

Jordan C. Brooks\*, Mark J. van der Laan, Daniel E. Singer, and Alan S. Go

# Targeted Minimum Loss-Based Estimation of Causal Effects in Right-Censored Survival Data with Time-Dependent Covariates: Warfarin, Stroke, and Death in Atrial Fibrillation

**Abstract:** Causal effects in right-censored survival data can be formally defined as the difference in the marginal cumulative event probabilities under particular interventions. Conventional estimators, such as the Kaplan-Meier (KM), fail to consistently estimate these marginal parameters under dependent treatment assignment or dependent censoring. Several modern estimators have been developed that reduce bias under both dependent treatment assignment and dependent censoring by incorporating information from baseline and time-dependent covariates. In the present article we describe a recently developed targeted minimum loss-based estimation (TMLE) algorithm for general longitudinal data structures and present in detail its application in right-censored survival data with time-dependent covariates. The treatment-specific marginal cumulative event probability is defined via a series of iterated conditional expectations in a time-dependent counting process framework. The TMLE involves an initial estimator of each conditional expectation and sequentially updates these such that the resulting estimator solves the efficient influence curve estimating equation in the nonparametric statistical model. We describe the assumptions required for consistent estimation of statistical parameters and additional assumptions required for consistent estimation of the causal effect parameter. Using simulated right-censored survival data, the mean squared error, bias, and 95% confidence interval coverage probability of the TMLE is compared with those of the conventional KM and the inverse probability of censoring weight estimating equation, conventional maximum likelihood substitution estimator, and the double robust augmented inverse probability of censoring weighted estimating equation. We conclude the article with estimation of the causal effect of warfarin medical therapy on the probability of “stroke or death” within a 1-year time frame using data from the ATRIA-1 observational cohort of persons with atrial fibrillation. Our results suggest that a fixed policy of warfarin treatment for all patients would result in 2% fewer deaths or strokes within 1-year as compared with a policy of withholding warfarin from all patients.

**Keywords:** targeted minimum loss-based estimation, survival analysis, dependent right-censoring, efficient influence curve, targeted maximum likelihood estimation

---

\*Corresponding author: Jordan C. Brooks, Life Expectancy Project, San Francisco, CA, USA,

E-mail: Brooks@LifeExpectancy.org

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

Daniel E. Singer, Clinical Epidemiology Unit, Massachusetts General Hospital, Boston, MA, USA, E-mail: DESINGER@partners.org

Alan S. Go, Division of Research, Kaiser Permanente of Northern California, Oakland, CA, USA, E-mail: Alan.S.Go@nsmtp.kp.org

## 1 Introduction

A commonly estimated causal effect parameter of survival data is the difference in marginal treatment-specific survival probabilities. That is, we wish to know the difference in the overall population survival proportion under two different treatment regimes. The estimation of treatment effects may be confounded by imbalances of certain covariates within each treatment group, and in most applications, right-censoring,

which may be dependent on an observed time-dependent covariate process, poses an additional challenge to estimation. Consistent estimation of causal effects from right-censored survival data therefore requires proper adjustment for both confounding due to covariate imbalances in treatment groups and confounding due to time-dependent censoring.

With complex medical data, it is often preferred to work under a nonparametric statistical model. In light of this, many authors in the medical literature construct the Kaplan-Meier (KM) [1] product limit estimator stratified within each treatment group. This procedure yields a consistent estimation of the treatment-specific marginal survival probabilities if treatment assignment is independent of covariates and if right-censoring is independent of survival time. While covariate imbalance may sometimes be addressed via a randomized study design, i.e. a randomized controlled trial (RCT), this is not always possible. Further, there is concern that causal effect estimates from RCTs may not generalize to the clinical patient population in which approved drugs are routinely prescribed because trial participants tend to be healthier, and medical care provided in the RCT may differ from typical practice. The assumption of independence of censoring and survival times may not hold in practice because sicker subjects are more likely than healthy subjects to drop out of a study earlier as it may be too difficult or unethical to continue.

In the last two decades several estimators have been proposed that, under certain assumptions, achieve consistent estimation of causal effects in the presence of covariate imbalances in treatment groups or right-censoring, which is explainable by baseline and time-dependent covariate processes. They are suitable for estimation of causal effects in RCTs or from observational data. As is the case with conventional estimators, the consistency of these estimators requires that survival time and censoring time be independent given the covariate process. In addition, their consistency relies on the assumptions that (1) the distribution of treatment assignment and censoring time can be estimated consistently as a function of the observed covariate process, and (2) that the distribution of the outcome can be estimated consistently as a function of the observed covariate process. For example, the consistency of the Inverse Probability of Censoring Weighted (IPCW) estimator [2], which is constructed as the solution to the IPCW estimating equation, relies on (1). The consistency of Maximum Likelihood Estimation (MLE) substitution estimators relies on (2). Double robust inverse probability of censoring weighted estimators (DRIPCW) [2, 27] and Targeted Minimum Loss Estimators (TMLE) [4, 5] are doubly robust in the sense that they are consistent under (1) or (2) and achieve the efficiency bound of a nonparametric model, i.e. have the highest possible precision amongst all regular estimators, if both (1) and (2).

In this article, we implement the recent TMLE algorithm developed in van der Laan and Gruber [6] for estimation of intervention-specific marginal means in general longitudinal data structures. This algorithm was based on key ideas from the double robust estimating equation methodology of Bang and Robins [7]. In brief, the algorithm computes the estimator by exploiting the iterative law, i.e. the tower property, of conditional expectation. The implementation here is specific to right-censored survival data with time-dependent covariates. The marginal mean parameter corresponds to the cumulative probability of a terminal event by a specific time point. The right-censored survival data model imposes particular constraints on the probability distribution, which allows a more computationally efficient implementation than that for the general algorithm presented in van der Laan and Gruber [6]. We apply the TMLE to the ATRIA-1 research study to estimate the additive causal effect of sustained warfarin therapy versus no warfarin therapy on the marginal cumulative probability of stroke or death by 1 year.

## 2 Data, model, and parameter

### 2.1 Data structure, $O$

Consider a right-censored survival data structure in which subject covariates are assessed at a baseline time point. After this assessment, the subjects are assigned a treatment and are then followed up until they

experience some terminal event of interest or are censored, where censoring events may include being lost to follow-up, study drop out, or administrative censoring at the end of the study. Time-dependent covariates are collected throughout the follow-up period. These may be predictive of the survival/death outcome or may be predictive of censoring. It is convenient to provide a formal representation of this data as a random variable in a counting process framework. Assume that time points are discrete and are indexed by the time counter,  $t \in 0, 1, \dots, K+1$ , that starts at 0 and increases by one unit until some ultimate time point of interest  $K+1$ .

Let  $(L_1(t) \in \{0, 1\} : t = 0, \dots, K+1)$  be the terminal event counting process. In brief, it is a random variable that takes value 0 until an event is observed. When an event is observed, the value switches to 1 and remains fixed for subsequent time points. We assume that  $L_1(0) = 0$ , i.e. each subject is alive at the beginning of follow-up. For reasons that will become apparent later we define  $Y = L_1(K+1)$  to be the value of the event process at the ultimate time point of interest. Let  $(L_2(t) \in \mathbb{R}^d : t = 0, \dots, K)$  be the covariate process, which includes both fixed and time-dependent covariates and is observed until the penultimate time  $K$ . We refer to the event and covariate processes collectively as  $L(t) = (L_1(t), L_2(t))$ , and for notational convenience we define  $\bar{L}(t) = (\bar{L}_1(t), \bar{L}_2(t))$  to be the histories of these processes up to time point  $t$ .

Let  $(A_1(t) \in \{0, 1\} : t = 0, \dots, K)$  be the treatment process. It is a random variable that takes value 0 if the subject is not treated at time  $t$  and 1 if he is treated. Let  $(A_2(t) \in \{0, 1\} : t = 0, \dots, K)$  be the counting process for censoring. Like the event counting process, this is a random variable that takes value 0 before the subject is censored and then becomes fixed to value 1 at the time when censoring is observed. The treatment and censoring processes, like the covariate process, are only observed up to the penultimate time  $K$ . We refer to the treatment and censoring processes collectively as  $A(t) = (A_1(t), A_2(t))$ , and we define  $\bar{A}(t) = (\bar{A}_1(t), \bar{A}_2(t))$  to be the histories of these processes.

By convention, we assume the ordering  $L(t) \rightarrow A(t)$  and when factorizing the likelihood at a particular time  $t$ , we use the ordering  $L_1(t) \rightarrow L_2(t) \rightarrow A_1(t) \rightarrow A_2(t)$ . Our data can now be written as the longitudinal data structure  $O = (Y, L(t), A(t) : t = 0, \dots, K)$ .

## 2.2 Factorization of the likelihood, $P_0$

We assume that our data is a random variable with a particular likelihood,  $O \sim P_0$ . This likelihood is the joint density of the random variables that make up  $O$ . As will be seen, it is useful to factorize  $P_0$  into orthogonal components indexed by the time-ordering  $t$ . Start by noting that the joint density of  $O$  can be factorized into the product of the time-specific conditional densities of the event, covariate, treatment, and censoring processes given the history.

$$P_0(O) = \prod_{t=0}^{K+1} P_0(L_1(t) | Pa(L_1(t))) \prod_{t=0}^K P_0(L_2(t) | Pa(L_2(t)))$$

$$\prod_{t=0}^K P_0(A_1(t) | Pa(A_1(t))) P_0(A_2(t) | Pa(A_2(t)))$$

where  $Pa(X)$  denotes the parents or all random variables that precede  $X$ . Note that:

$$Pa(L_1(t)) = (\bar{L}(t-1), \bar{A}(t-1))$$

$$Pa(L_2(t)) = (L_1(t), \bar{L}(t-1), \bar{A}(t-1))$$

$$Pa(A_1(t)) = (\bar{L}(t), \bar{A}(t-1))$$

$$Pa(A_2(t)) = (\bar{L}(t), A_1(t), \bar{A}(t-1))$$

From here on out we will refer to the conditional densities of the event and covariate processes as  $Q$ -factors and the conditional densities of the treatment and censoring processes as  $g$ -factors of the likelihood.

$$P_0(O) = \prod_{t=0}^{K+1} Q_{0,L_1(t)}(Pa(L_1(t))) \prod_{t=0}^K Q_{0,L_2(t)}(Pa(L_2(t)))$$

$$\prod_{t=0}^K g_{0,A_1(t)}(Pa(A_1(t))) g_{0,A_2(t)}(Pa(A_2(t)))$$

We denote the conditional expectations of the  $L_1(t)$  and  $A(t)$  processes given the history with the bar notation,  $\bar{Q}$  and  $\bar{g}$ .

$$\bar{Q}_{0,L_1(t)}(Pa(L_1(t))) = E_0[L_1(t)|Pa(L_1(t))]$$

$$\bar{g}_{0,A_1(t)}(Pa(A_1(t))) = E_0[A_1(t)|Pa(A_1(t))]$$

$$\bar{g}_{0,A_2(t)}(Pa(A_2(t))) = E_0[A_2(t)|Pa(A_2(t))]$$

Note that these conditional expectations are functions of the histories.

### 2.3 Statistical model, $\mathcal{M}$

At this point we must commit to a statistical model,  $\mathcal{M}$ , which is a set of assumptions we wish to impose on  $P_0$ . Formally,  $\mathcal{M}$  represents a collection of probability distributions (or likelihoods) of which  $P_0$  is the particular member that generates our data,  $O$ .

$$O \sim P_0 \in \mathcal{M}$$

In general we would like to avoid strong distributional assumptions on  $P_0$ , thus allowing our statistical model,  $M$ , to be as nonparametric as possible. There are, however, several constraints that are implicit in right-censored survival data. These are related to the fact that the time-dependent processes become degenerate after censoring or after an event has been observed.

After an event is observed:

- The event process remains fixed at 1.

$$\bar{Q}_{0,L_1(t)}(L_1(t-1) = 1, \bar{L}(t-1), \bar{A}(t-1)) = 1$$

- Censoring cannot be observed and remains fixed at 0.

$$\bar{g}_{0,A_2(t)}(L_1(t) = 1, \bar{L}(t), A_1(t), \bar{A}(t-1)) = 0$$

- The covariate and treatment processes can no longer be observed and are therefore fixed to a degenerate value, e.g. 0.

After censoring is observed:

- The event cannot be observed and remains fixed at 0.

$$\bar{Q}_{0,L_1(t)}(\bar{L}(t-1), A_2(t-1) = 1, \bar{A}(t-1)) = 0$$

- The censoring process remains fixed at 1.

$$\bar{g}_{0,A_2(t)}(\bar{L}(t-1), A_2(t-1) = 1, A_1(t), \bar{A}(t-1)) = 1$$

- The covariate and treatment processes can no longer be observed and are therefore fixed to a degenerate value, e.g. 0.

## 2.4 Statistical parameter $\psi_0 = \Psi(P_0)$

Our statistical parameter of interest,  $\psi_0$ , is the marginal cumulative event probability (or cumulative incidence) at time point  $K + 1$  according to a specific intervention,  $\bar{A}(K) = \bar{a}$ . Note that “intervention” here corresponds to both treatment and censoring processes. We define our parameter as a mapping applied to the probability distribution in the statistical model,  $\psi_0 = \Psi(P_0), \Psi : \mathcal{M} \rightarrow \mathbb{R}$ .

Note that this parameter is only well-defined under the positivity assumption,  $P_0(\bar{A}(t) = \bar{a} | Pa(\bar{A}(t))) > 0$  *a.e.* for all  $t$ . If some patient subgroup will never receive a particular treatment then a marginal intervention-specific outcome that includes this subgroup would be nonsensical. Similarly, if a certain subgroup is always censored before the ultimate time point then again a marginal parameter is not well-defined. If positivity holds then we may proceed in identifying the parameter.

The parameter mapping may be identified as a function of the  $Q$ -factors of  $P_0$ , i.e.  $\psi_0 = \Psi(P_0) = \Psi(Q_0)$ . This mapping is given below. For notational convenience here we drop the 0 subscript from the  $Q$ -factors. Start by recalling the definition  $Y = L_1(K + 1)$ .

By the law of iterated conditional expectation:

$$\begin{aligned}\bar{Q}_{Y|\bar{L}(K-0),\bar{a}} &= E[Y|\bar{L}(K-0), \bar{A}(K-0) = \bar{a}] \\ \bar{Q}_{Y|\bar{L}(K-1),\bar{a}} &= E[\bar{Q}_{Y|\bar{L}(K-0),\bar{a}}|\bar{L}(K-1), \bar{A}(K-1) = \bar{a}] \\ \bar{Q}_{Y|\bar{L}(K-2),\bar{a}} &= E[\bar{Q}_{Y|\bar{L}(K-1),\bar{a}}|\bar{L}(K-2), \bar{A}(K-2) = \bar{a}] \\ &\vdots \\ \bar{Q}_{Y|\bar{L}(1),\bar{a}} &= E[\bar{Q}_{Y|\bar{L}(2),\bar{a}}|\bar{L}(1), A(0) = \bar{a}] \\ \bar{Q}_{Y|L(0),\bar{a}} &= E[\bar{Q}_{Y|\bar{L}(1),\bar{a}}|L(0)] \\ \psi &= E[\bar{Q}_{Y|L(0),\bar{a}}] \\ \psi &= \Psi(Q) = \Psi(\bar{Q}_{Y|L(0),\bar{a}}, \bar{Q}_{Y|L(1),\bar{a}}, \dots, \bar{Q}_{Y|L(K),\bar{a}})\end{aligned}$$

As noted above, the true parameter,  $\psi_0$ , is simply this same mapping applied to the true  $Q_0$  distribution, i.e.  $\psi_0 = \Psi(Q_0)$ . This representation makes it clear that our parameter is only a function of the  $Q$ -factors through iterated conditional means at each time point. This is very convenient from an estimation perspective. In general, the estimation of conditional means is much easier than the estimation of conditional densities.

## 2.5 Efficient influence curve $D^*$

The efficient influence curve,  $D^* = D^*(P_0)(O)$ , is a functional of the density of the data that plays a central role in robust estimation in nonparametric statistical models [8]. The mean of  $D^*$  is 0, and the variance of  $D^*$  divided by the sample size defines the efficiency bound, i.e. the minimal variance that may be achieved by regular asymptotically linear estimators.

Let

$$g_{0,t} = \prod_{j=0}^t g_{0,A_1(j)} g_{0,A_2(j)}$$

be the probability of treatment and censoring by time  $t$  conditional on the history. The efficient influence curve for parameter mapping of the previous section is then:

$$D_{Q_0, g_0}^*(O) = \frac{I(\bar{A}(K) = \bar{a}(K))}{g_{0,K}} (Y - \bar{Q}_{0,Y|\bar{L}(K)}) + \sum_{t=0}^{K-1} \frac{I(\bar{A}(t) = \bar{a}(t))}{g_{0,t}} (\bar{Q}_{0,Y|\bar{L}(t+1)} - \bar{Q}_{0,Y|\bar{L}(t)}) + \bar{Q}_{0,Y|L(0)} - \psi_0$$

This efficient influence curve was originally presented by Bang and Robins [7] in the context of estimating equations and was derived in the context of TMLE in the recent article of van der Laan and Gruber [6]. When  $g_0$  and  $Q_0$  factors of the likelihood are estimated consistently, the TMLE solves an efficient (targeted minimum loss) score equation and also solves the efficient influence curve estimating equation,  $\sum_i D^*(O_i) = 0$ .

## 2.6 Causal treatment-specific mean parameter, $\psi^{\bar{a}}$

We now move from the statistical parameter to the causal parameter. Here we assume the existence of a counterfactual data distribution indexed by a particular intervention and then define the causal parameter as a marginal mean over this counterfactual distribution [9].

Consider the intervention,  $\bar{a} = \bar{a}(K)$ , which corresponds with a fixed intervention on the treatment and censoring processes. For example, we may choose

$$\bar{a} = (A_1(0) = A_1(1) = A_1(K) = 1, \quad A_2(0) = A_2(1) = A_2(K) = 0)$$

corresponding to a fixed treatment and no censoring up to time point  $K$ . Setting the values of all intervention nodes according to this intervention results in a data structure with a new probability density,  $P_0^{\bar{a}}$ . Our causal parameter is now defined as a mapping applied to the intervention-specific probability distribution  $P_0^{\bar{a}}$ .

$$\psi^{\bar{a}} = \Psi(P_0^{\bar{a}}) = \Psi(Q_0^{\bar{a}}) = E_0[Y^{\bar{a}}]$$

Our observed data  $O$  is not generated by the counterfactual distribution  $P_0^{\bar{a}}$  of the data under intervention  $\bar{a}$ , but rather from the observed data distribution  $P_0$ . Fortunately, we are under certain assumptions able to identify the causal quantity  $\psi_0^{\bar{a}}$  as the mapping  $\Psi$  applied to  $P_0$  described in the previous section. Positivity is required for the same reasons discussed previously. Also required for the identification of  $E[Y^{\bar{a}}]$  from  $P_0$  is the consistency assumption,  $O = (A, L_A)$ , i.e. that we observe the counterfactual corresponding with the treatment we observe, and the sequential randomization assumption:  $A(t) \perp L^{\bar{a}} | Pa(A(t))$  for all  $t$ . Note that the latter encompasses the usual assumption that censoring time is independent of survival time given covariates, and the heuristic translation is that there should be no unmeasured confounding variables. This assumption is untestable in practice, and therefore represents a point of controversy in causal effect estimation. However, even if this assumption is not valid, we may still arrive at an interesting statistical parameter, best described as the association of the intervention with the outcome adjusted for a subset of confounders. If these assumptions are met then  $\psi_0 = \psi_0^{\bar{a}}$  and consistent estimation of the statistical parameter is equivalent to consistent estimation of the causal parameter.

## 3 Targeted minimum loss estimation

TMLE was first introduced by van der Laan and Rubin [5]. TMLE is a two-stage substitution estimator for a finite-dimensional parameter,  $\psi_0 = \Psi(Q_0)$ , represented as a mapping from the density of the data to a vector of real numbers.

The first stage of the TMLE is to construct an estimate,  $Q_n$ , of the relevant portion of the density of the data,  $Q_0$ . An increasingly popular approach here is the use of loss-based machine learning algorithms that aim to minimize a global loss function for the (typically high-dimensional) density. The Super Learner [3, 10] is an example of a loss-based machine learning algorithm that relies on cross validation. This methodology

has been explored and utilized in a number of studies [11–14] and has been made available as an R package [15]. Substitution of  $Q_n$  for  $Q_0$  yields a typical minimum loss-based estimator  $\psi_n = \Psi(Q_n)$ .

The second stage of the TMLE is the “targeting” stage. Recall that our initial estimator was based on the minimization of the global loss function for the (typically high-dimensional) density of the data. As such, this estimator is unlikely to be optimal for and will almost certainly be biased our finite-dimensional target parameter. Thus the goal of the targeting step is to update our initial density estimate in such a way that we reduce bias and minimize the loss for our parameter of interest.

TMLE achieves this by fitting a least favorable parametric submodel [8] that takes the initial density estimator  $Q_n$  as an offset and a (possibly multivariate) covariate. Heuristically, the covariate defines a direction in which we must stretch our initial density estimator in order to remove bias for the parameter of interest, and is therefore often called the “clever covariate.” Typically, the clever covariate involves the treatment assignment or censoring mechanisms  $g_0$ . Thus the construction of the clever covariate usually requires an estimator of these factors of the likelihood. The parameter of the submodel represents the magnitude of the stretching required to reduce the bias in our initial estimate. The parameter can usually be estimated with standard regression software. The submodel estimate defines an updated density fit,  $Q_n^*$ , and the TMLE is defined by the substitution,  $\psi_n^{\text{TMLE}} = \Psi(Q_n^*)$ .

TMLEs have been derived for several parameters including marginal means and causal effects parameters in point treatment data [5, 16–18, 26], right-censored survival data [19, 20], longitudinal data structures with time-dependent covariates [6, 16, 17], and case-control settings [21]. TMLE has also been used to estimate variable importance measures [22] and for prediction calibration [23].

The TMLE presented here is a special case of the general algorithmic template given in van der Laan and Gruber [6] for the estimation of intervention-specific means in general longitudinal data structures. The latter was inspired in part by the estimating equation methods proposed by Bang and Robins [7]. The particular TMLE implementation discussed here is appropriate for right-censored survival data structures with a baseline treatment intervention and a covariate process that may include both fixed and time-dependent covariates. It allows for both treatment and censoring mechanisms that are dependent on either fixed or time-dependent covariates or on time itself.

### 3.1 TMLE algorithm

Van der Laan and Gruber [6] introduced the TMLE algorithm for the estimation of marginal means in the nonparametric statistical model for general longitudinal data structures. Here we outline a particular implementation of this algorithm that exploits additional information in the nonparametric statistical model for right-censored survival data.

#### 3.1.1 Step $K+1$

**Initial Estimator:** Propose an initial estimator,  $\bar{Q}_{n,Y|\bar{L}(K)}^{\bar{a},0}$ , of the innermost conditional expectation,  $\bar{Q}_{0,Y|L(K)}^{\bar{a}}$ . While one may be tempted to fit a regression straight away, some simple algebra allows more insight. Note that

$$\begin{aligned}\bar{Q}_{n,Y|\bar{L}(K)}^{\bar{a},0} &= I(L_1(K) = 1) \\ &\quad + I(L_1(K) = 0)E_n[Y^{\bar{a}}|\bar{L}_2^{\bar{a}}(K), L_1^{\bar{a}}(K) = 0]\end{aligned}$$

The first term reflects our knowledge that if the failure occurred at or before time  $K$  then event counting process is fixed at value 1 for all subsequent time points with probability 1. This implies that no update is required for observations with  $I(L_1(K) = 1) = 1$ . The second term is the conditional expectation of the event



process for persons who have not yet had the event and (due to the intervention settings) are not censored. In other words, this term corresponds to the intensity or hazard at time  $t = K + 1$ . A reasonable approach to constructing the initial estimator here might include estimation of intensity/hazard by smoothing over all time points and substituting the estimate at time  $t = K$ .

**Parametric submodel:**  $\text{Logit}(\bar{Q}_{n,Y|\bar{L}(K)}^{\bar{a},0}(\varepsilon_{K+1})) = \text{Logit}(\bar{Q}_{n,Y|\bar{L}(K)}^{\bar{a},0}) + \varepsilon_{K+1}H_{K+1}$  where the Logit transform of our initial estimator,  $\text{Logit}(\bar{Q}_{n,Y|\bar{L}(K)}^{\bar{a},0})$ , is used as an offset. The score of this submodel at  $\varepsilon_{K+1} = 0$  is  $H_{K+1}(Y - \bar{Q}_{n,Y|\bar{L}(K)}^{\bar{a},0}(\varepsilon_{K+1}))$  where  $H_{K+1} = \frac{I(\bar{A}(K)=\bar{a}(K))}{g_{n,K}}$ , i.e. the inverse probability of censoring weights constructed from our estimate of the  $g$ -factors of the likelihood.  $H_{K+1}$  is called the “clever covariate” because it is constructed so that the above score at  $\varepsilon_{K+1} = 0$  spans the  $(K + 1)$ th component of the efficient influence curve for  $\Psi(P_0)$ . The parameter estimate  $\varepsilon_{n,K+1}$  is computed through minimization of the log-loss function,  $\varepsilon_{n,K+1} = \underset{\varepsilon_{K+1}}{\text{argmin}} \frac{1}{n} \sum_{i=1}^n \mathcal{L}(\bar{Q}_{n,Y|\bar{L}(K)}^{\bar{a},0}(\varepsilon_{K+1}))(O_i)$ , where

$$\mathcal{L}(\bar{Q}_{n,Y|\bar{L}(K)}^{\bar{a},0}(\varepsilon_{K+1})) = -\log \left\{ \left( \bar{Q}_{n,Y|\bar{L}(K)}^{\bar{a},0}(\varepsilon_{K+1}) \right)^Y \left( 1 - \bar{Q}_{n,Y|\bar{L}(K)}^{\bar{a},0}(\varepsilon_{K+1}) \right)^{1-Y} \right\}$$

Standard software for logistic regression is used to fit this submodel.

**Targeted update:** The targeted update is then computed by plugging in our estimate,  $\varepsilon_{n,K+1}$ , into the parametric submodel.

$$\bar{Q}_{n,Y|\bar{L}(K)}^{\bar{a},*} = \bar{Q}_{n,Y|\bar{L}(K)}^{\bar{a},0}(\varepsilon_{n,K+1})$$

This solves the  $(K + 1)$ th component of the efficient influence curve estimating equation.

### 3.1.2 Step $K$ :

**Initial Estimator:** Now we must construct an initial estimator,  $\bar{Q}_{n,Y|\bar{L}(K-1)}^{\bar{a},0}$ , of the iterated conditional expectation, i.e. the conditional expectation (at time  $K$ ) of the conditional expectation at time  $K + 1$  given the history up to time  $K - 1$ . Note that

$$\begin{aligned} \bar{Q}_{n,Y|\bar{L}(K-1)}^{\bar{a},0} &= E_n[\bar{Q}_{n,Y|\bar{L}(K)}^{\bar{a},*} | \bar{L}^{\bar{a}}(K-1)] \\ &= I(L_1(K-1) = 1) \\ &\quad + I(L_1(K-1) = 0) E_n[\bar{Q}_{n,Y|\bar{L}(K)}^{\bar{a},*} | \bar{L}_2^{\bar{a}}(K-1), L_1^{\bar{a}}(K-1) = 0] \end{aligned}$$

Here again the first term reflects our knowledge that if the failure occurred at or before time  $K - 1$  then event counting process is fixed at value 1 for all subsequent time points with probability 1. We do not need to update this value in the targeted maximum likelihood step. Like the previous step, the second term here corresponds to people who have not yet had the event. Unlike the previous step, however, it is no longer the hazard or intensity at time  $K$ . It can be thought of as a random variable that consists of probabilities. This initial estimator is constructed from a regression of the previously targeted  $\bar{Q}_{n,Y|\bar{L}(K)}^{\bar{a},*}$  fit on the history up to  $K - 1$ .

**Parametric submodel:**  $\text{Logit}(\bar{Q}_{n,Y|\bar{L}(K-1)}^{\bar{a},0}(\varepsilon_K)) = \text{Logit}(\bar{Q}_{n,Y|\bar{L}(K-1)}^{\bar{a},0}) + \varepsilon_K H_K$  where  $\text{Logit}(\bar{Q}_{n,Y|\bar{L}(K-1)}^{\bar{a},0})$  is an offset, and  $H_K = \frac{I(\bar{A}(K-1)=\bar{a}(K-1))}{g_{K-1}}$  is the covariate. The score of this submodel at  $\varepsilon_K = 0$  is  $H_K(\bar{Q}_{n,Y|\bar{L}(K)}^{\bar{a},*} - \bar{Q}_{n,Y|\bar{L}(K-1)}^{\bar{a},0}(\varepsilon_K))$ . The parameter estimate  $\varepsilon_{n,K}$  is computed through minimization of the log-loss function,  $\varepsilon_{n,K} = \underset{\varepsilon_K}{\text{argmin}} \frac{1}{n} \sum_{i=1}^n \mathcal{L}(\bar{Q}_{n,Y|\bar{L}(K-1)}^{\bar{a},0}(\varepsilon_K))(O_i)$ , where

$$\mathcal{L}(\bar{Q}_{n,Y|\bar{L}(K-1)}^{\bar{a},0}(\varepsilon_K)) = -\log \left\{ \left( \bar{Q}_{n,Y|\bar{L}(K-1)}^{\bar{a},0}(\varepsilon_K) \right)^{\bar{Q}_{n,Y|\bar{L}(K)}^{\bar{a},*}} \left( 1 - \bar{Q}_{n,Y|\bar{L}(K-1)}^{\bar{a},0}(\varepsilon_K) \right)^{1-\bar{Q}_{n,Y|\bar{L}(K)}^{\bar{a},*}} \right\}$$



Most standard software packages for logistic regression can be applied to data with outcomes in the range  $[0,1]$  and can therefore be used to fit this submodel.

**Targeted update:** The targeted update is then given by

$$\bar{Q}_{n,Y|\bar{L}(K-1)}^{\bar{a},*} = \bar{Q}_{n,Y|\bar{L}(K-1)}^{\bar{a},0}(\varepsilon_{n,K})$$

This solves the  $(K)$ th component of the efficient influence curve estimating equation.

### 3.1.3 Steps $K - 1, \dots, 1$

The process for Step  $K$  is repeated for every time point, by conditioning on the history at time points  $t = K - 2, \dots, 0$ .

### 3.1.4 The final step

The final step is to take the empirical mean of the conditional expectations for all observations in the data set. Note that the empirical distribution of the pre-intervention variables is already targeted towards the estimation of our parameter, and therefore does not require any update. More specifically, given the empirical distribution as initial estimator the maximum likelihood estimator of the parameter  $\varepsilon_{t=0}$  in the parametric submodel is exactly equal to 0 [4]. We therefore arrive at the TMLE:

$$\psi_n^* = E_n \left[ \bar{Q}_{n,Y|L(0)}^{\bar{a},*} \right] = \frac{1}{n} \sum_{i=1}^n \bar{Q}_{n,Y|L(0)}^{\bar{a},*}$$

## 3.2 Statistical properties

The TMLE estimator is a substitution estimator based on targeted fits of functions of  $Q$ -factors of the likelihood,  $\psi_n^* = \Psi(\bar{Q}_{n,Y|L(t=0)}^{\bar{a},*}, \bar{Q}_{n,Y|\bar{L}(t=1)}^{\bar{a},*}, \dots, \bar{Q}_{n,Y|\bar{L}(t=K)}^{\bar{a},*})$ . Each  $t$ -specific targeted  $Q$ -factor solves the score equation for the  $t$ -specific parametric submodel. The submodel is a generalized linear model, and its score is therefore the “clever covariate” times the residual. Thus, the TMLE iteratively solves the score equation of each  $t$ -specific submodel and also iteratively solves each  $t$ -specific component of the efficient influence curve estimating equation. Upon completion of the algorithm the TMLE solves the entire influence curve estimating equation, which is the sum over all  $(K + 1)$  components, in a single pass through the data.

$$\sum_{i=1}^n D_{Q_n^*, g_n}^*(O_i) = \sum_{i=1}^n \sum_{t=0}^K D_{Q_n^*, Y|\bar{L}(t), g_n}^*(O_i) = 0$$

This implies that the TMLE is doubly robust in the sense that it will be consistent if either the intervention distribution  $g_n$  or the initial estimators of the relevant  $Q$ -factors,  $(\bar{Q}_{n,Y|L(0),\bar{a}}^0, \bar{Q}_{n,Y|\bar{L}(1),\bar{a}}^0, \bar{Q}_{n,Y|\bar{L}(K),\bar{a}}^0)$ , are estimated consistently. If both are estimated consistently then the TMLE,  $\psi_n^*$  has a normal distribution centered at the truth with variance that can be estimated with

$$\frac{1}{n} \sum_{i=1}^n \left( D_{Q_n^*, g_n^*}^*(O_i) \right)^2$$

The bootstrap can be used for inference in the absence of assumptions regarding the consistency of  $g_n$  or  $Q_n$ .

## 4 Simulation results

This section investigates the performance of the TMLE on a simulated dataset with a known distribution. Its performance is compared with that of the KM survival curve, the IPCW estimating equation estimator, the MLE substitution estimator, and the DRPCW estimating equation estimator. The data consisted of fixed covariates  $W_1, W_2, W_3, W_4$  and fixed treatment  $A_1$  drawn from the following distributions:

$$W_1, W_2 \sim \mathcal{N}(0, 1), W_3, W_4 \sim \text{Ber}(0.5)$$

$$A_1 \sim \text{Ber}(\text{expit}(0.001W_1 + 0.01W_2 - 0.5W_3 + 0.5W_4 - 0.2W_3W_1))$$

We also included a time-dependent covariate  $L_2(t)$  that was allowed to jump from value 0 to 1 once and then stay fixed over the course of follow-up. This covariate along with the censoring indicator  $A_2(t)$ , and event indicator  $L_1(t)$  at each time point of follow-up were drawn from the distributions implied by the following hazard/intensity functions. Note that for convenience we use the abuse of notation,  $R(t) = 1$ , to indicate that the subject is in the “risk set”, i.e. they have not had the event or been censored, by time  $t$ .

$$E[L_2(t)|Pa(L_2(t)), R(t) = 1] = \text{expit}(-3 + 0.01W_1 - 0.5A_1 + 0.01t)$$

$$E[A_2(t)|Pa(A_2(t)), R(t) = 1] = \text{expit}(-7 + 0.001W_1 - 0.002W_2 + 0.5W_3 \\ - 3W_4 + 0.25A_1 + 4L_2(t-1) - 0.01t)$$

$$E[L_1(t)|Pa(L_1(t)), R(t) = 1] = \text{expit}(-5 + 0.01W_1 - 0.002W_2 + 2W_3 - 3W_4 \\ - A_1 + A_1W_3W_4 + 2L_2(t-1) + 0.01t)$$

If we denote  $Y = L_1(20)$  our data can be written as  $O = (Y, W, A_1, L_1(t), L_2(t), A_2(t) : t = 0, \dots, 19) \sim P_0$ . The parameter of interest is the marginal cumulative event probability at time  $t = 20$ . This is given by  $E[Y^{\bar{a}}]$  under specific interventions:  $\bar{a}^0 = (A_1 = 0, A_2(0) = A_2(1) = \dots = A_2(19) = 0)$  for “no treatment without censoring,” and  $\bar{a}^1 = (A_1 = 1, A_2(0) = A_2(1) = \dots = A_2(19) = 0)$  for “treatment without censoring.” The true parameter value was computed using an empirical sample consisting of 1,000,000 independent draws from  $P^{\bar{a}}$  under each intervention.

The KM, IPCW, MLE, DRPCW, and TMLE were computed for 1,000 samples of size  $n = 1,000$  drawn from  $P_0$ . Bias was computed as the difference between the mean of the sample estimates and the true values. Mean squared error was computed in the usual way. For the KM, IPCW, DRPCW, and TMLE, we computed standard errors and coverage probabilities of 95% confidence intervals based on the variance of the their respective influence curves. Influence curve-based inference is not available for the MLE, and this explains the “na” values for the coverage probabilities for this estimator in the tables of our results. Data simulations and estimators were implemented in SAS software v 9.2.

### 4.1 Performance under positivity violations

Positivity is a requirement for the identification of both statistical and causal effect parameters. As a heuristic example, if a treatment or intervention of interest is absolutely never given to a particular subgroup of patients, perhaps because it is against medical guidelines or the law, then the average causal effect (of treatment) parameter for such patients is nonsensical. Mathematically, a violation of positivity corresponds with a parameter whose efficient influence curve is unbounded. For the parameter considered here it is easy to see that a positivity violation means that we get a  $D^*$  comprised of at least one fraction with a 0 in the denominator, which means  $D^* = \inf$ .

In practice, we often encounter data that exhibits near-positivity violations. For example, in the ATRIA-1 study, it was uncommon for persons with a history of bleeding events and falls to be prescribed warfarin. The case where treatment is particularly rare for certain subgroups is generally known as a practical positivity violation. For estimators whose influence curves have factors that involve the inverse probability of treatment, practical positivity violations are a concern. Because such estimators behave statistically like the empirical mean of their respective influence curves, a very small probability of treatment in the denominator can lead to a very large outlying value. The effect of this type of outlier may be likened to the effect of an outlier in the typical estimation of an empirical mean. The TMLE, IPCW, and DRIPCW are all estimators of this type, and a careful investigation of their respective influence curves can give insight into their behaviors under practical positivity violations.

Given the true  $g_0$ , the TMLE has influence curve

$$\begin{aligned} D_{Q^*, g_0}^{\text{TMLE}}(O) &= \frac{I(\bar{A}(K) = \bar{a}(K))}{g_{0,K}} \left( L_1(K+1) - \bar{Q}_{Y|\bar{L}(K)}^* \right) \\ &\quad + \sum_{t=0}^{K-1} \frac{I(\bar{A}(t) = \bar{a}(t))}{g_{0,t}} \left( \bar{Q}_{Y|\bar{L}(t+1)}^* - \bar{Q}_{Y|\bar{L}(t)}^* \right) \\ &\quad + \bar{Q}_{Y|\bar{L}(0)}^* - \psi_0 \end{aligned}$$

where  $Q^*$  is the asymptotic limit of our estimator  $Q_n^*$ . When  $g_0$  is estimated with  $g_n$  then the influence curve of the TMLE is the above influence curve minus its projection on the tangent space [8] of the model for  $g_0$ . Confidence intervals may be based on the empirical variance of  $D_{Q_n^*, g_n}$ . The DRIPCW influence curve is similar except that it is based on the initial estimator of the  $Q$ -factors of the likelihood, i.e.  $D_{Q_n, g_n}$ .

Given the true  $g_0$ , the IPCW estimator influence curve is

$$D_{g_0}^{\text{IPCW}}(O) = \frac{I(\bar{A}(\tilde{t}) = \bar{a}(\tilde{t}))}{g_{\tilde{t},0}} (L_1(\tilde{t}+1)) - \psi_0$$

where  $\tilde{t} = \min(K, \arg\max R(t) = 0)$ .

Note that the KM<sup>f</sup> can be expressed as an IPCW where the estimator of the  $g$ -factors of the likelihood does not depend on the covariates. Its influence curve follows the same form as that of the IPCW above.

Under regularity conditions, including when  $D^*$  is bounded, the TMLE is efficient when  $Q_n^*$  and  $g_n$  both converge to true  $Q_0$  and  $g_0$ , and remains consistent if one of them is correct. Despite this, its performance relative to the IPCW under positivity violations is not obvious. The reason is that the TMLE influence curve involves a sum of fractions with  $g_0$  in the denominator, whereas the IPCW influence curve only has one such fraction. In the TMLE, the numerator of the fraction is always a number bounded by  $[-1, 1]$  while the IPCW is exactly equal to 0 or 1.

The simulation study presented here investigates the performance of each estimator under positivity violations, introduced by scaling the linear component of the treatment assignment mechanism by constant factors, 1, 10, and 20 corresponding, respectively, to “no”, “substantial”, and “extreme” positivity violations. In the “no positivity” scenario, no truncation of the  $g$ -factors was required. In the “substantial” positivity scenario, 28% of the treatment group, for whom  $A = 1$ , required truncation, and 4% of the control group, for whom  $A = 0$ , required truncation. In the “extreme” case, 41% of the treatment group and 12% of the controls required truncation.

For the TMLE, IPCW, and DRIPCW the treatment assignment and censoring mechanisms were estimated with consistent parametric logistic regressions. To reflect common practice, the estimated  $g$ -factors at every time point were truncated to the interval  $[0.01, 0.99]$ . To carry out the TMLE, MLE, and DRIPCW, we first stratified our data set according to each treatment intervention, and then fit initial estimators of the intervention-specific  $Q$ -factors at every time point with Super Learners [10–14]. In brief, the Super Learner at each time point is a convex weighted average of the predictions from a library of five candidate

**Table 1** Simulation: performance under positivity violations

| Parameters:                      | $\psi_0^1 = 0.104$               | $\psi_0^0 = 0.217$    | $\psi_0^1 - \psi_0^0 = -0.113$ |
|----------------------------------|----------------------------------|-----------------------|--------------------------------|
| Estimator                        | MSE (Bias) – 95% CI coverage (%) |                       |                                |
| No positivity violation          |                                  |                       |                                |
| KM                               | 0.0010 (–0.0128) – 61            | 0.0101 (0.0613) – 49  | 0.0088 (–0.0741) – 30          |
| IPCW                             | 0.0003 (–0.0002) – 95            | 0.0003 (0.0001) – 97  | 0.0005 (–0.0003) – 97          |
| MLE                              | 0.0005 (–0.0160) – na            | 0.0004 (–0.0091) – na | 0.0006 (–0.0069) – na          |
| DRIPCW                           | 0.0003 (0.0010) – 93             | 0.0003 (–0.0016) – 95 | 0.0005 (0.0026) – 94           |
| TMLE                             | 0.0003 (0.0011) – 94             | 0.0003 (–0.0016) – 95 | 0.0005 (0.0028) – 94           |
| Substantial positivity violation |                                  |                       |                                |
| KM                               | 0.0045 (–0.0654) – 1             | 0.0636 (0.2030) – 0   | 0.0914 (–0.2685) – 0           |
| IPCW                             | 0.0036 (–0.0342) – 46            | 0.0004 (–0.0018) – 95 | 0.0039 (–0.0324) – 55          |
| MLE                              | 0.0028 (–0.0357) – na            | 0.0003 (–0.0069) – na | 0.0025 (–0.0288) – na          |
| DRIPCW                           | 0.0030 (–0.0107) – 59            | 0.0004 (–0.0020) – 93 | 0.0033 (–0.0087) – 69          |
| TMLE                             | 0.0036 (–0.0023) – 49            | 0.0004 (–0.0020) – 93 | 0.0039 (–0.0003) – 59          |
| Extreme positivity violation     |                                  |                       |                                |
| KM                               | 0.0049 (–0.0688) – 1             | 0.0571 (0.2028) – 0   | 0.0877 (–0.2717) – 0           |
| IPCW                             | 0.0054 (–0.0683) – 12            | 0.0004 (–0.0043) – 94 | 0.0052 (–0.0640) – 17          |
| MLE                              | 0.0037 (–0.0402) – na            | 0.0003 (–0.0067) – na | 0.0035 (–0.0335) – na          |
| DRIPCW                           | 0.0034 (–0.0269) – 26            | 0.0004 (–0.0010) – 95 | 0.0037 (–0.0260) – 46          |
| TMLE                             | 0.0039 (–0.0100) – 24            | 0.0004 (–0.0009) – 94 | 0.0042 (–0.0091) – 41          |

Note: Influence curve-based inference is not available, “na”, for MLE.

estimators. The convex weights are estimated from the data to minimize the expectation of the 5-fold cross validated negative Bernoulli loglikelihood. In the simulations here, the library of candidate estimators included (1) the unconditional mean; (2) a logistic regression for continuous outcomes in the interval [0,1]; (3) ordinary least squares regression; (4) neural network 3-layer perceptron with two hidden units; and (5) a recursive partitioning decision tree.

Results are shown in Table 1. In the absence of positivity violations, the KM is heavily biased due to the fact that both the treatment and censoring probabilities depend on covariate values. The KM also had the highest mean squared error. The MLE substitution estimator performed somewhat better, but was still more biased and had higher standard error than the remaining estimators. The IPCW, DRIPCW, and TMLE had the lowest mean squared errors. The biases for the TMLE and DRIPCW were higher than that of the IPCW, though their magnitudes were negligible with respect to statistical inference and 95% confidence intervals based on their respective influence curves had reasonably good coverage probabilities.

The bias and mean squared error of the IPCW increased with higher levels of positivity violation. The DRIPCW and TMLE performed better than the IPCW under extreme positivity violations, though the best estimator with respect to overall mean squared error in this scenario was the MLE. Interestingly, the performance of the MLE, which does not rely on estimation of the  $g$ -factors of the likelihood, also suffered with increasing positivity.

## 4.2 Performance under model misspecification

Model misspecification is a major concern for estimation in nonparametric statistical models. As noted previously the IPCW is consistent in the nonparametric model only if the  $g$ -factors of the likelihood are estimated consistently. The TMLE and DRIPCW are double robust in that they achieve consistency if either the  $g$ -factors or the relevant  $Q$ -factors are estimated consistently. The TMLE and DRIPCW are also locally efficient in that they achieve the nonparametric efficiency bound when both the  $g$  and  $Q$ -factors are

estimated consistently. This section investigates these properties via simulation. In order to avoid redundancy with the section on positivity violations above, the linear part of the treatment assignment mechanism was scaled by a factor of 2.

We assess the performance of the TMLE, IPCW, and DRIPCW estimators built upon parametric estimators of correctly specified  $g$ -factors,  $g_{\text{cor}}$ , or misspecified  $g$ -factors,  $g_{\text{mis}}$ , respectively. Model misspecification was specific to the treatment assignment probabilities. Both  $g_{\text{cor}}$  and  $g_{\text{mis}}$  used correctly specified logistic regressions model for the censoring hazard function. With regard to the treatment assignment,  $g_{\text{cor}}$  used the correctly specified logistic regression model,  $E_{\text{cor}}[A_1|W] = \text{expit}(\beta_0 + \beta_1 W_1 + \beta_2 W_2 + \beta_3 W_3 + \beta_4 W_4 - \beta_5 W_3 W_1)$ . The misspecified logistic regression model,  $E_{\text{mis}}[A_1|W] = \text{expit}(\beta_0 + \beta_1 W_1 + \beta_2 W_4)$ , was used for  $g_{\text{mis}}$ .

The gains in efficiency for TMLE and DRIPCW are investigated by comparing three different approaches to the estimation of the  $Q$ -factors. The first approach used the NULL model in which each  $\bar{Q}_{Y|\bar{L}(k),\bar{a}} : k = 0, K$  was assumed to be independent of the covariates. The second approach estimated the  $Q$ -factors with logistic regression functions of the baseline covariates and the most recent time-dependent covariate. The third approach constructed Super Learners for each  $Q$ -factor, as in the previous section.

Results are shown in Table 2. Again the KM, which assumes independence of both the treatment and censoring processes, is quite biased. In this sense, the KM is in fact an IPCW estimator based on a highly misspecified estimator of the  $g$ -factors of the likelihood. The IPCW is largely unbiased under  $g_{\text{cor}}$ , but quite biased under  $g_{\text{mis}}$ . The TMLEs and DRIPCWs presented here had very similar bias and mean squared error under all model specification settings. The performance of both the TMLE and DRIPCW was better than or equal to that of the IPCW and KM in terms of both bias and MSE. The TMLE and DRIPCW, based on the NULL estimators of the  $Q$ -factors, behave very similar to the IPCW. The double robustness property of the TMLE and DRIPCW are best illustrated in simulations where the  $g$ -factors are misspecified. This is evidenced

**Table 2** Simulation: performance under model misspecification

| Parameters:                       | $\psi_0^1 = 0.104$               | $\psi_0^0 = 0.217$    | $\psi_0^1 - \psi_0^0 = -0.113$ |
|-----------------------------------|----------------------------------|-----------------------|--------------------------------|
| Estimator                         | MSE (Bias) – 95% CI coverage (%) |                       |                                |
| KM:                               | 0.0015 (–0.0291) – 25            | 0.0205 (0.0992) – 7   | 0.0235 (–0.1283) – 0           |
| IPCW:                             |                                  |                       |                                |
| $g_{\text{mis}}$                  | 0.0007 (–0.0219) – 68            | 0.0009 (0.0242) – 84  | 0.0027 (–0.0461) – 59          |
| $g_{\text{cor}}$                  | 0.0004 (–0.0005) – 95            | 0.0003 (–0.0003) – 97 | 0.0006 (–0.0002) – 98          |
| MLE:                              |                                  |                       |                                |
| $Q_{\text{NULL}}$                 | 0.0012 (–0.0335) – na            | 0.0022 (0.0428) – na  | 0.0063 (–0.0763) – na          |
| $Q_{\text{GLM}}$                  | 0.0003 (0.0011) – na             | 0.0003 (–0.0023) – na | 0.0005 (0.0034) – na           |
| $Q_{\text{SL}}$                   | 0.0008 (–0.0207) – na            | 0.0003 (–0.0079) – an | 0.0007 (–0.0128) – na          |
| DRIPCW:                           |                                  |                       |                                |
| $g_{\text{mis}}, Q_{\text{NULL}}$ | 0.0007 (–0.0219) – 67            | 0.0009 (0.0244) – 81  | 0.0027 (–0.0463) – 52          |
| $g_{\text{mis}}, Q_{\text{GLM}}$  | 0.0003 (0.0018) – 88             | 0.0003 (–0.0018) – 95 | 0.0005 (0.0036) – 93           |
| $g_{\text{mis}}, Q_{\text{SL}}$   | 0.0003 (–0.0045) – 85            | 0.0003 (–0.0004) – 95 | 0.0005 (–0.0040) – 92          |
| $g_{\text{cor}}, Q_{\text{NULL}}$ | 0.0004 (–0.0005) – 94            | 0.0003 (–0.0003) – 96 | 0.0006 (–0.0002) – 97          |
| $g_{\text{cor}}, Q_{\text{GLM}}$  | 0.0003 (0.0016) – 94             | 0.0003 (–0.0020) – 94 | 0.0006 (0.0036) – 95           |
| $g_{\text{cor}}, Q_{\text{SL}}$   | 0.0004 (0.0007) – 94             | 0.0003 (–0.0020) – 94 | 0.0006 (0.0027) – 95           |
| TMLE:                             |                                  |                       |                                |
| $g_{\text{mis}}, Q_{\text{NULL}}$ | 0.0007 (–0.0219) – 67            | 0.0009 (0.0243) – 81  | 0.0027 (–0.0463) – 52          |
| $g_{\text{mis}}, Q_{\text{GLM}}$  | 0.0003 (0.0020) – 88             | 0.0003 (–0.0019) – 95 | 0.0005 (0.0038) – 92           |
| $g_{\text{mis}}, Q_{\text{SL}}$   | 0.0003 (–0.0010) – 86            | 0.0003 (–0.0009) – 95 | 0.0005 (0.0000) – 92           |
| $g_{\text{cor}}, Q_{\text{NULL}}$ | 0.0004 (–0.0005) – 95            | 0.0003 (–0.0003) – 97 | 0.0006 (–0.0002) – 97          |
| $g_{\text{cor}}, Q_{\text{GLM}}$  | 0.0003 (0.0016) – 94             | 0.0003 (–0.0020) – 94 | 0.0006 (0.0036) – 94           |
| $g_{\text{cor}}, Q_{\text{SL}}$   | 0.0004 (0.0009) – 94             | 0.0003 (–0.0020) – 94 | 0.0006 (0.0029) – 95           |

Note: Influence curve-based inference is not available, “na”, for MLE.

by the major bias reduction associated with using GLM or SL fits for the  $Q$ -factors and underscores the importance of estimating the  $Q$ -factors when the model for  $g$  is unknown. In practice, one might choose between different TMLEs indexed by different  $Q$ -factor estimators on the basis of the cross validated variance of the influence curve, though in the simulations presented here the observed differences in efficiency were modest at best. The MLE in some settings had lower mean squared error than both the TMLE and DRIPCW, though it appeared to have somewhat higher bias.

## 5 The causal effect of warfarin

Warfarin is a medication used in the prevention of thromboembolic stroke. However, its physiological mechanism of action also puts users at higher risk for adverse bleeding events, whose health consequences may be just as devastating as thromboembolic stroke. The efficacy of warfarin has been shown in randomized clinical trials, though some practitioners have questioned whether the extent of the benefits is the same in the general atrial fibrillation population. The reason is that typical patients have more comorbidities and may not be as healthy as trial participants. One of the major research objectives of the ATRIA-1 cohort study was to estimate the causal effect of warfarin on adverse outcomes in the general atrial fibrillation patient population at Kaiser Permanente Northern California.

Like most drugs, warfarin prescription decisions are based on the patient's medical history. In particular those who are thought to be at high risk for stroke are more likely to receive the drug. The risk factors for stroke, however, are largely the same as those for falls, bleeds, and death by any cause. The warfarin prescription decision must therefore weigh the potential benefits of stroke prevention against the potential risks of adverse outcomes. Thus treatment is highly dependent on clinical covariates, and estimation of causal effects must take this into account.

Here we estimate intervention-specific marginal probabilities of "stroke or death" within a 1-year time frame. The first intervention consists of sustained warfarin usage and no censoring, while the second intervention withholds warfarin therapy and enforces no censoring. The causal effect of warfarin is then defined as the difference in these two intervention-specific marginal probabilities.

### 5.1 ATRIA-1 cohort data

Using the same counting process framework outlined in Section 2.1.  $L_1(t)$  will denote the counting process for the stroke or death event at some time point  $t$ , and we are interested in its value at the 1-year (365 days) time point  $Y = L_1(K + 1 = 365)$ ;  $L_2(t)$  will denote the (fixed and time-dependent) covariate processes;  $A_1 = A_1(t) : t = 0, \dots, 364$  will denote the fixed baseline warfarin treatment; and  $A_2(t)$  will denote the censoring counting process, where censoring events include disenrollment from Kaiser, end of study, or switching warfarin treatment. We regard the 13,599 as an empirical sample drawn from a probability distribution,  $O_1, \dots, O_{13,599} \sim P_0$ , and denote the empirical distribution  $P_n$ .

At the beginning of the study, 5,289 subjects were taking warfarin and 8,270 were not. After the baseline assessment, 2,267 subjects started using warfarin and 1,228 stopped using warfarin within the first year. These individuals were censored on the day when the treatment was switched. One sixty eight subjects were censored due to disenrollment from the Kaiser system, and none were administratively censored within the 1-year time frame. Among those who were not censored, 171 had a thromboembolic stroke and 485 died within the first year of follow-up. The remaining 9,240 were observed to be alive, stroke-free, and had not switched treatments at the end of the first year of study. Under the data structure described above the working data set contained  $13,559 \text{ persons} \times 365 \text{ days} = 4,949,035$  rows, though after the timing of censoring and terminal events were taken into account the dataset consisted of 3,915,087 person-days at risk.



The  $L_2(t)$  covariate process for each individual included fixed indicators of gender, race, educational attainment, and income along with a time-dependent covariate for age at time  $t$ . Other time-dependent covariates included indicators for diagnoses of prior stroke, hypertension, intracranial hemorrhage, other bleeds, coronary heart disease, congestive heart failure, diabetes mellitus, retinopathy, dialysis, history of fall, dementia, seizure, peripheral artery disease, or hepatitis. The values of these covariates were set to 0 for persons with no history of the diagnoses and were fixed to 1 on the date of diagnosis and all subsequent time points. Time-dependent values for lab measurements of total hemoglobin, HbA1c, international normalized ratio, glomerular filtration rate, and proteinuria abstracted from Kaiser's administrative databases. These were measured under standard medical protocols and the values between measurements were imputed with the last observation carried forward.

## 5.2 Causal effect of warfarin parameter

Our outcome here is defined as “stroke or death” within a 1-year time frame, and our parameter of interest is the intervention-specific marginal mean of this outcome:  $\psi^{\bar{a}} = \Psi(Q^{\bar{a}}) = E[Y^{\bar{a}}]$ . Equivalently, this is the marginal cumulative probability of stroke or death within 1 year, or the proportion of the population who would experience stroke or death within 1 year under the specified intervention. The fixed interventions of interest are no warfarin for 1-year ( $\bar{a} = 0$ ) versus sustained warfarin therapy for 1-year ( $\bar{a} = 1$ ). Both interventions also prevent early censoring. We define the additive causal effect parameter as the difference between the two intervention-specific marginal means.

$$\psi^{\bar{a}=1} - \psi^{\bar{a}=0}$$

where

$$(\bar{a} = 1) : = (A_1(0) = \dots = A_1(364) = 1, A_2(0) = \dots = A_2(364) = 0)$$

$$(\bar{a} = 0) : = (A_1(0) = \dots = A_1(364) = 0, A_2(0) = \dots = A_2(364) = 0)$$

## 5.3 TMLE implementation

The  $g$ -factor corresponding to the fixed treatment assignment mechanism was given by the conditional expectation of warfarin treatment given the medical history at baseline,  $g_{A_1,0} = E_0[A_1(0)|L(0)]$ . This was estimated with a Super Learner [10, 13] for the negative Bernoulli loglikelihood risk. This Super Learner was fitted on a dataset with 13,559 rows – i.e. one per subject. The  $g$ -factor corresponding to the time-dependent censoring mechanism was given by the conditional intensity of censoring given the history up to the most recent time point,  $g_{A_2(t),0} = (E_0[A_2(t)|\bar{L}(t), \bar{A}(t-1)] : t = 0, \dots, 365)$ . This was estimated with a Super Learner for the risk defined as the expectation of the sum of the negative Bernoulli loglikelihood loss at each time point.

The initial fit for the terminal  $Q$ -factor corresponding to the time-dependent event intensity mechanism at day 365 was estimated with a Super Learner for the risk defined as the expectation of the sum of the negative Bernoulli loglikelihood at each time point. To improve precision we used the entire ATRIA-1 dataset (even time points beyond day 365) to fit this Super Learner. The initial fits for the remaining  $Q$ -factors were fit with Super Learners for the negative Bernoulli loglikelihood risk. The Super Learner for each time point specific  $Q$ -factor was fit using only those observations that had not yet been censored or experienced the event.

All of the Super Learners (both for the  $g$  and  $Q$ -factors) used 5-fold cross validation and a library with the following candidates:

- MEAN – the unconditional mean, i.e. the NULL model
- LOGIT or LOGIT\_CTS01 – logistic regression or logistic regression for continuous outcomes in  $[0,1]$

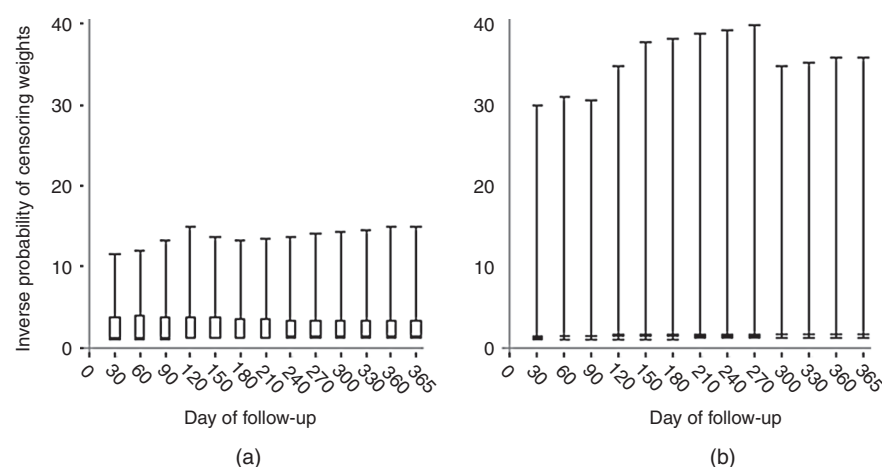


- LDA – linear discriminant analysis (not used for the  $Q$ -fits with continuous outcomes)
- NN2–neural network 3-layer perceptron with two hidden units
- TREE – recursive partitioning decision tree

The Super Learners and the TMLE algorithm were implemented in SAS v 9.2.

## 5.4 Results

The Super Learner for the fixed treatment assignment mechanism was a weighted average of the unconditional mean (9%), logistic regression (37%), linear discriminant analysis (12%), neural network (39%), and decision tree (3%). The Super Learner for the time-dependent censoring mechanism only involved the neural network (6%), and decision tree (94%). The fitted  $g$ -factors did not suggest major practical positivity violations, and we did not truncate inverse probability weights. The inverse probability weights are plotted as a function of the follow-up time for both the warfarin and the no warfarin interventions are plotted below (Figure 1).



**Figure 1** Inverse probability of censoring weights. (a) Warfarin. (b) No Warfarin

The Super Learner for the event intensity was a weighted average of logistic regression (66%), linear discriminant analysis (4%), neural network (20%), and decision tree (10%). The weights for the Super Learner for remaining  $Q$ -factors (364 in total) are too numerous to describe in detail, though generally the higher weight was given to the logistic regression candidate estimator. The initial estimator of the  $Q$ -factors for every time point, days 0 to 365, were updated with TMLE to solve the efficient score equation.

Table 3 presents the KM, IPCW, MLE, DRIPCW and TMLE estimators and their standard errors based on their respective influence curves. All three estimators yielded point estimates that indicate the sustained warfarin intervention decreased the proportion of stroke or death over the no warfarin intervention. The KM estimator suggests that the marginal stroke or death probability under the sustained warfarin intervention was 3.5% less than that under the no warfarin intervention. However, because warfarin treatment decisions are made based on the patient's medical covariate history, the KM is not consistent as an estimator of the marginal probability. It is instead a stratified estimator of the probability of stroke or death amongst persons on warfarin and amongst those who were not, respectively. Further, the KM is subject to bias when there is informative censoring, which appears to have been the case here.

**Table 3** Causal effect of warfarin on the probability of stroke or death within 1 year

| Estimator | Warfarin | Nowarfarin | Difference | S.E.   | <i>p-value</i> |
|-----------|----------|------------|------------|--------|----------------|
| KM        | 0.0398   | 0.0751     | −0.0352    | 0.0046 | 0.0001         |
| IPCW      | 0.0492   | 0.0696     | −0.0204    | 0.0070 | 0.0019         |
| MLE       | 0.0372   | 0.0646     | −0.0274    | na     | na             |
| DRIPCW    | 0.0519   | 0.0716     | −0.0197    | 0.0068 | 0.0020         |
| TMLE      | 0.0530   | 0.0719     | −0.0189    | 0.0068 | 0.0028         |

Note: Influence curve-based inference is not available, “na”, for MLE.

The IPCW, MLE, DRIPCW, and TMLE are of more interest here because these may be consistent causal effect estimators. The range of the estimates of the  $g$ -factors of the likelihood used in both TMLE and IPCW suggested that there was no major positivity violation for the estimation of the parameter of interest. Thus, if our estimates of the treatment assignment and the censoring mechanisms are consistent, then the TMLE, IPCW, and DRIPCW are consistent for the statistical parameter of interest. If the fits of the  $Q$ -factors are consistent then the TMLE, MLE, and DRIPCW are consistent for the statistical parameter of interest. Under the additional (untestable) assumption that there are no unmeasured confounders, then these estimators are also consistent for the causal effect parameter. The IPCW suggests a reduction of 2.0%, the MLE 2.7%, the DRIPCW 2.0% and the TMLE 1.9%. As expected, the influence curves of these estimators indicate that the TMLE and DRIPCW were slightly more precise than the IPCW.

It is noteworthy that the TMLE, DRIPCW, MLE, and IPCW all yielded higher “stroke or death” probability estimates than KM under the warfarin treatment and lower estimates under the no warfarin treatment. Because warfarin is prescribed to patients who are considered to be at higher risk for stroke, one might anticipate the opposite finding, i.e. that the KM estimator should make the drug look less beneficial. On the other hand, warfarin is often withheld from patients with a high risk of death for safety reasons. For example, recent falls and bleeding events are common contraindications for warfarin.

We explored this apparent confounding by regressing our Super Learner fit for the “stroke or death” event intensity on the baseline warfarin treatment and found that a higher event intensity at baseline was associated with a lower probability of having received warfarin. Recall that individuals who started warfarin therapy later and those who stopped taking it were censored. As an informal exploration, we also fit an estimator of the censoring mechanism as a function of the event intensity, separately for those who started out on warfarin and those who did not. Amongst persons who started the study on warfarin, higher event intensities were positively associated with censoring, which suggests that when a patient’s underlying risk of “stroke or death” increased they were subsequently taken off warfarin.

In summary, an intervention defined by giving warfarin to all atrial fibrillation patients significantly reduced the marginal probability of stroke or death by 2% relative to the null intervention (no warfarin at all). This effect estimate was smaller than the estimate that ignored confounding, seemingly introduced by clinician’s apparent reluctance to prescribe warfarin to patients whose medical history implied a higher underlying risk of death.

## 6 Discussion

Recently developed estimators, e.g. IPCW, MLE, DRIPCW, or TMLE, extend the analysis of causal effects to observational data settings. This is remarkable in that it allows the estimation of causal effects directly from the patient populations in which therapies are typically prescribed, as opposed to trial populations that tend to be comprised of healthier subjects. Increasing awareness of such estimators has led to academic debate

over which estimators are more useful in theory and practice. The TMLE combines attractive theoretical features of the estimating equation approaches with those of minimum loss-based estimation and avoids several of their respective shortcomings. Estimating equation approaches require that the parameter of interest be defined as the solution to an estimating equation, and this is not always possible. Even when it is, the estimating equation may have multiple solutions without a clear-cut optimality criterion. Further, inconsistent estimation of the  $g$ -factors of the likelihood can lead to parameter estimates that do not obey the constraints of the statistical model. For example, it is possible to compute IPCW confidence intervals for probabilities that lie outside the range  $[0, 1]$ . Because the TMLE is a substitution estimator, parameter estimates never escape the bounds of the statistical model. The parameter is defined as a mapping applied to the probability distribution of the data and therefore extends to cases where it cannot be written as the solution to an estimating equation. The minimum loss, or maximum likelihood, principle provides a clear-cut theoretical optimality criterion for use in practical estimation problems.

In this article we implemented a TMLE algorithm [4] that can be applied in either the RCT or observational scenarios to estimate the marginal intervention-specific cumulative event probability for right-censored survival data with time-dependent covariates. The TMLE solves the efficient influence curve estimating equation for the marginal intervention-specific cumulative event probability and is doubly robust in the sense that it is consistent if either the intervention distribution or the relevant components of the event process distribution are estimated consistently. If both are estimated consistently, it achieves the efficiency bound in the nonparametric model. The implementation here assumes that complete histories of covariate, treatment, event and censoring processes are observed for each individual at each time point in a discrete set of time points. The benefits of double robust estimators TMLE and DRIPCW were evident in our simulation studies. These two estimators, however, performed very similarly with regard to bias, mean squared error, and coverage probabilities for 95% confidence intervals. Although not readily apparent in the settings studied here, the work of Stitleman et al. [20] documented notable gains in the TMLE as compared to the DRIPCW in settings where the solution to the DRIPCW estimating fell outside the bounds of the model.

In applied work, time-dependent covariate measures are nearly always measured at different times for different subjects, and the researcher must make some plausible imputation of the histories at each time point. Our analysis of the atrial fibrillation data worked with the standard “last observation carried forward” imputation. Imputation of this type does not cause bias, because it does not come at any loss in information. It is important to note here that this imputation is not an attempt to control for unmeasured confounders (such as time-specific values of biomarkers that were not observed). We are merely adjusting for the measured covariate processes and assume the sequential randomization assumption with respect to these measured processes in order to interpret our statistical estimands as a causal effect. Loss in information (and possible bias) only occurs if we use a too coarse time scale, such that the actual observed data must be somehow truncated or otherwise modified.

The more important assumption in this discrete framework is that the intervention nodes (censoring and treatment) are assumed to be correctly defined at each time point. As noted above for the covariate process, an overly coarse discrete time scale may induce bias with respect to the desired causal effect. On the other hand, an overly fine discrete time scale may result in practical violations of the positivity assumption in cases where the treatment process is allowed to change at any point in time. In our case, we were able to use a very fine time scale (daily) since the treatment process and censoring process only jump once. Methods to estimate the causal effects of treatment processes that can jump in continuous time at various time points, instead of enforcing the data to adhere to the discrete time scale, are the subject of ongoing research.

The results of the data analysis indicated that a medical policy of treating all patients with warfarin versus no treatment would have reduced 1-year stroke or death probability from 7% to 5%. It must be restated that these interventions were chosen to illustrate the TMLE method. Such stringent interventions, i. e. sustained warfarin for all patients or no warfarin for any, are unrealistic in practice. Future work applied work in this area should focus on comparisons of realistic treatment rules.

Finally, we reiterate that the TMLE approach presented here is quite general and may be adapted to a variety of other problems. The same method may be used to estimate the causal effects of dynamic treatment rules, in which treatment decisions are made in response to the evolution of the time-dependent covariate process. Extension to the estimation of counterfactual survival curves (i.e. survival probabilities at multiple time points) is straightforward and simultaneous confidence bands may be computed as described in Quale and van der Laan [24]. Further extensions that involve marginal structural models may be used to estimate continuous treatment effects or treatment effects that depend on a function of baseline covariates [25].

## References

1. Kaplan EL, Meier P. Nonparametric estimation from incomplete observations. *J Am Stat Assoc* 1958;53:457–81.
2. Robins JM, Rotnitzky A. *AIDS epidemiology – methodological issues*. Boston, MA: Birkhäuser, 1992.
3. van der Laan MJ, Dudoit S. Unified cross-validation methodology for selection among estimators and a general cross-validated adaptive epsilon-net estimator: finite sample oracle inequalities and examples. Technical report, Division of Biostatistics, University of California, 2003.
4. van der Laan MJ, Rose S. *Targeted learning: causal inference for observational and experimental data*. New York: Springer, 2011.
5. van der Laan MJ, Rubin D. Targeted maximum likelihood learning. *Int J Biostat* 2006;2. Available at: <http://www.bepress.com/ijb/vol2/iss1/11>.
6. van der Laan MJ, Gruber S. Targeted minimum loss-based estimation of an intervention specific mean outcome. Technical Report 290, Division of Biostatistics, University of California, 2011.
7. Bang H, Robins JM. Doubly robust estimation in missing data and causal inference models. *Biometrics* 2005;61:962–72.
8. Bickel PJ, Klaassen CA, Ritov Y, Wellner J. *Efficient and adaptive estimation for semiparametric models*. New York: Springer-Verlag, 1993.
9. Pearl J. *Causality: models, reasoning, and inference*, 2nd edn. New York: Cambridge University Press, 2009.
10. van der Laan MJ, Polley EC, Hubbard AE. Super learner. *Stat Appl Genet Mol Biol* 2007;6. Available at: <http://www.bepress.com/sagmb/vol6/iss1/art25>.
11. Chaffee P, Hubbard AE, van der Laan MJ. Permutation-based pathway testing using the super learner algorithm. Technical Report 263, Division of Biostatistics, University of California, 2010. Available at: <http://www.bepress.com/ucbbiostat/paper263>.
12. Diaz Munoz I, van der Laan MJ. Super learner based conditional density estimation with application to marginal structural models. *Int J Biostat* 2011;7. Available at: <http://www.bepress.com/ijb/vol7/iss1/38>.
13. Polley EC, van der Laan MJ. Super learner in prediction. Technical Report 266, Division of Biostatistics, University of California, 2010. Available at: <http://www.bepress.com/ucbbiostat/paper266>.
14. Sinisi SE, Polley EC, Petersen ML, Rhee S-Y, van der Laan MJ. Super learning: an application to the prediction of Hiv-1 drug resistance. *Stat Appl Genet Mol Biol* 2007;6. Available at: <http://www.bepress.com/sagmb/vol6/iss1/art7>.
15. Polley E, van der Laan M. SUPERLEARNER: super learner prediction, 2011. Available at: <http://CRAN.R-project.org/package=SuperLearner>, r package version 2.0-4.
16. van der Laan MJ. Targeted maximum likelihood based causal inference part i. *Int J Biostat* 2010;6. Available at: <http://www.bepress.com/ijb/vol6/iss2/2>.
17. van der Laan MJ. Targeted maximum likelihood based causal inference part ii. *Int J Biostat* 2010;6. Available at: <http://www.bepress.com/ijb/vol6/iss2/3>.
18. Porter KE, Gruber S, van der Laan MJ, Sekhon JS. The relative performance of targeted maximum likelihood estimators. *Int J Biostat* 2011;7. Available at: <http://www.bepress.com/ijb/vol7/iss1/31>.
19. Stitleman OM, van der Laan MJ. Collaborative targeted maximum likelihood for time to event data. *Int J Biostat* 2010;6. Available at: <http://www.bepress.com/ijb/vol6/iss1/21>.
20. Stitleman OM, Wester CW, De Gruttola V, van der Laan MJ. Targeted maximum likelihood estimation of effect modification parameters in survival analysis. *Int J Biostat* 2011;7. Available at: <http://www.bepress.com/ijb/vol7/iss1/19>.
21. van der Laan MJ. Estimation based on case-control designs with known prevalence probability. *Int J Biostat* 2008;4. Available at: <http://www.bepress.com/ijb/vol4/iss1/17>.
22. Tuglus C, van der Laan MJ. Repeated measures semiparametric regression using targeted maximum likelihood methodology with application to transcription factor activity discovery. *Stat Appl Genet Mol Biol* 2011;10. Available at: <http://www.bepress.com/sagmb/vol10/iss1/art2>.

23. Brooks J, van der Laan MJ, Go AS. Targeted maximum likelihood estimation for prediction calibration. *Int J Biostat* 2012; 8(1). DOI: 10.1515/1557-4679.1385.
24. Quale CM, van der Laan MJ. Inference with bivariate truncated data. *Lifetime Data Anal* 2000;6:391–408.
25. Petersen ML, Schwab J, Gruber S, Blaser N, Schomaker M, van der Laan MJ. Targeted maximum likelihood estimation for dynamic and static longitudinal marginal structural working models. Technical Report 312, Division of Biostatistics, University of California, 2013.
26. van der Laan MJ, Gruber S. Collaborative Double robust targeted maximum likelihood estimation. *Int J Biostat* 2010;6. Available at: <http://www.bepress.com/ijb/vol6/iss1/17>.
27. van der Laan MJ, Robins JM. Unified methods for censored longitudinal data and causality. New York: Springer-Verlag, 2003.