by the first two binary variables because when $CD_4(t, 1)$ and $CD_4(t, 2)$ equal zero the third indicator function, $CD_4(t, 3)$, will equal 1. The binary indicators, $VL(t, 1)$, $VL(t, 2)$ and $VL(t, 3)$, for the viral load process are constructed in the same way with break points for the levels of 400, and 2,000.

Incorporating these ordered categorical variables into the likelihood involves constructing a hazard formulation of the ordered variables within each time point

for each intermediate continuous variable. In the statistical literature a hazard refers to the probability an event will happen at a given time given that it has not happened yet. It has a similar meaning as used here for the ordered categorical variables. It is the probability that a subject will belong to a particular level given that the subject did not belong to any of the previous levels, as defined by the ordering. So the hazard for $CD_4(t, 2)$ is the probability that $CD_4(t, 2)$ equals 1 given $CD_4(t, 1)$ equals zero.

After casting the intermediate variables in the hazard formulation discussed above, the temporal ordering may be re-expressed as: $W \rightarrow \mathbf{A} \rightarrow Y(1) \rightarrow CD_4(1, 1) \rightarrow CD_4(1, 2) \rightarrow VL(1, 1) \rightarrow VL(1, 2) \rightarrow \mathbf{A}(\mathbf{1}) \dots \rightarrow Y(k-1) \rightarrow CD_4(k-1, 1) \rightarrow CD_4(k-1, 2) \rightarrow VL(k-1, 1) \rightarrow VL(k-1, 2) \rightarrow \mathbf{A}(\mathbf{k}-\mathbf{1}) \rightarrow Y(k)$. Notice that $CD_4(t, 3)$ and $VL(t, 3)$ are not in the time ordering because they may be written in terms of the other two variables at the particular time point. Also, recall that the bolded variables are those on which we wish to intervene.

Now each subject's contribution of $P(CD_4(t)|Pa(CD_4(t)))$ to the likelihood expressed in equation (1) may be written in terms of the hazard distributions of $CD_4(t, 1)$ and $CD_4(t, 2)$ in the following way:

$$P(CD_4(t, 1)|Pa(CD_4(t)))^{CD_4(t, 1)} \\ \{1 - P(CD_4(t, 1)|Pa(CD_4(t)))\}^{CD_4(t, 2)} \{P(CD_4(t, 2)|Pa(CD_4(t)), CD_4(t, 1) = 0)\}^{CD_4(t, 2)} \\ \{1 - P(CD_4(t, 1)|Pa(CD_4(t)))\}^{CD_4(t, 3)} \{1 - P(CD_4(t, 2)|Pa(CD_4(t)), CD_4(t, 1) = 0)\}^{CD_4(t, 3)},$$

where $CD_4(t, 3) = 1 - CD_4(t, 1) - CD_4(t, 2)$. The same is true for the $P(VL(t)|Pa(VL(t)))$ contribution to the likelihood. Furthermore, the $Pa(\cdot)$ of all the other conditional distributions now includes $CD_4(t, 1)$ and $CD_4(t, 2)$ rather than $CD_4(t)$ and $VL(t, 1)$ and $VL(t, 2)$ rather than $VL(t)$. This same approach may be expanded to more levels of each intermediate variable.

An alternative notation for the same parameterization of the random variables is more convenient for expressing the binary likelihood and the TMLE algorithm we will propose. However, this notation is not as intuitive in terms of what each variable represents. For this reason we have chosen to introduce the variables as we have done and now explain how those variables correspond with this alternative binary notation. We will now define a new set of random variables $L(t, j, l)$. $L(t, j, l)$ are binary random variables that include all of the binary processes that Q_0 is composed of. Let t be a discrete time index as before, j an index for the process (e.g. Death, VL, or CD_4 process), and l an index for the level of the process. These random binary variables may be rolled up into $L(t, j)$, a set of variables defining a particular process at a particular time, and $L(t)$ the set of all processes at a particular time. Also, let $L(0)$ equal W .

Specifically, $L(t)$ is decomposed as follows:

$$L(t) = (L(t, j) : j = 1, \dots, n(t)), \quad (3)$$

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 analysis $n(t)$ is equal to 3, the number of processes in Q_0 . $L(t, j)$ are further decomposed into binary variables $L(t, j, l)$ in the following way:

$$L(t, j) = (L(t, j, l) : l = 1, \dots, n(t, j)), \quad (4)$$

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 and in the implementation as described above $n(t, 2)$ is equal to 3, the number of levels the CD4 process was divided into. Notice that $CD_4(t, 1)$ is equal to $L(t, 1, 1)$ in our new notation and $CD_4(t, 1)$ is equal to $L(t, 1, 2)$.

The g-comp formula may be expressed in our new notation in the following way:

$$P_{a,0}(L) = \overbrace{P(L(0))}^{Q_{L_0}} \prod_{t=1}^k \prod_{j=1}^{n(t)} \prod_{l=1}^{n(t,j)-1} \overbrace{P(L(t, j, l) | Pa(L(t, j, l)), A = a, \bar{A}(t-1) = 0)}^{Q_{L(t,j,l),a,0}}$$

Now that the g-comp formula is expressed in terms of binary variables, the components of the parameter of interest, $P_{0,A=1,A^c(t)=0}(Y(k) = 0)$, may be written in closed form. In our new notation this closed form solution of the treatment specific survival curve at time k is:

$$\Psi(Q_0) = E_{L(0)} \sum_{L(k,1,1)=0} \prod_{t=1}^k \prod_{j=1}^{n(t)} \prod_{l=1}^{n(t,j)-1} \overbrace{P(L(t, j, l) | Pa(L(t, j, l)), A = a, A(t-1) = 0)}^{Q_{L(t,j,l),a,0}}, \quad (5)$$

where the sum is taken over all possible combinations of variables ending in $L(k, 1, 1) = 0$.

Furthermore, this binary factorization allows us to estimate all of the conditional distributions, including the hazard distributions, using the multitude of data adaptive machine learning algorithms available for binary outcomes. Each factor of

the binary factorization has a corresponding contribution to the likelihood $Q_{L(t,j,l)}$ which is the conditional distribution of that binary variable given its parents and may be estimated using standard tools for estimating binary conditional distributions.

Once the binary factorization is established and the initial estimates of Q_0 are constructed, the next step in the algorithm is to update each component of Q_n to the targeted maximum likelihood estimate, Q_n^* , of Q_0 . Thus, each estimate $Q_{L(t,j,l),n}$ needs to be updated to $Q_{L(t,j,l),n}^*$, the targeted estimate for that binary factor. In van der Laan (2010a) the proposed approach is through a backward passing algorithm that starts with the last factor in the likelihood and works toward the first point in time. At each step the estimate $Q_{L(t,j,l),n}$ is updated through the use of a logistic fluctuation of the initial estimate. The parametric fluctuation used is of the following form:

$$\text{logit}Q_{(t,j,l),n}^*(\varepsilon) = \text{logit}Q_{(t,j,l),n} + \varepsilon_{t,j,l}C_{t,j,l}(Q_n, g_n),$$

where $C_{t,j,l}(Q_n, g_n)$ is referred to as the clever covariate for $Q_{L(t,j,l)}$ and is chosen such that the update solves the associated component of the efficient influence curve equation. It is established in van der Laan (2010a) that the clever covariates, $C_{t,j,l}(Q_n, g_n)$, necessary for the targeting step may be expressed in the same way for each node in Q , indexed by t, j and l as: $C_{tjl}(Q, g) = C_{tjl}(Q)C_{tjl}(g)$, where

$$C_{tjl}(Q) = \{P[L_{a,0}(K+1, 1, 1) = 0 \mid L(t, j, l) = 1, Pa(L(t, j, l))] - P[L_{a,0}(K+1, 1, 1) = 0 \mid L(t, j, l) = 0, Pa(L(t, j, l))]\}, \quad (6)$$

and

$$C_{tjl}(g) = \frac{I(A = a)I(C > t_-)}{g_{A(0)}(1 \mid L(0)) \prod_{s=1}^{t-1} g_{A^c(s)}(0 \mid Pa(A^c(s)))}. \quad (7)$$

Once the final factor is updated, the clever covariate for the subsequent factor, moving backwards, is calculated using the updated terminal node and then that node is updated. This process is continued until the first factor, $Q_{L(1,1,1),n}$, is reached. Once all $Q_{L(t,j,l),n}$ are updated to $Q_{L(t,j,l),n}^*$ the the set of estimates is the targeted estimate of Q_0 . Now $\Psi(Q_0)$ may be estimated by substituting $Q_{L(t,j,l),n}^*$ into equation (5) and evaluating the right side of the equation. Note that the factorization of $C_{tjl}(Q, g)$ into $C_{tjl}(Q)$ and $C_{tjl}(g)$ is not stressed in van der Laan (2010a) but is presented here stressing the factorization because it plays an important role in an adjusted algorithm we propose below.

It turns out that as the number of intermediate variables and the number of binary factors for each variable increases, the computational complexity of this algorithm also increases exponentially and makes it impossible to implement with even a reasonable amount of these factors. In the next section we propose adjustments to this algorithm that allows the complexity of the algorithm to grow linearly with the number of binary factors while retaining the robustness and efficiency properties of a TMLE.

3 Computationally Feasible TMLE

The method described in the preceding section, though theoretically correct, is impractical and in most cases computationally impossible to implement. In this section we explore why the presented approach is infeasible and provide adjustments to the algorithm that make implementation of TMLE in this setting possible. In fact, the adjustments proposed make creating an augmented IPCW estimator in the flavor of those presented in van der Laan and Robins (2003) also possible. Our solution involves multiple mathematical tricks and simplifications that are designed to increase feasibility without sacrificing efficiency and unbiasedness and ultimately result in an algorithm that is extremely fast even relative to our initial algorithms that only incorporated baseline variables (Stitelman and van der Laan (2011), Stitelman and van der Laan (2010)).

The TMLE algorithm presented in the previous section requires evaluating an integral over a very large number of possible combinations of the binary factors in both the evaluation of $\Psi_a(Q_n^*)$ and the clever covariates. Moreover, since each $Q_{L(t,j,l),n}$ and $g_{A^c(t),n}(0 | Pa(A^c(s)))$ are a function of the entire past history of parents, each possible combination of history must be evaluated for every clever covariate and the overall evaluation of $\Psi(Q_n^*)$. Various numerical integration procedures may be employed, including Monte Carlo simulation. However, these methods are computationally expensive and as k and the number of intermediate variables increases, the number of computations needed increases exponentially. This exponential increase in computational complexity makes the algorithm computationally expensive and not possible to implement in many situations.

In order to make the evaluation procedure more feasible, we make some simplifying assumptions. These assumptions were specifically chosen to retain the efficiency and consistency characteristics of TMLE while creating an algorithm that is possible to implement given reasonable computational computing restraints. This feasible algorithm is based on three major adjustments to the algorithm presented in the preceding section:

1. Data adaptive simplifying Markov assumptions.
2. Backward passing algorithm using Iterative conditional expectations.
3. Weighted logistic regression updates.

We will now explain in detail how each of these adjustments are implemented and how they contribute toward making the algorithm computationally practical.

3.1 Markov property for initial estimate to speed up algorithm

Even though the binary factorization of the likelihood allows us to write the clever covariates and the parameter of interest as a closed form solution of the expected value of the terminal node over all paths through the binary variables, the computational complexity of such an evaluation is immense. As a result, without any further adjustments to our methodology, our evaluation is still limited to either a Monte Carlo approach or by integration over all paths in the G -computation formula. Both of these approaches are computationally costly since the number of paths up until 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 binary variable in Q . At first glance it appears as though the binary factorization of the variables has not bought much in terms of feasibility. However, by combining the binary factorization with several other simplification assumptions, we will be able to produce an algorithm that is computationally feasible. One way to simplify the amount of necessary computations is to enforce that the conditional distribution of each $L(t, j, l)$ 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 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 $C_{t,j,l}(Q_n)$ is linear in the number of binary variables.

For the purpose of our HIV analysis, we enforce a Markov property on all the conditional distributions expressed in the likelihood. In each case we assume that each $L(t, j, l)$ is a function of the time dependent covariates through the most recently observed levels. Thus, each conditional distribution is a function of not being censored yet ($\bar{A}(t) = 0$), not having the event yet ($\bar{Y} = 0$, equivalently, $L(t, 1, 1) = 0$), the baseline variables (W), and the most recent observed values of the time dependent covariates. So the conditional distributions in the likelihood,

rather than being a function of the full past, will be defined as a function of the reduced past in the following way:

$$\begin{aligned} P(Y(t)|Pa(Y(t))) &= P(Y(t)|\bar{A}(t-1)=0, \bar{Y}(t)=0, W, CD_4(t-1), VL(t-1)) \\ P(CD_4(t)|Pa(CD_4(t))) &= P(CD_4(t)|\bar{A}(t-1)=0, \bar{Y}(t)=0, W, CD_4(t-1), VL(t-1)) \\ P(VL(t)|Pa(VL(t))) &= P(VL(t)|\bar{A}(t-1)=0, \bar{Y}(t)=0, W, CD_4(t), VL(t-1)) \end{aligned}$$

The distributions of W , $A(0)|W$, and $A^c(t)|Pa(A^c(t))$ are unaffected by the Markov assumption. We deal with the large dimension of pasts for $A^c(t)|Pa(A^c(t))$ with another modification to the TMLE algorithm using regression weights, that is explained below. In other analyses the Markov 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.

3.2 Backward passing algorithm using Iterative conditional expectations

The algorithm described in the previous section uses a backward updating scheme to make sure that the updated distribution P_n^* solves the efficient influence curve equation. We also employ a TMLE using a *backwards* updating algorithm that first updates the last factor, and then proceeds backwards which converges to P_n^* in one set of updates through the M nodes in the Q factors of the causal graph. The reason such a method is viable is that the clever covariate at each node is only a function of P_n through the nodes that follow it in the time ordering. As a result of the backward passing algorithm, those nodes are already updated when the algorithm arrives at any particular node. Our modification to the backward passing algorithm takes advantage of the the fact that the clever covariate at each subsequent step may be written as an iterative conditional expectation of the preceding step.

Recall that \mathcal{L}_{g_n} is a loss function indexed by the estimate g_n . The loss function we used for our analysis was the negative log-likelihood loss function. However, any appropriate loss function may be used. Let

$$\epsilon_n^M = \arg \min_{\epsilon^M} P_n \mathcal{L}_{g_n}(Q_{n,M}^0(\epsilon^M)).$$

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

$$\epsilon_n^{M-1} = \arg \min_{\epsilon^{M-1}} P_n \mathcal{L}_{g_n}(Q_{n,M-1}^1(\epsilon^{M-1})).$$

This yields a second update Q_n^2 obtained by updating the $M - 1$ -th factor $Q_{n,M-1}^1$ with $Q_{n,M-1}^1(\epsilon_n^{M-1})$. This updating process is continued through all the nodes resulting in a sequence $Q_n^1, Q_n^2, \dots, Q_n^M$ of M 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 , $J = 1, \dots, M$. The last M -th update involves the update of the first factor and all the updated other factors. This set of all updated Q_n is Q_n^* the fully targeted estimate of Q_0 . The TMLE of $\psi_0 = \Psi(Q_0)$ is the corresponding substitution estimator $\psi_n^* = \Psi(Q_n^*)$. The way we 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 $O(l)$, $l = 0, \dots, L$, where $O(0) = W$, $O(1) = A$, and several of subsequent $O(l)$ correspond with $A^c(t)$, 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

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

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 $A^c(t)$ -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

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

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 $O(k-2)$.

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. It should also be stated that the estimate of $P(W)$, $P_n(W)$, does not need to be updated because we use the empirical distribution as the estimate. The empirical distribution is the non-parametric efficient estimate of $P(W)$, and as a result, fluctuating the estimate would result in the same estimate (see section 5.2.4 of van der Laan and Rose (2011) for more details about this).

3.3 TMLE with weighted-log-likelihood loss function

Even if we enforce this Markov assumption on the initial estimate Q_n , note that $C_{t,j,l}(g_n)$ is still a function of the full history $Pa(A^c(t))$ through $g_{A^c(t)}$, 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 $C_{t,j,l}(g_n)$ from being a factor of the clever covariate to being a weight in the log-likelihood loss function. This is why we stressed the fact that $C_{tjl}(Q, g)$ may be factorized into $C_{tjl}(Q)$ and $C_{tjl}(g)$. We use a weighted logistic regression for each update with weights equal to $C_{t,j,l}(g_n)$, and a new clever covariate $C_{t,j,l}(Q_n)$ instead of $C_{t,j,l}(Q_n, g_n)$. This corresponds with using a weighted-log-likelihood loss $L_g(Q) = -\log Q_{L(0)} - \sum_{tjl} \{ \log Q_{tjl} \} C_{tjl}(g)$ 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 :

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

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, this weighted-log-likelihood and submodel also satisfies the condition 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. Furthermore it is not required that the censoring process be estimated using the factorized intermediate variables and thus it may be estimated using the continuous intermediate variables, $CD_4(t)$ and $VL(t)$ in the case presented here.

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) which only adjusted for baseline covariates.

In this section we provided a computationally feasible method for constructing the TMLE estimate of the treatment specific survival curve that incorporates time dependent covariates. This method relied on several adjustments to previously used approach to constructing such estimators. These adjustments were specifically implemented to make the approach computationally feasible while preserving the double robust and locally efficient properties of a TMLE. In the following sections we will present a simulation study that compares the TMLE presented here to other alternative methods for estimating the treatment specific survival curve and will also include an analysis of an HIV data set.

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: 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 abbreviation for an estimating equation based estimator that 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 here). The EE involves representing the efficient influence curve for the longitudinal data structure as an estimating function in the target parameter ψ_0 . Then defining the EE estimator as the solution to 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

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

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

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

for $t = 1, \dots, K+1$. 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 , $A^c(t)$ 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:

$$\begin{aligned} W_4(t) &= .2A(0) + .5W_1(0) - .4W_2(0) - .4W_3(0) + 2W_4(t-1) + 2W_5(t-1) + U_4 \\ W_5(t) &= .1A(0) + .1W_1(0) + .1W_2(0) - .4W_3(0) + 2W_4(t) + 2W_5(t-1) + U_5, \end{aligned}$$

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, $A^c(t)$, 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:

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

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 specific scenario where a biased estimator happens to produce the right answer. This has to do with the fact that the covariates that strongly affect 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 were at levels of W which were almost completely predictive of an event and being uncensored in the first scenario are now at levels of W that are still almost completely predictive of an event but also completely predictive of censoring.

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

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

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

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 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

In this section we repeat the simulation study but with independent censoring. The data generating distribution is as in the previous section except the censoring mechanism is modified to be independent of both the baseline variables and time dependent covariates. 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

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 Using Modified Censoring Process: Mean of Estimates and Mean Square Error for All Six Estimators

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 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 more efficient than 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 importance of being a substitution estimator. This gain is most likely due

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

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.

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 more details concerning these confidence intervals see Stitelman, De Gruttola, and van der Laan (2011).

4.3 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 which 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 4 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. 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 possess this quality. However, the Mean SE which is based on the mis-specified 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 .

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 of 650 subjects 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. For more details about the study see Wester et al. (2010). 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)).

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 4: Simulation Results For Both Models Mis-specified: Mean of Estimates and Mean Square Error for All Six Estimators

Here we extend this TMLE to account for potential bias due to informative censoring by time-dependent covariates, CD4 and viral load, which have an affect on both the time to drop-out and the time to event of interest. We will directly compare results using the TMLE that only incorporates the baseline covariates to the TMLE that accounts for time dependent confounding. 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 difference in additive risk by gender at 36 months after randomization to cART therapy. We will estimate this parameter using the six estimators examined in the simulation analysis in the previous section. 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 5 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. 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, a 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. 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.

Table 6 shows the results for treatment effect modification by gender for the death outcome. In this table, 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 difference between TD TMLE and BASE TMLE is

Risk Difference @ 36 Months						
	Time Dependent			Baseline		
	TMLE	DR-EE	IPCW	TMLE	DR-EE	IPCW
Est	20.0%	20.1%	18.3%	18.9%	18.9%	18.3%
SE	5.0%	4.9%	10.3%	4.9%	4.9%	10.3%
p	<0.001	<0.001	0.074	<0.001	<0.001	0.074

Table 5: Gender Effect Modification on TLOVR

Risk Difference @ 36 Months						
	Time Dependent			Baseline		
	TMLE	DR-EE	IPCW	TMLE	DR-EE	IPCW
Est	6.3%	6.5%	5.2%	5.1%	5.1%	5.2%
SE	2.3%	2.3%	12.5%	2.4%	2.4%	12.5%
p-value	0.005	0.004	0.680	0.029	0.030	0.680

Table 6: Gender Effect Modification on Death

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 at 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 make it unreliable in practice.

6 Discussion

This article represents the first implementation of TMLE to estimate the causal effect of a multiple time point intervention that is subject to time-dependent confounding. In this particular case, the multiple time point intervention is represented




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

by a point treatment at baseline, and a time-dependent process that can only jump 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 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




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

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.

The results of our simulations and data analyses 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. The Cox-proportional hazards analysis is known to be biased because it ignores both baseline and time-dependent covariates and is 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. One particular limitation of TMLE is that it is a more complex method than IPCW estimators and this may make practitioners less likely to adopt TMLE. However, hopefully the advantages of TMLE in terms of stability and robustness will outweigh the complexity and computational costs needed for implementing TMLE.

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 estimates of the censoring and treatment mechanisms based on the log-likelihood of the resulting TMLE (van der Laan and Gruber (2010)). In the Appendix to our earlier technical report 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 (Stitelman et al. (2011)).

References

- Bang, H. and J. Robins (2005): “Doubly robust estimation in missing data and causal inference models,” *Biometrics*, 61, 962–972.
- 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.
- Kang, J. and J. Schafer (2007): “Demystifying double robustness: A comparison of alternative strategies for estimating a population mean from incomplete data (with discussion),” *Statistical Science*, 22, 523–39.
- 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., *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 (2011): “The relative performance of targeted maximum likelihood estimators,” *UC Berkeley Division of Biostatistics Working Paper Series*, 279.
- 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”,” *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*, Birkhäuser.
- Robins, J., M. Sued, Q. Lei-Gomez, and A. Rotnitzky (2007): “Comment: Performance of double-robust estimators when “inverse probability” weights are highly variable,” *Statistical Science*, 22, 544–559.
- 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.
- 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., V. De Gruttola, and M. van der Laan (2011): “A general implementation of tmle for longitudinal data applied to causal inference in survival analysis,” *UC Berkeley Division of Biostatistics Working Paper Series*, 281.

- 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 effect modification parameters in survival analysis,” *The International Journal of Biostatistics*, 7.
- 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 targeted maximum likelihood estimation,” *The International Journal of Biostatistics*, 6.
- 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.
- Wester, C. et al. (2010): “Non-nucleoside reverse transcriptase inhibitor outcomes among combination antiretroviral therapy-treated adults in botswana,” *AIDS*, 24, S27–S36.

A Appendix A: The Clever Covariate Under The Binary Factorization

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

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

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

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

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

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

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

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

and

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