UC Berkeley
UC Berkeley Previously Published Works

Title
Doubly Robust and Efficient Estimation of Marginal Structural Models for the Hazard Function

Permalink
https://escholarship.org/uc/item/5w11w6kr

Journal
The International Journal of Biostatistics, 12(1)

ISSN
2194-573X

Authors
Zheng, Wenjing
Petersen, Maya
van der Laan, Mark J

Publication Date
2016-05-01

DOI
10.1515/ijb-2015-0036

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

HHS Public Access

Author manuscript

Int J Biostat. Author manuscript; available in PMC 2016 August 01.

Published in final edited form as:

Int J Biostat. 2016 May 1; 12(1): 233–252. doi:10.1515/ijb-2015-0036.

Doubly Robust and Efficient Estimation of Marginal Structural Models for the Hazard Function

Wenjing Zheng[*], **Maya Petersen**, and **Mark van der Laan**

School of Public Health, University of California, Berkeley

Abstract

In social and health sciences, many research questions involve understanding the causal effect of a longitudinal treatment on mortality (or time-to-event outcomes in general). Often, treatment status may change in response to past covariates that are risk factors for mortality, and in turn, treatment status may also affect such subsequent covariates. In these situations, Marginal Structural Models (MSMs), introduced by Robins (1997), are well-established and widely used tools to account for time-varying confounding. In particular, a MSM can be used to specify the *intervention-specific counterfactual hazard function*, i.e. the hazard for the outcome of a subject in an ideal experiment where he/she was assigned to follow a given intervention on their treatment variables. The parameters of this hazard MSM are traditionally estimated using the Inverse Probability Weighted estimation (IPTW, van der Laan and Petersen (2007), Robins et al. (2000b), Robins (1999), Robins et al. (2008)). This estimator is easy to implement and admits Wald-type confidence intervals. However, its consistency hinges on the correct specification of the treatment allocation probabilities, and the estimates are generally sensitive to large treatment weights (especially in the presence of strong confounding), which are difficult to stabilize for dynamic treatment regimes. In this paper, we present a pooled targeted maximum likelihood estimator (TMLE, van der Laan and Rubin (2006)) for MSM for the hazard function under longitudinal dynamic treatment regimes. The proposed estimator is semiparametric efficient and doubly robust, hence offers bias reduction and efficiency gain over the incumbent IPTW estimator. Moreover, the substitution principle rooted in the TMLE potentially mitigates the sensitivity to large treatment weights in IPTW. We compare the performance of the proposed estimator with the IPTW and a non-targeted substitution estimator in a simulation study.

1 Introduction

In social and health sciences, many research questions involve understanding the causal effect of a longitudinal treatment on a time-to-event outcome (say mortality). Often, treatment status may change in response to past covariates that are risk factors for mortality, and in turn, treatment status may also affect such subsequent covariates. Moreover, right censoring may also be present in a study of this nature, typically in response to past covariates and treatment. In these situations, conventional time-dependent regression methods, which predict mortality rate at each time based on history up to that time, may fail to properly account for this time-varying confounding of the treatment effect (e.g. Robins et

[*] wenjing.zheng@berkeley.edu.

Author Manuscript Author Manuscript Author Manuscript Author Manuscript

al. (2000b)). Marginal Structural Models (MSMs), introduced by Robins (1997), are well-established and widely used tools in this setting.

To define the causal effect of a longitudinal treatment on mortality, we can use a formal causal framework (e.g. Robins (1986), Pearl (2009)). In it, we formulate the research question of interest in terms of an ideal experiment that randomizes interventions on the treatment variables of interest. The comparative effect of the different interventions is assessed by comparing the distribution of the would-be (*counterfactual*) outcome processes under said interventions. The MSMs specify the marginal distribution of this *intervention-specific counterfactual outcome processes*.

In survival analysis, the effect of treatment interventions on mortality can be assessed by studying how the hazard of the intervention-specific counter-factual mean outcome changes as a function of time and the intervention, possibly conditional on a set of baseline covariates. In many applications, this hazard function may be high-dimensional, in which case a MSM for the hazard can be used to summarize its key features. This *hazard MSM* can provide adequate dimension reduction over a potentially complicated hazard function, and may also mitigate near positivity violations (due to lack of experimental support for an intervention) by extrapolation.

Depending on the research questions of interest, the interventions can be *static* — treatment option is the same for all subjects, or they can be *dynamic* — treatment option is a function of past covariates. For example, among HIV+ patients who failed their first line antiretroviral regimen, we wish to evaluate how their survival depends on when the regimen modification takes place. On the one hand, we may be interested in comparing the lengths of delay in regimen modification: Petersen et al. (2008) and Petersen et al. (2014b) defined static interventions by fixed lengths of delay and used a hazard MSM to assess how the survival changes as a function of the length of delay. On the other hand, we can take a patient-centered approach (perhaps more inline with clinical practice) and conceptualize the delay in regimen modification as a function of past CD4 T-cell counts: in (van der Laan and Petersen, 2007), dynamic interventions prescribe regimen modification only after CD4 counts dropped below a given threshold, and a hazard MSM assesses how the survival changes as a function of such CD4 count thresholds. Other examples of hazard MSMs include the study of the effect of “when to initiate” antiretroviral therapy on survival, under static interventions (Hernan et al. (2000)), using Case-Cohort designs (Cole et al. (2012)), or under CD4-based dynamic interventions (Cain et al. (2010) and HIV-CAUSAL-Collaboration et al. (2011)).

The parameters of an MSM are traditionally estimated using Inverse Probability of Treatment Weighted estimation (IPTW, Robins et al. (2000b), Robins (1999), van der Laan and Petersen (2007), Robins et al. (2008)). This estimator is intuitive, can be implemented using standard software, and admits influence curve based variance estimates and confidence intervals. However, its consistency hinges on correct specification of the conditional treatment probabilities. Moreover, in the presence of strong confounding, the inverse probability weights can become unwieldily large, thus producing very unstable estimates. In the case of static intervention, this instability can be attenuated by introducing marginal

Author Manuscript Author Manuscript Author Manuscript Author Manuscript

kernel weights that down-weight treatment options with little data support; but this solution has limited applicability for dynamic interventions prescribed by time-varying history.

To address the sensitivity to model misspecification of IPTW, a doubly robust and efficient Augmented-IPTW (A-IPTW) estimator for a MSM under static interventions was proposed in Robins and Rotnitzky (1992), Robins (2000) and Robins et al. (2000a). Under this framework, estimators are defined as solutions to the estimating equation given by the efficient influence curve, which also depends on nuisance parameters orthogonal to the treatment probabilities. Robins (2000), Robins (2002) and Bang and Robins (2005) offered an innovative insight that the corresponding efficient influence curves for the MSM for the intervention-specific mean can be expressed in terms of sequential conditional expectations. This observation allowed for construction of an A-IPTW with estimation of minimal nuisance parameters beyond the treatment mechanism. While A-IPTW estimator provides efficiency gain and bias reduction over a misspecified IPTW estimator, as an estimation equation-based method, it still suffers from the same general sensitivity to large inverse probability weights, and it may involve solving estimating equations that have no unique solution.

To improve the stability of the estimates, the targeted maximum likelihood estimation (TMLE, van der Laan and Rubin (2006) and Gruber and van der Laan (2010)) provides a general doubly robust and efficient estimator using the plug-in principle, which incorporates global information encoded in the parameter map and the model. A TMLE estimator for longitudinal static MSMs using a stratified approach was proposed by Schnitzer et al. (2014). A stratified TMLE uses the longitudinal TMLE for mean outcomes (van der Laan and Gruber (2012)) to separately estimate each intervention-specific mean; these means are then used to fit MSMs for the hazard and survival functions. This estimator readily improves upon IPTW in both robustness and efficiency. However, in applications with numerous interventions of interest, the stratified TMLE is vulnerable to insufficient support for certain interventions, as it does not directly target the MSM parameters to take advantage of the extrapolation across interventions. Most recently, Petersen et al. (2014a), building on the results from Robins (2000), Robins (2002) and Bang and Robins (2005), presented a pooled TMLE estimator for a longitudinal dynamic MSM for intervention-specific means (in particular, for survival functions). This estimator directly targets the parameters of the MSM, and pools over all interventions of interest in the updating step, potentially weakening the data support needed to achieve efficiency and double robustness.

This paper builds upon the work of Petersen et al. (2014a) to present a pooled TMLE for the MSM for the hazard function under longitudinal dynamic interventions. The efficient influence curve of the hazard MSM requires that a plug-in estimator be updated iteratively until convergence, but we show that, in practice, a stopping rule based on the relative sizes of the empirical mean of efficient score equation and the estimated standard error will suffice to produce MSM estimates with the desired theoretical properties. The proposed estimator is efficient and doubly robust, hence offers an improvement over traditional IPTW estimator; it directly targets the MSM parameters by pooling across interventions to update the initial nuisance parameter estimates, hence is less sensitive to data support issues than the stratified TMLE.

Author Manuscript Author Manuscript Author Manuscript Author Manuscript

In practice, when the research question of interest directly concerns the survival function and one can correctly specify its MSM, using a TMLE for the survival MSM, as proposed in Petersen et al. (2014a), would be a more straightforward approach than using our proposed estimator for the hazard MSM. However, in many applications the hazard function itself (or failure rate, as coined in some disciplines) is of substantive interest, since it describes how the rate of mortality changes with time. Moreover, in some cases, one may have more domain knowledge to specify the hazard MSM than to specify the survival MSM. In particular, the survival function is cumulative and hence must be monotonic in time, whereas the hazard function can take different shapes, and hence can capture more domain information (such as a decreasing hazard, or a U-shaped hazard with spikes at the beginning and end time points). In applications where the research question is best served (or the domain knowledge is best used) by studying the hazard function, our proposed estimator which directly estimates the parameters of a hazard MSM would be a more targeted approach. In addition to these application considerations, the hazard MSM also allows for the use of more stable weights that only need to be defined among subjects that are still alive.

Organization of paper

This paper is organized as follows. Section 2, describes the data structure and defines the parameter using a nonparametric causal framework. Section 3 first reviews the non-targeted substitution estimator (a.k.a. G-computation estimator) and the IPTW estimator for the parameters of interest, and then presents the efficient influence curve, and describes the proposed TMLE estimator. Section 4 evaluates the performance of these three estimators in a simulation study mimicking an observational cohort study. This paper concludes with a summary.

2 Parameters of Interest

Consider a longitudinal data structure

$$O=(L_0, A_0, L_1, A_1, ..., L_t, A_t, ..., L_K, A_K, L_{K+1}),$$

where $K+1$ is a user-supplied maximum follow-up time of interest. Here, L_0 denotes the baseline covariates; A_t encodes the treatment variable A_t^1 and the censoring indicator A_t^2, where $A_t^2=1$ indicates the subject was right censored at or before time t; and L_t denotes all the time-varying covariates measured between A_{t-1} and A_t. In particular, L_t includes the counting outcome process $Y_t \subset L_t$, where $Y_t = 1$ indicates event (say death) has occurred at or before time t. For notational convenience, after death or censoring, all variables are encoded by carrying forward the last observation. We shall use the boldface notation $\mathbf{L}_t \equiv (L_0,\ldots,L_t)$, $\mathbf{L}_{j,t} \equiv (L_j,\ldots,L_t)$ and $\mathbf{L}_{-1} \equiv \emptyset$; similarly for their realizations $\ell_t \equiv (\ell_0,\ldots,\ell_t)$, $\ell_{j,t} \equiv (\ell_j,\ldots,\ell_t)$. Analogous notations apply to the vector $\mathbf{A}_t$. The observed data consist of n independent and identically distributed (i.i.d.) copies of O drawn from a distribution P_0. Let $\mathcal{M}$ be a statistical model for P_0; the assumptions on this statistical model are limited to true user knowledge, in particular, we avoid strong and restrictive parametric assumptions.

Author Manuscript Author Manuscript Author Manuscript Author Manuscript

To illustrate these notations and the subsequent concepts, let us consider an example from HIV research (van der Laan and Petersen (2007)). In HIV+ patients with virologic failure on first line antiretroviral therapy (ART), current guidelines (DHHS (2014)) recommend prompt switch to a second line regimen. However, delayed switch is common in resource limited settings. Suppose we wish to assess the effect of delayed switch on mortality (e.g. van der Laan and Petersen (2007), (Petersen et al. (2008), Petersen et al. (2014b)). Our sample consists of n individuals drawn from a target population of HIV+ patients with confirmed virologic failure on their first line ART. The baseline $t = 0$ is the time of confirmed failure, and data is collected on a monthly basis thereafter. Variable L_t encompasses time varying covariates at time t, including CD4+ T cell counts and Y_t, an indicator of death by time t. Variable L_0 consists of the baseline values of these covariates at $t = 0$, as well as time-independent variables such as patient demography and history prior to virologic failure. The treatment variable A_t^1 is the indicator of switching to second line regimen by time t. For simplicity sake, for now we assume no right censoring in this example.

2.1 Causal Model and Causal Parameters

The causal assumptions implied in the notation of O can be made explicit by a nonparametric structural equations model (NPSEM, Pearl (2009)):

$$L_t = f_{L_t}(\mathbf{L}_{t-1}, \mathbf{A}_{t-1}, U_{L_t}); A_t = f_{A_t}(\mathbf{L}_t, \mathbf{A}_{t-1}, U_{A_t}), \quad \text{for } t = 0, \ldots, K+1.$$

This causal model assumes that each variable X in the observed data structure is an unknown deterministic function f_X of certain observed variables, which we refer to as the *parents of X* and denote by $Pa(X)$, and some unmeasured exogenous random factors U_X. This causal model defines a random variable with distribution $P_{O,U}$ on a unit.

In practice, additional exclusion restriction assumptions informed by subject matter knowledge can be represented by limiting the variables in $Pa(X)$. In the HIV example, at each time t, the investigators may specify that CD4 counts, death, and regimen modification all depend on the patient's entire observed past and unmeasured factors. But if they know that the decision to switch regimen is based only on the most recent CD4 measurement, then $Pa\left(A_t^1\right)$ may be restricted to exclude all earlier CD4 measurements.

As we alluded to in section 1, research questions can be formalized in terms of intervention rules on the treatment variables. An *intervention rule d* is a function that deterministically assigns treatment at time t based on covariate history according to $A_t = d(\mathbf{L}_t)$. This rule may be static — assigning the same treatment option a_t regardless of covariate history, i.e. $A_t = d(\mathbf{L}_t) \equiv a_t$; or it may be dynamic — assigning different treatment options to different covariate histories. We will use the boldface notation $\mathbf{A}_t = \mathbf{d}(\mathbf{L}_t)$ to denote the vector $(A_k = d(\mathbf{L}_k))_{k=0,\ldots,t}$; this means that each of the variables A_k, from $k = 0$ to t, was assigned value according to rule d and its covariate history. In the HIV example, a static rule would be to always switch regimen at m months after virologic failure, and a dynamic rule would be to

Author Manuscript Author Manuscript Author Manuscript Author Manuscript

switch ART at the first time t when the patient's CD4 counts drop below a pre-specified threshold θ.

Given a set $\mathscr{D}$ of intervention rules of interest, investigators are often concerned about their comparative causal effect on the outcome process Y_t. To be more precise, consider an ideal experiment where all subjects are assigned treatment under an intervention rule d and right censoring is also prevented; the covariates, on the other hand, take the value that they may in response to rule d. We call the variables A_t the *intervention variables*, as they are the ones that are subject to manipulation. This ideal experiment can be formalized in the NPSEM by setting the equations for A_t to $A_t = d(\mathbf{L}_t)$, and replacing the input $\mathbf{A}_{t-1}$ in f_{L_t} with $\mathbf{d}(\mathbf{L}_{t-1})$, which are treatment assignments from time 0 to $t-1$ under rule d and history $\mathbf{L}_{t-1}$. As a result, the only random endogenous variables of the system are the covariates; we use $L_t(d)$ to denote the time varying covariates that result under the intervention rule d, in particular $Y_t(d)$ is the indicator of death under such a regimen. The comparative causal effect of rule d_1 vs rule d_2 can be assessed by comparing the distributions of the outcome processes $Y_t(d_1)$ vs $Y_t(d_2)$.

Suppose we wish to understand how the outcome process $Y_t(d)$ changes as a function of d, t and some baseline covariate $V \subset L_0$. To this end, we study the *intervention-specific hazard function*

$$\lambda(d,t,V) \equiv P(Y_t(d)=1 | Y_{t-1}(d)=0, V) \quad (1)$$

on the space $\mathscr{D} \times \tau \times \mathscr{V}$, where $\mathscr{D}$ is the set of rules we wish to compare, $\tau = \{1,\ldots K+1\}$ contains all the follow-up times of interest, and $\mathscr{V}$ is the outcome space of V.

Studying the entire function $(d,t,V) \mapsto \lambda(d,t,V)$ is often difficult due to feasibility of computation, interpretation of results, and other challenges associated with high-dimensionality. Instead, we can study a more tractable, simplified model/summary of this function which captures how λ changes as a function of (d,t,V) in only a few summarizing parameters. More specifically, following the *marginal structural working model* framework in Neugebauer and van der Laan (2007), let $m_\psi(d,t,V)$ be a working MSM for $\lambda(d,t,V)$, parameterized by $\psi \in \mathbb{R}^J$. To be concrete, from here on, we will use a generalized linear hazard MSM with logit link,

$$m_\psi(d,t,V) = \text{expit}\,(\psi \cdot \phi(d,t,V)),$$

where $\phi(d,t,V)$ is the vector of linear predictors. However, it is important to note that the methods presented here can be generalized to other working MSM.

In addition to specifying the working MSM, we also specify a user-supplied kernel weight function $h(d,t,V)$, and the standard log-likelihood loss. Formally, the *causal parameter of interest* is defined as

Author Manuscript Author Manuscript Author Manuscript Author Manuscript

$$\Psi(P_{O,U}) \equiv \arg\min_{\psi} E_0 \left\{ \sum_{t\in\tau, d\in\mathscr{D}} h(d,t,V) I(Y_{t-1}(d)=0) \times \log \left(m_\psi(d,t,V)^{Y_t(d)} (1 - m_\psi(d,t,V))^{1-Y_t(d)} \right) \right\}. \\ = \arg\min_{\psi} E_0 \sum_{t\in\tau, d\in\mathscr{D}} h(d,t,V) P(Y_{t-1}(d)=0|L_0) \times \log \left(m_\psi(d,t,V)^{\lambda(d,t,L_0)} (1 - m_\psi(d,t,V))^{1-\lambda(d,t,L_0)} \right). \quad (2)$$

In words, our causal parameter of interest is the vector of coefficients of the MSM weighted by h. Under the working MSM framework, this parameter can be understood as the *best* (assessed by the loss function) *weighted* (by h) *approximation* (given by the working model m_ψ) of the hazard function λ. Embedded in the definition of the parameter are the stabilizing weights $h(d,t,V)$. To mitigate positivity violations, $h(d,t,V)$ can be chosen to down weight a rule d which has little data support at time t within strata of marginal covariates V. But the effectiveness of this stabilization will be limited when d uses time-varying tailoring variables.

Continuing our HIV example, suppose we wish to assess how delay in switching regimen affect mortality. Let d_m denote the rule that dictates switching regimen at m months after virologic failure. That is, $A_t = d_m(\mathbf{L}_t) = 0$ for $t < m$ and $A_t = d_m(\mathbf{L}_t) = 1$ for $t \geq m$, given $Y_t = 0$. Let $\mathscr{D}$ be the set of all possible switching times we wish to compare, i.e. $\mathscr{D} = \{d_m : m = 0, \ldots, K, \infty\}$. If we are only interested in the marginal hazard, then we set $V = \emptyset$. Or, if we wish to assess the hazard function stratified by CD4 at the time of virologic failure, then we can set $V = \mathbf{1}\{CD4_0 < 500 \text{ cells}/\mu\ell\}$. For simplicity, let us continue this example with the former option, and let the weights be $h(d,t) = 1$. For this example, we choose the logistic working model to be $m_\psi(d_m,t,V) = \text{expit}(\psi_0 + \psi_1 t + \psi_2 max(t-m,0))$, where $max(t-m,0) = 0$ if the patient has not switched by time $t-1$, and $max(t-m,0) = t-m$ encodes how long the patient has been on his second line regimen if he has already switched.

2.2 Statistical Parameter

Thus far, we have used the NPSEM to formulate the parameter of interest $\Psi(P_{O,U})$ in (2). This parameter is a function of the distributions of $L_t(d)$ and $Y_t(d)$, which are generated within an ideal experiment. Unfortunately, such ideal experiments are not always possible in real life. Then, what are the sufficient assumptions on the data-generating process under which the parameter $\Psi(P_{O,U})$ can be estimated using the observed data?

To answer this question, we review the expression in (2) and note that $P(Y_{t-1}(d)=0|L_0) = 1 - E(Y_{t-1}(d)|L_0)$ and

$$\lambda(d,t,L_0) = \frac{P(Y_t(d)=1,L_0) - P(Y_t(d)=1, Y_{t-1}(d)=1, L_0)}{P(Y_{t-1}(d)=0,L_0)} \\ = \frac{E(Y_t(d)|L_0) - E(Y_{t-1}(d)|L_0)}{1 - E(Y_{t-1}(d)|L_0)}. \quad (3)$$

Author Manuscript Author Manuscript Author Manuscript Author Manuscript

Consequently, to identify $\Psi(P_{O,U})$ as a function of the data generating distribution P_0, it suffices to establish the identification of the time-dependent causal dose-response curve $\{E(Y_t(d)|L_0):d \in \mathscr{D}, t \in \tau\}$.

To this end, we make the *sequential randomization assumption* (Robins (1986)):

$$A_k \perp \mathbf{L}(d)|Pa(A_k), \quad (4)$$

and the *positivity assumption*

$$P_0(A_k=d(\mathbf{L}_k)|\mathbf{L}_k, \mathbf{A}_{k-1}=\mathbf{d}(\mathbf{L}_{k-1}))>0 \; a.e., \quad (5)$$

for every $d \in \mathscr{D}$, $k \in \tau$, and $V \in \mathscr{V}$ with $h(d,t,V) > 0$. Assumption (4) specifies that at each k, the variable A_k is randomized conditional on its parent variables. In our HIV example, this can be satisfied if the parent variables of A_k contain all common determinants of mortality and the decision to switch regimen. Assumption (5) requires that under P_0, for each rule d and its compatible covariate history for which the weight $h(d,k,V)$ is non-zero, there is non-zero probability that the subject's treatment will continue to follow this rule. In our HIV example with $h(d,t,V) = 1$, given one has not yet switched regimen, there should be non-zero probability that the patient will switch at any given time m in $\mathscr{D}$. Assumption (5) would be violated if, say, $\mathscr{D}$ contains the rule to switch at 6 months but no patient in the study population actually switches at 6 months post-failure. This assumption would also be violated if, say, all patients with switch regimen at month 1.

Under assumptions (4) and (5), we obtain identification of the intervention-specific mean:

$$E(Y_t(d)|L_0)=\sum_{\ell_{1,t}} y_t \prod_{j=1}^{t} P_0(L_j=\ell_j|L_0, \ell_{1,j-1}, \mathbf{A}_{j-1}=\mathbf{d}(L_0, \ell_{1,j-1})) \equiv Q_1^{d,t}(P_0)(L_0). \quad (6)$$

We are reminded that $\boldsymbol{\ell}_{1,j-1} \equiv (\ell_1, \ell_{j-1})$. This is generally known as the G-computation formula (Robins (1986)) for the intervention-specific mean. This in turn identifies the intervention-specific hazard function:

$$\lambda(d,t,L_0)=\frac{Q_1^{d,t}(P_0)(L_0) - Q_1^{d,t-1}(P_0)(L_0)}{1 - Q_1^{d,t-1}(P_0)(L_0)}. \quad (7)$$

Consequently, the causal parameter of interest $\Psi(P_{O,U})$ in (2) is identified as a function of the observed data distribution given by

Author Manuscript Author Manuscript Author Manuscript Author Manuscript

$$\psi_0 \equiv \Psi(P_0) \equiv \underset{\psi}{\operatorname{argmin}} \left\{ E_{P_0} \sum_{t\in\tau, d\in\mathscr{D}} h(d,t,V) \left(1 - Q_{1,0}^{d,t-1}(L_0)\right) \right.$$
$$\left. \times \left(\frac{Q_{1,0}^{d,t}(L_0) - Q_{1,0}^{d,t-1}(L_0)}{1 - Q_{1,0}^{d,t-1}(L_0)} \log m_\psi(d,t,V) + \frac{1 - Q_{1,0}^{d,t}(L_0)}{1 - Q_{1,0}^{d,t-1}(L_0)} \log\left(1 - m_\psi(d,t,V)\right) \right) \right\}. \quad (8)$$

Note that the sequential randomization assumption (4) is an untestable assumption on the data-generating process under which we can claim the causal parameter (2) equals the statistical parameter (8). It provides a guiding principle to judge the causal interpretation of the statistical parameter. By contrast, the positivity assumption (5) is a testable assumption on the data-generating distribution P_0 under which the statistical parameter (8) is well-defined.

The remainder of this paper focuses on statistical inference for ψ_0.

3 Estimators for ψ_0

Recall that the observed data consist of n i.i.d. copies of $O \sim P_0 \in \mathscr{M}$. Before we introduce the proposed efficient and doubly robust estimator, we will review two available estimators for ψ_0: the non-targeted substitution G-computation estimator (Robins (1986), Taubman et al. (2009)) and the aforementioned IPTW estimator. The proposed targeted maximum likelihood estimator uses either of these to obtain an initial estimator of ψ_0. But before we proceed, we shall agree on the following notation.

3.1 Notations

We use P_n to denote the empirical distribution of n i.i.d. copies of $O \sim P_0$. Given a function $O \mapsto f(O)$, $P_n f$ denotes the empirical mean $P_n f \equiv \frac{1}{n}\sum_{i=1}^n f(O_i)$. More generally, for any $P \in \mathscr{M}$, $Pf \equiv E_P f(O)$.

For a generic $P \in \mathscr{M}$. We use Q^{L_0} to denote the marginal distribution of L_0. Generalizing the functional in (6), for a given $t \in \tau$ and $d \in D$, denote at $P \in \mathscr{M}$.

$$Q_k^{d,t}(P)(\mathbf{L}_{k-1}) \equiv \sum_{\ell_{k,t}} y_t \prod_{j=k}^{t} P(L_j = \ell_j | \mathbf{L}_{k-1}, \ell_{k,j-1}, \mathbf{A}_{j-1} = \mathbf{d}(\mathbf{L}_{k-1}, \ell_{k,j-1})), \text{ for } k \le t, \quad (9)$$

and $Q_{t+1}^{d,t}(\cdot) \equiv Y_t$. In the above notation, the superscript t signals the expectant outcome variable Y_t and the subscript k signals the length of the conditioning covariate history. Under assumptions (4) and (5), $Q_k^{d,t}(P_0)(\mathbf{L}_{k-1}) = E(Y_t(d) | \mathbf{L}_{k-1}, \mathbf{A}_{k-1} = \mathbf{d}(\mathbf{L}_{k-1}))$, which is the conditional mean of $Y_t(d)$ given an observed past that has followed the rule d up to time $k-1$. For simplicity, we may sometimes write $Q_k^{d,t}$ instead of $Q_k^{d,t}(P)$ when referring to the functional evaluated at a generic $P \in \mathscr{M}$, and $Q_{k,0}^{d,t}(\mathbf{L}_{k-1})$ instead of $Q_k^{d,t}(P_0)(\mathbf{L}_{k-1})$ when

Author Manuscript Author Manuscript Author Manuscript Author Manuscript

the functional is evaluated at $P_0 \in \mathscr{M}$. Bang and Robins (2005) made a key observation that these functionals satisfy the relation

$$Q_k^{d,t}(\mathbf{L}_{k-1}) = E_P\left[Q_{k+1}^{d,t}(\mathbf{L}_k)|\mathbf{L}_{k-1}\right]. \quad (10)$$

For a counting outcome process, these functionals are also monotonic in t, i.e. for a given k, $Q_k^{d,t} \leq Q_k^{d,t+1}$ for all $t \geq k$. We denote the components corresponding to the non-intervention variables as $Q \equiv \left(Q^{L_0}, \left\{Q_k^{d,t}:, t \in \tau, k \leq t, d \in \mathscr{D}\right\}\right)$. For the intervention variables, we denote the treatment allocation probabilities $P(A_k/\mathbf{L}_k, \mathbf{A}_{k-1})$ as $g(A_k/\mathbf{L}_k, \mathbf{A}_{k-1})$, and the product $\prod_{k=1}^{t} g(A_k|\mathbf{L}_k, \mathbf{A}_{k-1})$ as $\mathbf{g}(\mathbf{A}_t/\mathbf{L}_t)$. For our purposes, the couple (Q,g) readily specifies a distribution $P \in \mathscr{M}$, so sometimes we may abuse notation and write $P = (Q,g)$. At the data generating distribution P_0, we adopt the subscripts Q_0 and g_0.

3.2 G-computation Estimator

The G-computation formula in (6) readily delivers a non-targeted substitution estimator (as opposed to the targeted substitution estimator that is TMLE), which is generally known as the parametric G-computation (Gcomp) estimator. More precisely, using the notations in section 3.1, the statistical parameter $\Psi(P_0)$ in (8) can be expressed as $\Psi(Q_0)$. Therefore, fitting the hazard MSM using a non-targeted estimator Q_n of Q_0 yields a non-targeted substitution estimator $\Psi(Q_n)$ of $\Psi(Q_0)$.

Recall that Q_0 consists of the marginal distribution of L_0, $Q^{L_0}(P_0)$, and the conditional means of $Y_t(d)$, $Q_k^{d,t}(P_0)$ defined in (9). To estimate the marginal distribution $Q^{L_0}(P_0)$, we can use the empirical distribution of L_0, denoted $Q_n^{L_0}$. To estimate the conditional means $Q_k^{d,t}(P_0)$, one approach is to estimate the conditional densities for each L_t given its parents and use the definition in (9). While this density-based approach ensures that the monotonicity in t of $Q_k^{d,t}$ is preserved, the dimension of the nuisance parameter and the computational cost grows with the dimension of L_t and the number of time points. One way to implement this approach is to use simplifying parametric modeling assumptions on these nuisance parameters. We refer to Taubman et al. (2009) and Young et al. (2011) for expositions of this technique.

Another approach to estimate the conditional means $Q_k^{d,t}(P_0)$ is to minimize estimation of nuisance parameters by exploiting the recursive relation (10) noted by Bang and Robins (2005), which allows one to use the many available regression techniques (parametric or data-adaptive) in the literature. We can implement this regression-based approach by running the following two-level algorithm. Recall that in our notation $Q_k^{d,t}(P_0)$, d is the intervention of interest, t is the time of the expectant outcome, and k is the length of the conditioning history.

Author Manuscript Author Manuscript Author Manuscript Author Manuscript

1. Starting at $t = K+1$, estimate the conditional means $\left\{Q_k^{d,t}(P_0): k \leq t, d \in \mathscr{D}\right\}$ as follows:

   a. Initiate at $k = t$. Recall that $Q_t^{d,t}(P_0) \equiv E_{P_0}(Y_t | \mathbf{L}_{t-1}, \mathbf{A}_{t-1} = \mathbf{d}(\mathbf{L}_{t-1}))$ and $Q_{k+1}^{d,t} \equiv Y_t$. We obtain an estimator $Q_{t,n}^{d,t}$ of $Q_t^{d,t}(P_0)$ by regressing Y_t on the observed values $\mathbf{L}_{t-1}$ and $\mathbf{A}_{t-1}^1$ among observations that remain uncensored by time $t-1$, and evaluate the fitted function at the observed $\mathbf{L}_{t-1}$ and the intervened values $\mathbf{A}_{t-1}^1 = \mathbf{d}(\mathbf{L}_{t-1})$, for these uncensored observations. For a subject that has died before time t, the value is deterministically set to 1. We can organize the data by having one row for each patient i and each regimen d.

   b. At each subsequent $k = t-1, \ldots, 1$: At the previous step, we have thus far obtained an estimator $Q_{k+1,n}^{d,t}$ of $Q_{k+1}^{d,t}(P_0)$. To obtain an estimator $Q_{k,n}^{d,t}$ of $Q_k^{d,t}(P_0)$, we regress $Q_{k+1,n}^{d,t}(\mathbf{L}_k)$ on the observed values $\mathbf{L}_{k-1}$ and $\mathbf{A}_{k-1}^1$ among those observations that remained uncensored by time $k-1$, and evaluate the fitted function at the observed $\mathbf{L}_{k-1}$ and the intervened values $\mathbf{A}_{k-1}^1 = \mathbf{d}(\mathbf{L}_{k-1})$ on these uncensored observations. For observations that had died before time k, the value is deterministically set to 1.

   c. After iterating step (b) in order of decreasing k, we have obtained a sequence of means $\left\{Q_{k,n}^{d,t}: k \leq t, d \in \mathscr{D}\right\}$.

2. Repeat step 1 in order of decreasing t, from $t = K$ to $t = 1$. At the end, we will have obtained estimators $\left\{Q_{k,n}^{d,t}: t \in \tau, k \leq t, d \in \mathscr{D}\right\}$

For each history k, monotonicity in $t \geq k$ of the expectations, i.e. $Q_{k,n}^{d,t} \leq Q_{k,n}^{d,t+1}$, can be enforced in step 1.b with respect to $Q_{k,n}^{d,t+1}$ obtained at the previous $t+1$ level through a simple sequential truncation. To do this, for subjects that have not died or censored, we set $Q_{k,n}^{d,t} = Q_{k,n}^{d,t+1}$ whenever $Q_{k,n}^{d,t} > Q_{k,n}^{d,t+1}$. Other more sophisticated approaches are available, but they are outside the scope of this paper.

At the end of this algorithm, we have the sequence of conditional means $\left\{Q_{1,n}^{d,t}(L_0): t \in \tau, d \in \mathscr{D}\right\}$ for each of the n observations. Pooling together these estimates $Q_{1,n}^{d,t}(L_0)$ over all d and t, we have one row per patient i, rule $d \in \mathscr{D}$ and $t \in \tau$. The G-computation estimator $\psi_n^{Gcomp} \equiv \Psi(Q_n)$ is obtained by fitting a weighted logistic regression of $\dfrac{Q_{1,n}^{d,t}(L_0) - Q_{1,n}^{d,t-1}(L_0)}{1 - Q_{1,n}^{d,t-1}(L_0)}$ according to the MSM, with weights $h(d,t,V)(1 - Q_{1,n}^{d,t-1}(L_0))$.

Author Manuscript Author Manuscript Author Manuscript Author Manuscript

Consistency of ψ_n^{Gcomp} relies on consistency of Q_n. In the density-based approach, this means consistent estimation of all the conditional densities of L_t given its parents; in the regression approach, this means consistent estimation of the conditional means $Q_k^{d,t}(P_0)$. In either case, correct specification of Q_0 under a finite dimensional parametric model is rarely possible in practice. Alternatively, we may use machine learning algorithms, such as Super Learner. This option is more enticing, especially when used with the regression-based approach, since there are more data-adaptive techniques available to estimate a conditional mean via regression. However, unlike the Inverse Probability of Treatment Weighted Estimator (section 3.3) and the Targeted Maximum Likelihood Estimator (section 3.4) which satisfy a central limit theorem under appropriate empirical process conditions, analogous results for ψ_n^{Gcomp} are not available, regardless of model size. In addition, a non-targeted estimator Q_n of Q_0 is obtained by minimizing a global loss function for Q_0, not for $\Psi(Q_0)$. This means, in particular, that the bias-variance tradeoff in Q_n is optimized for the high-dimensional nuisance parameter Q_0, instead of a much lower-dimensional parameter of interest $\Psi(Q_0)$. The proposed targeted estimator in section 3.4 aims to address these two issues by providing a substitution estimator that is asymptotically linear (under appropriate regularity conditions), and optimizes the bias-variance tradeoff of Q_n towards $\Psi(Q_0)$ via an updating step.

3.3 Inverse Probability of Treatment Weighted Estimator

Inverse probability of treatment weighted estimation is the standard methodology for estimating the parameters of a marginal structural model in the presence of time-varying confounding (Robins et al. (2000b)), due to its ease of implementation and its asymptotic linearity, which allows for construction of Wald-type confidence intervals.

To begin, we first note that the parameter in (8) can be rewritten in an IPTW form:

$$\begin{aligned}\Psi(P_0) &= \Psi^{IPTW}(g_0)\\ &\equiv \underset{\psi}{\operatorname{argmin}} \left\{ E_{P_0} \sum_{t\in\tau, d\in\mathscr{D}} h(d,t,V) \frac{I(\mathbf{A}_{t-1}=\mathbf{d}(\mathbf{L}_{t-1}))}{\mathbf{g}_0(\mathbf{A}_{t-1}=\mathbf{d}(\mathbf{L}_{t-1})|\mathbf{L}_{t-1})}(1-Y_{t-1}) \right.\\ &\quad \left. \times \log\left(m_\psi(d,t,V)^{Y_t}(1-m_\psi(d,t,V))^{1-Y_t} \right) \right\}. \end{aligned} \tag{11}$$

The IPTW estimator ψ_n^{IPTW} is obtained by fitting a weighted logistic regression of Y_t according to the MSM, with weights $h(d,t,V)\frac{I(\mathbf{A}_{t-1}=\mathbf{d}(\mathbf{L}_{t-1}))}{\mathbf{g}_n(\mathbf{A}_{t-1}=\mathbf{d}(\mathbf{L}_{t-1})|\mathbf{L}_{t-1})}(1-Y_{t-1})$, where g_n is an estimator for g_0.

The asymptotic theory of the IPTW estimator is well understood in the literature. We refer the reader to Robins (1999), van der Laan and Robins (2003) and van der Laan and Petersen (2007), where the last reference specifically addresses dynamic intervention rules. In summary, ψ_n^{IPTW} described above satisfies the estimating equation $P_n D^{IPTW}(\psi, g_n) = 0$, where

Author Manuscript Author Manuscript Author Manuscript Author Manuscript

$$D^{IPTW}(\psi, g) = \sum_{t\in\tau, d\in\mathscr{D}} \tilde{h}(d,t,V)\frac{I(\mathbf{A}_{t-1}=\mathbf{d}(\mathbf{L}_{t-1}))}{\mathbf{g}(\mathbf{A}_{t-1}=\mathbf{d}(\mathbf{L}_{t-1})|\mathbf{L}_{t-1})}(1-Y_{t-1})(Y_t - m_\psi(d,t,V)), \quad (12)$$

with $\tilde{h}(d,t,V) \equiv h(d,t,V)\phi(d,t,V)$. The estimator ψ_n^{IPTW} is well-defined if it is a unique minimizer, and it is a consistent estimator of ψ_0 provided g_n is a consistent estimator of g_0. Moreover, ψ_n^{IPTW} is asymptotically linear with influence curve

$$IC_{IPTW}(g_0) = M^{IPTW}(\psi_0, g_0, P_0)^{-1} D^{IPTW}(\psi_0, g_0), \quad (13)$$

where

$$M^{IPTW}(\psi', g, P) \equiv -\frac{\partial\, PD^{IPTW}(\psi,g)}{\partial\,\psi}|_{\psi=\psi'}$$
$$= E_P\left\{ \sum_{t\in\tau, d\in\mathscr{D}} h(d,t,V)\phi(d,t,V)\phi(d,t,V)^T \frac{I(\mathbf{A}_{t-1}=\mathbf{d}(\mathbf{L}_{t-1}))}{\mathbf{g}(\mathbf{A}_{t-1}=\mathbf{d}(\mathbf{L}_{t-1})|\mathbf{L}_{t-1})}(1-Y_{t-1}) \times m_{\psi'}(d,t,V)(1-m_{\psi'}(d,t,V))\right\}.$$

The sample covariance matrix of the estimated influence curve is given by $\sum_n^{IPTW} \equiv M^{IPTW}(\psi_n^{IPTW}, g_n, P_n)^{-1} D^{IPTW}(\psi_n^{IPTW}, g_n)$. The variance of $\sqrt{n}(\psi_n^{IPTW} - \psi_0)$ can be estimated using $\sum_n^{IPTW}$. Consequently, we can construct Wald confidence intervals of level $(1-\alpha)$ as $\left[\psi_n^{IPTW} \pm \xi_{1-\alpha/2}(\sum_n^{IPTW}/n)^{1/2}\right]$, where $\xi_{1-\alpha/2}$ is the $(1-\alpha/2)$-quantile of the standard normal distribution. These confidence intervals ought to be conservative when g_0 is estimated using a correctly specified parametric model. But no such theoretical guarantee for data-adaptive estimation of g. Note that this influence curve based variance estimate assumes that the weights $h(d,t,V)$ are known functions; when these weights are estimated, this variance estimate should be interpreted as estimating the variance of an estimator of MSM parameters defined by the estimated weights.

Though both G-computation estimator and IPTW estimator properly account for time-varying confounding, the popularity of IPTW over G-computation estimator is apparent from its straightforward implementation and its theoretical validity for confidence intervals. Moreover, the treatment probabilities g_0 may arguably be easier to specify correctly than the sequential conditional means $Q_k^{d,t}(P_0)$. However, the IPTW estimator may be generally more susceptible to near positivity violations due to the inverse probability weighting in (11). Often, the kernel functions $h(d,t,V)$ are chosen to stabilize these inverse probability weights. This remedy, however, is less effective when the rules d are dynamic and use time varying covariates as tailoring variables, since the marginal function $h(d,t,V)$ cannot depend on the time varying covariates L_t.

Author Manuscript Author Manuscript Author Manuscript Author Manuscript

3.4 Targeted Maximum Likelihood Estimator

As discussed earlier, consistency of the IPTW estimator relies on consistency of g_n, while consistency of the G-computation estimator relies on consistency of Q_n. In this section, we propose a semiparametric efficient estimator that is robust against misspecification of either Q_0 or g_0. These theoretical promises hinge on the use of one important ingredient — the efficient influence curve for ψ_0.

3.4.1 The Efficient Influence Curve—Central to our methods is viewing the parameter of interest ψ_0 as the value of the map $\Psi : \mathscr{M} \to \mathbb{R}^J$ evaluated at P_0, where $\Psi(P)$ is given by equation (8), with the corresponding functionals at $P \in \mathscr{M}$. It is also straightforward to note from (8) that $\Psi(P) = \Psi(Q)$. From its definition, we see that $\psi = \Psi(P)$ satisfies the characterizing equation

$$\begin{aligned} 0 &= u(\psi, Q, P) \\ &\equiv E_P \sum_{t\in\tau, d\in\mathscr{D}} \tilde{h}(d,t,V)\left(1 - Q_1^{d,t-1}(L_0)\right)\left(\frac{Q_1^{d,t}(L_0) - Q_1^{d,t-1}(L_0)}{1 - Q_1^{d,t-1}(L_0)} - m_\psi(d,t,V)\right). \end{aligned} \tag{14}$$

The mapping Ψ is pathwise differentiable on $\mathscr{M}$, its efficient influence curve (EIC) sheds light on the asymptotic properties of all regular and asymptotically linear estimators of $\Psi(P_0)$. The latter statement is formalized in the following lemma. We refer the reader to Bickel et al. (1997), van der Laan and Robins (2003) van der Vaart and Wellner (1996) for definitions and proofs about properties of efficient influence curves in general. The efficient influence curve for the mapping $P \mapsto \Psi(P)$ in (8) can be derived using the characterizing equation (14) via the functional delta method; this derivation can be found in Petersen et al. (2014a) and Schnitzer et al. (2014), we provide it in the appendix for completeness sake.

Lemma 1 (Efficient influence curve for Ψ): Suppose the mapping Ψ: $\mathscr{M} \to \mathbb{R}^J$ is well-defined at P, in the sense that it is a unique minimizer and hence characterized by the equation $U(\Psi(P), Q, P) = 0$. Then, its efficient influence curve at $P = (Q, g)$, denoted $D^*(Q, g)$, is given by $D^*(\Psi(Q), Q, g)$, where

$$D^*(\psi, Q, g)(O) \equiv M(\psi, Q, P)^{-1} \times \left\{ \sum_{t\in\tau} \sum_{d\in\mathscr{D}} D^{d,t}(\psi, Q, g)(O) + D^{L_0}(\psi, Q)(O) \right\} \tag{15}$$

with

Author Manuscript Author Manuscript Author Manuscript Author Manuscript

$$M(\psi, Q, P) \equiv -\frac{\partial U(\psi', Q, P)}{\partial \psi'}\Big|_{\psi'=\psi}$$

$$= E_P \left\{ \sum_{t\in\tau, d\in\mathscr{D}} h(d,t,V)\phi(d,t,V)\phi(d,t,V)^T \left(1 - Q_1^{d,t-1}(L_0)\right) m_\psi(d,t,V)(1 - m_\psi(d,t,V)) \right\},$$

$$D^{d,t}(\psi, Q, g)(O) \equiv J_\psi(d,t,V) \sum_{k=1}^{t} \frac{I(\mathbf{A}_{k-1}=\mathbf{d}(\mathbf{L}_{k-1}))}{\mathbf{g}(\mathbf{A}_{k-1}=\mathbf{d}(\mathbf{L}_{k-1})|\mathbf{L}_{k-1})} \left(Q_{k-1}^{d,t}(\mathbf{L}_k) - Q_k^{d,t}(\mathbf{L}_{k-1})\right),$$

$$D^{L_0}(\psi, Q)(O) \equiv \sum_{t\in\tau, d\in\mathscr{D}} \tilde{h}(d,t,V) \left(1 - Q_1^{d,t-1}(L_0)\right) \left(\frac{Q_1^{d,t}(L_0) - Q_1^{d,t-1}(L_0)}{1 - Q_1^{d,t-1}(L_0)} - m_\psi(d,t,V) \right).$$

with $J_\psi(d,t,V) = \tilde{h}(d,t,V) - I(t \le K)\tilde{h}(d,t+1,V)(1 - m_\psi(d,t+1,V))$.

The variance $\mathrm{Var}_P D^*(\Psi(Q),Q,g)(O)$ is a generalized Cramér-Rao lower bound for the asymptotic variance of any regular and asymptotically linear estimator of $\Psi(P)$.

Moreover, if either $Q = Q_0$ or $g = g_0$, then $P_0 D^*(\Psi(Q),Q,g) = 0$ implies that $\Psi(Q) = \Psi(Q_0)$.

Proof: See appendix for derivation of (15) and proof of double robustness; see chapters 2 and 3 of Bickel et al. (1997), chapters 1 and 2 of van der Laan and Robins (2003) and chapter 3 of Tsiatis (2006) for the statement regarding variance bounds. □

Now, we are ready to describe the implementation of a TMLE estimator using D^*.

3.4.2 The Loss Function, the Fluctuation Model, and the Algorithm—In a glimpse, our strategy consists in targetedly updating given initial estimators Q_n of Q_0 by minimizing a pre-specified loss along a least favorable (with respect to ψ_0) submodel through Q_n; iterate this updating procedure until the estimating equation $P_n D^*(Q_n^*, g_n) = 0$ is solved at some final targeted estimate Q_n^* of Q_0, and then evaluate Ψ at this Q_n^*.

More specifically, for each $d \in \mathscr{D}$, $t = K + 1, \ldots, 1$ and $k \le t$, the relation in (10) allows us to consider $Q_k^{d,t}$ as a conditional expectation of $Q_{k+1}^{d,t}$, therefore we specify the quasi negative log likelihood loss for $Q_k^{d,t}$, indexed by its expectant $Q_{k+1}^{d,t}$:

$$\mathscr{L}\left(Q_k^{d,t}\right)(O) \equiv -I(\mathbf{A}_{k-1}=\mathbf{d}(\mathbf{L}_{k-1})) \log \left(Q_k^{d,t}(\mathbf{L}_{k-1})^{Q_{k+1}^{d,t}(\mathbf{L}_k)} \left(1 - Q_k^{d,t}(\mathbf{L}_{k-1})\right)^{1 - Q_{k+1}^{d,t}(\mathbf{L}_k)} \right). \quad (16)$$

We suppressed the indexing by $Q_{k+1}^{d,t}$ in the notation. The corresponding least favorable submodel through $Q_k^{d,t}$, parametrized as $\left\{ Q_k^{d,t}(\varepsilon) : \varepsilon \in \mathbb{R}^J \right\}$ (recall that J is the dimension of the parameter ψ_0), is chosen to satisfy the score condition

Author Manuscript Author Manuscript Author Manuscript Author Manuscript

$$\sum_{k=1}^{t} \frac{\partial \mathscr{L}\left(Q_k^{d,t}(\varepsilon)\right)}{\partial \varepsilon}\Big|_{\varepsilon=0} = D^{d,t}(\psi, Q, g).$$

In particular, we can choose

$$Q_k^{d,t}(\varepsilon) = \text{expit}\left(\text{logit}\, Q_k^{d,t} + \varepsilon\, H_k^{d,t}(\psi, g)\right), \quad (17)$$

where $H_k^{d,t}(\psi, g)(\mathbf{L}_{k-1}) \equiv \frac{J_\psi(d,t,V)}{\mathbf{g}(\mathbf{A}_{k-1}=\mathbf{d}(\mathbf{L}_{k-1})|\mathbf{L}_{k-1})}$. Note that the dependency of $Q_k^{d,t}(\varepsilon)$ on ψ and g are suppressed in the notation.

Before describing the algorithm, we make the following observation. Due to the form of the efficient influence curve, the direction of fluctuation $H_k^{d,t}(\psi, g)$ depends on the parameter ψ itself; consequently, the implementation of the TMLE for the hazard MSM conceals more subtleties than its counterpart for the survival MSM in Schnitzer et al. (2014) and Petersen et al. (2014a). For one, in addition to non-targeted initial estimates Q_n and g_n, an initial estimate of ψ_0 is also needed to perform the first update for Q_n. One should choose a consistent initial estimator (albeit not doubly robust), such as the G-computation or the IPTW estimator. Since iterations will be performed, choosing either one of these should have the same asymptotic implications. However, it may make a difference in finite sample performance when Q_n is misspecified, since the fluctuation serves to provide bias reduction over such misspecified Q_n. The second subtlety is that this TMLE estimator requires iterative updates, wherein each iteration uses the previously obtained updated estimate of ψ_0 to steer the direction of fluctuation. The goal of these iterations is to produce targeted estimators Q_n^* of Q_0 that satisfy the efficient score equation $P_n D^*(Q_n^*, g_n) = 0$ in order to achieve proper bias reduction. However, once the residual term $P_n D^*(Q_n^*, g_n)$ becomes smaller than the standard error of the estimator, computational efforts spent to further minimize it will only yield diminishing returns on bias reduction, as can be evidenced by studying the linear expansion of the proposed estimator. Therefore, we use as stopping rule $|P_n D^*(Q_n^*, g_n)|/\widehat{SE}_n < 1/\sqrt{n}$, where $\widehat{SE}$ is estimated using $\sqrt{\widehat{\text{Var}}(D^*(Q_n^*, g_n))/n}$.

With these issues in mind, we are now ready to describe the algorithm.

1. Obtain initial estimators g_n of g_0 and Q_n of Q_0. For the latter, estimate the marginal distribution of L_0 using the empirical distribution, and estimate the conditional means $Q_k^{d,t}(P_0)$ using the regression-based technique of Bang and Robins (2005) described in section 3.2. Obtain initial estimator ψ_n of ψ_0 using either the IPTW estimator $\Psi^{IPTW}(g_n)$ or the G-computation estimator $\Psi(Q_n)$.

Author Manuscript Author Manuscript Author Manuscript Author Manuscript

2. Given initial estimators ψ_n, g_n and $Q_n = \left(Q_n^{L_0}, \left\{Q_{k,n}^{d,t}, t \in \tau, k \leq t, d \in \mathscr{D}\right\}\right)$, we sequentially update the conditional means $Q_{k,n}^{d,t}$ in a two-level algorithm as follows:

a. (a) Starting at $t = K+1$, estimate the sequence $\left\{Q_k^{d,t}(P_0) : k \leq t, d \in \mathscr{D}\right\}$ in order of decreasing k as follows (recall that $Q_{k+1}^{d,t} \equiv Y_t$):

i. At each $k = t, \ldots, 1$: We have an initial estimator $Q_{k,n}^{d,t}$ of $Q_k^{d,t}(P_0)$, and at the previous step, we have obtained an updated estimator $Q_{k+1,n}^{d,t,*}$, of which $Q_{k+1}^{d,t}(P_0)$ which $Q_k^{d,t}(P_0)$ is the expectation. The updated estimator $Q_{k,n}^{d,t,*}$ is given by $Q_{k,n}^{d,t,*} \equiv Q_{k,n}^{d,t}(\varepsilon_k^t)$ where

$$\varepsilon_{k,n}^t \equiv \arg\min P_n \sum_{d \in \mathscr{D}} \frac{\mathscr{L}\left(\text{expit}\left(\text{logit}\, Q_{k,n}^{d,t} + \varepsilon\, J_{\psi_n}(d,t,V)\right)\right)}{\mathbf{g}(\mathbf{A}_{k-1} = \mathbf{d}(\mathbf{L}_{k-1}) | \mathbf{L}_{k-1})}.$$

We can obtain this coefficient by regressing the expectant $Q_{k+1,n}^{d,t,*}(\mathbf{L}_k)$ on $J_{\psi_n}(d,t,V)$ with offset logit $Q_{k,n}^{d,t}(\mathbf{L}_{k-1})$ and weights

$$I\left(\mathbf{A}_{k-1}^1 = \mathbf{d}(\mathbf{L}_{k-1})\right) / \mathbf{g}\left(\mathbf{A}_{k-1} = \mathbf{d}(\mathbf{L}_{k-1}) | \mathbf{L}_{k-1}\right),$$

among observations that remained uncensored by time $k-1$. For observations that had died before time k, the value of the updated estimator is deterministically set to 1.

ii. After iterating step (i) in order of decreasing k, we have obtained updated estimators $\left(Q_{k,n}^{d,t,*} : k \leq t, d \in \mathscr{D}\right)$.

b. Repeat step (a) in order of decreasing t, from $t = K$ to $t = 1$.

3. At the end of the two-level algorithm in step (2) above, we have obtained targeted estimators $Q_n^* = \left(Q_n^{L_0}, \left\{Q_{k,n}^{d,t,*}(\cdot) : t \in \tau, k \leq t, d \in \mathscr{D}\right\}\right)$. In particular, we have $\left\{Q_{1,n}^{d,t,*}(L_0) : t \in \tau, d \in \mathscr{D}\right\}$ for each of the n observations. Pool together these estimates $Q_{1,n}^{d,t,*}(L_0)$ over all d and t, and we have one row per patient i, rule $d \in \mathscr{D}$ and $t \in \tau$. We can then update the parameter estimate using $\psi_n^* \equiv \Psi(Q_n^*)$; this can be implemented by fitting a weighted logistic

Author Manuscript Author Manuscript Author Manuscript Author Manuscript

regression of $\frac{Q_{1,n}^{d,t,*}(L_0) - Q_{1,n}^{d,t-1,*}(L_0)}{1 - Q_{1,n}^{d,t-1,*}(L_0)}$ according to the MSM, with weights $\tilde{h}(d,t,V)(1 - Q_{1,n}^{d,t-1,*}(L_0))$.

4. Repeat step 2 and 3 using ψ_n^* and Q_n^* as initial estimators. Iterate this procedure until the stopping criteria $P_n D^*(Q_n^*, g_n)/\widehat{SE}_n < 1/\sqrt{n}$, where $\widehat{SE}_n = \sqrt{\widehat{\mathrm{Var}}(D^*(Q_n^*, g_n))/n}$, is satisfied.

5. The final updates Q_n^* and ψ_n^* are the targeted maximum likelihood estimators for Q_0 and ψ_0, respectively.

3.4.3 Statistical Inference of TMLE—As we alluded to earlier, given the pre-specified loss function (16), the corresponding least favorable submodel (17) is chosen so that, by design, the TMLE estimator Q_n^* satisfies $|P_n D^*(Q_n^*, g_n)| \approx 0$. From this property stems the doubly robust and locally efficient properties of ψ_n^*.

Specifically, under regularity and empirical process conditions (e.g. van der Laan and Rose (2011)), if both Q_n^* and g_n are consistent estimators, then ψ_n^* is asymptotically linear with influence curve $D^*(Q_0,g_0)$; if g_n converges to g_0, but Q_n^* converges to some Q^* (which may be correctly specified or otherwise), then the influence curve of ψ_n^* equals $D^*(Q_0,g_0)$ minus its projection onto the tangent space of $\mathscr{M}$ for g_0. In either case, $\sqrt{n}(\psi_n^* - \psi_0)$ converges weakly to a normal distribution with each diagonal element of the covariance matrix equal or greater than its counterpart on the covariance matrix of $D^*(\psi_0,Q^*,g_0)$. Consequently, the asymptotic variance of $\sqrt{n}(\psi_n^* - \psi_0)$ can be conservatively estimated using $\sum_n^*$, the sample covariance matrix of $D^*(\psi_n^*, Q_n^*, g_n)$. We can construct asymptotically conservative confidence intervals of level $(1-\alpha)$ as $\left[\psi_n^* \pm \xi_{1-\alpha/2}(\sum_n^*/n)^{1/2}\right]$. The same comments as for IPTW regarding the estimated kernel weights also apply here.

4 Simulation Study

In this section, we evaluate the relative performances of the G-computation estimator, IPTW estimator and TMLE estimator for the parameters of a marginal structural model for the hazard function. For each estimator, we assess the bias, variance, mean squared error (MSE) and coverage estimates for the influence-curve based 95% confidence intervals.

4.1 Data Generating Process

We use a data generating process that resembles the running example, but now including right censoring. A sample consists of n i.i.d. copies of

$$O=(W, CD4_0, A_0^2, A_0^1, Y_1, CD4_1, A_1^2, A_1^1, \ldots, Y_K, CD4_K, A_K^2, A_K^1, Y_{K+1}),$$

Author Manuscript Author Manuscript Author Manuscript Author Manuscript

with $K + 1 = 5$. The baseline variable W encodes sex, the baseline age and disease stage, the time varying covariate $CD4_t$ encodes the most recent CD4 count measurement at time t, and Y_t is the indicator of death by time t. The intervention variables are $A_t=(A_t^1, A_t^2)$, where the treatment variable A_t^1 is the indicator of having switched regimen by time t, and the censoring variable A_t^2 indicates loss to follow up by time t.

We briefly summarize the data generating process here, and defer the details of the data generating distributions are given in the appendix. First, we draw the baseline time independent covariates W. At each time $t \geq 0$, if still alive and uncensored by $t-1$, then first draw mortality status Y_t based on W, past CD4 counts $CD4_{t-1}$ and past switching status A_{t-1}^1 (note at time $t = 0$, $Y_t = 0$ by default). If still alive, then draw the CD4 count measurement for time t based on baseline covariate W, prior CD4 counts $CD4_{t-1}$, and treatment status at previous time point A_{t-1}^1. After drawing CD4, draw censoring by loss to followup A_t^2 based on W and current CD4 counts. If still uncensored and have not switched regimen, draw indicator of whether to switch regiment at time t based on W and current CD4 counts $CD4_t$.

4.2 Target Parameter

An intervention rule of interest d_m switches regimen at m months after confirmed virologic failure and prevents right censoring, given the subject is still alive. The set of regimens of interest are indexed by switching times of interest, $\mathscr{D}=\{d_m : m=0, \ldots, K+1, \infty\}$. Note that each d_m is in fact a static rule assigning $A_t^1=a_t^m$, where $a_t^m=0$ at $t < m$ and $a_t^m=1$ at $t \geq m$. We summarize the hazard function $\lambda(d_m, t) = P(Y_t(d_m) = 1 \mid Y_{t-1}(d_m) = 0)$ as

$$m_\psi(d_m, t)=\text{expit}(\psi_0+\psi_1 t+\psi_2 max(t-m, 0)),$$

and use kernel weights

$$h(d_m, t)=P_0(\mathbf{A}_{t-1}=\mathbf{a}_{t-1}^m) I(t<t^*),$$

where t^* is the first time point where all subjects have either died or censored. Both $P_0(\mathbf{A}_{t-1}=\mathbf{a}_{t-1}^m)$ and t^* are estimated from data, but are treated as given in our parameter definition. This weight mitigates near positivity violation by down-weighting rules with little support in data.

Under the given data-generating distribution, the probabilities of regimen switch are bounded between 0.15 and 0.57. The probabilities of right censoring are bounded between 0.11 and 0.2. The theoretical bounds for the cumulative function **g** for an individual across all regimens of interest are between 0.0069 and 0.22.

4.3 Estimators

We use the regression-based approach in section 3.2 to estimate the initial estimator Q_n of Q_0. A correctly specified $Q_{k,n}^{d,t}$ is obtained by regressing $Q_{k+1,n}^{d,t}$ on W and the entire past

Author Manuscript Author Manuscript Author Manuscript Author Manuscript

covariate history $\mathbf{L}_{k-1} \equiv (CD4_0,\ldots,CD4_{k-1})$ and treatment history $\mathbf{A}_{k-1}^1$. A misspecified $Q_{k,n}^{d,t}$ uses an intercept model. A correctly specified g_n estimates the conditional censoring probabilities and the treatment allocation probabilities by adjusting for W and $CD4_t$. A misspecified g_n uses an intercept model. Both g_n and Q_n are fitted using a logistic regression model.

We implement the Gcomp and the IPTW estimators as described in their respective sections. To implement the TMLE estimator as described in 3.4, we use the IPTW estimator as an initial estimate of ψ_0. In both IPTW and TMLE, the product of g in the denominator is truncated below at 0.01. The iterations in TMLE stops at either $P_n D^*(Q_n^*, g_n)/\widehat{SE}_n < 1/\sqrt{n}$, or at the 17th iteration, whichever comes first. For TMLE and IPTW, the asymptotic variances are estimated by the sample covariance of their respective influence curves, $\sum_n^*$ and $\sum_n^{IPTW}$; consequently, finite sample variance of these estimators can be approximated using the sample variance of their influence curves divided by n. As we mentioned in section 3, there are no influence curve based variance estimators for the Gcomp. But in the spirit of exploration, we implemented a naive version of such variance estimators for the Gcomp by using the variance estimate for the TMLE but using the initial estimators for Q, instead of the updated ones. It is important to remember that the G-comp's IC-based variances estimates are not backed by empirical process theory.

4.4 Results

We assess bias, variance, MSE and coverage probability of the confidence intervals of the estimators over 500 samples. The true values of the target parameters are approximated by generating large datasets where the treatment variables are assigned according to the regimens of interest and the covariates and outcomes are drawn according to the specified data-generating distributions. We display here the results for the TMLE implemented with IPTW as initial estimator. While we had also implemented the TMLE with Gcomp as initial estimator. The results between the two TMLEs are similar up to the displayed digits of significance, so we omit the Gcomp version here for clarity of exposition.

The TMLE implementations in all simulations ended by the stopping rule after one or two iterations, and are deemed convergent for our purposes. In these tables "IC-varEst" indicates the average of the influence-curve based variance estimates, "Coverage" indicates the empirical coverage of the influence-curve based Wald confidence intervals, and "Coverage truevarCI" indicates the empirical coverage of the Wald confidence intervals given by the true sample variance of the estimator.

The simulation results are displayed in table 1 (for sample size $n = 500$) and table 2 (for sample size $n = 2000$). When both Q_0 and g_0 are correctly specified, theoretical results predict that Gcomp, IPTW and TMLE are all consistent. Indeed, the MSE is reduced at the rate of sample size increase for all three estimators. At both sample sizes, the efficiency of TMLE and IPTW are quite similar under this data generating distribution. When one of the nuisance parameters are misspecified, lemma 1 predicts that the double robustness properties of TMLE should provide bias reduction. Indeed, when g_n is misspecified and Q_n is correctly

Author Manuscript Author Manuscript Author Manuscript Author Manuscript

specified, TMLE provides significant bias reduction over the misspecified IPTW, with only slight increase in variance, hence leading to overall smaller MSE. The benefit of this bias-variance trade-off is more apparent across sample sizes, since the overall MSE of the TMLE estimators will decrease at the speed of sample size growth, whereas those of IPTW will stabilize. When Q_n is misspecified and g_n is correct, TMLE also provides substantial bias reduction over the misspecified Gcomp. Though this TMLE may have a larger variance than the biased Gcomp, the bias reduction is significant enough to yield smaller MSE for TMLE and ensure MSE converges at the rate of sample size increase.

In addition to parameter estimation, we are also interested in providing variance estimate and confidence intervals for our estimators. When both Q_0 and g_0 are correctly specified, both IPTW and TMLE are asymptotically normal, and hence one should achieve the desired coverage using Wald confidence intervals based on influence curve-based variance estimates ('IC-varEst'). Firstly we see that IC-varEst tends to underestimate the true variance of the estimators, but the approximation gets better with sample size. At sample size 500, the Wald 95% confidence intervals of TMLE have empirical coverage that is below the nominal coverage. Is this low coverage attributed to the IC-based variance estimates or to an underlying problem with achieving asymptotic normality? To answer this, we construct a confidence intervals using the true variance of the estimators (as estimated by the sample variance across the 500 simulations). The coverage of these confidence intervals is reported under "Coverage: true varCI". We see that this second set of confidence intervals have coverage close to the 95% confidence level, given the moderate sample size. This suggest that the suboptimal coverage of the IC-based confidence intervals is attributed to the performance of the IC-based variance estimator, not to the asymptotic normality. Indeed, at sample size 2000, the coverage of the IC-based confidence intervals of TMLE approximate the nominal coverage. This underscores the need for variance estimators with improved finite sample performance, possibly exploiting high-order derivatives of the parameter estimators, but this is beyond the scope of this paper. We remind the readers that we also report here a IC-based variance estimate and corresponding confidence intervals for the Gcomp, only for the sake of comparison, as the performance of these are not backed up by theoretical results.

5 Summary

In this paper, we have presented a doubly robust and efficient substitution estimator that directly targets a longitudinal dynamic (or static) marginal structural model for the hazard function. This work builds upon the pooled TMLE methodology in Petersen et al. (2014a) for the survival function, as well as earlier work by Robins (2000), Robins (2002) and Bang and Robins (2005).

Unlike TMLE for MSM for survival functions or means, TMLE for MSM for a hazard function is a bona fide iterated estimator and requires an initial estimator of the parameter of interest. There is no theoretical predictions on which initial estimator is more advantageous. The simulations have also offer little evidence in favor of either. This seems to suggests that the choice of initial estimator has little effect, at least in moderate and large sample sizes. In terms of computing costs, like the TMLE for survival MSM, the main computational

Author Manuscript Author Manuscript Author Manuscript Author Manuscript

resources for TMLE for hazard MSM are expended to obtain the initial estimators for $Q_k^{d,t}(P_0)$; these costs are heavily dependent on sample size and the algorithms used to fit these estimators, but only have to be performed once. Beyond those, each iteration of TMLE updates fit $O(K^2)$ logistical models of dimension J. TMLE for survival MSM only performs one iteration of such updates, whereas TMLE for hazard MSM may perform more than one (though in our simulation two iterations was enough). Further simulations are needed to assess the actual added computational burden in situations with big datasets, large number of outcome times, or that require higher number of iterations.

The proposed TMLE estimator offers theoretical advantages over the popular IPTW estimator and a non-targeted substitution Gcomp estimator: 1) it offers protection against model misspecifications, 2) it is locally efficient, 3) it may also provide protection against unstable probability weights. The bias reduction over a misspecified IPTW or Gcomp estimator is clear in the simulation studies even for a moderate sample size. We also see that the influence curve based variance estimate of IPTW and TMLE tend to be conservative estimates (underestimate) of the true variance, but this variance estimate can be improved with increased sample size. As mentioned in section 3.4, under certain empirical process conditions, these influence-curve based estimators provide conservative variances estimates when only one of the nuisance parameters are correctly specified. Such conditions depend on the rates of convergence for the nuisance parameter estimates and the complexity of the correctly specified model. In the case of IPTW, the consistencies of both the parameter estimate and the variance estimate hinge on correct specification of the treatment mechanisms. Influence curve based estimates provide a viable alternative or complement to bootstrapping when the computing burden of the latter is high. But their empirical process conditions also suggest that a more sophisticated theory-based variance estimates, as well as appropriate diagnostics, are needed for the TMLE and the IPTW estimator. Future research priorities should focus on variance estimation and inference methods that remain resilient in the face of moderate confounding and multiple time points.

Acknowledgments

We thank Josh Schwab for his help in the software implementation. This work was supported by R01 A1074345-07 (P.I. M. van der Laan).

6 Paper Appendix

6.1 Derivation of Efficient Influence Curve

In order to apply the functional delta method, we first rewrite the characterizing equation (14) as

Author Manuscript Author Manuscript Author Manuscript Author Manuscript

$$U(\psi,Q,P)=E_P\sum_{t\in\tau,d\in\mathscr{D}}\tilde{h}(d,t,V)\left(1-Q_1^{d,t-1}(L_0)\right)\left(\frac{Q_1^{d,t}(L_0)-Q_1^{d,t-1}(L_0)}{1-Q_1^{d,t-1}(L_0)}-m_\psi(d,t,V)\right)$$
$$=E_P\sum_{t\in\tau,d\in\mathscr{D}}\tilde{h}(d,t,V)\left(Q_1^{d,t}(L_0)-Q_1^{d,t-1}(L_0)-m_\psi(d,t,V)+m_\psi(d,t,V)Q_1^{d,t-1}(L_0)\right)$$
$$=E_P\sum_{t\in\tau,d\in\mathscr{D}}J_\psi(d,t,V)Q_1^{d,t}(L_0)-\tilde{h}(d,t,V)m_\psi(d,t,V).$$

(18)

Since 0 = $U(\Psi(P),Q,P)$, it follows from implicit differentiation that

$$\frac{d\Psi(P)}{dP}=-\frac{dU(\Psi(P),Q,P)^{-1}}{d\Psi(P)}\frac{dU(\psi,Q,P)}{dP}.$$

Straightforward computations shows that indeed

$$=E_P\left\{\sum_{d\in\mathscr{D}}\sum_{t=1}^{\tau}h(d,t,V)\left(1-Q_1^{d,t-1}(L_0)\right)m_\psi(d,t,V)(1-m_\psi(d,t,V))\phi(d,t,V)\phi(d,t,V)^T\right\}^{-\frac{dU(\psi,Q,P)}{d\psi}}.$$

This proves the expression for $M(\psi,Q,P)$ in (15).

Next, note that $\frac{dU(\psi,Q,P)}{dP}=\frac{dU(\psi,Q,P)}{dQ^{L_0}}\frac{dQ^{L_0}}{dp}+\sum_{t\in\tau}\sum_{d\in\mathscr{D}}\frac{dU(\psi,Q,P)}{dQ_1^{d,t}}\frac{dQ_1^{d,t}}{dP}$. can be obtained by a simple application of the functional delta method method, using the efficient influence functions for Q^{L_0} and $Q_1^{d,t}$ obtained in Petersen et al. (2014a). Therefore, we have

$$D^{L_0}(\psi,Q)(O)\equiv\sum_{t\in\tau}\sum_{d\in\mathscr{D}}\tilde{h}(d,t,V)\left(1-Q_1^{d,t-1}(L_0)\right)\left(\frac{Q_1^{d,t}(L_0)-Q_1^{d,t-1}(L_0)}{1-Q_1^{d,t-1}(L_0)}-m_\psi(d,t,V)\right),$$
$$D^{d,t}(\psi,Q,g)(O)\equiv J_\psi(d,t,V)\sum_{k=1}^{t}\frac{I(\mathbf{A}_{k-1}=\mathbf{d}(\mathbf{L}_{k-1}))}{\mathbf{g}_k(\mathbf{A}_{k-1}=\mathbf{d}(\mathbf{L}_{k-1})|\mathbf{L}_{k-1})}\left(Q_{k+1}^{d,t}(\mathbf{L}_k)-Q_k^{d,t}(\mathbf{L}_{k-1})\right).$$

Now, we show the robustness property. If $Q = Q_0$, the result is trivial by definition of $\Psi(Q)$ and Q_k^t. We only need to check the second case. When $g = g_0$, at each t, the sum from $k = 1$ to $k = t$ forms a telescopic sum, leaving only the first and the last term. Using expression in (18) to rewrite $D^{L0}(\Psi(Q),Q)$, we have, up to a constant normalizing matrix,

Author Manuscript Author Manuscript Author Manuscript Author Manuscript

$$
\begin{aligned}
0 = P_0 D^*(\Psi(Q), Q, g_0) &= P_0 \sum_{t=1}^{\tau} \sum_{d \in \mathscr{D}} J_{\Psi(Q)}(d,t,V) Q_{1,0}^{d,t}(L_0) - P_0 \sum_{t=1}^{\tau} \sum_{d \in \mathscr{D}} J_{\Psi(Q)}(d,t,V) Q_1^{d,t}(L_0) \\
&\quad + P_0 \sum_{t=1}^{\tau} \sum_{d \in \mathscr{D}} J_{\Psi(Q)}(d,t,V) Q_1^{d,t}(L_0) - \tilde{h}(d,t,V) m_{\Psi(Q)}(d,t,V) \\
&= P_0 \sum_{t=1}^{\tau} \sum_{d \in \mathscr{D}} J_{\Psi(Q)}(d,t,V) Q_{1,0}^{d,t}(L_0) - \tilde{h}(d,t,V) m_{\Psi(Q)}(d,t,V) \\
&= U(\Psi(Q), Q_0, P_0).
\end{aligned}
$$

Therefore, $\Psi(Q) = \Psi(Q_0)$.

6.2 Data Generating Distribution for the Simulation Study

The baseline covariate W consists of $W = (W_1, W_2, W_3, W_4)$, where $W_1 = I(30 \le age \le 39)$, $W_2 = I(age > 39)$, W_3 indicates sex, and W_4 = disease stage.

Let $\varepsilon_{1,t}$ be errors drawn from a standard normal distribution. The data generating distributions are given as follows:

$$
\begin{gathered}
W_1 \sim Bern(0.3),\ (W_2 | W_1 = 0) \sim Bern(0.5),\ W_3 \sim Bern(0.5),\ W_4 \sim Bern(0.3); \\
P(Y_t = 1 | W, CD4_{t-1}, A_{t-1}^1, Y_{t-1} = 0) = \\
\begin{cases} 0 \text{ if } t = 0; \\ \text{expit}(-2 - 0.5W_1 - 0.1W_2 + 0.2W_3 + 0.3W_4 - 0.7CD4_{t-1} - 0.9A_{t-1}^1), \text{else}. \end{cases} \\
CD4_t = \max(\min(\mu(W, CD4_{t-1}, A_{t-1}^1) + \varepsilon_{1,t}, 2), -2), \text{where } \mu(W, CD4_{t-1}, A_{t-1}^1) = \\
\begin{cases} -W_4 \text{ if } t = 0 \\ 0.1W_1 - 0.5W_2 - 0.1W_3 - 0.2W_4 + 0.4CD4_{t-1} + 0.5A_{t-1}^1 - 0.1CD4_{t-1}A_{t-1}^1, \text{else}; \end{cases} \\
P(A_t^2 = 1 | W, CD4_t, A_{t-2}^2 = 0, Y_t = 0) = \\
1 - \text{expit}(1.5 + 0.1W_1 + 0.2W_2 + 0.1W_3 + 0.1W_4 + 0.1CD4_t) \\
P(A_t^1 = 1 | A_{t-1}^1, W, CD4_t, A_t^2 = 0, Y_t = 0) = \\
\begin{cases} 1, \ \text{ if } t \ge 1 \text{ and } A_{t-1}^1 = 1; \\ \text{expit}\,(-0.5 - 0.1W_1 - 0.2W_2 - 0.2W_3 + 0.1W_4 - 0.4CD4_t), \text{else}. \end{cases}
\end{gathered}
$$

References

Bang H, Robins JM. Doubly robust estimation in missing data and causal inference models. Biometrics. 2005; 61:962–972. [PubMed: 16401269]

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

Cain L, Robins JM, Lanoy E, Logan R, Costagliola D, Hernan M. When to start treatment? a systematic approach to the comparison of dynamic regimes using observational data. International Journal of Biostatistics. 2010; 6 Article 18.

Cole S, Hudgens M, Tien P, Anastos K, Kingsley L, Chmiel J, Jacobson L. Marginal structural models for case-cohort study designs to estimate the association of antiretroviral therapy initiation with incident aids or death. Am J Epidemiol. 2012; 175:381–90. [PubMed: 22302074]

DHHS. Guidelines for the use of antiretroviral agents in hiv-1-infected adults and adolescents. department of health and human services. 2014. URL http://aidsinfo.nih.gov/contentfiles/lvguidelines/adultandadolescentgl.pdf

Gruber S, van der Laan M. A targeted maximum likelihood estimator of a causal effect on a bounded continuous outcome. International Journal of Biostatistics. 2010; 6 Article 26.

Author Manuscript Author Manuscript Author Manuscript Author Manuscript

Hernan M, Brumback B, Robins J. Marginal structural models to estimate the causal effect of zidovudine on the survival of HIV-positive men. Epidemiology. 2000; 11:561–570. [PubMed: 10955409]

HIV-CAUSAL-Collaboration. Cain L, Logan R, Robins J, Sterne J, Sabin C, Bansi L, Justice A, Goulet J, van Sighem A, de Wolf F, Bucher H, von Wyl V, Esteve A, Casabona J, del Amon J, Moreno S, Seng R, Meyer L, Perez-Hoyos S, Muga R, Lodi S, Lanoy E, Costagliola D, Hernan M. When to initiate combined antiretroviral therapy to reduce mortality and aids-defining illness in hiv-infected persons in developed countries: an observational study. Ann Intern Med. 2011; 154:509–15. [PubMed: 21502648]

Neugebauer R, van der Laan M. Nonparametric causal effects based on marginal structural models. Journal of Statistical Planning and Inference. 2007; 137:419–434.

Pearl, J. Causality: Models, Reasoning and Inference. 2nd. New York: Cambridge University Press; 2009.

Petersen M, Schwab J, Gruber S, Blaser N, Schomaker M, van der Laan M. Targeted maximum likelihood estimation for dynamic and static longitudinal marginal structural working models. J Causal Inference. 2014a; 2:147–185. [PubMed: 25909047]

Petersen M, Tran L, Geng E, Reynolds S, Kambugu A, Wood R, Bangsberg D, Yiannoutsos C, Deeks SG, Martin JN. Delayed switch of antiretroviral therapy after virologic failure associated with elevated mortality among hiv-infected adults in africa. AIDS. 2014b; 28:2097–107. [PubMed: 24977440]

Petersen M, van der Laan M, Napravnik S, Eron J, Moore R, Deeks SG. Long term consequences of the delay between virologic failure of highly active antiretroviral therapy and regimen modification. AIDS. 2008; 22:2097–2106. [PubMed: 18832873]

Robins J. A new approach to causal inference in mortality studies with sustained exposure periods - application to control of the healthy worker survivor effect. Mathematical Modelling. 1986; 7:1393–1512.

Robins, J. Statistical models in epidemiology: the environment and clinical trials. Springer-Verlag; 1999. Marginal structural models versus structural nested models as tools for causal inference; p. 95-134.

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

Robins J. Commentary on “using inverse weighting and predictive inference to estimate the effects of time-varying treatments on the discrete-time hazard. Stat Med. 2002; 210:1663–80.

Robins, J.; Rotnitzky, A. Recovery of information and adjustment for dependent censoring using surrogate markers. In: Jewell, N.; Dietz, K.; Farewell, V., editors. AIDS Epidemiology. Methodological issues, Bikhäuser; 1992. p. 297-331.

Robins J, Rotnitzky A, van der Laan M. Comment on “on profile Likelihood” by S.A. Murphy and A.W. van der Vaart. Journal of the American Statistical Association – Theory and Methods. 2000a; 450:431–435.

Robins JM. Marginal structural models. Proceedings of the Amer-ican Statistical Association Section on Bayesian Statistical Science. 1997:1–10.

Robins JM, Hernan MA, Brumback B. Marginal structural models and causal inference in epidemiology. Epidemiology. 2000b; 11:550–560. [PubMed: 10955408]

Robins JM, Orellana L, Rotnitzky A. Estimaton and extrapolation of optimal treatment and testing strategies. Statistics in Medicine. 2008; 27(23):4678–4721. [PubMed: 18646286]

Schnitzer ME, Moodie E, van der Laan MJ, Platt R, Klein MB. Modeling the impact of hepatitis c viral clearance on end-stage liver disease in an hiv co-infected cohort with targeted maximum likelihood estimation. Biometrics. 2014; 70:144–152. [PubMed: 24571372]

Taubman S, Robins J, Mittleman M, Hernan MA. Intervening on risk factors for coronary heart disease: an application of the parametric g-formula. Int J Epidemiol. 2009; 380:1599–611. [PubMed: 19389875]

Tsiatis, A. Semiparametric Theory and Missing Data. New York: Springer; 2006.

van der Laan M, Petersen M. Causal effect models for realistic individualized treatment and intention to treat rules. International Journal of Biostatistics. 2007; 3 Article 3.

Author Manuscript Author Manuscript Author Manuscript Author Manuscript

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

van der Laan, M.; Rose, S. Targeted Learning: Causal Inference for Observational and Experimental Data. first. Springer; 2011. Springer Series in Statistics

van der Laan M, Rubin D. Targeted maximum likelihood learning. The International Journal of Biostatistics. 2006; 2:1–40.

van der Laan MJ, Gruber S. Targeted minimum loss based estimation of causal effects of multiple time point interventions. The International Journal of Biostatistics. 2012; 8 Article 8.

van der Vaart, A.; Wellner, J. Weak Convergence and Empirical Processes. Springer-Verlag; New York: 1996.

Young JG, Cain LE, Robins JM, OReilly EJ, Hernán MA. Comparative effectiveness of dynamic treatment regimes: an application of the parametric g-formula. Statistics in biosciences. 2011; 3:119–143. [PubMed: 24039638]

Author Manuscript Author Manuscript Author Manuscript Author Manuscript

Table 1

Simulation Results for 500 simulations with $n = 500$ and $K + 1 = 5$. "IC-varEst": average of the influence curve based variance estimates across 500 simulations; "Covarage": coverage of the Wald 95% confidence interval using IC-based variance estimate. "true var. CI": coverage of the Wald confidence interval using true variance (approximated through the 500 simulations).

	Q_n and g_n correct			correct Q_n, misspec. g_n		misspec Q_n, correct. g_n	
	Gcomp	**IPTW**	**TMLE**	**IPTW**	**TMLE**	**Gcomp**	**TMLE**
Bias							
$\hat{\psi}_0$	−1.05e−02	−7.30e−03	−1.02e−02	−6.43e−02	−4.96e−02	−9.31e−02	6.93e−03
$\hat{\psi}_1$	−1.35e−02	−4.39e−03	−6.35e−04	−2.17e−01	2.89e−02	−2.55e−02	−1.90e−02
$\hat{\psi}_2$	9.19e−03	−4.52e−03	−9.63e−03	4.16e−01	−2.55e−02	7.93e−02	1.33e−03
Variance							
$\hat{\psi}_0$	4.57e−02	5.43e−02	5.76e−02	3.66e−02	6.33e−02	4.65e−02	4.79e−02
$\hat{\psi}_1$	1.08e−02	1.57e−02	1.75e−02	6.29e−03	2.12e−02	1.19e−02	1.21e−02
$\hat{\psi}_2$	1.02e−02	1.45e−02	1.46e−02	2.75e−04	1.69e−02	1.08e−02	1.16e−02
MSE							
$\hat{\psi}_0$	4.58e−02	5.43e−02	5.76e−02	4.06e−02	6.56e−02	5.50e−02	4.79e−02
$\hat{\psi}_1$	1.09e−02	1.57e−02	1.74e−02	5.33e−02	2.20e−02	1.25e−02	1.25e−02
$\hat{\psi}_2$	1.03e−02	1.45e−02	1.47e−02	1.73e−01	1.75e−02	1.71e−02	1.16e−02
IC-varEst							
$\hat{\psi}_0$	4.98e−02	5.20e−02	4.75e−02	7.90e−02	4.99e−02	4.72e−02	4.09e−02
$\hat{\psi}_1$	1.25e−02	1.34e−02	1.13e−02	2.56e−02	1.15e−02	1.14e−02	9.86e−03

Author Manuscript Author Manuscript Author Manuscript Author Manuscript

Int J Biostat. Author manuscript; available in PMC 2016 August 01.

Table 2

Simulation Results for 500 simulations with $n = 2000$ and $K + 1 = 5$. "IC-varEst": average of the influence curve based variance estimates across 500 simulations; "Covarage": coverage of the Wald 95% confidence interval using IC-based variance estimate. "true var. CI": coverage of the Wald confidence interval using true variance (approximated through the 500 simulations)

	Q_n and g_n correct			correct Q_n, misspec. g_n		misspec Q_n, correct. g_n	
	Gcomp	**IPTW**	**TMLE**	**IPTW**	**TMLE**	**Gcomp**	**TMLE**
Bias							
$\hat{\psi}_0$	1.92e−03	−1.44e−02	−1.36e−02	−7.49e−02	−1.88e−02	−1.00e−01	4.25e−03
$\hat{\psi}_1$	−1.06e−02	5.81e−03	5.46e−03	−2.10e−01	9.12e−03	−1.79e−02	−1.19e−02
$\hat{\psi}_2$	2.04e−03	−4.74e−03	−4.47e−03	4.17e−01	−6.12e−03	8.07e−02	1.83e−03
Variance							
$\hat{\psi}_0$	1.14e−02	1.28e−02	1.26e−02	9.18e−03	1.32e−02	1.19e−02	1.20e−02
$\hat{\psi}_1$	2.52e−03	3.30e−03	3.25e−03	1.56e−03	3.45e−03	2.78e−03	2.78e−03
$\hat{\psi}_2$	2.64e−03	3.44e−03	3.34e−03	7.09e−05	3.44e−03	2.76e−03	2.83e−03
MSE							
$\hat{\psi}_0$	1.14e−02	1.30e−02	1.28e−02	1.48e−02	1.36e−02	2.19e−02	1.19e−02
$\hat{\psi}_1$	2.63e−03	3.32e−03	3.28e−03	4.55e−02	3.53e−03	3.10e−03	2.92e−03
$\hat{\psi}_2$	2.63e−03	3.45e−03	3.35e−03	1.74e−01	3.47e−03	9.26e−03	2.83e−03
IC-varEst							
$\hat{\psi}_0$	1.25e−02	1.29e−02	1.22e−02	1.97e−02	1.25e−02	1.17e−02	1.03e−02
$\hat{\psi}_1$	3.19e−03	3.33e−03	3.04e−03	6.28e−03	3.05e−03	2.77e−03	2.50e−03