

Maya Petersen*, Joshua Schwab, Susan Gruber, Nello Blaser, Michael Schomaker and Mark van der Laan

Targeted Maximum Likelihood Estimation for Dynamic and Static Longitudinal Marginal Structural Working Models

Abstract: This paper describes a targeted maximum likelihood estimator (TMLE) for the parameters of longitudinal static and dynamic marginal structural models. We consider a longitudinal data structure consisting of baseline covariates, time-dependent intervention nodes, intermediate time-dependent covariates, and a possibly time-dependent outcome. The intervention nodes at each time point can include a binary treatment as well as a right-censoring indicator. Given a class of dynamic or static interventions, a marginal structural model is used to model the mean of the intervention-specific counterfactual outcome as a function of the intervention, time point, and possibly a subset of baseline covariates. Because the true shape of this function is rarely known, the marginal structural model is used as a working model. The causal quantity of interest is defined as the projection of the true function onto this working model. Iterated conditional expectation double robust estimators for marginal structural model parameters were previously proposed by Robins (2000, 2002) and Bang and Robins (2005). Here we build on this work and present a pooled TMLE for the parameters of marginal structural working models. We compare this pooled estimator to a stratified TMLE (Schnitzer et al. 2014) that is based on estimating the intervention-specific mean separately for each intervention of interest. The performance of the pooled TMLE is compared to the performance of the stratified TMLE and the performance of inverse probability weighted (IPW) estimators using simulations. Concepts are illustrated using an example in which the aim is to estimate the causal effect of delayed switch following immunological failure of first line antiretroviral therapy among HIV-infected patients. Data from the International Epidemiological Databases to Evaluate AIDS, Southern Africa are analyzed to investigate this question using both TML and IPW estimators. Our results demonstrate practical advantages of the pooled TMLE over an IPW estimator for working marginal structural models for survival, as well as cases in which the pooled TMLE is superior to its stratified counterpart.

Keywords: dynamic regime, semiparametric statistical model, targeted minimum loss based estimation, confounding, right censoring

DOI 10.1515/jci-2013-0007

*Corresponding author: **Maya Petersen**, Division of Biostatistics, University of California, Berkeley, CA, USA, E-mail: mayaliv@berkeley.edu

Joshua Schwab, Division of Biostatistics, University of California, Berkeley, CA, USA, E-mail: joshuaschwab@yahoo.com

Susan Gruber, Department of Epidemiology, Harvard School of Public Health, Boston, MA, USA, E-mail: sgruber@hsph.harvard.edu

Nello Blaser, Institute of Social and Preventive Medicine (ISPM), University of Bern, Bern, Switzerland, E-mail: nblaser@ispm.unibe.ch

Michael Schomaker, Centre for Infectious Disease Epidemiology & Research, University of Cape Town, Cape Town, South Africa, E-mail: michael.schomaker@uct.ac.za

Mark van der Laan, Division of Biostatistics, University of California, Berkeley, CA, USA, E-mail: laan@stat.berkeley.edu

1 Introduction

Many studies aim to learn about the causal effects of longitudinal exposures or interventions using data in which these exposures are not randomly assigned. Specifically, consider a study in which baseline covariates, time-varying exposures or treatments, time-varying covariates, and an outcome of interest, such as death, are observed on a sample of subjects followed over time. The exposures of interest can both depend on past covariates and affect future covariates, as well as the outcome. Censoring may also occur, possibly in response to past treatment and covariates. Such data structures are ubiquitous in observational cohort studies. For example, a sample of HIV-infected patients might be followed longitudinally in clinic and data collected on antiretroviral prescriptions, determinants of prescription decisions including CD4+ T cell counts and plasma HIV RNA levels (viral loads), and vital status. Such data structures also occur in randomized trials when the exposure of interest is (non-random) compliance with a randomized exposure or includes non-randomized mediators of an exposure's effect.

The causal effects of longitudinal exposures (as well as the effects of single time point exposures when the outcome is subject to censoring) can be formally defined by contrasting the distribution of a counterfactual outcome under different “interventions” to set the values of the exposure and censoring variables. For example, the counterfactual survival curve of HIV-infected subjects following immunological failure of antiretroviral therapy might be contrasted under a hypothetical intervention in which all subjects were switched immediately to a new antiretroviral regimen versus an intervention in which all subjects remained on their failing therapy [1]. In the presence of censoring due to losses to follow up, these counterfactuals of interest might be defined under a further intervention to prevent censoring. Interventions such as these, under which all subjects in a population are deterministically assigned the same vector of exposure and censoring decisions (for example, do not switch and remain under follow up) are referred to as “static regimes.”

More generally, counterfactuals can be defined under interventions that assign a treatment or exposure level to each subject at each time point based on that subject's observed past. For example, counterfactual survival might be compared under interventions to switch all patients to second line antiretroviral therapy the first time their CD4+T cell count crosses a certain threshold, for some specified set of thresholds [2]. Such subject-responsive treatment strategies have been referred to as individualized treatment rules, adaptive treatment strategies, or “dynamic regimes” (see, for example, Robins [3]; Murphy et al. [4]; Hernan et al. [5]). Additional examples include strategies for deciding when to start antiretroviral therapy [6, 7] and strategies for modifying dose or drug choice based on prior response and adverse effects. Investigation of the effects of such dynamic regimes makes it possible to learn effective strategies for assigning an intervention based on a subject's past and is thus relevant to any discipline that seeks to learn how best to use past information to make decisions that will optimize future outcomes.

The static and dynamic regimes described above are longitudinal – they involve interventions to set the value of multiple treatment and censoring variables over time. For example, counterfactual survival under no switch to second line therapy corresponds to a subject's survival under an intervention to prevent a patient from switching at each time point from immunologic failure until death or the end of the study. A time-dependent causal dose–response curve, which plots the mean of the intervention-specific counterfactual outcome at time t as a function of the interventions through time t , can be used to summarize the effects of these longitudinal interventions. For example, a plot of the counterfactual survival probability as a function of time since immunologic failure, for a range of alternative CD4+ T cell count thresholds used to initiate a switch captures the effect of alternative switching strategies on survival.

Formal causal frameworks provide a tool to establish the conditions under which such causal dose–response curves can be identified from the observed data. Longitudinal static and dynamic regimes are often subject to time-dependent confounding – time-varying variables may confound the effect of future treatments while being affected by past treatment [8]. Traditional approaches to the identification of point treatment effects, which are based on selection of a single set of covariates for regression or stratification-based

adjustment, break down when such time-dependent confounding is present. However, the mean counterfactual outcome under a longitudinal static or dynamic regime may still be identified under the appropriate sequential randomization and positivity assumptions (reviewed in Robins and Hernan [9]).

Under these assumptions, causal dose–response curves can be estimated by generating separate estimates of the mean counterfactual outcome for each time point and intervention (or regime) of interest. For example, one could generate separate estimates of the counterfactual survival curve for each CD4-based threshold for switching to second line therapy. In this manner, one obtains fits of the time-dependent causal dose–response curve for each of a range of possible thresholds, which together summarize how the mean counterfactual outcome at time t depends on the choice of threshold.

A number of estimators can be used to estimate intervention-specific mean counterfactual outcomes. These include inverse probability weighted (IPW) estimators (for example, [3, 5, 10]), “G-computation” estimators (typically based on parametric maximum likelihood estimation of the non-intervention components of the data generating process) (for example, [7, 11, 12]), augmented-IPW estimators (for example, [13–16, 31]), and targeted maximum likelihood (or minimum loss) estimators (TMLEs) (for example, [17, 18]). In particular, van der Laan and Gruber [19] combine the targeted maximum likelihood framework [20, 21] with important insights and the iterated conditional expectation estimators established in Robins [3, 29] and Bang and Robins [22].

Both the theoretical validity and the practical utility of these estimators rely, however, on reasonable support for each of the interventions of interest, both in the true data generating distribution and in the sample available for analysis. For example, in order to estimate how survival is affected by the threshold CD4 count used to initiate an antiretroviral treatment switch, a reasonable number of subjects must in fact switch at the time indicated by each threshold of interest. Without such support, estimators of the intervention-specific outcome will be ill-defined or extremely variable. Although one might respond to this challenge by creating coarsened versions of the desired regimes, so that sufficient subjects follow each coarsened version, such a method introduces bias and leaves open the question of how to choose an optimal degree of coarsening.

Since adequate support for every intervention of interest is often not available, Robins [23] introduced marginal structural models (MSMs) that pose parametric or small semiparametric models for the counterfactual conditional mean outcome as a function of the choice of intervention and time. For example, static MSMs have been used to summarize how the counterfactual hazard of death varies as a function of when antiretroviral therapy is initiated [24] and when an antiretroviral regimen is switched [25]. The extrapolation assumptions implicitly defined by non-saturated MSMs make it possible to estimate the coefficients of the model, and thereby the causal dose–response curve, even when few or no subjects follow some interventions of interest.

While MSMs were originally developed for static interventions [8, 10, 23, 24] they naturally generalize to classes of dynamic (or even more generally, stochastic) interventions as shown in van der Laan and Petersen [2] and Robins et al. [26]. Dynamic MSMs have been used, for example, to investigate how counterfactual hazard of death varies as a function of CD4+ T cell count threshold used to initiate antiretroviral therapy [6] or to switch antiretroviral therapy regimens [2]. Because the true shape of the causal dose–response curve is typically unknown, we have suggested that MSMs be used as working models. The target causal coefficients can then be defined by projecting the true causal dose–response curve onto this working model [20, 27].

The coefficients of both static and dynamic MSMs are frequently estimated using IPW estimators [2, 8, 10, 26]. These estimators have a number of attractive qualities: they can be intuitively understood, they are easy to implement, and they provide an influence curve-based approach to standard error estimation. However, IPW estimators also have substantial shortcomings. In particular, they are biased if the treatment mechanism used to construct the weights is estimated poorly (for example, using a misspecified parametric model). Further, IPW estimators are unstable in settings of strong confounding (near or partial positivity violations) and the resulting bias in both point and standard error estimates can result in poor inference (for a review of this issue see Petersen et al. [28]). Dynamic MSMs can exacerbate this problem, as the options for effective weight stabilization are limited [6, 26].

Asymptotically efficient and double robust augmented-IPW estimators of the estimand corresponding to longitudinal static MSM parameters were developed by Robins and Rotnitzky [14], Robins [13], Robins et al. [16]. These estimators are defined as a solution of an estimating equation, and as a result may be unstable due to failure to respect the global constraints implied by the model and the parameter. Robins [13, 29] and Bang and Robins [22] introduced an alternative double robust estimating equation-based estimator of longitudinal MSM parameters based on the key insight that both the statistical target parameter and the corresponding augmented-IPW estimating function (efficient influence curve) for MSMs on the intervention-specific mean can be represented as a series of iterated conditional expectations. In addition, they proposed a targeted sequential regression method to estimate the nuisance parameters of the augmented-IPW estimating equation. This innovative idea allowed construction of a double robust estimator that relies only on estimation of minimal nuisance parameters beyond the treatment mechanism.

In this paper, we describe a double robust substitution estimator of the parameters of a longitudinal marginal structural working model. The estimator presented incorporates the key insights and prior estimator of Robins [13, 29] and Bang and Robins [22] into the TMLE framework. Specifically, we expand on this prior work in several ways. We propose a TMLE for marginal structural working models for longitudinal dynamic regimes, possibly conditional on pre-treatment covariates. The TMLE described is defined as a substitution estimator rather than as solution to an estimating equation and incorporates data-adaptive/machine learning methods in generating initial fits of the sequential regressions. Finally, we further generalize the TMLE to apply to a larger class of parameters defined as arbitrary functions of intervention-specific means across a user-supplied class of interventions.

TMLE for the parameters of a MSM for “point treatment” problems, in which adjustment for a single set of covariates known not to be affected by the intervention of interest is sufficient to control for confounding, including history-adjusted MSMs, have been previously described [30, 31]. However, the parameter of a longitudinal MSM on the intervention-specific mean under sequential interventions subject to time-dependent confounding is identified as a distinct, and substantially more complex, estimand than the estimand corresponding to a point treatment MSM, and thus requires distinct estimators. An alternative TMLE for longitudinal static MSMs, which we refer to as a stratified TMLE, was described by Schnitzer et al. [32]. The stratified TMLE uses the longitudinal TMLE of van der Laan and Gruber [19] for the intervention-specific mean to estimate each of a set of static treatments and combines these estimates into a fit of the coefficients of a static longitudinal MSM on both survival and hazard functions. The stratified TMLE [32] resulted in substantially lower standard error estimates than an IPW estimator in an applied data analysis and naturally generalizes to dynamic MSMs. However, it remains vulnerable when there is insufficient support for some interventions of interest. In contrast, the TMLE we describe here pools over the set of dynamic or static interventions of interest as well as optionally over time when updating initial fits of the likelihood. It thus substantially relaxes the degree of data support required to remain an efficient double robust substitution estimator.

In summary, a large class of causal questions can be formally defined using static and dynamic longitudinal MSMs, and the parameters of these models can be identified from non-randomized data under well-studied assumptions. This article describes a TMLE that builds on the work of Robins [13, 29] and Bang and Robins [22] in order to directly target the coefficients of a marginal structural (working) model for a user-supplied class of longitudinal static or dynamic interventions. The theoretical properties of the pooled TMLE are presented, its implementation is reviewed, and its practical performance is compared to alternatives using both simulated and real data. R code [33] implementing the estimator and evaluating it in simulations is provided in online supplementary materials and as an open source R library `ltmle` [34].

1.1 Organization of paper

In Section 2, we define the observed data and a statistical model for its distribution. We then specify a non-parametric structural equation model for the process assumed to generate the observed data. We define counterfactual outcomes over time based on static or dynamic interventions on multiple treatment and

censoring nodes in this system of structural equations. Our target causal quantity is defined using a marginal structural working model on the mean of these intervention-specific counterfactual outcomes at time t . The general case we present includes marginal structural working models on both the counterfactual survival and the hazard. We briefly review the assumptions under which this causal quantity is identified as a parameter of the observed data distribution. The statistical estimation problem is thus defined in terms of the statistical model and statistical target parameter.

Section 3 presents the TMLE defined by (a) representation of the statistical target parameter in terms of an iteratively defined set of conditional mean outcomes, (b) an initial estimator for the intervention mechanism and for these conditional means, (c) a submodel through this initial estimator and a loss function chosen so that the generalized score of the submodel with respect to this loss spans the efficient influence curve, (d) a corresponding updating algorithm that updates the initial estimator and iterates the updating till convergence, and (e) final evaluation of the TMLE as a plug-in estimator. We also present corresponding influence curve-based confidence intervals for our target parameter.

Section 4 illustrates the results presented in Section 3 using a simple three time point example and focusing on a marginal structural working model for counterfactual survival probability over time. This example is used to clarify understanding of notation and to provide a step-by-step overview of implementation of the pooled TMLE.

Section 5 compares the pooled TMLE described in this paper with alternative estimators for the parameters of longitudinal dynamic MSMs for survival. We provide a brief overview of the stratified TMLE [32], discuss scenarios in which each estimator may be expected to offer superior performance, and illustrate the breakdown of the stratified TMLE in a finite sample setting in which some interventions of interest have no support. As IPW estimators are currently the most common approach used to fit longitudinal dynamic MSMs, we also discuss two IPW estimators for these parameters.

Section 6 presents a simulation study in which the pooled TMLE is implemented for a marginal structural working model for survival at time t . Its performance is compared to IPW estimators and to the stratified TMLE for a simple data generating process and in a simulation designed to be similar to the data analysis presented in the following section, which includes time-dependent confounding and right censoring.

Section 7 presents the results of a data analysis investigating the effect of switching to second line therapy following immunologic failure of first line therapy using data from HIV-infected patients in the International Epidemiological Databases to Evaluate AIDS (IeDEA), Southern Africa. Throughout the paper, we illustrate notation and concepts using a simplified data structure based on this example.

Appendices contain a derivation of the efficient influence curve, further simulation details, an alternative TMLE, and reference table for notation. In online supplementary files, we present R code that implements the pooled TMLE, the stratified TMLE, and two IPW estimators for a marginal structural working model of survival. A corresponding publicly available R-package, `ltmle`, was released in May 2013 (<http://cran.r-project.org/web/packages/ltmle/>).

2 Definition of statistical estimation problem

Consider a longitudinal study in which the observed data structure O on a randomly sampled subject is coded as

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

where $L(0)$ are baseline covariates, $A(t)$ denotes an intervention node at time t , and $L(t)$ denotes covariates measured between intervention nodes $A(t-1)$ and $A(t)$. Assume that there is an outcome process $Y(t) \subseteq L(t)$ for $t = 1, \dots, K+1$, where $L(K+1) = Y(K+1)$ is the final outcome measured after the final treatment $A(K)$. The intervention node $A(t) = (A_1(t), A_2(t))$ has a treatment node $A_1(t)$ and a censoring indicator $A_2(t)$, where $A_2(t) = 1$ indicates that the subject is right censored by time t . We observe n

independent and identically distributed (i.i.d.) copies O_1, \dots, O_n of O , and we will denote the probability distribution of O with $P_{O,0}$, or more simply, as P_0 . Throughout, we use subscript 0 to denote the true distribution.

Running example. Here and in subsequent sections, we illustrate notation using an example in which n i.i.d. HIV-infected subjects with immunological failure on first line therapy are sampled from some target population. Here $t = 0$ denotes time of immunological failure. $L(t)$ denotes time-varying covariates, and includes CD4+ T cell count at time t and $Y(t)$, an indicator of death by time t . In addition to baseline values of these time-varying covariates, $L(0)$ includes non-time-varying covariates such as sex. The intervention nodes of interest are $A(t)$, $t = 0, \dots, K$, where $A(t)$ is defined as an indicator of switch to second line therapy by time t ; in our simplified example, we assume no right censoring. For notational convenience, after a subject dies all variables for that subject are defined as equal to their last observed value.

2.1 Statistical model

We use the notation $\bar{L}(k) = (L(0), \dots, L(k))$ to denote the history of time-dependent variable L from $t = 0, \dots, k$. Define the “parents” of a variable $L(k)$, denoted $Pa(L(k))$, as those variables that precede $L(k)$ (i.e., $Pa(L(k)) = \bar{L}(k-1), \bar{A}(k-1)$). Similarly, $\bar{A}(k) = (A(0), \dots, A(k))$ is used to denote the history of the intervention process and $Pa(A(k))$ to denote a specified subset of the variables that precede $A(k)$ such that the distribution of $A(k)$ given the whole past is equal to the distribution of $A(k)$ given its parents ($Pa(A(k)) \subseteq \bar{L}(k), \bar{A}(k-1)$). Under our causal model, which we introduce below, these parent sets $Pa(L(k))$ and $Pa(A(k))$ correspond to the set of variables that may affect the values taken by $L(k)$ and $A(k)$, respectively.

We use $Q_{L(k),0}$ to denote the conditional distribution of $L(k)$, given $Pa(L(k))$, and $g_{A(k),0}$ to denote the conditional distribution of $A(k)$, given $Pa(A(k))$. We also use the notation: $g_{0:k} \equiv \prod_{j=0}^k g_{A(j),0}$, $g_0 \equiv g_{0:K}$ and define $Q_{0:k} \equiv \prod_{j=0}^k Q_{L(j),0}$ and $Q_0 \equiv Q_{0:K+1}$. In our example, $Q_{L(k),0}$ denotes the joint conditional distribution of CD4 count and death at time k , given the observed past (including past CD4 count and switching history), and $g_{A(k),0}$ denotes the conditional probability of having switched to second line by time k given the observed past (deterministically equal to one for those time points at which a subject has already switched).

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

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

We consider a statistical model \mathcal{M} for P_0 that possibly assumes knowledge on the intervention mechanism g_0 . For example, the treatment of interest, such as switch time, may be known to be randomized, or to be assigned based on only a subset of the observed past. If \mathcal{Q} is the set of all values for Q_0 and \mathcal{G} the set of possible values of g_0 , then this statistical model can be represented as $\mathcal{M} = \{P = Qg : Q \in \mathcal{Q}, g \in \mathcal{G}\}$. In this statistical model, \mathcal{Q} puts no restrictions on the conditional distributions $Q_{L(k),0}$, $k = 0, \dots, K+1$.

2.2 Causal model and counterfactuals of interest

By specifying a structural causal model [35, 36] or equivalently, a system of non-parametric structural equations, it is assumed that each component of the observed longitudinal data structure (e.g. $A(k)$ or $L(k)$)

is a function of a set of observed parent variables and an unmeasured exogenous error term. Specifically, consider the non-parametric structural equation model (NPSEM) defined by

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

and

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

in terms of a set of deterministic functions ($f_{L(k)} : k = 0, \dots, K + 1$), ($f_{A(k)} : k = 0, \dots, K$), and a vector of unmeasured random errors or background factors $U = (U_{L(0)}, \dots, U_{L(K+1)}, U_{A(0)}, \dots, U_{A(K)})$ [35, 36].

To continue our HIV example, we might specify a causal model in which both time-varying CD4 count, death, and the decision to switch potentially depended on a subject's entire observed past, as well as unmeasured factors. Alternatively, if we knew that switching decisions were made only in response to a subject's most recent CD4 count and baseline covariates, the parent set of $A(k)$ could be restricted to exclude earlier CD4 count values.

This causal model represents a model \mathcal{M}^F for the distribution of (O, U) and provides a parameterization of the distribution of the observed data structure O in terms of the distribution of the random variables (O, U) modeled by the system of structural equations. Let $P_{O,U,0}$ denote the latter distribution. The causal model \mathcal{M}^F encodes knowledge about the process, including both measured and unmeasured variables, that generated the observed data. It also implies a model for the distribution of counterfactual random variables under specific interventions on (or changes to) the observed data generating process. Specifically, a post-intervention (or counterfactual) distribution is defined as the distribution that O would have had under a specified intervention to set the value of the intervention nodes $\bar{A} = (A(0), \dots, A(K))$.

The intervention of interest might be static, with $f_{A(k)}$, $k = 0, \dots, K$ replaced by some constant for all subjects. For example, an intervention to set $A(k) = 0$ for $k = 0, \dots, K$ corresponds to a static intervention to delay switching indefinitely for all subjects. Alternatively, the intervention might be dynamic, with $f_{A(k)}$, $k = 0, \dots, K$ replaced by some specified function $d_k(\bar{L}(k))$ of a subject's observed covariates. For example, an intervention could set $A(k)$ to 1 the first time a subject's CD4 count drops below some threshold. As static regimes are a special case of dynamic regimes, in the following sections we define the statistical estimation problem and develop our estimator for the more general dynamic case. Throughout, we use "rule" to refer to a specific intervention, static or dynamic, that sets the values of \bar{A} .

Given a rule d , the counterfactual random variable $L_d = (L(0), L_d(1), \dots, L_d(K + 1))$ is defined by deterministically setting all the $A(k)$ nodes equal to $d_k(\bar{L}(k))$ in the system of structural equations. The probability distribution of this counterfactual L_d is called the post-intervention or counterfactual distribution of L and is denoted with $P_{d,0}$. Causal effects are defined as parameters of a collection of post-intervention distributions under a specified set of rules. For example, we might compare mean counterfactual survival over time under a range of possible switch times.

2.3 Marginal structural working model

Our causal quantity of interest is defined using a marginal structural working model to summarize how the mean counterfactual outcome at time t varies as a function of the intervention rule d , time point t , and possibly some baseline covariate V that is a function of the collection of all baseline covariates $L(0)$. Specifically, given a class of dynamic treatment rules \mathcal{D} , we can define a true time-dependent causal dose-response curve $(E_{P_{d,0}}(Y_d(t)|V) : d \in \mathcal{D}, t \in \tau)$ for some subset $\tau \subseteq \{1, \dots, K + 1\}$. Note that choice of V (as well as choice of τ and \mathcal{D}) depends on the scientific question of interest. In many cases V will be defined as the empty set. In other cases, it may be of interest to estimate how the causal dose-response curve varies depending on the value of some subset of baseline variables.

We specify a working model $\Theta \equiv \{m_\beta : \beta\}$ for this true time-dependent causal dose-response curve. Our causal quantity of interest is then defined as a projection of the true causal dose-response curve onto this

working model, which yields a definition m_{β_0} representing this projection. For example, if $Y(t) \in [0, 1]$, we may use a logistic working model $\text{Logit } m_{\beta}(d, t, V) = \sum_{j=1}^J \beta_j \phi_j(d, t, V)$ for a set of basis functions, and define our causal quantity of interest as

$$\begin{aligned} \Psi^F(P_{O,U,0}) = \operatorname{argmin}_{\beta} & -E_0 \sum_{t \in \tau} \sum_{d \in \mathcal{D}} h(d, t, V) \{Y_d(t) \log m_{\beta}(d, t, V)\}, \\ & + (1 - Y_d(t)) \log(1 - m_{\beta}(d, t, V)), \end{aligned}$$

where $h(d, t, v)$ is a user-specified weight function. We discuss choice of $h(d, t, v)$ further below.

Such a $\Psi_0^F = \beta_0$ solves the equation

$$0 = E_0 \sum_{t \in \tau} \sum_{d \in \mathcal{D}} h(d, t, V) \frac{\frac{d}{d\beta_0} m_{\beta_0}(d, t, V)}{m_{\beta_0}(1 - m_{\beta_0})} (E_0(Y_d(t) | V) - m_{\beta_0}(d, t, V)).$$

This equation can be replaced by

$$0 = E_0 \sum_{t \in \tau} \sum_{d \in \mathcal{D}} h(d, t, V) \frac{\frac{d}{d\beta_0} m_{\beta_0}(d, t, V)}{m_{\beta_0}(1 - m_{\beta_0})} (E_0(Y_d(t) | L(0)) - m_{\beta_0}(d, t, V)),$$

which corresponds with

$$\begin{aligned} \Psi^F(P_{O,U,0}) = \operatorname{argmin}_{\beta} & -E_0 \sum_{t \in \tau} \sum_{d \in \mathcal{D}} h(d, t, V) \{E_0(Y_d(t) | L(0)) \log m_{\beta}(d, t, V) \\ & + (1 - E_0(Y_d(t) | L(0))) \log(1 - m_{\beta}(d, t, V))\}. \end{aligned}$$

In this case we have that $\frac{\frac{d}{d\beta_0} m_{\beta_0}(d, t, V)}{m_{\beta_0}(1 - m_{\beta_0})} = (\phi_j(d, t, V) : j = 1, \dots, J)$.

To be completely general, we will define our causal quantity of interest as a function f of $E(Y_d(t) | L(0))$ across $d \in \mathcal{D}, t \in \tau$ and the distribution $Q_{L(0)}$ of $L(0)$. Thus we define

$$\Psi^F(P_{O,U,0}) = f((E_0(Y_d(t) | L(0)) : d \in \mathcal{D}, t \in \tau), Q_{L(0),0}).$$

In addition to including the above example, this general formulation allows us to include marginal structural working models on continuous outcomes and on the intervention-specific hazard.

2.3.1 Choice of a weight function $h(d, t, V)$

Unless one is willing to assume that the MSM m_{β} is correctly specified, choice of the weight function changes the target quantity being estimated. Choice of the weight function should thus be guided by the motivating scientific question. For example, the simple weight function $h(d, t, V) = 1$ gives equal weight to all time points and switch times. Alternatively, choice of a weight function equal to the marginal probability of following rule d through time t within strata of V gives greater weight to those rule, time, and baseline strata combinations with more support in the data, and zero weight to values (d, t, v) without support. As discussed further below, choice of a weight function can thus also affect both identifiability of the target parameter and the asymptotic and finite sample properties of IPW and TML estimators.

2.3.2 Running example

Continuing our HIV example, recall that static regimes are a special case of dynamic regimes and define the set of treatment rules of interest \mathcal{D} as the set of possible switch times (switch at time 0, switch at time 1, \dots ,

never switch). We might focus on the marginal counterfactual survival curves under a range of switch times (with V defined as the empty set). Alternatively, we might investigate how survival under a specific switch time differs among subjects that have a CD4+ T cell count < 50 versus ≥ 50 cells/ μl at time of failure ($V = I(CD4(0) < 50)$ where $CD4(0) \subset L(0)$). For simplicity, for the remainder of the paper we use a running example in which we avoid conditioning on baseline covariates (i.e. $V = \{\}$).

The true time-dependent causal dose–response curve ($E_0(Y_d(t)) : d \in \mathcal{D}, t \in 1, \dots, K+1$) corresponds to the set of counterfactual survival curves (through time $K+1$) that would have been observed for the population as a whole under each possible switch time. In this example, each rule d implies a single vector \bar{a} ; we use $d(t)$ to refer to the value $a(t)$ implied by rule d and s_d to refer to the switch time assigned by rule d . One might then specify the following marginal structural working model to summarize how the counterfactual probability of death by time t varies as a function of t and assigned switch time:

$$\text{Logit } m_\beta(d, t) = \beta_0 + \beta_1 t + \beta_2(d(t-1)(t-s_d)),$$

where $d(t-1)(t-s_d)$ is time since switch for subjects who have switched by time $t-1$, and otherwise 0. For simplicity, we choose $h(d, t) = 1$ and define the target causal quantity of interest as the projection of ($E_0(Y_d(t)) : d \in \mathcal{D}, t \in 1, \dots, K+1$) onto $m_\beta(d, t)$ according to

$$\begin{aligned} \Psi^F(P_{0,U,0}) = \underset{\beta}{\operatorname{argmin}} & -E_0 \sum_{t \in \tau} \sum_{d \in \mathcal{D}} \{Y_d(t) \log m_\beta(d, t) \\ & + (1 - Y_d(t)) \log(1 - m_\beta(d, t))\}. \end{aligned} \quad (1)$$

2.4 Identifiability and definition of statistical target parameter

We assume the sequential randomization assumption [11]

$$A(k) \prod L_d | Pa(A(k)), k = 0, \dots, K \quad (2)$$

(noting that weaker identifiability assumptions are also possible; see, for example, Robins and Hernan [9]). In our HIV example, the plausibility of this assumption would be strengthened by measuring all determinants of the decision to switch to second line therapy that also affect mortality via pathways other than switch time.

We further assume positivity, informally an assumption of support for each rule of interest across covariate histories compatible with that rule of interest. Specifically, for each d, t, V for which $h(d, t, V) \neq 0$, we assume

$$P_0(A(k) = d_k(\bar{L}(k)) | \bar{L}(k), \bar{A}(k-1) = \bar{d}(\bar{L}(k-1))) > 0, k = 0, \dots, K \text{ – almost everywhere.} \quad (3)$$

In our HIV example, in which $h(d, t) = 1$, a subject who has not already switched should have some positive probability of both switching and not switching regardless of his covariate history. Under these assumptions, the counterfactual probability distribution of L_d is identified from the true observed data distribution P_0 and given by the G -computation formula P_0^d [11]:

$$P_0^d(l) = \prod_{k=0}^{K+1} Q_{L(k),0}^d(\bar{l}(k)), \quad (4)$$

where $Q_{L(k),0}^d(\bar{l}(k)) = Q_{L(k),0}(l(k) | \bar{l}(k-1), \bar{A}(k-1) = \bar{d}(\bar{L}(k-1)))$. Thus this G -computation formula P_0^d is defined by the product over all $L(k)$ -nodes of the conditional distribution of the $L(k)$ -node, given its parents, and given $\bar{A}(k-1) = \bar{d}(\bar{L}(k-1))$. If identifiability assumptions (2) and (3) hold for each rule $d \in \mathcal{D}$, then the time-dependent causal dose–response curve ($E_0(Y_d(t) | V) : d \in \mathcal{D}, t \in \tau$) is also identified from P_0 through the collection of G -computation formulas ($P_0^d : d \in \mathcal{D}$). For the remainder of the paper, we choose $\tau = 1, \dots, K+1$ and at times suppress the index set τ .

Let $L^d = (L(0), L^d(1), \dots, L^d(K+1))$ denote a random variable with probability distribution P_0^d , which includes as a component the process $Y^d = (Y^d(0), Y^d(1), \dots, Y^d(K+1))$. The above-defined causal quantities can now be defined as a parameter of P_0 . For example, if $Y(t) \in [0, 1]$ and the causal parameter of interest is a vector of coefficients in a logistic MSM, then we have

$$\begin{aligned} \Psi^F(P_{O,U,0}) &= \operatorname{argmin}_{\beta} -E_0 \sum_t \sum_{d \in \mathcal{D}} h(d, t, V) \{E_0(Y^d(t)|L(0)) \log m_{\beta}(d, t, V) \\ &\quad + (1 - E_0(Y^d(t)|L(0))) \log(1 - m_{\beta}(d, t, V))\} \\ &\equiv \Psi(P_0). \end{aligned}$$

The estimand $\psi_0 = \beta_0$ solves the equation

$$0 = E_0 \sum_t \sum_{d \in \mathcal{D}} h(d, t, V) \frac{\frac{d}{d\beta_0} m_{\beta_0}(d, t, V)}{m_{\beta_0}(1 - m_{\beta_0})} (E_0(Y^d(t)|L(0)) - m_{\beta_0}(d, t, V)).$$

The causal identifiability assumptions put no restrictions on the probability distribution P_0 so that our statistical model is unchanged, with the exception that we now also assume positivity (3). The statistical target parameter is now defined as a mapping $\Psi : \mathcal{M} \rightarrow \mathbb{R}^J$ that maps a probability distribution $P \in \mathcal{M}$ of O into a vector of parameter values $\Psi(P)$.

The statistical estimation problem is now defined: We observe n i.i.d. copies O_1, \dots, O_n of $O \sim P_0 \in \mathcal{M}$ and we want to estimate $\Psi(P_0)$ for a defined target parameter mapping $\Psi : \mathcal{M} \rightarrow \mathbb{R}^J$. For this estimation problem, the causal model plays no further role – even when one does not believe any of the causal assumptions, one might still argue that the statistical parameter $\Psi(P_0) = \psi_0$ represents an effect measure of interest controlling for all the *measured* confounders.

3 Pooled TMLE of working MSM for dynamic treatments and time-dependent outcome process

The TMLE algorithm starts out with defining the target parameter as a $\Psi(\bar{Q}, Q_{L(0)})$ for a particular choice \bar{Q} that is easier to estimate than the whole likelihood Q . It requires the derivation of the efficient influence curve $D^*(P)$ which can also be represented as $D^*(\bar{Q}, Q_{L(0)}, g)$. Subsequently, it defines a loss function $\mathcal{L}(\bar{Q}, Q_{L(0)})$ for $(\bar{Q}_0, Q_{L(0),0})$ and a submodel $((\bar{Q}(\epsilon, g), Q_{L(0)}(\epsilon_0)) : \epsilon, \epsilon_0)$ through $(\bar{Q}, Q_{L(0)})$ at $(\epsilon, \epsilon_0) = 0$, indexed by the intervention mechanism g , chosen so that $\frac{d}{d(\epsilon, \epsilon_0)} \mathcal{L}(\bar{Q}(\epsilon, g), Q_{L(0)}(\epsilon_0)) \Big|_{\epsilon=0}$ spans the efficient influence curve $D^*(\bar{Q}, Q_{L(0)}, g)$. Given these choices, it remains to define the updating algorithm which simply uses the submodel through the initial estimator to determine the update by fitting (ϵ, ϵ_0) with minimum loss based estimation (MLE), and this updating step is iterated till convergence at which point the MLE of (ϵ, ϵ_0) equals 0. By the fact that an MLE solves its score equation, it then follows that the final update $\bar{Q}_n^*, Q_{L(0),n}^*$ also solves the efficient influence curve equation $\sum_i D^*(\bar{Q}_n, Q_{L(0),n}^*, g_n)(O_i) = 0$, which provides the foundation for its asymptotic linearity and efficiency. The remainder of this section presents each of these steps in detail.

An estimator of ψ_0 is efficient among the class of regular estimators if and only if it is asymptotically linear with influence curve $D^*(Q_0, g_0)$ [37]. The efficient influence curve can thus be used as an ingredient for the construction of an efficient estimator. One approach is to represent the efficient influence curve as an estimating function $D^*(Q, g, \psi)$ and define an estimator ψ_n as the solution of $P_n D^*(Q_n, g_n, \psi) = 0$, given initial estimators Q_n, g_n . This is referred to as the estimating equation methodology for construction of locally efficient estimators [38]. Here, we instead use the efficient influence curve to define a targeted

maximum likelihood (substitution) estimator $\Psi(Q_n^*)$ that, as a by-product of the procedure, satisfies $P_n D^*(Q_n^*, g_n) = 0$ and thus also solves the efficient influence curve estimating equation. Under regularity conditions, one can now establish that, if $D^*(Q_n^*, g_n)$ consistently estimates $D^*(Q_0, g_0)$, then $\Psi(Q_n^*)$ is asymptotically linear with influence curve equal to the efficient influence curve, so that $\Psi(Q_n^*)$ is asymptotically efficient. In addition, robustness properties of the efficient influence curve are naturally inherited by the TMLE.

Robins [13, 29] and Bang and Robins [22] reformulate the statistical target parameter and corresponding efficient influence curve for longitudinal MSMs on the intervention-specific mean as a series of iterated conditional expectations. For completeness, and to generalize to dynamic marginal structural working models possibly conditional on baseline covariates, as well as to general functions of the intervention-specific mean across a user-supplied class of interventions, we present this reformulation of the statistical target parameter below. The corresponding efficient influence curve is given in Appendix B. We will use the common notation $Ph = \int h(O) dP(O)$ for the expectation of a function $h(O)$ with respect to P .

3.1 Reformulation of the statistical target parameter in terms of iteratively defined conditional means

For the case $Y_d(t) \in [0, 1]$ in Section 2.4 we defined $\Psi(P)$ as

$$\Psi(Q) = \operatorname{argmin}_{\beta} -E \sum_t \sum_{d \in \mathcal{D}} h(d, t, V) \left\{ \bar{Q}_{L(1)}^{d,t} \log m_{\beta}(d, t, V) + (1 - \bar{Q}_{L(1)}^{d,t}) \log(1 - m_{\beta}(d, t, V)) \right\}, \quad (5)$$

where $\bar{Q}_{L(1)}^{d,t} = E_P(Y^d(t)|L(0))$. Thus, $\Psi(P)$ only depends on P through $\bar{Q}_{L(1)} = (\bar{Q}_{L(1)}^{d,t} : d \in \mathcal{D}, t)$ and $Q_{L(0)}$. Therefore, we will also refer to the statistical target parameter $\Psi(P)$ as $\Psi(Q)$ where we redefine $Q \equiv (\bar{Q}_{L(1)}, Q_{L(0)})$. For each given t , we can use the following recursive definition of $E_P(Y^d(t)|L(0))$: for $k = t, t-1, \dots, 1$ we have

$$\begin{aligned} \bar{Q}_{L(k)}^{d,t} &= E\left(Y^d(t) | \bar{L}^d(k-1)\right) \\ &= E_{L(k)}\left(\bar{Q}_{L(k+1)}^{d,t} | \bar{L}(k-1), \bar{A}(k-1) = \bar{d}_{k-1}(\bar{L}(k-1))\right), \end{aligned}$$

where we define $\bar{Q}_{L(t+1)}^{d,t} = Y(t)$. This defines $\bar{Q}_{L(1)}^{d,t}$ as an iteratively defined conditional mean [22].

To obtain $\Psi(Q)$ we simply put $\bar{Q}_{L(1)} = (\bar{Q}_{L(1)}^{d,t} : d \in \mathcal{D}, t)$, combined with the marginal distribution of $L(0)$ into the above representation $\Psi(Q) = \Psi(\bar{Q}_{L(1)}, Q_{L(0)})$. As mentioned in the previous section, we have that $\Psi(Q)$ solves the score equations given by

$$\begin{aligned} 0 &= E \sum_t \sum_{d \in \mathcal{D}} h(d, t, V) \frac{\frac{d}{d\beta} m_{\beta}(d, t, V)}{m_{\beta}(1 - m_{\beta})} (E(Y^d(t)|L(0)) - m_{\beta}(d, t, V)) \\ &\equiv E \sum_t \sum_{d \in \mathcal{D}} h_1(d, t, V) (E(Y^d(t)|L(0)) - m_{\beta}(d, t, V)), \end{aligned}$$

where we defined

$$h_1(d, t, V) \equiv h(d, t, V) \frac{\frac{d}{d\beta} m_{\beta}(d, t, V)}{m_{\beta}(1 - m_{\beta})}.$$

The TMLE for the linear working model using the squared error loss function is obtained by simply redefining $h_1(d, t, V) \equiv h(d, t, V) \frac{d}{d\beta} m_{\beta}(d, t, V)$.

In general, the above shows that we can represent $\Psi(P) = f((E(Y^d(t)|L(0)) : t, d), Q_{L(0)})$ as a function $f(Q)$, where $Q = (\bar{Q}_{L(1)} = (E(Y^d(t)|L(0)) : d, t), Q_{L(0)})$ and we have an explicit representation of the derivative equation corresponding with f .

3.2 Estimation of intervention mechanism g_0

The log-likelihood loss function for g_0 is $-\log g$. Specifically, we can factorize the likelihood g_0 as

$$g_0 = \prod_{k=0}^K g_{1,k}(A_1(k)|Pa(A_1(k)))g_{2,k}(A_2(k)|Pa(A_2(k))),$$

where $(g_{1,k} : k)$ represents the treatment mechanism and $(g_{2,k} : k)$ represents the censoring mechanism. Both mechanisms can be estimated separately with a log-likelihood based logistic regression estimator, either according to parametric models, or preferably using the state of the art in machine learning. In particular, we can use the log-likelihood based super learner based on a library of candidate machine learning algorithms, which uses cross-validation to determine the best performing weighted combination of the candidate machine learning algorithms [39]. Use of such aggressive data-adaptive algorithms is recommended in order to ensure consistency of g_n .

If there are certain variables in the $Pa(A(k))$ that are known to be instrumental variables (variables that affect future Y nodes only via their effects on $A(k)$), then these variables should be excluded from our estimates of g_0 in the TMLE procedure. In that case our estimate of the conditional distribution of $A(k)$ is in fact not estimating the conditional distribution of $A(k)$ given its parents; however, for simplicity we do not make this explicit in our notation.

3.3 Loss functions and initial estimator of Q_0

We will alternate notation $\bar{Q}_k^{d,t}$ and $\bar{Q}_{L(k)}^{d,t}$. Recall that $\Psi(Q)$ depends on Q through $Q_{L(0)}$, and $\bar{Q} \equiv (\bar{Q}_k^{d,t} : d \in \mathcal{D}, t \in \tau, k = 1, \dots, t)$. Note $\bar{Q}_k^{d,t}$ is a function of $\bar{L}(k-1)$, $t = 1, \dots, K+1$, $k = 1, \dots, t$, $d \in \mathcal{D}$. We will use the following loss function for $\bar{Q}_k^{d,t}$:

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

This is an application of the log-likelihood loss function for the conditional mean of $\bar{Q}_{k+1}^{d,t}$ given past covariates and given that past treatment has been assigned according to rule d . For example, fitting a parametric logistic regression model of $\bar{Q}_{k+1}^{d,t}$ on past covariates among subjects with $\bar{A}(k-1) = \bar{d}_{k-1}(\bar{L}(k-1))$ would minimize the empirical mean of this loss function over the unknown parameters of the logistic regression model. Alternatively, one could use loss-based machine learning algorithms, such as loss-based super learning, with this loss function.

In this loss function, the outcome $\bar{Q}_{k+1}^{d,t}$ is treated as known. In implementation of our estimator, it will be replaced by an estimate; we thus refer to $\bar{Q}_{k+1}^{d,t}$ as a nuisance parameter in this loss function. The collection of loss functions from $k = 1, \dots, t$ implies a sequential regression procedure where one starts at $k = t$ and sequentially fits $\bar{Q}_k^{d,t}$ for $k = t, \dots, 1$. We describe this procedure in greater detail in the next subsection, for a sum-loss function that sums the above loss function over a collection of rules $d \in \mathcal{D}$.

By summing over $d \in \mathcal{D}$, the time points t , and $k = 1, \dots, t$, we obtain the loss function

$$\mathcal{L}_{\bar{Q}}(\bar{Q}) \equiv \sum_t \sum_{k=1}^t \sum_{d \in \mathcal{D}} \mathcal{L}_{d,t,k,\bar{Q}_{k+1}^{d,t}}(\bar{Q}_k^{d,t})$$

for the whole $\bar{Q} = (\bar{Q}_k^{d,t} : d \in \mathcal{D}, k = 1, \dots, t, t = 1, \dots, K + 1)$.

We will use the log-likelihood loss $\mathcal{L}(Q_{L(0)}) = -\log Q_{L(0)}$ as loss function for the distribution $Q_{0,L(0)}$ of $L(0)$, but this loss will play no role since we will estimate $Q_{0,L(0)}$ with the empirical distribution function $Q_{L(0),n}$. To conclude, we have presented a loss function for all components of $(\bar{Q}, Q_{L(0)})$ our target parameter depends on, and the sum-loss function $\mathcal{L}_Q(\bar{Q}, Q_{L(0)}) \equiv \mathcal{L}_{\bar{Q}}(\bar{Q}) - \log Q_{L(0)}$ is a valid loss function for $(\bar{Q}, Q_{L(0)})$ as a whole.

3.4 Non-targeted substitution estimator

These loss functions imply a sequential regression methodology for fitting each of the required components of $(\bar{Q}, Q_{L(0)})$. These initial fits can then be used to construct a non-targeted plug-in estimator of the target parameter ψ_0 . As noted, we estimate the marginal distribution of $L(0)$ with the empirical distribution. We now describe how to obtain an estimator $\bar{Q}_{k,n}^{d,t}$, $d \in \mathcal{D}$, $k = 1, \dots, t$, for any given $t = 1, \dots, K + 1$. We define $\bar{Q}_{t+1}^{d,t} = Y(t)$ for all d , and recall that $\bar{Q}_{t,n}^{d,t}$ is the regression of $Y(t)$ on $\bar{A}(t-1) = \bar{d}_{t-1}(\bar{L}(t-1))$ and $\bar{L}(t-1)$. This latter regression can be carried out conditional on $\bar{A}_1(t-1), \bar{L}(t-1)$, stratifying only on not being censored through time $t-1$ (i.e. $\bar{A}_2(t-1) = 0$). The resulting fit for all $\bar{A}_1(t-1)$ values can then be evaluated at $\bar{A}(t-1) = \bar{d}_{t-1}(\bar{L}(t-1))$. In this manner, if certain rules have little support, one can still obtain an initial estimator that smoothes across all observations.

Given the regression fit $\bar{Q}_{t,n}^{d,t}$, for a $d \in \mathcal{D}$, we regress $\bar{Q}_{t,n}^{d,t}$ onto $\bar{A}(t-2), \bar{L}(t-2)$ and evaluate it at $\bar{A}(t-2) = \bar{d}_{t-2}(\bar{L}(t-2))$, giving us $\bar{Q}_{t-1,n}^{d,t}$. This is carried out for each $d \in \mathcal{D}$, giving us $\bar{Q}_{t-1,n}^{d,t}$ for each $d \in \mathcal{D}$. Again, given this regression $\bar{Q}_{t-1,n}^{d,t}$, we regress this on $\bar{A}(t-3), \bar{L}(t-3)$, and evaluate it at $\bar{A}(t-3) = \bar{d}_{t-3}(\bar{L}(t-3))$, giving us $\bar{Q}_{t-2,n}^{d,t}$. We carry this out for each $d \in \mathcal{D}$, giving us $\bar{Q}_{t-2,n}^{d,t}$, for each $d \in \mathcal{D}$. This process is iterated until we obtain an estimator of $\bar{Q}_{1,n}^{d,t}(L(0))$ for each $d \in \mathcal{D}$. Since this process is carried out for each $t = 1, \dots, K + 1$, this results in an estimator $\bar{Q}_{1,n}^{d,t}$ for each $d \in \mathcal{D}$ and $t = 1, \dots, K + 1$. We denote this estimator of $\bar{Q}_{1,0} = (\bar{Q}_{1,0}^{d,t} : d, t)$ with $\bar{Q}_{1,n} = (\bar{Q}_{1,n}^{d,t} : d, t)$. Note that a plug-in estimator $\Psi(\bar{Q}_{1,n}, Q_{L(0),n})$ of $\psi_0 = \Psi(\bar{Q}_{1,0}, Q_{L(0)})$ is now obtained by regressing $\bar{Q}_{1,n}^{d,t}$ onto d, t, V according to the working marginal structural model using weighted logistic regression based on the pooled sample $(\bar{Q}_{1,n}^{d,t}(L_i(0)), V_i, d, t)$, $d \in \mathcal{D}, i = 1, \dots, n, t = 1, \dots, K + 1$, with weight $h(d, t, V_i)$.

The pooled TMLE presented below utilizes this same sequential regression algorithm and makes use of these initial fits of Q_0 . In order to provide a consistent initial estimator of Q_0 and thereby improve the efficiency of the TMLE, use of an aggressive data-adaptive algorithm such as super learning [39] when generating the initial regression fits is recommended. These initial fits are then updated to remove bias in a series of targeting steps that rely on the fit g_n of g_0 . The updating steps involve submodels whose score spans the efficient influence curve.

3.5 Loss function and least favorable submodel that span the efficient influence curve

Recall that we use the notation $g_{0:k} = \prod_{j=0}^k g_{A(j)}$ for the cumulative product of conditional intervention distributions. Consider the submodel $\bar{Q}_k^t(\epsilon, g) = (\bar{Q}_k^{d,t}(\epsilon, g) : d \in \mathcal{D})$ with parameter ϵ defined by

$$\text{Logit } \bar{Q}_{L(k)}^{d,t}(\epsilon, \mathbf{g}) = \text{Logit } \bar{Q}_{L(k)}^{d,t} + \epsilon \frac{h_1(d, t, V)}{g_{0:k-1}}, \quad k = t, \dots, 1.$$

This parameter ϵ is of same dimension as β and h_1 . This defines a submodel $\bar{Q}^t(\epsilon, \mathbf{g})$ with parameter ϵ through $\bar{Q}^t = (\bar{Q}_k^{d,t} : d \in \mathcal{D}, k = 1, \dots, t)$. Note that

$$\left. \frac{d}{d\epsilon} \mathcal{L}_{d,t,k,\bar{Q}_{k+1}^{d,t}}(\bar{Q}_k^{d,t}(\epsilon, \mathbf{g})) \right|_{\epsilon=0} = h_1(d, t, V) \frac{I(\bar{A}(k-1) = \bar{d}_{k-1}(\bar{L}(k-1)))}{g_{0:k-1}} (\bar{Q}_{k+1}^{d,t} - \bar{Q}_k^{d,t}).$$

This shows that

$$\begin{aligned} & \frac{d}{d\epsilon} \sum_{t=1}^{K+1} \sum_{d \in \mathcal{D}} \sum_{k=1}^t \mathcal{L}_{d,t,k,\bar{Q}_{k+1}^{d,t}}(\bar{Q}_k^{d,t}(\epsilon, \mathbf{g})) \Big|_{\epsilon=0} \\ &= \sum_{t=1}^{K+1} \sum_{d \in \mathcal{D}} \sum_{k=1}^t h_1(d, t, V) \frac{I(\bar{A}(k-1) = \bar{d}_{k-1}(\bar{L}(k-1)))}{g_{0:k-1}} (\bar{Q}_{k+1}^{d,t} - \bar{Q}_k^{d,t}) \\ &= c(Q)[D^*(P) - D_{L(0)}^*(Q)] \end{aligned}$$

where $D^*(P)$ is the efficient influence curve as presented in Corollary (1), Appendix (B), and we define

$$c(Q) \equiv E_{Q_{L(0)}} \sum_{t,d} h_1(d, t, V) \frac{d}{d\beta} m_\beta(d, t, V),$$

giving

$$D_{L(0)}^*(Q) = c(Q)^{-1} \sum_{t,d} h_1(d, t, V) (\bar{Q}_{L(1)}^{d,t} - m_\beta(d, t, V)).$$

In other words, the sum-loss function

$$\mathcal{L}_{\bar{Q}}(\bar{Q}) = \sum_{t=1}^{K+1} \sum_{d \in \mathcal{D}} \sum_{k=1}^t \mathcal{L}_{d,t,k,\bar{Q}_{k+1}^{d,t}}(\bar{Q}_k^{d,t})$$

and submodel $\bar{Q}(\epsilon, \mathbf{g}) = (\bar{Q}_k^{d,t}(\epsilon, \mathbf{g}) : k, d, t)$ through $\bar{Q} = (\bar{Q}_k^{d,t} : k, d, t)$ generates the component $D^*(P) - D_{L(0)}^*(Q)$ of the efficient influence curve $D^*(P)$.

Consider also a submodel $Q_{L(0)}(\epsilon_0)$ of $Q_{L(0)}$ with score $D_{L(0)}^*(Q)$, but this submodel and loss will play no role in the TMLE algorithm since we will estimate $Q_{L(0)}$ with its NPMLE, the empirical distribution of $L_i(0)$, $i = 1, \dots, n$, so that the MLE of ϵ_0 will be equal to zero. This defines our submodel $(Q_{L(0)}(\epsilon_0), \bar{Q}(\epsilon, \mathbf{g}) : \epsilon_0, \epsilon)$. The sum-loss function $\mathcal{L}_{\bar{Q}}(Q_{L(0)}, \bar{Q}) = \mathcal{L}(\bar{Q}) - \log Q_{L(0)}$ and this submodel satisfy the condition that the generalized score spans the efficient influence curve:

$$D^*(Q, \mathbf{g}) \in \left\langle \frac{d}{d(\epsilon, \epsilon_0)} \mathcal{L}_{\bar{Q}}(Q_{L(0)}(\epsilon_0), \bar{Q}(\epsilon, \mathbf{g})) \Big|_{(\epsilon, \epsilon_0)=0} \right\rangle. \quad (6)$$

3.6 Pooled TMLE

We now describe the TMLE algorithm based on the above choices of (1) the representation of $\Psi(P)$ as $\Psi(\bar{Q}, Q_{L(0)})$, (2) the loss function for $(\bar{Q}, Q_{L(0)})$, and (3) the least favorable submodels $((\bar{Q}(\epsilon, \mathbf{g}) : \epsilon), (Q_{L(0)}(\epsilon_0) : \epsilon_0))$ through $(\bar{Q}, Q_{L(0)})$ at $(\epsilon, \epsilon_0) = 0$ for fluctuating these parameters $(\bar{Q}, Q_{L(0)})$. We utilize the same sequential regression approach described in Section 3.4, but now incorporate sequential targeted updating of the initial regression fits. We assume an estimator g_n of g_0 . We first specify where in the algorithm updating occurs and then describe the updating process.

Recall that we define $\bar{Q}_{t+1}^{d,t} = Y(t)$ for all d and that $\bar{Q}_{t,n}^{d,t}$ is the regression of $Y(t)$ on $\bar{A}(t-1) = \bar{d}_{t-1}(\bar{L}(t-1)), \bar{L}(t-1)$. For any given $t = 1, \dots, K+1$, the initial estimator $\bar{Q}_{t,n}^{d,t}$ is first updated to $\bar{Q}_{t,n}^{d,t,*}$ using a logistic regression fit of our least favorable submodels, as described below. For a $d \in \mathcal{D}$, we then regress the updated regression fit $\bar{Q}_{t,n}^{d,t,*}$ onto $\bar{A}(t-2), \bar{L}(t-2)$, and evaluate it at $\bar{A}(t-2) = \bar{d}_{t-2}(\bar{L}(t-2))$, giving us $\bar{Q}_{t-1,n}^{d,t}$. This is carried out for each $d \in \mathcal{D}$, giving us $\bar{Q}_{t-1,n}^{d,t}$ for each $d \in \mathcal{D}$. The regressions $\bar{Q}_{t-1,n}^{d,t}$ are then updated for each $d \in \mathcal{D}$, as described below, giving us $\bar{Q}_{t-1,n}^{d,t,*}$ for each $d \in \mathcal{D}$. For a $d \in \mathcal{D}$, we then regress the updated regression fit $\bar{Q}_{t-1,n}^{d,t,*}$ on $\bar{A}(t-3), \bar{L}(t-3)$ and evaluate it at $\bar{A}(t-3) = \bar{d}_{t-3}(\bar{L}(t-3))$, giving us $\bar{Q}_{t-2,n}^{d,t}$. We again carry this out for each $d \in \mathcal{D}$, giving us $\bar{Q}_{t-2,n}^{d,t}$ for each $d \in \mathcal{D}$ and again update the resulting regressions, giving us $\bar{Q}_{t-2,n}^{d,t,*}$, for each $d \in \mathcal{D}$. This process is iterated until we obtain an updated estimator $\bar{Q}_{1,n}^{d,t,*}(L(0))$ for each $d \in \mathcal{D}$. Since this process is carried out for each $t = 1, \dots, K+1$, this results in an estimator $\bar{Q}_{1,n}^{d,t,*}$ for each $d \in \mathcal{D}$ and $t = 1, \dots, K+1$. We denote this estimator of $\bar{Q}_{1,0} = (\bar{Q}_{1,0}^{d,t} : d, t)$ with $\bar{Q}_{1,n}^* = (\bar{Q}_{1,n}^{d,t,*} : d, t)$.

The updating steps are implemented as follows: for each $t \in \{1, \dots, K+1\}$, and for $k = t$ to $k = 1$, we compute

$$\epsilon_{k,n} \equiv \underset{\epsilon_k}{\operatorname{argmin}} P_n \sum_{d \in \mathcal{D}} \mathcal{L}_{d,t,k,\bar{Q}_{k+1,n}^{d,t,*}} \left(\bar{Q}_{k,n}^{d,t}(\epsilon_k, \mathbf{g}_n) \right),$$

and compute the corresponding update $\bar{Q}_{k,n}^{d,t,*} = \bar{Q}_{k,n}^{d,t}(\epsilon_{k,n}, \mathbf{g}_n)$, for all $d \in \mathcal{D}$. Note that

$$\begin{aligned} \epsilon_{k,n} &= \underset{\epsilon}{\operatorname{argmin}} \sum_{d \in \mathcal{D}} \mathcal{L}_{d,t,k,\bar{Q}_{k+1,n}^{d,t,*}} \left(\bar{Q}_{k,n}^{d,t}(\epsilon, \mathbf{g}_n) \right) \\ &= \underset{\epsilon}{\operatorname{argmin}} \sum_{d \in \mathcal{D}} \sum_{i=1}^n I(\bar{A}_i(k-1) = \bar{d}_{k-1}(\bar{L}_i(k-1))) \\ &\quad \left\{ \bar{Q}_{k+1,n}^{d,t,*}(\bar{L}_i(k)) \log \bar{Q}_{k,n}^{d,t}(\epsilon, \mathbf{g}_n)(\bar{L}_i(k-1)) \right. \\ &\quad \left. + (1 - \bar{Q}_{k+1,n}^{d,t,*}(\bar{L}_i(k))) \log(1 - \bar{Q}_{k,n}^{d,t}(\epsilon, \mathbf{g}_n)(\bar{L}_i(k-1))) \right\} \\ &k = 1, \dots, K+1. \end{aligned}$$

Thus $\epsilon_{k,n}$ can be obtained by fitting a logistic regression of the outcome $\bar{Q}_{k+1,n}^{d,t,*}(\bar{L}_i(k))$ with offset Logit $\bar{Q}_{k,n}^{d,t}$ on multivariate covariate

$$h_1(d, t, V_i) I(\bar{A}_i(k-1) = \bar{d}_{k-1}(\bar{L}_i(k-1))) / g_{0:k-1}(O_i),$$

using a data set pooled across $i = 1, \dots, n, d \in \mathcal{D}$ (consisting of $n \times |\mathcal{D}|$ observations).

This defines the TMLE $\bar{Q}_n^* = (\bar{Q}_{k,n}^{d,t,*} : d \in \mathcal{D}, t, k = 1, \dots, t)$. In particular, $\bar{Q}_{1,n}^* = (\bar{Q}_{1,n}^{d,t,*} : d \in \mathcal{D}, t)$ is the TMLE of $\bar{Q}_{1,0} = (E_0(Y^d(t)|L(0)) : d \in \mathcal{D}, t)$. This defines now the TMLE $(Q_{L(0),n}, \bar{Q}_n^*)$ of $(Q_{L(0),0}, \bar{Q}_0)$, where $Q_{L(0),n}$ is the empirical distribution of $L(0)$.

The TMLE of ψ_0 is the plug-in estimator corresponding with $\bar{Q}_{1,n}^*$ and $Q_{L(0),n}$:

$$\psi_n^* = \Psi(\bar{Q}_{1,n}^*, Q_{L(0),n}).$$

This plug-in estimator $\Psi(\bar{Q}_{1,n}^*, Q_{L(0),n})$ of $\psi_0 = \Psi(\bar{Q}_{1,0}, Q_{L(0),0})$ is obtained by regressing $\bar{Q}_{1,n}^{d,t,*}$ onto d, t, V according to the marginal structural working model in the pooled sample $(\bar{Q}_{1,n}^{d,t,*}(L_i(0)), V_i, d, t)$, $d \in \mathcal{D}, i = 1, \dots, n, t = 1, \dots, K+1$, using weights $h(d, t, V_i)$.

An alternative pooled TMLE that only fits a single ϵ to compute the update is described in Appendix C.

3.7 Statistical inference for pooled TMLE

By construction, the TMLE solves the efficient influence curve equation $0 = P_n D^*(\bar{Q}_n^*, g_n, \Psi(\bar{Q}_n^*, Q_{L(0),n}))$, thereby making it a double robust locally efficient substitution estimator under regularity conditions, and positivity (3) (van der Laan [20], theorem 8.5, appendix A.18). Here, we provide standard error estimates and thereby confidence intervals for the case that g_n is a maximum likelihood estimator for g_0 using a correctly specified semiparametric model for g_0 .

Specifically, if g_n is a maximum likelihood estimator of g_0 according to a correctly specified semiparametric model for g_0 , and \bar{Q}_n^* converges to some possibly misspecified \bar{Q} , then under regularity conditions the TMLE ψ_n^* is asymptotically linear with an influence curve given by $D^*(\bar{Q}, g_0, \psi_0)$ minus its projection onto the tangent space of this semiparametric model for g_0 . As a consequence, the asymptotic variance of $\sqrt{n}(\psi_n^* - \psi_0)$ is more spread-out or equal to the covariance matrix $\Sigma_0 = P_0 D^*(\bar{Q}, g_0, \psi_0)^2$. A consistent estimator of this asymptotic variance is given by

$$\Sigma_n = P_n \{D^*(\bar{Q}_n^*, g_n, \psi_n^*)\}^2.$$

As a consequence, $\psi_n^*(j) \pm 1.96 \frac{\sqrt{\Sigma_n(jj)}}{\sqrt{n}}$ is an asymptotically conservative 95% confidence interval for $\psi_0(j)$, and we can also use this multivariate normal limit result, $\psi_n^* \sim N(\psi_0, \Sigma_0/n)$, to construct a simultaneous confidence interval for ψ_0 and to test null hypotheses about ψ_0 . This variance estimator treats weight function h as known. If h is estimated, then this variance estimator still provides valid statistical inference for the statistical target parameter defined by the estimated h .

In the case that g_n is a data-adaptive estimator converging to g_0 , we suggest (without proof), that this variance estimator will still provide an asymptotically conservative confidence interval under regularity conditions. However, ideally the data-adaptive estimator g_n should also be targeted [40]. An approach to valid inference in the case where g_n is inconsistent but Q_n is consistent is also discussed in van der Laan [40]; however, it remains to be generalized to the parameters in this paper.

4 Implementation of the pooled TMLE

The previous section reformulated the statistical parameter in terms of iteratively defined conditional means and described a pooled TMLE for this representation. In this section, we illustrate notation and implementation of this TMLE to estimate the parameters of a marginal structural working model on counterfactual survival over time.

4.1 The statistical estimation problem

We continue our motivating example, in which the goal is to learn the effect of switch time on survival. For illustration, focus on the two time point case where $K = 1$. Let the observed data consist of n i.i.d. copies O_1, \dots, O_n of $O_i = (L_i(0), A_i(0), L_i(1), A_i(1), Y_i(2)) \sim P_0$. Let $L(t) = (Y(t), CD4(t))$, where $Y(t)$ is an indicator of death by time t , and $CD4(t)$ is CD4 count at time t . Assume all subjects are alive at baseline ($Y(0) = 0$). As above, $A(t)$ is an indicator of switch to second line by time t . We assume no right censoring so that all subjects are followed until death or the end of the study (for convenience define variable values after death as equal to their last observed value). We specify a NPSEM such that each variable may be a function of all variables that precede it and an independent error, and assume the corresponding non-parametric statistical model for P_0 .

Define the set of treatment rules of interest \mathcal{D} as the set of all possible switch times $\{0, 1, 2\}$ (where 2 corresponds to no switch). Each rule d implies a single vector $\bar{a} = (a(0), a(1))$; we use $d(t)$ to refer to the

value $a(t)$ implied by rule d , and s_d to refer to the switching time implied by rule d . We specify the following marginal structural working model for counterfactual probability of death by time t under rule d :

$$\text{Logit } m_\beta(d, t) = \beta_0 + \beta_1 t + \beta_2(d(t-1)(t-s_d)). \quad (7)$$

The target causal parameter is defined as the projection of $(E_0(Y_d(t)) : d \in \mathcal{D}, t \in \{1, 2\})$ onto $m_\beta(d, t)$ according to eq. (1).

Under the sequential randomization (2) and positivity (3) assumptions, $E(Y_d(t)|L(0)) = E(Y^d(t)|L(0))$. The target statistical parameter is defined as the projection of $(E(Y^d(t)|L(0)) : d \in \mathcal{D}, t \in \{1, 2\})$, onto the marginal structural working model m_β , according to eq. (5) with $h(d, t) = 1$.

4.1.1 Reformulation of the statistical target parameter

Note that $E(Y^d(t=2)|L(0))$ for rule d (denoted $\bar{Q}_1^{d,2}$) can be expressed in terms of iteratively defined conditional means:

$$E_{L(1)}(E_{Y(2)}(Y(2)|L(1), L(0), A(1) = d(1), A(0) = d(0))|L(0), A(0) = d(0)),$$

while $E(Y^d(t=1)|L(0))$ (denoted $\bar{Q}_1^{d,1}$) equals $E(Y(1)|L(0), A(0) = d(0))$. The statistical target parameter $\Psi(Q)$ is defined by plugging $(\bar{Q}_1^{d,1}, \bar{Q}_1^{d,2}) : d \in \mathcal{D}$, and the marginal distribution of $L(0)$ into eq. (5).

4.2 Estimator implementation

We begin by describing implementation of a simple plug-in estimator of $\Psi(Q)$.

4.2.1 Non-targeted substitution estimator

1. For each rule of interest $d \in \mathcal{D}$, corresponding to each possible switch time, generate a vector $\bar{Q}_{1,n}^{d,2}$ of length n for $t = 2$:
 - (a) Fit a logistic regression of $Y(2)$ on $L(1), L(0), A(1), A(0)$ and generate a predicted value for each subject by evaluating this regression fit at $A(1) = d(1), A(0) = d(0)$. Note $E_0(Y(2)|Y(1) = 1) = 1$, so the regression need only be fit and evaluated among subjects who remain alive at time 1. This gives a vector $\bar{Q}_{2,n}^{d,2}$ of length n .
 - (b) Fit a logistic regression of the predicted values generated in the previous step on $L(0), A(0)$. Generate a new predicted value for each subject by evaluating this regression fit at $A(0) = d(0)$. This gives a vector $\bar{Q}_{1,n}^{d,2}$ of length n .
2. For each rule of interest $d \in \mathcal{D}$ generate a vector $\bar{Q}_{1,n}^{d,1}$ of length n for $t = 1$: Fit a logistic regression of $Y(1)$ on $L(0), A(0)$ and generate a predicted value for each subject by evaluating this regression fit at $A(0) = d(0)$.
3. The previous steps generated $\bar{Q}_{1,n} = (\bar{Q}_{1,n}^{d,t} : d \in \mathcal{D}, t \in \{1, 2\})$. Stack these vectors to give a single vector with length equal to the number of subjects n times the number of rules $|\mathcal{D}|$ times the number of time points $(n \times 3 \times 2)$. Fit a pooled logistic regression of $\bar{Q}_{1,n}$ on (d, t) according to model m_β (eq. 7), with weights given by $h(d, t)$ (here equal to 1). This gives an estimator of the target parameter $\Psi(Q)$.

We now describe how the pooled TMLE modifies this algorithm to update the initial estimator $\bar{Q}_{1,n}$. In the following section, we compare the pooled TMLE to this non-targeted substitution estimator and with other available estimators.

4.2.2 Pooled TMLE

1. Estimate $P_0(A(1)|A(0), L(1), L(0))$ and $P_0(A(0)|L(0))$. Denote these estimators $g_{1,n}$ and $g_{0,n}$, respectively, and let $g_{0:1,n} = g_{0,n}g_{1,n}$ denote their product. In our example, this step involves estimating the conditional probability of switching at time 0 given baseline CD4 count, and estimating the conditional probability of switching at time 1, given a subject did not switch at time 1, did not die at time 1, and CD4 count at times 0 and 1.
2. Generate a vector $\bar{Q}_{2,n}^{d,2,*}$ of length $n \times |\mathcal{D}|$ for $t = 2, k = 2$:
 - (a) Fit a logistic regression of $Y(2)$ on $L(1), L(0), A(1), A(0)$. Generate a predicted value for each subject and each $d \in \mathcal{D}$ by evaluating this regression fit at $A(1) = d(1), A(0) = d(0)$. Note that $E_0(Y(2)|Y(1) = 1) = 1$, so the regression need only be fit and evaluated among subjects who remain alive at time 1. This gives a vector of initial values $\bar{Q}_{2,n}^{d,2}$ of length $n \times |\mathcal{D}|$.
 - (b) For each subject, $i = 1, \dots, n$, create a vector consisting of one copy of $Y_i(2)$ for each $d \in \mathcal{D}$. Stack these copies to create a single vector of length $n \times |\mathcal{D}|$, denoted $\bar{Q}_{3,n}^{d,2,*}$.
 - (c) For each subject $i = 1, \dots, n$ and each $d \in \mathcal{D}$, create a new multidimensional weighted covariate:

$$h(d, t = 2) \frac{\frac{d}{d\beta} m_\beta(d, t = 2) I(\bar{A}_i = d)}{m_\beta(1 - m_\beta) g_{0:1,n}(O_i)}.$$

In our example, $h(d, t) = 1$, and $\frac{d}{d\beta} m_\beta(d, t)$ equals 1, t , and $d(t - 1)(t - s_d)$ for the derivative taken with respect to β_0, β_1 , and β_2 , respectively. The following 3×3 matrix would thus be generated for each subject i , with rows corresponding to switch at time 0, time 1, or do not switch:

$$\begin{pmatrix} \frac{1 \times I(A_i(0) = 1, A_i(1) = 1)}{g_{0:1,n}(O_i)} & \frac{2 \times I(A_i(0) = 1, A_i(1) = 1)}{g_{0:1,n}(O_i)} & \frac{2 \times I(A_i(0) = 1, A_i(1) = 1)}{g_{0:1,n}(O_i)} \\ \frac{1 \times I(A_i(0) = 0, A_i(1) = 1)}{g_{0:1,n}(O_i)} & \frac{2 \times I(A_i(0) = 0, A_i(1) = 1)}{g_{0:1,n}(O_i)} & \frac{1 \times I(A_i(0) = 0, A_i(1) = 1)}{g_{0:1,n}(O_i)} \\ \frac{1 \times I(A_i(0) = 0, A_i(1) = 0)}{g_{0:1,n}(O_i)} & \frac{2 \times I(A_i(0) = 0, A_i(1) = 0)}{g_{0:1,n}(O_i)} & \frac{0 \times I(A_i(0) = 0, A_i(1) = 0)}{g_{0:1,n}(O_i)} \end{pmatrix}.$$

Stack these matrices to create a matrix with $n \times |\mathcal{D}|$ rows and one column for each component of β (here, $3n \times 3$).

- (d) Among those subjects still alive at the previous time point ($Y(1) = 0$), fit a pooled logistic regression of $\bar{Q}_{3,n}^{d,2,*}$ (the $Y(2)$ vector) on the weighted covariates created in the previous step, suppressing the intercept and using as offset $\text{Logit} \bar{Q}_{2,n}^{d,2}$, the logit of the initial predicted values for $t = 2$ and $k = 2$. This gives a fit for multivariate ϵ_2 . Denote this fit $\epsilon_{2,n} = (\epsilon_{2,n}^{\beta_0}, \epsilon_{2,n}^{\beta_1}, \epsilon_{2,n}^{\beta_2})$.
- (e) Generate $\bar{Q}_{2,n}^{d,2,*}$ by evaluating the logistic regression fit in the previous step at each $d \in \mathcal{D}$ among those subjects for whom $Y(1) = 0$. For subject i and rule d , evaluate

$$\begin{aligned} & \text{Expit} \left(\text{Logit}(\bar{Q}_{2,n}^{d,2}(\bar{L}_i, d)) + \frac{\epsilon_{2,n}^{\beta_0}}{g_{0,n}(d(0)|L_i(0))g_{1,n}(d(1)|L_i(0), d(0), L_i(1))} \right. \\ & \quad + \frac{\epsilon_{2,n}^{\beta_1} \times 2}{g_{0,n}(d(0)|L_i(0))g_{1,n}(d(1)|L_i(0), d(0), L_i(1))} \\ & \quad \left. + \frac{\epsilon_{2,n}^{\beta_2} \times d(1)(2 - s_d)}{g_{0,n}(d(0)|L_i(0))g_{1,n}(d(1)|L_i(0), d(0), L_i(1))} \right). \end{aligned}$$

For subjects with $Y(1) = 1$, $\bar{Q}_{2,n}^{d,2,*} = \bar{Q}_{2,n}^{d,2} = 1$. This gives an updated vector $\bar{Q}_{2,n}^{d,2,*}$ of length $n \times |\mathcal{D}|$.

3. Generate a vector $\bar{Q}_{1,n}^{d,2,*}$ of length $n \times |\mathcal{D}|$ for $t = 2, k = 1$:
- Fit a logistic regression of $\bar{Q}_{2,n}^{d,2,*}$ (generated in the previous step) on $L(0), A(0)$. Generate a predicted value for each subject and each $d \in \mathcal{D}$ by evaluating this regression fit at $A(0) = d(0)$. This gives a vector of initial values $\bar{Q}_{1,n}^{d,2}$ of length $n \times |\mathcal{D}|$.
 - For each subject $i = 1, \dots, n$ and each $d \in \mathcal{D}$, create the multidimensional weighted covariate as above, now for $k = 1$:

$$h(d, t = 2) = \frac{\frac{d}{d\beta} m_\beta(d, t = 2) I(A_i(0) = d(0))}{m_\beta(1 - m_\beta) g_{0,n}(O_i)}.$$

The following 3×3 matrix would thus be generated for each subject i , with rows corresponding to switch at time 0, time 1, or don't switch:

$$\begin{pmatrix} \frac{1 \times I(A_i(0) = 1)}{g_{0,n}(O_i)} & \frac{2 \times I(A_i(0) = 1)}{g_{0,n}(O_i)} & \frac{2 \times I(A_i(0) = 1)}{g_{0,n}(O_i)} \\ \frac{1 \times I(A_i(0) = 0)}{g_{0,n}(O_i)} & \frac{2 \times I(A_i(0) = 0)}{g_{0,n}(O_i)} & \frac{1 \times I(A_i(0) = 0)}{g_{0,n}(O_i)} \\ \frac{1 \times I(A_i(0) = 0)}{g_{0,n}(O_i)} & \frac{2 \times I(A_i(0) = 0)}{g_{0,n}(O_i)} & \frac{0 \times I(A_i(0) = 0)}{g_{0,n}(O_i)} \end{pmatrix}.$$

Stack these matrices to create a matrix with $n \times |\mathcal{D}|$ rows and one column for each dimension of β .

- Fit a pooled logistic regression of $\bar{Q}_{2,n}^{d,2,*}$ (the updated fit generated in step 2) on these weighted covariates, suppressing the intercept and using as offset $\text{Logit} \bar{Q}_{1,n}^{d,2}$, the logit of the initial predicted values for $t = 2$ and $k = 1$. This gives a fit for multivariate ϵ_1 . Denote this fit $\epsilon_{1,n} = (\epsilon_{1,n}^{\beta_0}, \epsilon_{1,n}^{\beta_1}, \epsilon_{1,n}^{\beta_2})$.
- Generate $\bar{Q}_{1,n}^{d,2,*}$ by evaluating the logistic regression fit in the previous step at each $d \in \mathcal{D}$. For subject i and rule d , evaluate

$$\text{Expit} \left(\text{Logit} \left(\bar{Q}_{1,n}^{d,2}(L_i(0), d(0)) \right) + \frac{\epsilon_{1,n}^{\beta_0}}{g_{0,n}(d(0)|L_i(0))} + \frac{\epsilon_{1,n}^{\beta_1} \times 2}{g_{0,n}(d(0)|L_i(0))} + \frac{\epsilon_{1,n}^{\beta_2} \times d(1)(2 - s_d)}{g_{0,n}(d(0)|L_i(0))} \right).$$

This gives an updated vector $\bar{Q}_{1,n}^{d,2,*}$ of length $n \times |\mathcal{D}|$.

4. Generate a vector $\bar{Q}_{1,n}^{d,1,*}$ of length $n \times |\mathcal{D}|$ for $t = 1, k = 1$:
- Fit a logistic regression of $Y(1)$ on $L(0), A(0)$. Generate a predicted value for each subject and each $d \in \mathcal{D}$ by evaluating this regression fit at $A(0) = d(0)$. This gives a vector of initial values $\bar{Q}_{1,n}^{d,1}$ of length $n \times |\mathcal{D}|$.
 - For each subject, $i = 1, \dots, n$, create a vector consisting of one copy of $Y_i(1)$ for each $d \in \mathcal{D}$. Stack these copies to create a single vector of length $n \times |\mathcal{D}|$, denoted $\bar{Q}_{2,n}^{d,1,*}$.
 - For each subject $i = 1, \dots, n$ and each $d \in \mathcal{D}$, create a new multidimensional weighted covariate, for $t = 1, k = 1$:

$$h(d, t = 1) = \frac{\frac{d}{d\beta} m_\beta(d, t = 1) I(A_i(0) = d(0))}{m_\beta(1 - m_\beta) g_{0,n}(O_i)}.$$

The following 3×3 matrix would thus be generated for each subject i , with rows corresponding to switch at time 0, time 1, or don't switch:

$$\begin{pmatrix} \frac{1 \times I(A_i(0) = 1)}{g_{0,n}(O_i)} & \frac{1 \times I(A_i(0) = 1)}{g_{0,n}(O_i)} & \frac{1 \times I(A_i(0) = 1)}{g_{0,n}(O_i)} \\ \frac{1 \times I(A_i(0) = 0)}{g_{0,n}(O_i)} & \frac{1 \times I(A_i(0) = 0)}{g_{0,n}(O_i)} & \frac{0 \times I(A_i(0) = 0)}{g_{0,n}(O_i)} \\ \frac{1 \times I(A_i(0) = 0)}{g_{0,n}(O_i)} & \frac{1 \times I(A_i(0) = 0)}{g_{0,n}(O_i)} & \frac{0 \times I(A_i(0) = 0)}{g_{0,n}(O_i)} \end{pmatrix}.$$

- Stack these matrices to create a matrix with $n \times |\mathcal{D}|$ rows and one column for each component of β .
- (d) Fit a pooled logistic regression of $\bar{Q}_{2,n}^{d,1,*}$ (the $Y(1)$ vector) on these weighted covariates, suppressing the intercept and using as offset $\text{Logit} \bar{Q}_{1,n}^{d,1}$, the logit of the initial predicted values for $t = 1$ and $k = 1$. This gives a fit for multivariate ϵ_1 . Denote this fit $\epsilon_{1,n} = (\epsilon_{1,n}^{\beta_0}, \epsilon_{1,n}^{\beta_1}, \epsilon_{1,n}^{\beta_2})$.
- (e) Generate $\bar{Q}_{1,n}^{d,1,*}$ by evaluating the logistic regression fit in the previous step at each $d \in \mathcal{D}$. For subject i and rule d , evaluate

$$\text{Expit} \left(\text{Logit} \left(\bar{Q}_{1,n}^{d,1}(L_i(0), d(0)) \right) + \frac{\epsilon_{1,n}^{\beta_0}}{g_{0,n}(d(0)|L_i(0))} + \frac{\epsilon_{1,n}^{\beta_1}}{g_{0,n}(d(0)|L_i(0))} + \frac{\epsilon_{1,n}^{\beta_2} \times d(0)(1 - s_d)}{g_{0,n}(d(0)|L_i(0))} \right).$$

This gives an updated vector $\bar{Q}_{1,n}^{d,1,*}$ of length $n \times |\mathcal{D}|$.

5. The previous steps generated $\bar{Q}_{1,n}^* = (\bar{Q}_{1,n}^{d,t,*} : d \in \mathcal{D}, t = 1, 2)$. Stack these vectors to give a single vector with length equal to the number of subjects n times the number of rules $|\mathcal{D}|$ times the number of time points $(n \times 3 \times 2)$. Fit a pooled logistic regression of $\bar{Q}_{1,n}^*$ on (d, t) according model m_β (eq. 7), with weights given by $h(d, t)$ (here equal to 1). This gives the pooled TMLE of the target parameter $\Psi(Q)$.

5 Comparison with alternative estimators

In this section we compare the TMLE described with several alternative estimators available for dynamic MSMs for survival: non-targeted substitution estimators, IPW estimators, and the stratified TMLE of Schnitzer et al. [32].

5.1 Non-targeted substitution estimator

The consistency of non-targeted substitution estimators of $\Psi(Q)$ relies entirely on consistent estimation of the Q portions of the observed data likelihood. For estimators based on the parametric G-formula this requires correctly specifying parametric estimators for the conditional distributions of all non-intervention nodes given their parents [7, 11, 12, 41]. For the non-targeted estimator described in Section (3.4), this requires consistently estimating the iteratively defined conditional means $\bar{Q} \equiv (\bar{Q}_k^{d,t} : d \in \mathcal{D}, t \in \tau, k = 1, \dots, t)$.

Correct *a priori* specification of parametric models for \bar{Q} in either case is rarely possible, rendering such non-targeted plug-in estimators susceptible to bias. Further, while machine learning methods, such as Super Learning, can be used to estimate Q non-parametrically, the resulting plug-in estimator has no theory supporting its asymptotic linearity, and will generally be overly biased for the target parameter β [21].

5.2 Inverse probability weighted estimators

The IPW estimator described in van der Laan and Petersen [2], Robins et al. [26] is commonly used to estimate the parameters of a dynamic MSM. In brief, this estimator is implemented by creating one data line for each subject i , for each t , and for each d for which $\bar{A}(t-1) = d(\bar{L}(t-1))$. Each data line consists of $Y_i(t)$, any functions of (d, t, V_i) included as covariates in the MSM, and a weight $\frac{h(d,t,V_i)I(\bar{A}_i(t-1)=d(\bar{L}_i(t-1)))}{\prod_{j=0}^{t-1} g_n(\bar{A}_i(j)|Pa(\bar{A}_i(j)))}$. A weighted logistic regression is then fit, pooling over time and rules d .

The parameter mapping presented here for the TMLE suggests an alternative IPW estimator for dynamic MSM – namely, implement an IPW estimator for $E(Y_d)$ (possibly within strata of V if V is discrete) for

$d \in \mathcal{D}, t$ and project these estimates onto m_β . The IPW estimator employed could be either the standard Horvitz–Thompson estimator:

$$\frac{1}{n} \sum_{i=1}^n \frac{I(\bar{A}_i(t-1) = d(\bar{L}_i(t-1)))}{\prod_{j=0}^{t-1} g_n(A_i(j)|Pa(A_i(j)))} Y_i,$$

or its bounded counterpart:

$$\sum_{i=1}^n \frac{I(\bar{A}_i(t-1) = d(\bar{L}_i(t-1)))}{\prod_{j=0}^{t-1} g_n(A_i(j)|Pa(A_i(j)))} Y_i / \sum_{i=1}^n \frac{I(\bar{A}_i(t-1) = d(\bar{L}_i(t-1)))}{\prod_{j=0}^{t-1} g_n(A_i(j)|Pa(A_i(j)))}$$

Robins and Rotnitzky [14].

The consistency of both IPW estimators relies on having a consistent estimator g_n of g_0 ; further, even if g_n is estimated consistently, neither will be asymptotically efficient. Both also suffer from the general sensitivity of IPW estimators to strong confounding (data sparsity or near positivity violations). Standard IPW estimators for dynamic MSMs are typically more susceptible to instability in such settings than their counterparts for static MSMs, due to the limited ability to stabilize weights (with stabilizing function restricted to $h(d, t, V)$ versus $h(\bar{A}_i(t-1), t, V)$).

5.3 Stratified targeted maximum likelihood estimator

Similar to the pooled TMLE, the stratified TMLE [32] also relies on reformulating the statistical target parameter in terms of iteratively defined conditional means and updating initial fits of these conditional means using covariates that are functions of an estimator g_n of the intervention mechanism. The stratified and pooled differ, however, in several respects. In particular, in the pooled TMLE the update step is accomplished by fitting a single multivariate ϵ for each time point t and non-intervention node k , pooling across all rules of interest $d \in \mathcal{D}$. In contrast, the stratified TMLE fits a separate ϵ for each time point t , non-intervention node k , and rule of interest $d \in \mathcal{D}$. Specifically, the stratified TMLE consists of implementing the longitudinal TMLE of van der Laan and Gruber [19] for $E(Y^d(t))$ separately for each time point t and each rule of interest $d \in \mathcal{D}$, and then combining these estimates into a fit of m_β .

Let \bar{Q}_n denote the initial estimator of the iteratively defined conditional means that forms the basis of the pooled and the stratified TMLEs. Let \bar{Q}_n^* denote the targeted update of \bar{Q}_n for the two estimators (noting that the update is accomplished differently for the pooled and stratified estimators). As long as their corresponding update Q_n^* solves the efficient influence curve equation $P_n D^*(Q_n^*, g_n)$, both the stratified and the pooled estimators will share the desirable asymptotic properties of a TMLE. Both estimators will be consistent if either \bar{Q}_n or g_n is a consistent estimator of Q_0 or g_0 , respectively. Further, both the pooled and the stratified estimators will be asymptotically efficient if both the initial estimator \bar{Q}_n and the estimator g_n are consistent.

The pooled and stratified estimators may nonetheless differ in both their asymptotic and finite sample performance. The stratified TMLE uses a more saturated model when updating \bar{Q}_n than does the pooled TMLE. Thus if the initial estimator \bar{Q}_n is misspecified, the update of this initial estimator will be more extensive (the update will be further from the initial misspecified estimator) for the stratified, as compared to the pooled, TMLE, resulting asymptotically in a Q_n^* that is closer to the true Q_0 and thus improving efficiency (recall that the efficiency bound is achieved at Q_0, g_0). The extent to which this asymptotic property translates into meaningful finite sample gains in settings where \bar{Q}_n is misspecified remains to be investigated.

On the other hand, in some cases, it is no longer clear how to implement the stratified TMLE. For example, the target parameter may be defined using a MSM $m_\beta(d, t, V)$, conditional on a subset of baseline

covariates V . If V is discrete with adequate finite sample support at each value, the stratified TMLE can be applied by estimating $E(Y^d(t)|V = v)$ within each stratum v . When V is continuous, however, or has levels which are not represented in a given finite sample, such an approach will break down for many choices of weight function $h(d, t, V)$. Similarly, whenever there are some rules $d \in \mathcal{D}$ with no support in a given finite sample, no data will be available to fit ϵ for some rules of interest, and the key update step will no longer be possible using the stratified TMLE. Note that such lack of support in finite samples can occur even when the assumption of positivity in the observed data distribution (3) is satisfied.

In cases where V is discrete and there is support for some but not all rules $d \in \mathcal{D}$ within each stratum of V , one option is to define a stratified quasi-TMLE using the initial fit $\bar{Q}_n^{d,t}$ for those rules, time points, and strata of V where no data are available to fit ϵ . The estimator, implemented in the simulations below, remains defined even when not all rules of interest are supported within each stratum of V in a given sample. However, in such cases, the initial estimator \bar{Q}_n is only partially updated, and thus Q_n^* may no longer solve the efficient influence curve equation $P_n D^*(Q_n^*, g_n)$, even if g_n is a consistent estimator of g_0 . If the initial estimator \bar{Q}_n is poor (for example, if it is a misspecified parametric model), the resulting estimator of β will be biased. In contrast, the pooled TMLE retains the ability to fit ϵ and thus update \bar{Q}_n by pooling over rules d . In other words, both estimators rely on the theoretical positivity assumption on P_0 (3) for identifiability; however, they may respond differently to practical positivity violations in finite samples.

5.3.1 Numerical illustration

We use a simple simulation to illustrate a setting in which the positivity assumption holds, but many rules of interest have no support in a given finite sample. In this setting, the stratified estimator will be biased if the initial estimator \bar{Q}_n is misspecified. In contrast, the pooled TMLE remains asymptotically linear if g_n is a consistent estimator of g_0 . Simulation studies comparing the performance of the pooled and stratified TMLEs as well as IPW estimators under more realistic scenarios are provided in the following section.

We implemented a simulation with observed data consisting of n i.i.d. copies of $O = (L(0), A = (A(0), \dots, A(6)), Y)$, where $L(0)$ is a baseline covariate, $A(t)$ is a binary treatment assigned randomly at seven time points ($t = 0, \dots, 6$), and Y is a binary outcome. The data for a given individual i were generated by drawing sequentially from the following distributions:

$$L(0) \sim N(0, 1),$$

$$A(t) \sim \text{Bern}(p = 0.6) \text{ for } t = 0, \dots, 6,$$

$$Y|L(0), A \sim \text{Bern}\left(p = \left(\text{expit}\left(L(0) - 4\left(\frac{1}{7}\sum_{k=0}^6 A(t)\right)^2\right)\right)\right).$$

The target parameter was defined as the projection of $E_0(Y_d : d \in \mathcal{D})$ onto marginal structural working model $m_\beta(d) = \beta_0 + \beta_1 \sum_{t=0}^6 d(t)$ according to eq. (5), with weight function $h(d) = 1$, \mathcal{D} equal to the $2^7 = 128$ possible values of A , and $d(t)$ denoting the treatment level $a(t)$ assigned by a given rule d at time t .

Stratified and pooled TMLEs for β were implemented using estimators g_n and \bar{Q}_n based on intercept only logistic regressions; thus \bar{Q}_n was an inconsistent estimator of Q_0 . Table 1 shows estimated bias, bias to standard error ratio, variance, mean squared error, and 95% confidence interval coverage (using the variance estimator described in Section 3.7) based on 500 samples of size $n = 128$. Note that at this sample size many of the 128 rules of interest have no support in the data in a given sample, while the remainder have few observations available to fit ϵ in the stratified TMLE. As predicted by its double robust property, the pooled TMLE remains without meaningful bias despite use of a poor initial estimator \bar{Q}_n . In contrast, the

Table 1 Breakdown of stratified TMLE when some rules $d \in \mathcal{D}$ have no support

	Pooled TMLE	Stratified TMLE
Bias		
$\hat{\beta}_0$	0.1674	-0.9346
$\hat{\beta}_1$	-0.0563	0.2039
Bias/SE		
$\hat{\beta}_0$	0.1821	-2.3037
$\hat{\beta}_1$	-0.2326	2.0955
Variance		
$\hat{\beta}_0$	0.8451	0.1646
$\hat{\beta}_1$	0.0586	0.0095
MSE		
$\hat{\beta}_0$	0.8714	1.0377
$\hat{\beta}_1$	0.0616	0.0510
95% confidence interval coverage		
$\hat{\beta}_0$	88%	88%
$\hat{\beta}_1$	22%	33%

Note: True parameter values: $\beta_0 = 0.6214$ $\beta_1 = -0.4452$.

stratified TMLE has bias of approximately double the magnitude of its standard error, posing a substantial threat to valid inference. The stratified TMLE also exhibits markedly lower variance than its pooled counterpart, explained by the fact that for those rules d without support in the data, the stratified estimator uses an intercept only model to estimate \bar{Q}^d . Although unbiased, the pooled TMLE provides anti-conservative confidence interval coverage; we return to this point below.

This breakdown of the stratified TMLE in settings with no support for some rules of interest will not occur if the function h is chosen to give a weight of 0 to any rule (or more generally, any (d, V, t) combination) without support in the sample. For example, in the illustration above we could have defined $h(d) = P_0(A = d)$ and estimated it using the empirical distribution. Unless one is willing to assume that the MSM m_β is correctly specified, however, choice of h changes the target parameter being estimated [27]. Further, even with this choice of weight function, the estimators may still exhibit different finite sample performance in setting with marginal data support. We investigate this possibility further in the following section.

6 Simulation study

6.1 Overview

In this section, we investigate the relative performance of the pooled TMLE, stratified TMLE, and IPW estimators for the parameters of a marginal structural working model. For each candidate estimator, we report bias, variance, MSE, and 95% confidence interval coverage estimates based on influence curve variance estimators. We note that our influence curve-based estimators assume the weight function $h(d, t)$ is known; if the weight function is estimated, the influence curve should be corrected for this additional estimated component. We investigate two basic data generating processes. Simulation 1 investigates a simple process, in which the effect of the longitudinal treatment (time to switch) is confounded by baseline variables only, the outcome is observed at a single time point, and there is no censoring. Simulation 2 introduces more realistic complexity, designed to resemble the data analysis presented in the following section.

6.2 Simulation 1: baseline confounding only

6.2.1 Data generating process

We implemented a simulation with observed data consisting of n i.i.d. copies of $O = (L(0), A = (A(0), \dots, A(K)), Y)$, where $L(0)$ is a baseline covariate, $A(t)$ is a binary treatment assigned at time point $t = 0, \dots, K$, and Y is a binary outcome. The data for a given individual were generated by drawing sequentially from the following distributions:

$$L(0) \sim N(0, 1),$$

$$A(t)|L(0) \sim \text{Bern}(p = \max(\min(L(0) + 0.5, 0.62), 0.38)) \text{ for } t = 0, \dots, K,$$

$$Y|L(0), A \sim \text{Bern}\left(p = \left(\text{expit}\left(L(0) - \frac{1}{K+1} \sum_{t=0}^K A(t)\right)\right)\right).$$

In order to investigate the impact of decreasing support in the data (practical violations or near violations of the positivity assumption), we considered two versions of this data generating process, with $K = 2$ (Simulation 1a, lower bound on g_0 of 0.05) and $K = 6$ (Simulation 1b, lower bound on g_0 of 0.001).

6.2.2 Target parameter

The target parameter was defined as the projection of $E_0(Y_d : d \in \mathcal{D})$ onto marginal structural working model $m_\beta(d) = \beta_0 + \beta_1 \sum_{t=0}^K d(t)$ according to eq. (5), with \mathcal{D} equal to the $2^{(K+1)}$ possible values of A , $d(t)$ equal to the treatment level assigned by a given rule d at time, t and weight function $h(d) = P_0(A = d)$. In the case that some rules of interest had no support in a given sample, this choice of weight function (when estimated as the empirical proportion of subjects that followed rule d) ensured that the IPW estimator remained defined and that the updated fit Q_n^* used by the stratified TMLE solved the efficient influence curve equation when g_n was a consistent estimator of g_0 .

6.2.3 Estimators

This is a static point treatment problem, and thus a number of additional estimators are available. However, we use this as a special case of longitudinal dynamic MSMs and investigate the relative performance of three estimators described in Section 5: the pooled TMLE, the stratified TMLE, and the standard IPW estimator. All estimators were implemented using two estimators of g_0 : an estimator based on a correctly specified model for the conditional distribution of A given $L(0)$ and an estimator using an intercept only model. The estimators g_n were bounded from below at 0.001. TMLEs were implemented using two estimators of \bar{Q}_0 : an estimator based on main terms logistic regression models using the correct set of parents for a given node as independent variables (in a slight abuse, we refer to these as “correctly specified”) and an estimator using intercept only models. Performance was evaluated across 500 samples of size $n = 500$; 95% confidence interval coverage was based on the variance estimator described in Section 3.7.

6.2.4 Results

Results for Simulation 1a are shown in Table 2. When both \bar{Q}_n and g_n were based on correctly specified models, all three estimators were unbiased, had similar variance, and achieved close to nominal coverage. Table 2 also demonstrates double robustness; when \bar{Q}_n and g_n were based on a misspecified model, both

Table 2 Simulation 1a

	<i>g</i> and <i>Q</i> correct			<i>g</i> correct and <i>Q</i> incorrect			<i>g</i> incorrect and <i>Q</i> correct		
	PI TMLE	Str TMLE	IPW	PI TMLE	Str TMLE	IPW	PI TMLE	Str TMLE	IPW
Bias									
$\hat{\beta}_0$	-0.0047	-0.0037	0.0094	-0.0260	-0.0352	0.0094	-0.0083	-0.0108	-0.4578
$\hat{\beta}_1$	0.0024	0.0020	-0.0038	0.0178	0.0238	-0.0038	0.0056	0.0074	0.2994
Bias/SE									
$\hat{\beta}_0$	-0.0258	-0.0204	0.0453	-0.1450	-0.1953	0.0453	-0.0485	-0.0637	-2.4870
$\hat{\beta}_1$	0.0226	0.0197	-0.0336	0.1725	0.2306	-0.0336	0.0578	0.0767	2.8280
Variance									
$\hat{\beta}_0$	0.0334	0.0326	0.0427	0.0321	0.0325	0.0427	0.0295	0.0289	0.0339
$\hat{\beta}_1$	0.0109	0.0106	0.0130	0.0106	0.0107	0.0130	0.0093	0.0093	0.0112
MSE									
$\hat{\beta}_0$	0.0333	0.0325	0.0427	0.0327	0.0336	0.0427	0.0295	0.0290	0.2434
$\hat{\beta}_1$	0.0109	0.0105	0.0130	0.0109	0.0112	0.0130	0.0094	0.0093	0.1008
95% CI Coverage									
$\hat{\beta}_0$	95%	95%	97%	96%	96%	97%	97%	96%	29%
$\hat{\beta}_1$	95%	95%	96%	96%	95%	96%	97%	97%	17%

Notes: $g_0 > 0.05$. True parameter values: $\beta_0 = 0.0012$ $\beta_1 = -0.2771$; PI TMLE: Pooled TMLE; Str TMLE: Stratified TMLE.

TMLEs remained without meaningful bias. In this simulation, both TMLEs continued to achieve close to nominal coverage even when g_n was an inconsistent estimator of g_0 . In contrast and as expected, the IPW estimator was substantially biased with poor coverage when g_n was based on an intercept only model.

Results for Simulation 1b, with both Q_n and g_n based on correctly specified models, are shown in Table 3. In this simulation, in which the lower bound for g_0 is 0.001, the IPW estimator was minimally biased and retained good 95% confidence interval coverage. Interestingly, the performance of the stratified TMLE suffered in this setting, with bias approximately equal to the standard error, and 95% confidence interval coverage of 83% and 80% for β_0 and β_1 , respectively. The pooled TMLE remained unbiased and retained good confidence interval coverage.

Table 3 Simulation 1b

	Pooled TMLE	Stratified TMLE	IPW
Bias			
$\hat{\beta}_0$	-0.0025	-0.2622	-0.0486
$\hat{\beta}_1$	-0.0007	0.0744	0.0161
Bias/SE			
$\hat{\beta}_0$	-0.0092	-1.0907	-0.1618
$\hat{\beta}_1$	-0.0103	1.1504	0.1991
Variance			
$\hat{\beta}_0$	0.0716	0.0578	0.0904
$\hat{\beta}_1$	0.0051	0.0042	0.0065
MSE			
$\hat{\beta}_0$	0.0715	0.1264	0.0926
$\hat{\beta}_1$	0.0051	0.0097	0.0068
95% CI coverage			
$\hat{\beta}_0$	98%	83%	97%
$\hat{\beta}_1$	97%	80%	96%

Notes: $g_0 > 0.001$; Both g and Q are correct; True parameter values: $\beta_0 = 0.0014$ $\beta_1 = -0.1187$.

6.3 Simulation 2: resembling data analysis

In this simulation, we used a data generating process designed to resemble the data analysis presented in the following section, in which the goal was to investigate the effect of delayed switch to a new antiretroviral regimen on mortality among HIV-infected patients who have failed first line therapy. The data generating process thus contains both baseline and time-dependent confounders of a longitudinal binary treatment (time to switch), a repeated measures binary outcome (survival over time), and informative right censoring due to two causes (database closure and loss to follow up).

6.3.1 Data generating process

We implemented a simulation with observed data consisting of n i.i.d. copies of

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

for $K = 9$. Here, $L(0) = (W, CD4(0))$ and $L(t) = (Y(t), CD4(t))$, where W was a non-time-varying baseline covariate ($W = (W_1, \dots, W_4)$, representing baseline age, sex, and disease stage), $CD4(t)$ was a time-varying covariate representing most recent measured CD4 count at time t (square root transformed), and $Y(t)$ was an indicator of death by time t . The intervention nodes for a given time point t were $A(t) = (C_1(t), C_2(t), A_1(t))$, where $C_1(t)$ was an indicator of database closure by time t , $C_2(t)$ was an indicator of loss to follow up by time t , and $A_1(t)$ was an indicator of having switched to second line therapy by time t . In brief, the data for a given individual were generated by first drawing baseline characteristics W , then for each time point t , for as long as the subject remained alive and uncensored,

1. Drawing a time updated CD4 count $CD4(t)$ given W , prior CD4 counts, and regimen status at the prior time point ($A_1(t-1)$)
2. Determining censoring due to database closure $C_1(t)$ using a Bernoulli trial with probability dependent on W .
3. If still uncensored, determining censoring due to loss to follow up $C_2(t)$ using a Bernoulli trial with probability dependent on W , prior CD4 count and regimen status at the prior time point ($A_1(t-1)$).
4. If still uncensored and not yet switched, determining switching using a Bernoulli trial with probability dependent on W and prior CD4 count.
5. Determining death using a Bernoulli trial with probability dependent on W , prior CD4 counts and regimen status $A(t)$

The full data generating process is provided in Appendix D. Coefficients in the data generating process were chosen to approximate the degree of censoring, treatment, and death in the analysis data set, as detailed in Appendix D, Tables 7 and 9. The data generating process results in a true intervention mechanism g_0 not bounded away from 0.

6.3.2 Target parameter

The target parameter was defined as the projection of the counterfactual survival curve for each switch time, $(E(Y_d(t)) : d \in \mathcal{D}, t = 1, \dots, 10)$ onto marginal structural working model $\text{Logit } m_\theta(d, t) = \beta_0 + \beta_1 t + \beta_2(d(t-1)(t-s_d))$ according to eq. (5), where we use $d(t)$ to denote the value $a(t)$ assigned by rule d , s_d to denote the switch time assigned by rule d , and with \mathcal{D} consisting of each possible switch time, $\{0, \dots, 10\}$ combined with an intervention to prevent censoring. We used the following weight function:

$$h(d, t) = \frac{P_0(\bar{A}(t-1) = \bar{a}(t-1))}{n_t} I(t < t^*),$$

where n_t was the number of unique values $\bar{a}(t-1)$ compatible with $d \in \mathcal{D}$, t^* is the first time point at which all subjects have either died or been censored, and both $P_0(\bar{A}(t-1) = \bar{a}(t-1))$ and t^* were estimated with their empirical counterparts. This weight function gave 0 weight to rules without support in a given sample. It further avoided up-weighting specific values $\bar{a}(t-1)$ proportional to the number of rules (assigned switch times) that they were compatible with.

6.3.3 Estimators

We implemented the pooled TMLE, the stratified TMLE, the IPW estimator based on estimating $E(Y^d(t) : d \in \mathcal{D})$ using the bounded Horvitz–Thompson estimator and projecting the resulting estimates onto the model m_β (referred to as “Stratified IPW”) and the standard IPW estimator for dynamic MSM (referred to as a “Standard IPW” estimator).

The intervention rules $d \in \mathcal{D}$ could be considered static, in that they assign a fixed vector of treatment decisions \bar{a} irrespective of a subject’s covariate values. However, when the target parameter is defined using a MSM on survival, a static IPW estimator cannot be implemented in the standard way (fitting a weighted pooled regression of $Y(t)$ on observed treatment history $\bar{A}(t-1)$) because the full $\bar{A}(t-1)$ is not observed for subjects who die before time t . As noted by Picciotto et al. [42], one option, adopted by us here, is to instead define the interventions of interest as dynamic (switch at time s_d if still alive).

All estimators were implemented using an estimator g_n based on a correctly specified parametric model for g_0 , but bounding the resulting estimates from below at 0.01 in order to ensure that the denominator in the covariate used in the updating step remained bounded away from 0. The TMLEs were implemented using estimators of \bar{Q}_0 based on main terms logistic regression models using the correct set of parents for a given node as independent variables (not equivalent to use of a correctly specified parametric model to estimate \bar{Q}_0). The performance of each estimator was evaluated across 500 samples of size $n = 2,627$, corresponding to the sample size in the data analysis. 95% confidence interval coverage was based on the variance estimator described in Section 3.7; calculation of non-parametric bootstrap-based coverage was computationally prohibitive.

6.3.4 Results

Results for Simulation 2 are shown in Table 4. In this simulation, both the pooled and the stratified TMLEs were essentially unbiased for all coefficients, and the two TMLEs had comparable MSEs. Both TMLEs exhibited less than nominal 95% confidence interval coverage when using influence curve-based variance estimators. The anti-conservative performance of the influence curve-based variance estimator is likely due to the presence of practical positivity violations and relatively rare outcomes; the fact that the weight function was treated as known may also make a small contribution. Further work is needed to develop improved diagnostics and variance estimators in these settings.

In contrast, both IPW estimators were substantially biased for β_2 , which reflected the treatment effect, despite use of an estimator g_n based on a correctly specified parametric model. Both IPW estimators also showed higher MSE for β_2 , and achieved 95% confidence interval coverage for β_2 substantially below that of the TMLEs. This finding is consistent with the known susceptibility of IPW estimators to positivity violations and data sparsity, exacerbated by the limited ability when using a dynamic regime formulation to choose an effectively stabilizing weight function. Across simulations, the median minimum value of g_n used by IPW

Table 4 Simulation 2: resembling data analysis

	PI TMLE	Str TMLE	Stan IPW	Str IPW
Bias				
$\hat{\beta}_0$	0.0000	0.0060	-0.0239	-0.0501
$\hat{\beta}_1$	-0.0041	-0.0040	0.0179	0.0309
$\hat{\beta}_2$	-0.0006	0.0107	0.0984	0.0880
Bias/SE				
$\hat{\beta}_0$	0.0000	0.0218	-0.0867	-0.1829
$\hat{\beta}_1$	-0.0849	-0.0827	0.4118	0.7046
$\hat{\beta}_2$	-0.0096	0.1654	2.0741	1.8708
Variance				
$\hat{\beta}_0$	0.0738	0.0750	0.0758	0.0751
$\hat{\beta}_1$	0.0024	0.0024	0.0019	0.0019
$\hat{\beta}_2$	0.0037	0.0042	0.0022	0.0022
MSE				
$\hat{\beta}_0$	0.0736	0.0749	0.0762	0.0775
$\hat{\beta}_1$	0.0024	0.0024	0.0022	0.0029
$\hat{\beta}_2$	0.0037	0.0043	0.0119	0.0099
95% CI coverage				
$\hat{\beta}_0$	92%	92%	93%	94%
$\hat{\beta}_1$	82%	80%	90%	89%
$\hat{\beta}_2$	80%	77%	35%	42%

Notes: True parameter values: $\beta_0 = -4.9525$ $\beta_1 = 0.2962$ $\beta_2 = -0.2168$; PI TMLE: Pooled TMLE; Str TMLE: Stratified TMLE; Stan IPW: Standard IPW; Str IPW: Stratified IPW.

prior to bounding at 0.01 was 0.000297; 1.55% of values of g_n used by IPW were less than 0.01. Tables 8 and 10, provide further details on data support and number of events.

7 Data analysis

We analyzed data from the International Epidemiological Databases – Southern Africa in order to investigate the effect of switching to second line therapy on mortality among HIV-infected patients with immunological failure on first line antiretroviral therapy. The data set and clinical question are described in detail in Gsponer et al. [1]. In brief, data were drawn from clinical care facilities in Zambia and Malawi, in which HIV-infected patients were followed longitudinally in clinic and data were collected on baseline demographic and clinical variables (sex, age, and baseline disease stage), time-varying CD4 count, and time-varying treatment, summarized here as switch to second line therapy. Death was independently reported. The 2,627 subjects meeting WHO immunological failure criteria were included in the current analysis beginning at time of immunologic failure. Following common practice and prior analysis, time was discretized into 3-month intervals; time updated CD4 count was coded such that CD4 count for an interval preceded switching decisions in that interval. Data on a subject were censored at time of database closure or after four consecutive intervals without clinical contact.

The data structure and target parameter were identical to those described in Simulation 2, with $W = (W_1, W_2, W_3, W_4)$, W_1 equal to sex, W_2 and W_3 representing two levels of a three-level categorical age variable (< 30, 30 – 39, and > 39), and W_4 equal to disease stage. The analysis was implemented under the causal model assumed for Simulation 2; in particular assuming that monitoring times did not affect the outcome other than via effects on switching. We acknowledge that this assumption may not hold; however, relaxing it introduces a number of additional complications, as described in the Appendix. We implemented

the estimators described in Simulation 2: pooled and stratified TMLEs and standard and stratified IPW estimators. Estimators of g_0 and \bar{Q}_0 were based on main term logistic regression models, analogous to Simulation 2. Given the results in Simulation 2 suggesting the poor performance of the influence curve-based variance estimator, we also estimated the variance using a non-parametric bootstrap.

Results are given in Table 5. The IPW estimates for the effect of switching on mortality (β_2) are close to zero and non-significant. Both TMLE point estimates suggest a 0.88 relative odds of death per 3-month earlier switch, and all except for the stratified TMLE combined with bootstrap-based variance estimation were significant at the $\alpha = 0.05$ level. Such a protective effect of switching is consistent with clinical knowledge. Interestingly, these results appear consistent with those of Simulation 2, which suggested that the IPW estimator was substantially positively biased, underestimating the harm of delayed switch, while both TMLEs performed well in terms of bias. In summary, our results in both the simulation and data analysis are consistent with the TMLEs controlling for measured confounders more completely than the corresponding IPW estimators.

Table 5 Data analysis

	Estimate	95% CI ¹	95% CI ²	95% CI ³	95% CI ⁴
Pooled TMLE					
$\hat{\beta}_0$	-4.8291	-5.2235	-4.4348	-5.2157	-4.5436
$\hat{\beta}_1$	0.1874	0.1390	0.2357	0.1499	0.2291
$\hat{\beta}_2$	-0.1246	-0.2275	-0.0217	-0.2259	-0.0496
Stratified TMLE					
$\hat{\beta}_0$	-4.8382	-5.2357	-4.4407	-5.3582	-4.5790
$\hat{\beta}_1$	0.1883	0.1396	0.2370	0.1560	0.2876
$\hat{\beta}_2$	-0.1262	-0.2303	-0.0221	-0.2110	0.0184
Standard IPTW					
$\hat{\beta}_0$	-4.9166	-5.3136	-4.5196	-5.2692	-4.6413
$\hat{\beta}_1$	0.2053	0.1587	0.2519	0.1673	0.2429
$\hat{\beta}_2$	0.0238	-0.0696	0.1171	-0.0575	0.0900
Stratified IPTW					
$\hat{\beta}_0$	-5.0811	-5.5020	-4.6601	-5.5260	-4.8252
$\hat{\beta}_1$	0.2538	0.1970	0.3105	0.2189	0.3274
$\hat{\beta}_2$	0.0153	-0.0539	0.0844	-0.0374	0.0784

Notes: CI¹: Lower 95% influence curve-based confidence interval. CI²: Upper 95% influence curve-based confidence interval. CI³: Lower 95% bootstrap confidence interval. CI⁴: Upper 95% bootstrap confidence interval.

The poor coverage observed in Simulation 2, despite absence of bias for the TMLEs, suggests that the influence curve-based variance estimators may be systematically underestimating the true variance in this analysis. While the non-parametric bootstrap offers an alternative approach, it is not expected to resolve the challenge of anti-conservative variance estimation in the setting of practical positivity violations. Intuitively, rare treatment/covariate combinations, despite being theoretically possible, may simply not occur in a given finite sample and as a result, the corresponding extreme weights implied by these combinations will not occur. Because the non-parametric bootstrap resamples from the same finite sample, it fails to address the underlying problem. Indeed, the bootstrap-based confidence intervals in the data analysis were slightly smaller than confidence intervals based on the influence curve. Thus in this realistic setting of rare outcomes and moderately strong confounding, our results caution against reliance on either approach to variance estimation, for either IPW or TML estimators, and suggest that additional work developing robust variance estimators in this setting is urgently needed.

In addition to the issues raised above, limitations of the analysis include the potential for unmeasured confounding by factors such as unmeasured health status and adherence, as well as bias due to incomplete

health reporting, resulting in censoring due to loss to follow up that directly depends on death [43, 44]. These results also do not contradict the previous published results of Gsponer et al. [1], which found a protective effect of switching using a static IPW estimator for a hazard MSM; such an IPW estimator might perform substantially better than the dynamic IPW estimator for a survival MSM implemented here.

8 Discussion

In summary, we have presented a pooled TMLE for the parameters of static and dynamic longitudinal MSMs that builds on prior work by Robins [13, 29] and Bang and Robins [22]. We evaluated the performance of this estimator using simulated data and applied it, together with alternatives, in a data analysis. Both theory and simulations suggest settings in which the pooled TMLE offers advantages over alternative estimators. Software implementing this estimator, together with competitors, is included in supplementary files and is publicly available as part of R library `ltmle` (<http://cran.r-project.org/web/packages/ltmle/>).

The pooled TMLE presented in this paper, together with corresponding open source software, provides a new tool for estimation of the parameters of static or dynamic MSMs. It has clear theoretical advantages over available alternatives. Unlike IPW and augmented-IPW estimators, it is a substitution estimator. Unlike IPW estimators, it is double robust and asymptotically efficient depending on the initial estimators of g and Q . Unlike the previously proposed stratified TMLE, it does not require support in the data for every rule of interest and remains defined in the case that the target parameter is defined using a marginal structural working model conditional on a continuous baseline covariate.

In settings where some subset of discrete intervention rules has adequate support in a given sample, an alternative approach is to compare only this subset of rules [2]. However, in many settings smoothing over a large number of poorly supported rules is appropriate. For example, for a set of rules indexed by a continuous or multiple level ordered variable, smoothing over this variable provides a way to define a causal effect of interest despite inadequate support to estimate the counterfactual outcome under any of these rules individually.

The TMLE presented in this paper was developed for a causal model in which the non-intervention variables may be a function of the entire observed past ($Pa(L(k)) = (\bar{A}(k-1), \bar{L}(k-1))$) and the intervention variables may be a function of some subset of the observed past ($Pa(A(k)) \subseteq (\bar{A}(k-1), \bar{L}(k))$) – in other words, for a model in which exclusion restrictions were assumed, if at all, only for the intervention variables. In some cases, a causal model that also restricts the parent set of the non-intervention variables to a subset of the observed past may be appropriate. This smaller model is included in the larger model assumed in the current paper. As a result, while the TMLE developed for the larger model will still be valid, it will no longer be efficient.

Our simulations suggest that both stratified and pooled TMLEs may outperform both the IPW estimator typically used for dynamic MSMs, as well as an alternative “stratified” IPW estimator, in some settings with sparse data/near positivity violations. However, further work is needed to confirm this preliminary observation. Although the theory in this paper was developed for the general case of dynamic MSMs, including models of the time-specific hazard or survival functions, we focused our software implementation, examples, and simulations on MSMs for survival. The practical performance of the pooled TMLE relative to alternative estimators also remains to be investigated for the case of static and dynamic MSMs on the hazard. It also remains to implement and evaluate the relative performance of the alternative pooled TMLE described in Appendix C, in which the updating step pools not only over all rules of interest but also over all time points. Finally, the relative performance of the TMLE compared to double robust efficient estimating equation-based estimators for longitudinal MSM parameters, including those of Robins [13, 29] and Bang and Robins [22], remains to be evaluated.

Importantly, our simulations and data analysis illustrate the need for improved variance estimators for both TMLE and IPW in settings with moderately strong confounding, multiple time points, and relatively

rare outcomes. Improved approaches to variance estimation and valid inference, as well as appropriate diagnostics to warn applied practitioners of settings in which coverage is likely to be poor, are crucial research priorities.

Acknowledgments: This work was supported by NIH Grant #U01AI069924 (NIAID, NICHD, and NCI) (PIs: Egger and Davies), the Doris Duke Charitable Foundation Grant #2011042 (PI: Petersen), and NIH Grant #R01AI074345-06 (NIAID) (PI: van der Laan).

Appendix A

Notation

Table 6 Partial list of notation

Variable	Description
Observed data	
$L(t)$	covariates at time t
$Y(t) \subseteq L(t)$	outcome at time t
$A(t) = (A_1(t), A_2(t))$	intervention at time t
$A_1(t)$	treatment at time t
$A_2(t)$	right censoring at time t
$\bar{A}(k) = A(0), \dots, A(k)$	history of intervention nodes $t = 0, \dots, k$
$\bar{L}(k) = L(0), \dots, L(k)$	history of non-intervention nodes $t = 0, \dots, k$
$O = (L(0), A(0), \dots, L(K), A(K), L(K+1))$	observed data structure
$Pa(L(k)) = \bar{L}(k-1), \bar{A}(k-1)$	parents of non-intervention nodes $L(k)$
$Pa(A(k)) \subseteq \bar{L}(k), \bar{A}(k-1)$	parents of intervention nodes $A(k)$
Counterfactuals	
\mathcal{D}	set of dynamic intervention rules
$d \in \mathcal{D}$	intervention rule
L_d	counterfactual L under rule d
Y_d	counterfactual Y under rule d
Statistical counterparts	
L^d	random variable with distribution P_0^d equal to L_d in distribution under sequential randomization assumption (SRA) (2)
Y^d	outcome-component of L^d
Marginal structural working model	
$m_{\beta}(d, t, V)$	working model for $(E_0(Y_d(t) V) : d \in \mathcal{D}, t \in \tau)$
τ	index set of time points (often $\tau = 1, \dots, K+1$)
V	“effect modifier” of interest, a function of $L(0)$
Distributions	
$P_{O,U,0}$	distribution of (O, U) defined by structural causal model (SCM)
P_0	distribution of O
$P_{d,0}$	distribution of counterfactual L_d
P_0^d	G -computation formula for post-intervention distribution of L under rule d equal to $P_{d,0}$ under SRA (2)
$Q_{L(k),0}$	distribution of $L(k) Pa(L(k))$
$Q_{0:k} \equiv \prod_{j=0}^k Q_{L(j),0}$	
$Q_0 \equiv Q_{0:K+1}$	non-intervention factor of $P_0 = Q_0 g_0$
$g_{A(k),0}$	distribution of $A(k) Pa(A(k))$

(continued)

Table 6 (Continued)

Variable	Description
$g_{0:k,0} \equiv \prod_{j=0}^k g_{A(j),0}$	
$g_0 \equiv g_{0:k}$	intervention mechanism factor of $P_0 = Q_0 g_0$
$Q_{L(0),n}$	empirical distribution of $L_i(0), i = 1, \dots, n$
Models	
\mathcal{M}^F	SCM for $P_{0,U,0}$ defines causal assumptions: $P_{0,U,0} \in \mathcal{M}^F$
\mathcal{M}	statistical model for P_0 defines statistical assumptions: $P_0 \in \mathcal{M}$
Functions and parameters	
$\Psi^F : \mathcal{M}^F \rightarrow \mathbb{R}^J$	causal target parameter mapping
$\Psi^F(P_{0,U,0})$	causal target parameter value (causal quantity)
$\Psi : \mathcal{M} \rightarrow \mathbb{R}^J$	statistical target parameter mapping
$\Psi(P_0) = \Psi(\bar{Q}_0, Q_{L(0),0}) = \psi_0 = \beta_0$	target parameter value equals $\Psi^F(P_{0,U,0})$ under SRA (2) and positivity (3)
$\bar{L}(k)$	structural equation $(Pa(L(k)), U_{L(k)}) \mapsto L(k)$
$\bar{A}(k)$	structural equation $(Pa(A(k)), U_{A(k)}) \mapsto A(k)$
d_k	dynamic intervention $\bar{L}(k) \mapsto A_k$
h	weight function $(d, t, V) \mapsto \mathbb{R}$
h_1	$h(d, t, V) \frac{\frac{d}{dp} m_\beta(d, t, V)}{m_\beta(1 - m_\beta)}$ for logistic working model m_β
h_1	$h(d, t, V) \frac{d}{dp} m_\beta(d, t, V)$ for linear working model m_β
Efficient influence curve	
$D^*(P) = D^*(\bar{Q}, Q_{L(0)}, g)$	efficient influence curve of Ψ at P equals influence curve of TMLE $\Psi(Q_n^*)$ when Q_n^*, g_n are both consistent
Least favorable submodels	
$\{\bar{Q}_k^{d,t}(\epsilon, g) : \epsilon\}$	submodel through $\bar{Q}_k^{d,t}$ at parameter value $\epsilon = 0$ used to update initial estimate of $\bar{Q}_k^{d,t}$ in TMLE
Loss functions	
$\mathcal{L}_{d,t,k,\bar{Q}_{k+1}^{d,t}}(\bar{Q}_k^{d,t})$	loss function for $\bar{Q}_k^{d,t}$ relies on estimator of previous $\bar{Q}_{k+1}^{d,t}$ used to fit ϵ of least favorable model $\bar{Q}_k^{d,t}(\epsilon, g)$
Others	
$\bar{Q}_{L(k)}^{d,t} = \bar{Q}_k^{d,t}$	$E_P(Y^d(t) \bar{L}^d(k-1))$
$\bar{Q} \equiv (\bar{Q}_k^{d,t} : d \in \mathcal{D}, t \in \tau, k = 1, \dots, t)$	
$\bar{Q}_1 = (\bar{Q}_{L(1)}^{d,t} : d \in \mathcal{D}, t \in \tau)$	$(E_P(Y^d(t) L(0)) : d \in \mathcal{D}, t \in \tau)$
Q	redefined as $(\bar{Q}_1, Q_{L(0)})$ when used in $\Psi(Q)$
Estimators	
g_n	estimator of g_0
Q_n	initial estimator of Q_0
\bar{Q}_n	initial estimator of \bar{Q}_0
$Q_n^* = (\bar{Q}_{1,n}^*, Q_{L(0),n})$	updated estimator of Q_0 , targeted at Ψ
$\Psi(Q_n^*) = \psi_n^*$	TMLE of ψ_0

Appendix B

The efficient influence curve of the statistical target parameter

We have the following theorem for a general parameter $\Psi(P) = f(\bar{Q}_{L(1)} = (\bar{Q}_{L(1)}^{d,t} : d, t), Q_{L(0)})$.

Theorem 1 Consider a parameter $\Psi : \mathcal{M} \rightarrow \mathbb{R}^J$ that can be represented as $\Psi(P) = f(\bar{Q}_{L(1)}, Q_{L(0)})$, where $\bar{Q}_{L(1)} = (\bar{Q}_{L(1)}^{d,t} = E(Y^d(t) | L(0)) : d, t)$. Let $Q = (\bar{Q}_{L(1)}, Q_{L(0)})$. Assume that $f(\cdot)$ is such that Ψ is pathwise differentiable. Let $f'_{d,t}(Q)(w) = \frac{d}{d\bar{Q}_{L(1)}^{d,t}(w)} f(\bar{Q}_{L(1)}, Q_{L(0)})$ be the partial derivative of f with respect to

$\bar{Q}_{L(1)}^{d,t}(w) = E(Y^d(t)|L(0) = w)$ at Q . Let $D_{L(0)}^*(Q)$ be the influence curve of $f(\bar{Q}_{L(1)}, Q_{L(0),n})$ as an estimator of $f(\bar{Q}_{L(1)}, Q_{L(0)})$, where $Q_{L(0),n}$ is the empirical distribution of $L_i(0)$, $i = 1, \dots, n$.

Then, the efficient influence curve of Ψ at P can be represented as follows:

$$D^*(P) = D_{L(0)}^*(Q) + \sum_{d,t} f'_{d,t}(L(0)) \sum_{k=1}^t \frac{I(\bar{A}(k) = \bar{d}_k(\bar{L}(k)))}{g_{0:k}} (\bar{Q}_{L(k+1)}^{d,t} - \bar{Q}_{L(k)}^{d,t})$$

Proof: The efficient influence curve of the parameter $E(Y^d(t)|L(0) = w)$ (assuming discrete random variable $L(0)$) is given by

$$D_{d,t,w}^* = \frac{I(L(0) = w)}{Q_{L(0)}(w)} \sum_{k=1}^t \frac{I(\bar{A}(k) = \bar{d}_k(\bar{L}(k)))}{g_{0:k}} (\bar{Q}_{L(k+1)}^{d,t} - \bar{Q}_{L(k)}^{d,t})$$

(appendix A3, van der Laan and Rose [20]). By the delta-method, the efficient influence curve of $f(\bar{Q}_{L(1)}, Q_{L(0)})$ is thus given by

$$\begin{aligned} D^* &= D_{L(0)}^* + \sum_{w,d,t} f'_{d,t}(w) D_{d,t,w}^* \\ &= D_{L(0)}^* + \sum_{w,d,t} f'_{d,t}(w) \frac{I(L(0) = w)}{Q_{L(0)}(w)} \sum_{k=1}^t \frac{I(\bar{A}(k) = \bar{d}_k(\bar{L}(k)))}{g_{0:k}} (\bar{Q}_{L(k+1)}^{d,t} - \bar{Q}_{L(k)}^{d,t}) \\ &= D_{L(0)}^* + \sum_{d,t} f'_{d,t}(L(0)) \sum_{k=1}^t \frac{I(\bar{A}(k) = \bar{d}_k(\bar{L}(k)))}{g_{0:k}} (\bar{Q}_{L(k+1)}^{d,t} - \bar{Q}_{L(k)}^{d,t}) \end{aligned}$$

(appendix A3, van der Laan and Rose [20]). This completes the proof. \square

In order to determine the partial derivative of the function f and $D_{L(0)}^*$ the following is useful. Suppose that, as in our examples, $f(Q) = \arg \max_{\beta} M(\beta, Q)$ for some function M , and suppose that $Q = (\bar{Q}, Q_{L(0)})$. Then $\beta(Q) = f(Q)$ solves the equation $0 = \frac{d}{d\beta} M(\beta, Q) \equiv U(\beta, Q)$. By the implicit function theorem we have that

$$\frac{d}{dQ} \beta(Q) = - \left\{ \frac{d}{d\beta} U(\beta, Q) \right\}^{-1} \frac{d}{dQ} U(\beta, Q).$$

In particular,

$$\frac{d}{dQ} \beta(\bar{Q}, Q_{L(0)}) = - \left\{ \frac{d}{d\beta} U(\beta, Q) \right\}^{-1} \frac{d}{dQ} U(\beta, Q),$$

and

$$\frac{d}{dQ_{L(0)}} \beta(Q) = - \left\{ \frac{d}{d\beta} U(\beta, Q) \right\}^{-1} \frac{d}{dQ_{L(0)}} U(\beta, \bar{Q}, Q_{L(0)}).$$

We can now apply our general Theorem 1 to the example with $Y(t) \in [0, 1]$ and the logistic regression working MSM in which $\beta(Q)$ solves the equation $U(\beta, \bar{Q}_{L(1)}, Q_{L(0)})$ with

$$U(\beta, \bar{Q}_{L(1)}, Q_{L(0)}) \equiv E_{Q_{L(0)}} \sum_t \sum_{d \in \mathcal{D}} h_1(d, t, V) (\bar{Q}_{L(1)}^{d,t}(L(0)) - m_{\beta}(d, t, V)).$$

The same equation applies for the linear working MSM but with the other definition of h_1 as mentioned above. Define

$$c(Q) \equiv E_{Q_{L(0)}} \sum_{t,d} h_1(d, t, V) \frac{d}{d\beta} m_{\beta}(d, t, V).$$

Note that $-\frac{d}{d\beta}U(\beta, Q) = c(Q)$. Thus,

$$D_{L(0)}^*(Q) = c(Q)^{-1} \sum_{t,d} h_1(d, t, V)(\bar{Q}_{L(1)}^{d,t} - m_\beta(d, t, V)),$$

and

$$f'_{d,t}(L(0)) = c(Q)^{-1} h_1(d, t, V).$$

This proves the following corollary [13, 22, 29].

Corollary 1 Consider the target parameter $\Psi(Q)$ defined by eq. (5). This target parameter is pathwise differentiable at P with efficient influence curve given by

$$\begin{aligned} D^*(P) &= c(Q)^{-1} \sum_{t,d} h_1(d, t, V)(\bar{Q}_{L(1)}^{d,t} - m_\beta(d, t, V)) \\ &\quad + c(Q)^{-1} \sum_t \sum_{k=1}^t \sum_d h_1(d, t, V) \frac{I(\bar{A}(k) = \bar{d}_k(\bar{L}(k)))}{g_{0:k}} (\bar{Q}_{L(k+1)}^{d,t} - \bar{Q}_{L(k)}^{d,t}) \end{aligned}$$

The efficient influence curve is double robust. In other words we have that $-P_0 D^*(Q, g_0) = \Psi(Q) - \psi_0$, so that, in particular, if $\Psi(Q) = \psi_0$, then $P_0 D^*(Q, g_0) = 0$. As a consequence, our TMLE $\Psi(Q_n^*)$ will be a consistent estimator of ψ_0 if either Q_n^* is consistent for Q_0 or g_n is consistent for g_0 [38].

Appendix C

An alternative pooled TMLE that only fits a single ϵ to compute the update

The TMLE described in the main text relies on a separate ϵ for each $k = 1, \dots, t$ and for each $t = 1, \dots, K+1$ resulting in a collection of $\sum_{t=1}^{K+1} t$ estimators of ϵ that define the TMLE. A nice feature of this TMLE is that it exists in closed form. The following alternative TMLE only relies on fitting a single ϵ , but in this case the updating needs to be iterated until convergence. First construct an initial estimator \bar{Q}_n^0 of $\bar{Q}_0 = (\bar{Q}_{k,0}^{d,t} : k = 1, \dots, t, d \in \mathcal{D}, t = 1, \dots, K+1)$ as described above. Now, consider the above-presented submodel $(\bar{Q}_n^0(\epsilon) = (\bar{Q}_{k,n}^{d,t}(\epsilon, g) : \epsilon))$ through \bar{Q}_n^0 at $\epsilon = 0$. Compute

$$\epsilon_n^0 = \arg \min_{\epsilon} \sum_{t=1}^{K+1} \sum_{d \in \mathcal{D}} \sum_{k=1}^t \mathcal{L}_{d,t,k, \bar{Q}_{k+1,n}^{d,t,0}}(\bar{Q}_{k,n}^{d,t,0}(\epsilon, g_n)),$$

where the nuisance parameters of the loss function are estimated with the initial estimator \bar{Q}_n^0 . Note that ϵ_n^0 can be fit with a pooled logistic regression as stated above. This yields an update $\bar{Q}_n^1 = \bar{Q}_n^0(\epsilon_n^0)$. In general, at the m th step, given the estimator \bar{Q}_n^m , we compute

$$\epsilon_n^m = \arg \min_{\epsilon} \sum_{t=1}^{K+1} \sum_{d \in \mathcal{D}} \sum_{k=1}^t \mathcal{L}_{d,t,k, \bar{Q}_{k+1,n}^{d,t,m}}(\bar{Q}_{k,n}^{d,t,m}(\epsilon, g_n)),$$

and the resulting update $\bar{Q}_n^{m+1} = \bar{Q}_n^m(\epsilon_n^m, g_n)$. This updating process is iterated until $\epsilon_n^m \approx 0$. The resulting final update is denoted with \bar{Q}_n^* and is the TMLE of \bar{Q}_0 . By construction, we have that this TMLE also solves the efficient influence curve equation $P_n D^*(Q_n^*, g_n) = 0$ with arbitrary precision. The TMLE of ψ_0 is now computed with the corresponding plug-in estimator $\Psi(\bar{Q}_n^*, Q_{L(0),n})$, as above. The potential advantage of this alternative TMLE is that it is able to smooth across all time points t and k when computing the update, while the closed form TMLE presented above only smoothes over the rules $d \in \mathcal{D}$.

Appendix D

Supplementary material, Simulation 2

Overview

Below we describe the true data generating process used in Simulation 2 in greater detail. We also provide a summary table comparing the support and number of events over time in the simulated data and in the real data set it was designed to resemble.

In our presentation of the simulation below we have altered our notation slightly from that presented in the paper to match notation in the accompanying R code. Specifically, $L(t)$ is used to refer to time varying CD4 count $CD4(t)$. The observed data generated on a given subject consisted of

$$O = (W, Y(0), L(0), C_1(0), C_2(0), A_1(0), \dots, Y(9), L(9), C_1(9), C_2(9), A_1(9), Y(10))$$

Further, the data generating process included a non-monotone monitoring process (denoted $M(t)$) designed to mimic when subjects come into clinic, have their CD4 counts measured, and have an opportunity to switch regimens. This adds several additional complexities. First, observed CD4 count, denoted here as $L(t)$ (and in the main text as $CD4(t)$), is only updated to reflect the true underlying CD4 count process when a patient is seen. Below, $\hat{L}(t)$ is used to denote the true underlying CD4 value. Subsequent CD4 values and death are functions of this true underlying value, while the intervention nodes are functions of the observed values only. Because both switching and monitoring are generated only in response to the observed past, however, this time-dependent non-monotone monitoring process is a multivariate instrumental variable, warranting its exclusion from the adjustment set. Further, its inclusion would both be expected to harm efficiency and introduce positivity violations (for example, subjects not seen in clinic at a given time point have zero probability of switch at that time point). The non-monotone monitoring process, while retained to mimic the data analysis, is thus omitted from the observed data and presentation in the main text in order to simplify discussion. Second, in accordance with common practice in clinical cohort data, censoring due to loss to follow up C_2 is defined deterministically based on not being seen in clinic for a certain number of consecutive time points. Finally, a subject can only switch treatment when seen ($A_1(t)$ is only at risk of jumping when $M(t) = 1$). As above, $W = (W_1, W_2, W_3, W_4)$ and $Y(t)$ denotes an indicator of death by time t .

Data generating process

Data were generated for a given individual according to the following process, where ϵ_1 and ϵ_2 are draws from a standard normal distribution, and all binary variables were drawn from a Bernoulli distribution with the conditional probabilities given below. Data for a given subject were drawn sequentially until either $Y(t)$ jumped to one, $C_1(t)$ jumped to 1, $C_2(t)$ jumped to 1, or $Y(10)$ was generated. Tables 7 and 8 compare the number of deaths in the simulated data (median of 501 samples) and actual data among patients following a given regime.

$$P(W_1 = 1) = 0.3$$

$$P(W_2 = 1 | W_1 = 0) = 0.5$$

$$P(W_3 = 1) = 0.5$$

$$P(W_4 = 1) = 0.3$$

$$P(Y(t) = 1|W, \hat{L}(t-1), A_1(t-1)) =$$

$$\begin{cases} 0, & \text{if } t = 0 \\ -5.8 - 0.1W_1 - 0.1W_2 + 0.1W_3 - 0.2W_4 - 0.7\hat{L}(t-1) - 0.9A_1(t-1), & \text{if } t > 0 \end{cases}$$

$$\hat{L}(t) =$$

$$\begin{cases} \max(\min(\epsilon_1(t) - W_4, 4), -4), & \text{if } t = 0 \\ \max(\min(\epsilon_1(t) + 0.1W_1 - 0.1W_2 - 0.1W_3 - 0.5W_4 + 0.9\hat{L}(t-1) + A_1(t-1), 4), -4), & \text{if } t > 0 \end{cases}$$

$$P(M(t) = 1|W, L(t-1), A_1(t-1)) =$$

$$\begin{cases} 1, & \text{if } t = 0 \\ \text{expit}(0.4 + 0.1W_1 - 0.2W_2 + 0.3W_3 + 0.1W_4 - 0.1L(t-1) + 0.2A_1(t-1)), & \text{if } t > 0 \end{cases}$$

$$L(t) = \begin{cases} \hat{L}(t), & \text{if } M(t) = 1 \\ L(t-1), & \text{if } M(t) = 0 \end{cases}$$

$$P(C_1(t) = 1|W, L(0)) = 1 - \text{expit}(2 + 0.1W_1 + 0.2W_2 + 0.1W_3 + 0.1W_4 + 0.1L(0))$$

$$C_2(t) = I(M(t-2) = 0 \text{ and } M(t-1) = 0 \text{ and } M(t) = 0)$$

$$P(A_1(t) = 1|M(t), A_1(t-1), W, L(t)) =$$

$$\begin{cases} 1, & \text{if } t > 0 \text{ and } A_1(t-1) = 1 \\ 0, & \text{if } t > 0 \text{ and } A_1(t-1) = 0 \text{ and } M(t) = 0 \\ \text{expit}(-5 + 0.1W_1 + 0.1W_2 + 0.2W_3 + 0.2W_4 - 1.5L(t) + \epsilon_2(t)), & \text{otherwise} \end{cases}$$

Table 7 Events in data analysis

Time	Switch time										
	0	1	2	3	4	5	6	7	8	9	10
1	0	13									
2	1	0	8								
3	0	1	0	11							
4	0	0	0	0	5						
5	1	0	0	0	0	3					
6	0	0	0	0	0	0	4				
7	0	0	0	0	0	0	0	1			
8	0	0	0	0	0	0	0	0	1		
9	0	0	0	0	0	0	0	0	0	2	
10	0	0	0	0	0	0	0	0	0	0	3

Table 8 Events in simulated data

Time	Switch time										
	0	1	2	3	4	5	6	7	8	9	10
1	0	9									
2	0	1	10								
3	0	0	1	9							
4	0	0	0	1	7						
5	0	0	0	0	1	5					
6	0	0	0	0	0	0	4				
7	0	0	0	0	0	0	0	3			
8	0	0	0	0	0	0	0	0	2		
9	0	0	0	0	0	0	0	0	0	2	
10	0	0	0	0	0	0	0	0	0	0	1

Tables 9 and 10 compare the number of patients in the simulated data and actual data who are uncensored and following a given regime. In the data analysis, for example, in the first 3-month interval following immunologic failure (time = 1) there were no patient deaths among the 137 uncensored patients who switched immediately (switch time = 0) and 13 deaths among the 2,285 uncensored patients who did not

Table 9 Support (# uncensored and following rule) in data analysis

Time	Switch time											
	0	1	2	3	4	5	6	7	8	9	10	
1	137	2,285										
2	120	88	1,962									
3	104	80	71	1,703								
4	88	68	67	54	1,453							
5	77	59	57	51	61	1,184						
6	54	44	50	48	57	57	951					
7	40	37	43	48	54	55	51	760				
8	29	33	40	46	50	51	51	49	602			
9	17	29	34	41	44	48	51	49	50	508		
10	10	28	34	39	44	46	51	48	49	53	432	

Table 10 Support (# uncensored and following rule) in simulated data

Time	Switch time											
	0	1	2	3	4	5	6	7	8	9	10	
1	121	2,229										
2	107	118	1,890									
3	95	107	131	1,582								
4	82	94	120	127	1,237							
5	70	83	110	117	112	992						
6	61	74	97	108	102	101	798					
7	53	65	88	98	96	95	89	645				
8	47	58	80	91	90	90	85	82	527			
9	40	51	72	83	83	84	82	79	76	433		
10	35	47	65	77	77	79	76	75	74	72	358	

switch immediately (switch time = 1). In the second 3-month interval following immunologic failure (time = 2), there was one patient death among the 120 uncensored patients who switched immediately (switch time = 0), no patient deaths among the 88 uncensored patients who switched during the first 3-month interval (switch time = 1) and 8 deaths among the 1,962 uncensored patients who did not switch immediately or during the first 3-month interval (switch time = 2).

References

1. Gsponer T, Petersen M, Egger M, Phirid S, Maathuis M, Boule A, et al., and O. Keiser for IeDEA Southern Africa. The causal effect of switching to second line ART in programmes without access to routine viral load monitoring. *AIDS* 2012;26:57–65.
2. van der Laan MJ, Petersen ML. Causal effect models for realistic individualized treatment and intention to treat rules. *Int J Biostat* 2007;3(1):Article 3.
3. Robins JM. Analytic methods for estimating HIV treatment and cofactor effects. In: Ostrow DG, Kessler R, editors. *Methodological issues of AIDS Mental Health Research*. New York: Plenum Publishing, 1993;213–90.
4. Murphy SA, van der Laan MJ, Robins JM. Marginal mean models for dynamic regimes. *J Am Stat Assoc* 2001;96:1410–23.
5. Hernan MA, Lanoy E, Costagliola D, Robins JM. Comparison of dynamic treatment regimes via inverse probability weighting. *Basic Clin Pharmacol* 2006;98:237–42.
6. Cain LE, Robins JM, Lanoy E, Logan R, Costagliola D, Hernn MA. When to start treatment? A systematic approach to the comparison of dynamic regimes using observational data. *Int J Biostat* 2010;6(2):Article 18.
7. Schomaker M, Egger M, Ndirangu J, Phiri S, Moultrie H, Technau K, et al.. When to start antiretroviral therapy in children aged 2–5 years: A collaborative causal modelling analysis of cohort studies from Southern Africa. *PLoS Med* 2013;10: e1001555.
8. Robins JM, Hernan MA, Brumback B. Marginal structural models and causal inference in epidemiology. *Epidemiology* 2000;110:550–60.
9. Robins JM, Hernan MA. Estimation of the causal effects of time-varying exposures. In: Fitzmaurice G, Davidian M, Verbeke MG, Molenberghs G, editors. *Advances in longitudinal data analysis*. New York: Chapman and Hall/CRC Press, 2009:553–599.
10. Robins JM. Marginal structural models versus structural nested models as tools for causal inference. In: Halloran ME, Berry D, editors. *Statistical models in epidemiology, the environment, and clinical trials*. New York: Springer, 2000:95–133.
11. Robins JM. A new approach to causal inference in mortality studies with sustained exposure periods – application to control of the healthy worker survivor effect. *Math Model* 1986;7:1393–512.
12. Taubman SL, Robins JM, Mittleman MA, Hernan MA. [Intervening on risk factors for coronary heart disease: an application of the parametric g-formula](#). *Int J Epidemiol* 2009;38:1599–611.
13. Robins JM. Robust estimation in sequentially ignorable missing data and causal inference models. In: *Proceedings of the American Statistical Association on Bayesian Statistical Science, 1999, 2000:6–10*.
14. Robins JM, Rotnitzky A. Recovery of information and adjustment for dependent censoring using surrogate markers. In: Nicholas P, Jewell, Klaus Dietz, Vernon T. Farewell, editors, *AIDS epidemiology*. Boston: Birkhäuser, 1992:297–331.
15. Robins JM, Rotnitzky A. Comment on the Bickel and Kwon article, “Inference for semiparametric models: some questions and an answer”. *Stat Sin* 2001;110:920–36.
16. Robins JM, Rotnitzky A, van der Laan MJ. Comment on “On profile likelihood”. *J Am Stat Assoc* 2000;450:431–5.
17. Rosenblum M, van der Laan MJ. Simple examples of estimating causal effects using targeted maximum likelihood estimation. UC Berkeley Division of Biostatistics Working Paper Series, Working Paper 209. Available at: <http://biostats.bepress.com/jhubiostat/paper209>, 2011.
18. Stitelman OM, De Gruttola V, van der Laan MJ. A general implementation of TMLE for longitudinal data applied to causal inference in survival analysis. UC Berkeley Division of Biostatistics Working Paper Series, Working Paper 281. Available at: <http://biostats.bepress.com/ucbbiostat/paper281>, 2011.
19. van der Laan MJ, Gruber S. Targeted minimum loss based estimation of causal effects of multiple time point interventions. *Int J Biostat* 2012;8(1):Article 8.
20. van der Laan MJ, Rose S. *Targeted learning: causal inference for observational and experimental data*. Berlin/Heidelberg/New York: Springer, 2011.
21. van der Laan MJ, Rubin D. Targeted maximum likelihood learning. *Int J Biostat* 2006;6(2):Article 2.
22. Bang H, Robins JM. [Doubly robust estimation in missing data and causal inference models](#). *Biometrics* 2005;61:962–72.
23. Robins JM. Marginal structural models. In: 1997 Proceedings of the American Statistical Association, Section on Bayesian Statistical Science, 1998:1–10.
24. Hernan MA, Brumback B, Robins JM. Marginal structural models to estimate the causal effect of zidovudine on the survival of HIV-positive men. *Epidemiology* 2000;110:561–70.

25. Petersen ML, van der Laan MJ, Napravnik S, Eron J, Moore R, Deeks S. Long term consequences of the delay between virologic failure of highly active antiretroviral therapy and regimen modification: a prospective cohort study. *AIDS* 2008;22:2097–106.
26. Robins JM, Orellana L, Rotnitzky A. Estimation and extrapolation of optimal treatment and testing strategies. *Stat Med* 2008;27:4678–721.
27. Neugebauer R, van der Laan M. Nonparametric causal effects based on marginal structural models. *J Stat Plann Inference* 2007;137:419–34.
28. Petersen ML, Porter KE, Gruber S, Wang Y, van der Laan MJ. Diagnosing and responding to violations in the positivity assumption. *Stat Methods Med Res* 2012;21:31–54.
29. Robins JM. Commentary on “using inverse weighting and predictive inference to estimate the effects of time-varying treatments on the discrete-time hazard. *Stat Med* 2002;21:1663–80.
30. Rosenblum M, van der Laan MJ. Targeted maximum likelihood estimation of the parameter of a marginal structural model. *Int J Biostat* 2010;6(2):Article 19.
31. Scharfstein DO, Rotnitzky A, Robins JM. Adjusting for nonignorable drop-out using semiparametric nonresponse models (with discussion and rejoinder). *J Am Stat Assoc* 1999;94:1096–1120 (1121–1146).
32. Schnitzer ME, Moodie EM, van der Laan MJ, Platt RW, 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(1):144–52.
33. R Core Team. R: a language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria, 2013. Available at: <http://www.R-project.org/>. ISBN 3-900051-07-0.
34. Schwab J, Lendle S, Petersen M, van der Laan M. LTMLE: longitudinal targeted maximum likelihood estimation, 2013. Available at: <http://cran.r-project.org/web/packages/ltmlle/>. R package version 0.9.3.
35. Pearl J. Causal diagrams for empirical research. *Biometrika* 1995;82:669–710.
36. Pearl J. *Causality: models, reasoning, and inference*. Cambridge: Cambridge University Press, 2000.
37. Bickel PJ, Klaassen CA, Ritov Y, Wellner JA. *Efficient and adaptive estimation for semiparametric models*. Baltimore, MD: Johns Hopkins University Press, 1993.
38. van der Laan MJ, Robins JM. *Unified methods for censored longitudinal data and causality*. Berlin/Heidelberg/New York: Springer, 2003.
39. van der Laan MJ, Polley E, Hubbard A. Super learner. *Stat Appl Genet Mol Biol* 2007;6:Article 25.
40. van der Laan MJ. Statistical inference when using data adaptive estimators of nuisance parameters. UC Berkeley Division of Biostatistics Working Paper Series, Working Paper 302, 2012. Available at: <http://www.bepress.com/ucbbiostat/paper302>.
41. Young JG, Cain LE, Robins JM, O’Reilly EJ, Hernan MA. Comparative effectiveness of dynamic treatment regimes: an application of the parametric g-formula. *Stat Biosci* 2011;30:119–43.
42. Picciotto S, Hernan M, Page J, Young J, Robins J. Structural nested cumulative failure time models to estimate the effects of interventions. *J Am Stat Assoc* 2012;107:866–900.
43. Geng EH, Glidden D, Bangsberg DR, Bwana MB, Musinguzi N, Metcalfe J, et al. Causal framework for understanding the effect of losses to follow-up on epidemiologic analyses in clinic based cohorts: the case of HIV-infected patients on antiretroviral therapy in Africa. *Am J Epidemiol* 2012;175:1080–7.
44. Schomaker M, Gsponer T, Estill J, Fox M, Boulle A. Non-ignorable loss to follow-up: correcting mortality estimates based on additional outcome ascertainment. *Stat Med* 2014;33:129–42.