The International Journal of Biostatistics

Volume 7, Issue 1

2011

Article 19

Targeted Maximum Likelihood Estimation of Effect Modification Parameters in Survival Analysis

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

C. William Wester, *Vanderbilt University School of
Medicine, Vanderbilt Institute for Global Health, and
Harvard School of Public Health*

Victor De Gruttola, *Harvard School of Public Health*

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

Recommended Citation:

Stitelman, Ori M.; Wester, C. William; De Gruttola, Victor; and van der Laan, Mark J. (2011)
"Targeted Maximum Likelihood Estimation of Effect Modification Parameters in Survival
Analysis," *The International Journal of Biostatistics*: Vol. 7: Iss. 1, Article 19.

DOI: 10.2202/1557-4679.1307

Available at: <http://www.bepress.com/ijb/vol7/iss1/19>

©2011 Berkeley Electronic Press. All rights reserved.

Targeted Maximum Likelihood Estimation of Effect Modification Parameters in Survival Analysis

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

Abstract

The Cox proportional hazards model or its discrete time analogue, the logistic failure time model, posit highly restrictive parametric models and attempt to estimate parameters which are specific to the model proposed. These methods are typically implemented when assessing effect modification in survival analyses despite their flaws. The targeted maximum likelihood estimation (TMLE) methodology is more robust than the methods typically implemented and allows practitioners to estimate parameters that directly answer the question of interest. TMLE will be used in this paper to estimate two newly proposed parameters of interest that quantify effect modification in the time to event setting. These methods are then applied to the Tshepo study to assess if either gender or baseline CD4 level modify the effect of two cART therapies of interest, efavirenz (EFV) and nevirapine (NVP), on the progression of HIV. The results show that women tend to have more favorable outcomes using EFV while males tend to have more favorable outcomes with NVP. Furthermore, EFV tends to be favorable compared to NVP for individuals at high CD4 levels.

KEYWORDS: causal effect, semi-parametric, censored longitudinal data, double robust, efficient influence curve, influence curve, G-computation, Targeted Maximum Likelihood Estimation, Cox-proportional hazards, survival analysis

Author Notes: Portions of the work in this article were supported in part by the Russell M. Grossman Endowment, NIEHS grant ES015493, and NIH grants R01 AI074345-04 and R01 AI 51164.

1 Introduction

Current methods used to evaluate effect modification in time to event data, such as the Cox proportional hazards model or its discrete time analogue, the logistic failure time model, posit highly restrictive parametric models and attempt to estimate parameters which are specific to the model proposed (Cox (1972)). These methods, as a result of their parametric nature, tend to be biased and force practitioners to estimate parameters out of computational convenience rather than estimate quantities most directly related to questions of primary scientific interest. The targeted maximum likelihood estimation (TMLE) methodology, originally proposed by van der Laan and Rubin (2006), and applied to time to event outcomes by Moore and van der Laan (2009) improves on the currently implemented methods in both robustness, its ability to provide unbiased estimates, and flexibility, allowing practitioners to estimate parameters that directly address their question of interest. Two new parameters of interest designed to quantify causal effect modification in time to event studies will be presented here as well as extensions of methods presented by Moore and van der Laan (2009) for their estimation. Finally, as an example, the presented methods will be used to re-evaluate the statistical questions addressed by Wester et al. (2010) in the recently completed Adult Antiretroviral Treatment and Drug Resistance (*Tshepo*) study; an open-label, randomized, 3x2x2 factorial design HIV study conducted at Princess Marina Hospital in Gaborone, Botswana to evaluate the efficacy, tolerability, and development of drug resistance of six different first-line combination antiretroviral therapy (cART) regimens.

Effect modification by a particular variable occurs when the effect of a particular treatment, or intervention, is different at different levels of the variable. So, for example, if men react differently to a drug than women, then it is said that gender modifies the effect of the drug. Within the statistical literature effect modification is also referred to as interaction. That is the treatment variable interacts with the effect modifier on the outcome of interest. The reason that we use the term effect modification, and more specifically causal effect modification, is to acknowledge that we are in fact estimating causal parameters. There has been a lot of literature devoted to estimating interaction and effect modification across a wide range of subject matters. In fact, almost all introductory statistics books and books on experimental design address interaction (see e.g. Stone (1996) and Box, Hunter, and Hunter (2005)). Effect modification is also an extremely relevant topic in the causal literature and specifically the causal literature that focuses on robust parameter estimation (see e.g. Vansteelandt et al. (2008) and Tchetgen Tchetgen (2010)). In this article we address causal effect modification using the TMLE framework. Specifically, we examine several causal parameters that address causal effect modification on time-to-event outcomes. Though these parameters are posed in the

time-to-event framework similar causal effect parameters of interest may be used to quantify causal effect modification in different data structures and for different outcomes.

2 Data Structure

Time to event data captures information about the amount of time, T , required for a subject to experience a particular event. Usually one is interested in assessing the effect of a particular treatment, A , on the amount of time, T , until the occurrence of the event of interest.

Figure 1: Data Structures.

For the methods presented here we are only concerned with binary levels of treatment, $A \in \{0, 1\}$, and vector of covariates, W , as they are observed at baseline. We assume T is discrete and takes on the values $\{1, \dots, \tilde{T}\}$, where $\tilde{T}$ is the last possible time the subjects are monitored. In the case that T is not discrete one may discretize time into fine enough cut points as to preserve any signal in the data. The censoring time, C , is the last time that a subject is observed, which might be marked by the end of the study or an earlier censoring event (i.e. lost to follow-up). Thus, the observed data consists of n i.i.d. replicates of, $O = (A, W, \tilde{T}, \Delta)$, where $\tilde{T} = \min(T, C)$ and $\Delta = I(T \leq C)$. So, $\tilde{T}$ is the last time point observed for a particular individual and Δ is an indicator variable which denotes whether or not the event was observed at that time point. In the observed data each subject would account for one line in the data set and we will refer to this data structure as the short form. Data structure I in Fig. 1 is a sample data set of four subjects from this observed data structure. $W1$ and $W2$ are two sample covariates which were measured at baseline.

An alternative data structure, which we will refer to as the long form, is a more appropriate way of thinking of the data for our purposes. Additionally, define $N_1(t) = I(\tilde{T} \leq t, \Delta = 1)$ as the counting process which denotes whether an event has occurred or not and $N_2(t) = I(\tilde{T} \leq t, \Delta = 0)$ as the counting process which codes right-censoring events. Thus $dN_1(t) = 0$ for all time points up until there is an observed failure time event, and at the time of the observed event $dN_1(t) = 1$. After a censoring event $dN_1(t)$ can never jump from 0 to 1 since the event can no longer be observed. Similarly, $dN_2(t)$ remains 0 for all times at which a particular observation is uncensored, and $dN_2(t) = 1$ at the time at which the observation becomes censored. Thus, the observed data may be represented in its long form as n i.i.d observations of $O = (A, W, dN_1(t), dN_2(t) : t = 1, \dots, \tilde{T}) \sim P_0$ where P_0 denotes the probability distribution of the observed data structure, O . In this representation of the observed data structure each subject contributes a line of data for each time point at which the subject was observed up until the occurrence of an event or censoring. Furthermore, the values measured at baseline, A and W , are repeated at their initial values for each time point. Data structure II in Fig. 1 shows the exact same observations from Data Structure I, but in their long form.

In Fig. 1 the subject with ID equal to 1 has $\tilde{T} = 3$ and $\Delta = 1$. So this individual was observed for three time periods and during the third, the subject had the event of interest. As a result subject 1 contributes three lines of data to the long form, as can be seen in Data Structure II in Fig. 1. The subject's event process, $N_1(t)$, remains zero up until time 3 when it jumps to 1, and the subject's censoring process, $N_2(t)$, remains 0 for all time points since the subject is never right-censored. All covariates measured at baseline remain the same for all time points. Subject 2, on the other hand, was right-censored at time point 2, resulting

in $\tilde{T} = 2$ and $\Delta = 0$. Thus, Subject 2 contributes two lines to the data in its long form and $N_1(t)$ is zero for all t , since the event was not observed, and the censoring-counting process, $N_2(t)$, jumps to 1 at time 2 since the subject was censored at that time. A , W_1 and W_2 remain constant at every time point. Contributions from subjects 3 and 4 to the data sets in both the short and long form are included in Fig. 1 as additional examples. The data in its long form may be thought of as being conditional on the failure time event (implying that right-censoring did not occur before time t). So each observation contributes a line in the long form of the data if that subject has not yet experienced the event or been censored.

3 Cox Proportional Hazards and Logistic Failure Time Models

The Cox proportional hazards model for the estimation of survival curves is based on the proportional hazards assumption, and assumes that the effect of treatment and covariates on the hazard follows a particular parametric form. This assumption implies that survival curves for different strata must have hazard functions whose ratio is constant in time. Despite the importance of the resulting theory and generalized multiplicative intensity models for modeling intensity of general counting processes (Andersen, Borgan, Gill, and Keiding (1993)), its stringent assumptions make these models subject to similar limitations as parametric regression models. More than thirty five years after the introduction of the model the proportional hazard assumption is still assumed to hold for the majority of survival data analyses. Fortunately, not all semi-parametric methods in survival analysis are susceptible to the shortcomings of the proportional hazards model. In this section, we will discuss the flaws of the Cox proportional hazards model as well as its discrete time analogue.

It is common in analyses that intend to test whether or not a treatment or exposure, A , has an effect on a particular time to event outcome, to assume an *a priori* specified model for the conditional hazard; in such models one tests if the coefficient on A in the specified model differs from zero. As described above, a Cox model is often employed for continuous time; and a logistic failure time model for discrete time failure times. It is common practice in both cases to fit time as flexibly as possible and to model the effect of the treatment, and for covariates, if adjustment for them is desired, with a linear parametric model. Regardless of whether one does or does not adjust for baseline covariates, W , these methods are usually biased due to their dependence on highly restrictive parametric models that are not typically representative of the data generating distribution.

Another common practice in implementing these methods is to test the validity of the assumptions after implementing the method (see e.g. Chapter 4 of Collett (1991)). Many tests, both graphical and formal have been devised in order to test whether the proportional hazards assumption holds. In cases when the proportional hazards assumption is rejected some practitioners may adjust the model so that the assumption is not rejected. One way this is done is by introducing time-dependent covariates to the regression model. However, even if these tests are implemented, adjustments are made if the assumption is violated, and one proceeds as if the model is correct once the test is not rejected, the subsequent inference ignores that the model was data adaptively selected. Thus, there is no theoretical reason to rely on the p-values or confidence intervals produced by such an approach. Moreover, data adaptively selecting a model in this fashion does not result in a correctly specified model.

The underlying parameter being estimated is a major issue for this approach as well. Interpretation of the parameters of these models may be compromised, even if the models are correct, by the fact that they were chosen for convenience of estimation rather than because they are natural parameters generally applicable to less restrictive models of the hazard. The parameters in the Cox model may be represented in terms of an average over time of log hazard ratios or the log of the ratio of the log conditional survival probabilities. These two representations are very different parameters for most conditional distributions of T , given (A, W) , with very different interpretations, but they happen to be equal to each other for distributions that satisfy the constraints of the Cox proportional hazards model. The complexity of these parameters may make them difficult to interpret when trying to quantify the effect of A on the outcome. Thus, the parameters in these models may not directly address the questions practitioners are interested in answering. For this reason methods have been proposed that estimate more easily interpretable parameters of interest in time to event studies such as risk differences, or a closely related measure, the number needed to treat (see e.g. Altman and Andersen (1999) and Laubender and Bender (2010)). However, these methods are based on using the Cox proportional hazards model for estimating the underlying conditional hazard, and thus rely on the model being specified correctly to create unbiased estimates.

4 Defining the Parameter of Interest and Causal Effects

In this section we examine the advantages of defining the parameter of interest as a function of the data generating distribution as well as introduce several parameters

of interest in time to event studies. One can define the parameter of interest as a function of the data generating distribution, $\Psi(P_0)$, which permits estimation of exactly the feature of the data generating distribution of interest. The meaning of such parameters does not depend on an assumption of proportional hazards.

The treatment specific survival curve at a particular time point, t_k , is a simple example of such a parameter, $Pr(T_a > t_k)$, where T_a is the counterfactual event time, T , one would have observed had an individual's treatment been set, possibly contrary to fact, to treatment level a . The outcome which would have been observed had the subject been treated by a different level of A than that actually received is known as a counterfactual outcome. By formulating a set of causal assumptions through the use of a non-parametric structural equation model (NPSEM) one may define the distribution of the counterfactual outcomes indexed by an intervention on some treatment and censoring nodes in the NPSEM. The assumptions of NPSEM may be visualized through the use of a causal graph, another common tool for positing causal assumptions in the causal inference literature. Then one may define a causal parameter as a difference between the counterfactual distributions of the outcome for different interventions on the treatment and censoring nodes in the causal graph.

In order to link the NPSEM to the observed data, the distribution of the observed data is assumed to be implied by the causal graph/NPSEM. The NPSEM includes the distribution of the error/exogenous nodes, and the deterministic functions that define each endogenous node as a function of its parents and an error/exogenous node. Typically, one assumes that the observed data structure corresponds with a subset of the nodes, such as all the endogenous nodes. Under certain causal assumptions one can then show that the causal effect of interest can be identified as a parameter of the distribution of the observed data, $\Psi(P_0)$.

Fig. 2 posits a set of causal assumptions in the form of a causal graph/NPSEM in which our data structure corresponds with the displayed nodes. Exogenous nodes are not shown in the graph. This is common practice for displaying that each displayed node has an error/exogenous node with an arrow going into it and no arrows going into any other nodes.

Necessary conditions typically imposed to make causal parameters identifiable may be made through the use of this causal graph or nonparametric structural equation model, i.e. the consistency assumption and coarsening at random (CAR) or sequential randomization assumption (SRA). The consistency assumption states that the observed data consist of the counterfactual outcome corresponding with the intervention actually observed. This assumption is a direct consequence of defining the observed data structure as particular nodes in the causal graph/NPSEM. The strong sequential randomization assumption on the intervention nodes of the causal graph assumes that, for each intervention node, the conditional density of an inter-

Figure 2: Causal Assumptions

vention node, given the collection of all counterfactuals and the intervention nodes that are parents of the intervention node, is a function only of the observed parents of the intervention node. This is directly implied by the causal graph, as the only arrows into the treatment and censoring process variables are from their observed ancestors and no other nodes (including no error nodes). For the treatment variable this assumption is sometimes referred to as the randomization of treatment assumption or no unmeasured confounders assumption. Since in the analysis presented below the treatment assignment is in fact randomized there is no arrow from W to A in the causal graph. Note, that simple unobserved nodes that effect only treatment, A , or the censoring process, $dN_2(t)$, but not the future outcome process, $dN_1(t)$, do not violate these assumptions. This is because such nodes do not produce unblocked backdoor paths to the future outcome process, $dN_1(t)$. For an in depth account of causality and graphical models, see Pearl (2008).

One other assumption is necessary in order for the parameters of interest to be identifiable from the observed data; however, this assumption can not be expressed within the causal graph. This assumption is known as the experimental treatment assumption (ETA) (Neugebauer and van der Laan (2005)), or, the positivity assumption (Hernan and Robins (2006)). The assumption states, in particular, that there is no level of W that is completely predictive of treatment level, A . However, since treatment is randomized the design of the experiment in the *Tshepo* study assures that this assumption is met. In addition, since we are only interested in the censoring nodes set at value 0, the assumption states that for all level of covariates W there is a positive probability of being uncensored, or that $dN_2(t) = 0$.

The causal graph in Fig. 2 suggest the following factorization of the likelihood:

$$P_0(O) = \underbrace{P(W)}_{Q_{10}} \underbrace{P(A)}_{g_{10}} \prod_{t=1}^{\dot{T}} \underbrace{P(dN_1(t) | \bar{N}_1(t-1), \bar{N}_2(t-1), A, W)}_{Q_{20}} \prod_{t=1}^{\dot{T}} \underbrace{P(dN_2(t) | \bar{N}_1(t), \bar{N}_2(t-1), A, W)}_{g_{20}},$$

where, the bar notation is used to represent the past history of the particular process (e.g. $\bar{N}_1(t) = \{dN_1(1) \dots dN_1(t)\}$). Thus, the likelihood is factorized into a portion Q_0 , corresponding to the conditional distributions of the non-intervention nodes, and a portion g_0 , corresponding to the conditional distributions of the censoring and treatment nodes. Q_0 is composed of Q_{10} , the distribution of the baseline covariates, and Q_{20} the conditional distribution of the binary indicators $dN_1(t)$, given its parents, including the treatment, A , and the baseline covariates, W . g_0 is further factorized into the treatment mechanism, g_{10} , and censoring mechanism, g_{20} , which involves the conditional distributions of the binary censoring indicators $dN_2(t)$, given its parents, including the treatment A , and the baseline covariates, W . We note that Q_{20} and g_{20} are identified by the discrete-intensities

$$E_0(dN_1(t) | A, W, N_1(t-1), N_2(t-1)) \equiv I(\tilde{T} \geq t) Q_{20}(t, A, W)$$

$$E_0(dN_2(t) | A, W, N_1(t), N_2(t-1)) \equiv I(N_1(t) = 0, N_2(t-1) = 0) g_{20}(t, A, W).$$

Note that these two intensities of the event process N_1 and censoring process N_2 indeed equal an indicator of being at risk times a function of t, A, W . Under the coarsening at random assumption, it follows that $Q_{20}(t, A, W) = P_0(T = t | T \geq t, A, W)$ equals the conditional hazard of the failure time T . Let's also define $S_0(t_k | A, W) = Pr(T > t_k | A, W)$ which is the conditional survival of the event of interest, which can be expressed as a function of the conditional hazard $Q_{20}(t, A, W)$, under the coarsening at random assumption, as follows:

$$S_{0,a}(t_k | W) = \prod_{t=1}^{t_k} [1 - P(dN_1(t) = 1 | \bar{N}_1(t-1) = 0, \bar{N}_2(t-1) = 0, A = a, W)]$$

By intervening on the treatment A and censoring process $dN_2(t)$ in the NPSEM described by Fig. 2, by setting treatment A equal to the desired treatment, a , and setting $dN_2(t)$ equal to 0, or, equivalently, no censoring, for all t , one obtains the distribution of the event process under the desired treatment and without censoring. This counterfactual distribution of the data structure under an intervention can be identified from the observed data under the stated assumptions, and the resulting formula is known as the G-computation formula (Robins (1986)). All of the nodes which are intervened upon in the causal graph are set to their intervened upon level in the likelihood, and all the conditional distributions for those nodes are removed from the likelihood since they are no longer random variables. The following is the resulting G-computation formula, or distribution of the data under the intervention:

$$P(W) \prod_{t=1}^T P(dN_1(t) | \bar{N}_1(t-1) = 0, \bar{N}_2(t-1) = 0, A = a, W).$$

Based on these causal assumptions and the G-computation formula we can now write the marginal treatment specific survival probability, $P(T_a > t_k)$, in terms of the data generating distribution, P_0 :

$$\Psi_a(P_0)(t_k) = Pr(T_a > t_k) = E_W(S_{0,a}(t_k | W)).$$

Parameters which combine $\Psi_1(P_0)(t_k)$ and $\Psi_0(P_0)(t_k)$ allow one to quantify the effect of a change of treatment A on survival, T . Three examples of these types of parameters are the marginal additive difference in the probability of survival (or risk difference), the log relative risk of survival, and the marginal log relative hazard of survival:

$$\Psi_{RD}(P_0)(t_k) = \Psi_1(P_0)(t_k) - \Psi_0(P_0)(t_k), \quad (1)$$

$$\Psi_{RR}(P_0)(t_k) = \log \left(\frac{\Psi_1(P_0)(t_k)}{\Psi_0(P_0)(t_k)} \right), \quad (2)$$

$$\Psi_{RH}(P_0)(t_k) = \log \left(\frac{\log(\Psi_1(P_0)(t_k))}{\log(\Psi_0(P_0)(t_k))} \right). \quad (3)$$

For ease of interpretation we prefer the first two of these parameters; however, for completeness we show that the third parameter, which represents the parameter targeted by a marginal structural Cox Proportional Hazard model $P(T_a = t |$

$T_a \geq t) = \lambda_0(t) \exp(\beta a)$, may also be estimated through the methods presented here. However, the parameter expressed in equation (3) is undefined at t_k for which $P(T_a > t_k) = 1$. The above parameters quantify the effect of A at a particular time point, t_k ; averages of the above parameters over a set of time points may also be of interest.

One possible method for estimating the mean counterfactual outcome as expressed in the parameters proposed above is maximum likelihood estimation (MLE). The MLE approach to estimating $\Psi(Q_0)$ is to obtain a maximum likelihood estimate, Q_n , of Q_0 and construct the substitution estimator $\Psi(Q_n) = \psi_n^{MLE}$ according to the mapping $\Psi(\cdot)$.¹ Under the counterfactual distribution provided by the G-computation formula above under Q_n , ψ_n^{MLE} may be constructed. Thus, it is necessary to estimate both the marginal distribution of W , $Q_{10}(W)$, and the conditional hazard, $Q_{20}(t, A, W)$, and consequently its corresponding conditional survival function, $S_0(t_k | A, W)$. ψ_n^{MLE} is consistent when both of these distributions are estimated consistently. The marginal distribution of W , $Q_{10}(W)$, is estimated with the non-parametric maximum likelihood estimator, the empirical distribution of the baseline covariates in the sample. In practice several models have been used for $Q_{20}(t, A, W)$. Many practitioners *a priori* specify a parametric model, often using main terms logistic regression and obtain the maximum likelihood estimate within this highly restrictive model. This approach almost certainly leads to biased estimates of $Q_{20}(t, A, W)$, and thus a biased estimate of $\Psi(Q_0)$, since it is likely that the true Q_{20} is not contained in this highly restrictive model. If the model does not contain the true Q_{20} then the estimates of Q_{20} and $\Psi(Q_0)$ will generally not be consistent. For this reason a non-parametric model for Q_{20} often represents the available knowledge about Q_{20} . Since this nonparametric model is very large, sieve-based (data adaptive) maximum likelihood estimation (i.e., loss-based machine learning), involving fine-tuning of the amount of smoothing used becomes necessary. A detailed account of loss based learning and cross validation may be seen in van der Laan and Dudoit (2003). In addition, van der Laan, Dudoit, and Keles (2004) focus specifically on likelihood-based cross-validation and illustrate its asymptotic optimality. Many different methods have been developed for estimating the effect of covariates on a binary outcome in the non-parametric model including regression trees, DSA, and k-nearest neighbors. Rudser, LeBlanc, and Emerson (2008) use one of these methods, a non-parametric tree estimator, for estimating the survival distribution and subsequently $\Psi(Q_0)$. Even though a single estimator for these distributions in the non-parametric model may yield consistent estimates van der Laan, Polley, and Hubbard (2007) show that super learner, which combines many differ-

¹Since the parameter of interest is only a function of Q_0 , and not the g -part of the likelihood it is written here and in some instances below as $\Psi(Q_0)$ as opposed to $\Psi(P_0)$.

ent learners, will converge to the true distribution at a faster rate (van der Laan et al. (2007)). Super Learner is a data-adaptive machine learning algorithm which uses log-likelihood-based cross validation to choose among a library of different learners and is a recommended method for estimating Q_{20} because of its advantages in finite samples.

Once these estimates are obtained the MLE estimate of the treatment specific survival, $\Psi_a(P_0)(t_k)$, may be expressed as:

$$\hat{\Psi}_a^{MLE}(P_n)(t_k) = \frac{1}{n} \sum_{i=1}^n S_{n,a}(t_k | W_i).$$

The marginal distribution of W is estimated using the empirical distribution; this is expressed in the above $\hat{\Psi}_a^{MLE}$ as an empirical mean of the estimated survival probabilities. Fig. 3 is a diagram that shows how to map an estimate Q_{2n} of Q_{20} and the empirical distribution of W , Q_{1n} , into $\hat{\Psi}_1(Q_n)$, an estimate of the treatment specific survival curve at $A=1$. The MLE estimate of the parameters presented in equations (1)-(3) may be generated by combining MLE estimates of the treatment specific survival curves under alternative treatments.

5 Parameters of Interest which Quantify Causal Effect Modification

Parameters which quantify the level of effect modification due to another baseline variable V may also be of interest in studies aimed to assess the effect of a treatment, A , at different levels of V . However, we must now consider a causal graph with an explicit node for the effect modifier V , instead of collapsing it in single node W . First, consider V , which occurs after the baseline covariates, W , and before the rest of the censoring and event processes. This would correspond to a causal graph like in Fig. 2 but with an additional node V with an arrow coming into it from W and from V to the future censoring and event process. The causal graph just described states the causal assumptions regarding the potential effect modifier, V . This causal graph implies the following factorization of the likelihood:

$$\begin{aligned} P_0(O) &= P(W)P(V | W)P(A | W, V) \\ &\quad \prod_{t=1}^T P(dN_1(t) | \bar{N}_1(t-1) = 0, \bar{N}_2(t-1) = 0, A, V, W) \\ &\quad \prod_{t=1}^T P(dN_2(t) | \bar{N}_1(t) = 0, \bar{N}_2(t-1) = 0, A, V, W), \end{aligned}$$

Figure 3: How to map Q_{2n} and the empirical distribution of the baseline covariates, Q_{1n} , into $\hat{\Psi}(Q_n)$.

and corresponding G-computation formula for the intervention $A = a$, $V = v$, and all censoring nodes equal to zero:

$$P(W) \prod_{t=1}^T P(dN_1(t) \mid \bar{N}_1(t-1) = 0, \bar{N}_2(t-1) = 0, A = a, V = v, W). \quad (4)$$

The new parameter being targeted is the causal effect of the joint intervention (A, V) on T . As before, the G-computation estimate only requires an estimate of the marginal distribution of W , and the conditional hazard, $Q_{20}(t, A, W)$. Thus, we will define the following parameters of interest to measure the causal effect modification when the effect modifier occurs after W :

$$\Psi_{RD}^{CEM} = \sum_{t_k} w(t_k) [(\Psi_{11}(P_0)(t_k) - \Psi_{01}(P_0)(t_k)) - (\Psi_{10}(P_0)(t_k) - \Psi_{00}(P_0)(t_k))] \quad (5)$$

$$\Psi_{LR}^{CEM} = \sum_{t_k} w(t_k) \left[\log\left(\frac{\log(\Psi_{11}(P_0)(t_k))}{\log(\Psi_{01}(P_0)(t_k))}\right) - \log\left(\frac{\log(\Psi_{10}(P_0)(t_k))}{\log(\Psi_{00}(P_0)(t_k))}\right) \right] \quad (6)$$

where,

- $\Psi_{11}(P_0)(t_k)$ is counterfactual survival at t_k setting V to 1 and A to 1
- $\Psi_{01}(P_0)(t_k)$ is counterfactual survival at t_k setting V to 1 and A to 0
- $\Psi_{10}(P_0)(t_k)$ is counterfactual survival at t_k setting V to 0 and A to 1
- $\Psi_{00}(P_0)(t_k)$ is counterfactual survival at t_k setting V to 0 and A to 0

That is, the first subscript corresponds to treatment, A , and the second, to the level of the potential effect modifier, V . Here $w(t_k)$ are time varying weights. These weights may be dependent on the question of interest, may be set equal to the reciprocal of an estimate of the variance of the parameter estimate at the particular time, or simply set equal to 1. For the remainder of this article we will weight each time point specific estimator by the reciprocal of its estimated variance. This is done to emphasize time points for which more information is available, such as those with less prior censoring.

In situations where the effect modifier, V , is realized before the baseline covariates, W , a different causal graph must be considered. In this causal graph an additional node V is added to Fig. 2 with an arrow going from V into W , and all of the event and censoring process nodes. This causal graph implies the following factorization of the likelihood:

$$\begin{aligned}
 P_0(O) &= P(V)P(W | V)P(A | W, V) \\
 &\quad \prod_{t=1}^T P(dN_1(t) | \bar{N}_1(t-1) = 0, \bar{N}_2(t-1) = 0, A, V, W) \\
 &\quad \prod_{t=1}^T P(dN_2(t) | \bar{N}_1(t) = 0, \bar{N}_2(t-1) = 0, A, V, W),
 \end{aligned}$$

and resulting G-computation formula:

$$P(W | V = v) \prod_{t=1}^T P(dN_1(t) | \bar{N}_1(t-1) = 0, \bar{N}_2(t-1) = 0, A = a, V = v, W).$$

Characteristics such as sex and gender, which are set at birth, are examples of variables for which V is realized before W . In order to assess the level of effect modification due to these types of variables, an estimate of the conditional distribution of W , given V , $P_0(W | V = v)$, is needed instead of the marginal distribution $P_0(W)$ of W . The conditional hazard, $Q_{20}(t, A, W)$ may be estimated as before. The empirical distribution among $V = v$ will be used to estimate $P(W | V = v)$. We will refer to these causal effect modification parameters as stratified effect modification (SEM) parameters, as they can be expressed as follows:

$$\Psi_{RD}^{SEM} = \left[\sum_{t_k} w(t_k)(S_{1|1}(t_k) - S_{0|1}(t_k)) \right] - \left[\sum_{t_k} w(t_k)(S_{1|0}(t_k) - S_{0|0}(t_k)) \right] \quad (7)$$

$$\Psi_{LR}^{SEM} = \sum_{t_k} w(t_k) \left[\log \left(\frac{\log(S_{1|1}(t_k))}{\log(S_{0|1}(t_k))} \right) - \log \left(\frac{\log(S_{1|0}(t_k))}{\log(S_{0|0}(t_k))} \right) \right], \quad (8)$$

where

- $S_{1|1}(t_k)$ denotes survival at t_k for individuals with $V = 1$ and treatment set to 1
- $S_{0|1}(t_k)$ denotes survival at t_k for individuals with $V = 1$ and treatment set to 0
- $S_{1|0}(t_k)$ denotes survival at t_k for individuals with $V = 0$ and treatment set to 1
- $S_{0|0}(t_k)$ denotes survival at t_k for individuals with $V = 0$ and treatment set to 0.

These are also counterfactual survival probabilities consistent with setting $A = a$ and $V = v$; however, we have expressed them as $S_{a|v}(t_0)$ to emphasize the fact that they are different parameters of interest corresponding to the alternative causal graph, and corresponding factorization of the likelihood, in which V occurs before W .

The two examples above illustrate how the causal parameter of interest is directly related to the causal assumptions made by the NPSEM/causal graph. As a result the parameter being estimated is different depending on the causal assumptions concerning the effect modifier of interest. This illustrates how crucial it is to make valid causal assumptions. If the time ordering is unknown, then the correct NPSEM may not be specified, and the appropriate parameter of interest cannot be identified. In such a situation causal inference will typically not be applicable. However, If time ordering is not known, say it is unclear if W precedes A in chronological time, then one can still proceed if one knows that W is not affected by A , i.e $W(a) = W$.

The parameters proposed above are for a binary effect modifier, V . Similar parameters based on ordered categorical values are also possible. The difficulty with effect modification parameters that account for many levels of V or a continuous V is that even if they are valid there may not be enough support in the data to estimate them efficiently. If there are a few categorical levels of V a direct extension of the parameters proposed here is possible (e.g. a parameter of interest may be targeted that contrasts the risk difference of treatment across levels of V). An approach for estimating effect modification parameters based on a continuous

V or for V with a large number of categories is addressed in the discussion. Moreover, effect modification parameters of interest may be more difficult to interpret for many levels of V or a continuous V .

6 Targeted Maximum Likelihood Estimation

This section explores how TMLE may be used to estimate the parameters of interest introduced in the two previous sections. The TMLE of the causal parameters presented above improves on the MLE estimate in that it is consistent when either the treatment, $g_{10}(A, W)$ and censoring mechanism, $g_{20}(N_2(t), A, W)$, are estimated consistently, or Q_0 is estimated consistently. Thus, the method is double robust. In the TMLE algorithm, an initial hazard is estimated and the algorithm updates it by iteratively adding to it a time-dependent clever covariate. The latter is chosen specifically to reduce bias in the estimate of the parameter of interest; it is a function of time, treatment, and the baseline covariates, and requires an estimate of the treatment and censoring mechanisms.

Consider an initial estimator P_n^0 of P_0 consisting of estimates $Q_{2n}^0(t, A, W)$, $g_{1n}^0(A, W)$, $g_{2n}^0(t, A, W)$, and the empirical distribution of $W_1, \dots, W_n$. The 0 on the top indicates that it is the estimate prior to any iterations of the targeted maximum likelihood process. By updating the initial estimate $Q_{2n}^0(t, A, W)$ through addition of the term εh such that the score for the parametric fluctuation model at $\varepsilon = 0$ is equal to the corresponding component of the efficient influence curve, an updated distribution targeted toward the parameter of interest may be constructed. H^* is the time dependent clever covariate and will be written as $H^*(t, A, W)$ to stress its dependence on time. In practice this may be done by implementing a univariate logistic regression to fluctuate the initial estimate of the hazard by regressing the binary outcome $dN_1(t)$ on $H^*(t, A, W)$ using the initial estimate $Q_{2n}^0(t, A, W)$ as an offset (this can be implemented using a statistical programming language such as R).

In Moore and van der Laan (2009) the following time dependent covariates, $H^*(t, A, W)$, were shown to yield the appropriate score for targeting parameters $\Psi_1(P_0)(t_k)$ and $\Psi_0(P_0)(t_k)$ respectively:

$$H_1^*(t, A, W) = - \frac{I(A = 1)}{g_{10}(1 | W) \prod_{i=1}^{t-1} [1 - g_{20}(i | A, W)]} \frac{S_0(t_k | A, W)}{S_0(t | A, W)} I(t \leq t_k)$$

and

$$H_0^*(t, A, W) = - \frac{I(A = 0)}{g_{10}(0 | W) \prod_{i=1}^{t-1} [1 - g_{20}(i | A, W)]} \frac{S_0(t_k | A, W)}{S_0(t | A, W)} I(t \leq t_k).$$

The target parameters $\Psi_1(P_0)(t_k)$ and $\Psi_0(P_0)(t_k)$ have the following efficient influence curves:

$$D_1^*(P_0) = \sum_{t \leq t_k} H_1^*(t, A, W) [I(\tilde{T} = t, \Delta = 1) - I(\tilde{T} \geq t) Q_{20}(t, A = 1, W)] \\ + S_0(t_k | A = 1, W) - \Psi_1(P_0)(t_k),$$

and

$$D_0^*(P_0) = \sum_{t \leq t_k} H_1^*(t, A, W) [I(\tilde{T} = t, \Delta = 1) - I(\tilde{T} \geq t) Q_{20}(t, A = 0, W)] \\ + S_0(t_k | A = 0, W) - \Psi_0(P_0)(t_k).$$

Since the time-dependent clever covariates for fluctuating Q_{2n} are unknown due to their dependence on g_0 and Q_0 , estimates, $H_n^*(t, A, W)$, are constructed using the estimates g_n , and Q_n . So the clever covariate for estimating the treatment specific survival under treatment $A = 1$ at $t_k = 2$ for the first line of data structure II in Fig. 1 would be expressed as:

$$H_n^*(1, A = 1, W = (11, 3)) = -\frac{1}{g_{1n}(1 | W = (11, 3))} \frac{S_n(2 | A = 1, W = (11, 3))}{S_n(1 | A = 1, W = (11, 3))}$$

The clever covariate for the second row for the same individual is:

$$H_n^*(2, A = 1, W = (11, 3)) = -\frac{1}{g_{1n}(1 | W = (11, 3)) [1 - g_{2n}(1 | A = 1, W = (11, 3))]}$$

and at time point three $H_n^*(3, A = 1, W = (11, 3))$ will be zero since when $t > t_k$ the function $I(t \leq t_k)$ will evaluate to zero. For the subject represented by ID 2, the fourth and fifth row of data structure II, the clever covariate, $H_n^*(t, A, W)$ will equal zero since $I(A = 1)$ will evaluate to zero since that subject was treated at $A = 0$. Once the clever covariates are constructed, one targeting step may be achieved by specifying the following fluctuation of Q_{2n}^0 according to a parametric logistic regression model:

$$\text{logit} [Q_{2n}^1(\epsilon)(t, A, W)] = \text{logit} [Q_{2n}^0(t, A, W)] + \epsilon H_n^*(t, A, W).$$

This update is implemented by fitting a univariate logistic regression model of the event process, $N_1(t)$, on the clever covariate with the initial fit $Q_{2n}^0(t, A, W)$ as an offset. $Q_{2n}^1(\epsilon_n^1)(t, A, W)$ is the first step TMLE of Q_{20} , where ϵ_n^1 is the fitted regression coefficient. Since the time dependent covariates, $H^*(t, A, W)$, are functions of $Q_{20}(t, A, W)$ it is necessary to iterate this updating step. This process is detailed in Fig.4. The estimate $Q_{2n}^k(\epsilon_n^1)(t, A, W)$ plays the role of initial estimator, and a new update of this new initial estimator is obtained following the above procedure, resulting in the next $Q_{2n}^{k+1}(\epsilon_n^2)(t, A, W)$. This process is iterated until ϵ_n^k converges to zero. The resulting estimate after convergence is Q_{2n}^* . The empirical distribution

Figure 4: One iteration of the TMLE algorithm.

of W does not need to be updated since it is already a nonparametric maximum likelihood estimator: the updating process would not result in any updates. The resulting $(g_n, Q_n^* = (Q_{1n}, Q_{2n}^*))$ represents a density targeted toward estimating the specific parameters of interest. Finally, the targeted maximum likelihood estimates of $\Psi_a(P_0)(t_k)$ are the substitution estimators based on plugging in the targeted estimator Q_n^* :

$$\hat{\Psi}_a(Q_n^*)(t_k) = \frac{1}{n} \sum_{i=1}^n S_{n,a}^*(t_k | W_i), \quad (9)$$

where $S_n^*(t_k | A = a, W_i)$ is the survival probability corresponding with Q_{2n}^* .

$\hat{\Psi}_a(Q_n^*)(t_k)$ may be constructed as shown in Fig. 3. The above process should be implemented once for $\hat{\Psi}_1(P_n)(t_k)$ and a second time for $\hat{\Psi}_0(P_n)(t_k)$ since each parameter of interest has a different $H^*(t, A, W)$. Alternatively, two covariates can be added to the hazard, one for $\hat{\Psi}_1(P_n)(t_k)$ and one for $\hat{\Psi}_0(P_n)(t_k)$, and be updated simultaneously with ε_1 and ε_0 until they both converge toward zero.

The TMLE of $\Psi_{av}(t_k)$ needed for constructing the parameter estimates of equations (5) and (6) may be constructed in the same way as for one treatment variable, A , but by treating A and V together as a treatment variable. The treatment mechanism must be changed to estimate the joint probability of $A = a$ and $V = v$ and the resulting clever covariate is:

$$H^*(t, A, V, W) = - \frac{I(A = a, V = v)}{g_{10}(A = a, V = v | W) \prod_{i=1}^{t-1} [1 - g_{20}(i | A, V, W)]} \frac{S(t_k | A = a, V = v, W)}{S(t | A = a, V = v, W)} I(t \leq t_k)$$

and the TMLE is:

$$\hat{\Psi}_{av}(Q_n)(t_k) = \frac{1}{n} \sum_{i=1}^n S_{n,a,v}^*(t_k | W_i),$$

where $S_{n,a,v}^*$ is now the targeted maximum likelihood estimate of the treatment specific survival probability at $A = a$ and $V = v$. Thus, we need to adjust for confounders of the bivariate treatment (A, V) . The set of potential confounders for the bivariate treatment are included in W . Note however, that this requires modeling A given W and V given W . Different covariates from the set W may be included in these different treatment mechanisms.

The TMLE of $S_{a|v}(t_k)$ needed for constructing the parameter estimates of equations (7) and (8) may be constructed in the same way as for one treatment variable, A , but done separately for each level of V . So the TMLE of $S_{a|v}(t_k)$ is the same as the TMLE for $\Psi_a(t_k)$ just estimated on a data set that only includes individuals with $V = v$.

Since the TMLE $P_n^* = (Q_n^*, g_n)$ solves the efficient influence curve estimating equation $0 = P_n D^*(Q_n^*, g_n)$ under regularity conditions, the TMLE $\Psi(Q_n^*)$ of $\Psi(Q_0)$ is both double robust and locally efficient. As mentioned above, double robustness means that the parameter of interest is consistent if either the conditional hazard of T , Q_{20} , or the treatment mechanism $g_{10}(A | W)$, and censoring mechanism $g_{20}(t, A, W)$ are estimated consistently. The maximum likelihood based substitution

estimators presented above rely on the fact that Q_0 is estimated consistently. Local efficiency implies that the estimate achieves the semi-parametric efficiency bound when $Q_{20}(t, A, W)$, $g_{10}(A, W)$, and $g_{20}(t, A, W)$ are estimated consistently. For a more in-depth treatment of these properties, and the efficient influence curve see van der Laan and Robins (2003).

There are a few clear advantages of TMLE for time-to-event analysis that are a consequence of the double robustness and local efficiency of the method. First, if there is independent censoring and randomized treatment, the estimate of the parameter of interest is guaranteed to be asymptotically unbiased. This is because the treatment mechanism $g_{10}(A | W)$ is known and the censoring mechanism $g_{20}(t, A, W) = g_{20}(t)$ is known to be independent. Thus, the censoring mechanism may be consistently estimated with the Kaplan-Meier estimator. So by definition g_n is a consistent estimate of g and by the double robustness property the TMLE is consistent. However, it has also been shown that by estimating these mechanisms even when they are known one can improve the efficiency of the estimates through adjustment for empirical confounding. This is done by positing a model for these mechanisms that contains the truth, so that these estimates will converge to the truth. For a general account of how estimating the treatment and censoring mechanism can improve efficiency see van der Laan and Robins (2003) section 2.3.7. Second, with informative censoring and/or observational treatment the double robustness allows one to reduce bias due to the initial estimate of $Q_{20}(N_1(t), A, W)$ by estimating the treatment, $g_{10}(A, W)$, and censoring mechanism, $g_{20}(N_2(t), A, W)$ as well as possible.

The TMLE also improves on estimating equation based techniques, in which case ψ_0 is estimated with the closed form solution of the efficient influence curve estimating equation $P_n D^*(Q_n^0, g_n^0, \psi) = 0$. This is due to the fact that TMLE is a substitution estimator which, in particular, obeys the proper bounds of a survival probability. Estimating equation based results do not obey these bounds and may even result in estimates of probabilities that don't fall between zero and one. In cases where the treatment mechanism gets very close to zero, corresponding to practical violations of the ETA assumption stated above, the estimating equation based methods tend to become very unstable. When violations in the ETA assumption are a problem, the estimating equation based approaches may not only return estimates which are not probabilities, but may also suffer drastically in efficiency and therefore not approach the semi-parametric efficiency bound. However, TMLE estimates are more stable and may still achieve the semi-parametric efficiency bound

in these situations. Moreover, estimating equation based approaches that are unbounded can give undue influence to observations with very large weights, heavily influencing the final estimate and in some cases creating estimates that do not obey the proper bounds of the parameter space. In some situations a single observation could have such severe inverse probability weights that it results in an estimate that is larger than 1 or less than 0 when estimating a probability. This can never happen in TMLE since it is a substitution estimator that by definition obeys the proper bounds of the parameter space.

There has been some concern expressed in the literature that double robust methods that rely on a regression on a “clever” covariates which involve inverse weights in sparse data situations are unstable. In particular, this concern was expressed in Kang and Schafer (2007) regarding a method proposed in Bang and Robins (2005). Furthermore, Robins in his response to that paper expressed similar concerns and suggested alternative double robust methods that obeyed the proper bounds of the parameter space and relied on solving an estimating equation (Robins, Sued, Lei-Gomez, and Rotnitzky (2007)). As explained in Gruber and van der Laan (2010) the instability they observed in simulation is only a concern if one is estimating the effect of a treatment on a *continuous* outcome. The TMLE of the additive causal effect on a binary outcome is a robust substitution estimator. Moreover, Gruber and van der Laan show that the instability *only* occurs for continuous outcomes if one is using a least squares loss function and a linear fluctuation of the initial distribution. They then go on to create a stable TMLE that estimates the causal effect of treatment on a continuous outcome in terms of quasi-log-likelihood loss and logistic fluctuation model. Their paper fully explains this approach and demonstrates the stability of TMLE relative to the non-robust TMLE and estimating equation based estimators with a thorough simulation study. In addition, in Porter, Gruber, van der Laan, and Sekhon (forthcoming 2011), this TMLE of the additive causal effect on a continuous outcome is compared to estimating equation based double robust methods (including those that obey the proper bounds of the parameter space), as well as other causal methods, and TMLE is shown to be more stable than the other methods across a wide range of data generating distributions.

As immediately relevant to the TMLE presented in this article in Stitelman and van der Laan (2010) a simulation study is presented that exhibits the stability of TMLE for estimating the causal effect of treatment on survival. In that paper they compare TMLE and an extension, Collaborative-TMLE (C-TMLE), to other causal methods, including double robust estimating equation based approaches. Again, in that study the TMLE is shown to be more stable than the other methods. This increased stability may be attributed to the fact that TMLE is a substitution estimator, so that no single observation can throw off the estimator by more than $1/n$ (whereas,

a single observation can have a very large impact on an estimating equation based estimator since it is unbounded).

An additional concern for practitioners may be the difficulty and computational complexity involved in implementing the iterative method proposed here. However, in addition to the advantages expressed, in the preceding paragraphs, TMLE, in general, may be easily extended to more complicated data structures (including data structures that account for time-dependent covariates). Moreover, the instability of estimating equation based approaches may be exaggerated in these more complex data structures. So a major benefit of this additional complexity is that TMLE will be able to address a wider range of questions of interest. In fact, TMLE has begun to be used to estimate causal parameters in longitudinal data structures; simulation results as well as applied analysis will be presented in a forthcoming article Stitelman and van der Laan (2011). Furthermore, the algorithm used in that forthcoming article is not iterative and makes use of a backwards solving approach that requires less computational resources. Such an approach could also be used to estimate the parameters discussed here that do not account for time dependent covariates. Lastly, extensions that address causal parameters with time varying treatments are also currently in development.

R code, and an example of its implementation on a simulated data set, for the TMLE algorithm discussed here may be found online.² How to create confidence intervals for these estimates is addressed in Appendix A.

7 Data Example: The Tshepo Study

The *Tshepo* study was a three year randomized study using a 3x2x2 factorial design comparing efficacy, tolerability, and the development of drug resistance among different drug regimens. The Tshepo study is the first clinical trial evaluating the long-term efficacy and tolerability of efavirenz (EFV) vs. nevirapine (NVP) based cART among adults in Botswana. The study included 650 adults ranging in age from 20 to 64. Baseline characteristics, W , that were collected in the study and adjusted for in the following analyses included: age, gender, body mass index (BMI), HIV-1 RNA, CD4+ cell count, and WHO clinical stage. For a complete list of variables and descriptive statistics see Wester et al. (2010). Wester et al. (2010) performed an analysis of the *Tshepo* study using Cox proportional hazards to address the questions of whether there is a causal effect of cART treatment (individuals were randomized to one of two NNRTI-based cART therapies: EFV or NVP) on virologic failure and whether gender modifies this effect. They concluded that there was no significant difference by assigned NNRTI on time to virological failure and

²<http://www.stat.berkeley.edu/~laan/Software/index.html>

also showed a trend that women receiving NVP-based cART had higher virological failure rates when compared with EFV-treated women that did not reach statistical significance at $p=0.05$. Furthermore, they concluded, at a high level of statistical significance, that individuals treated with NVP had shorter times to treatment modification than did individuals treated with EFV. Using TMLE we will also address the causal effect of NNRTI treatment, and effect modification of gender as well as whether or not baseline CD4+ cell count modifies the effect of NNRTI treatment. However, for the purpose of this analysis we will focus on a slightly different outcome, TLOVR (Time to Loss of Virological Response), which is the FDA's preferred outcome for assessing the safety and efficacy of a drug: the minimal time to virologic failure, death, or treatment modification. The only censoring event for this outcome of interest is the end of study, implying the appropriateness of assuming independent censoring. A forthcoming article will focus on the medically relevant results of a similar analysis by examining alternative outcomes as well as other statistical questions of interest.

7.1 Causal Effect of NNRTI

Table 1 presents the estimates that address whether or not there is a causal effect of taking EFV versus NVP. The Cox proportional hazard estimates in the first column are the estimates of the coefficient in front of the treatment indicator unadjusted for any other variables. That is a standard analysis performed in assessing the effect of a randomized control trial on a time to event outcome. The targeted maximum likelihood estimates presented are the mean marginal additive difference in the probability of survival (equation (1)) and the mean marginal log relative hazard (equation (3)). The mean is taken over the first 34 months after randomization and the weights are based on the variance of the influence curve. The marginal difference of survival (equation (1)) at the final time point, 34 months, is also presented. The mean log relative hazard is the same target parameter as in cox proportional hazards regression but uses the TMLE method described here to estimate it. Thus, it does not rely on the Cox assumptions being correct. The Cox results are included in the table presented here, as well in the following two sections, since they represent current practice and it is interesting to contrast them to TMLE estimates of effect modification. When treatment is equal to NVP the treatment variable, A , equals 1 and when treatment equals EFV the treatment variable, A , equals 0. Thus, a positive mean log relative hazard and a negative risk difference corresponds with longer times until the specified outcome for the individuals treated with EFV compared to NVP.

For the TLOVR, which includes treatment modification, there is a highly significant causal effect of taking EFV versus NVP. The TMLE estimates for this

outcome are known to be unbiased since treatment is randomized and censoring is independent. Furthermore, the results are consistent with the results presented by Wester et al. (2010), who concluded that individuals treated with NVP tended to modify treatment sooner than individuals treated with EFV due to the toxicity of NVP.

Fig. 5 presents the TMLE estimates of the treatment specific survival curves for the TLOVR outcome. Examining the parameter estimates in conjunction with Fig. 5 reveals the difficulty in interpreting the Cox proportional hazards parameter (and our TMLE analogue) compared to the marginal additive difference in the probability of survival. The mean additive difference is $-.072$ which may be interpreted as implying that on average, the EFV specific survival probability is 7.2 percent higher than the NVP specific survival probability. A quick examination of Fig. 5 verifies this difference in the survival curves. The Cox proportional hazard estimate is 0.358 , and would approximate the mean of equation (3) over all time points according to the Cox proportional hazard model. This value has no easily interpretable meaning since it is the average of the log of the ratio of log survival probabilities. Alternatively, one can interpret it as an average of the log relative hazards—an interpretation that requires the user to fully understand the definition of a conditional hazard. It is clear that when it is positive the EFV specific survival curve is larger; however, there is no intuitive meaning of the size of the value. In addition, we can see that the TMLE estimates are more highly statistically significant than the Cox proportional hazard estimates (p-values of 0.022 vs 0.006 and 0.003). Efficiency theory specific to the TMLE and simulation results suggests that this gain in significance is due to a reduction in bias from using TMLE (van der Laan and Rubin (2006), Moore and van der Laan (2009)).

Table 1: Parameter Estimates To Assess Causal Effect of NNRTI Assignment.

	Parametric	Targeted Maximum Likelihood		
	Cox-PH	Mean RH	Mean RD	RD at $t = 34$
Estimate	0.358	0.451	-0.072	-0.060
SE	0.156	0.165	0.025	0.034
p-value	0.022	0.006	0.003	0.072

Figure 5: EFV and NVP specific survival curves.

7.2 Causal Effect Modification by Baseline-CD4+ Cell Count

Table 2 presents the estimates that address whether or not there is a causal effect modification due to baseline CD4+ cell count (High (201-350) versus Low (≤ 200)) on the effect of cART regimen (NVP/EFV). This is a situation where it is appropriate to assume that the effect modifier comes after the baseline covariates. The first column is the estimate from the Cox-proportional hazards model. All main terms for the baseline covariates, W , were included in the model as well as the cross term between A , EFV/NVP, and the effect modifier, V , CD4+ cell count level. The estimate presented is the β in front of the cross term, $A * V$ in the Cox proportional hazards model. This is the commonly implemented approach to assess effect modification using a Cox model. The targeted maximum likelihood estimates presented are the causal effect modification parameters that quantify the difference in the risk difference and difference in the marginal log relative hazard, equations (5) and (6), respectively. The mean is taken over the first 34 months after randomization and the weights are based on the variance of the influence curve. The difference in the risk difference at the final time point, 34 months, is also presented. The treatment variable, A , is set as for the analysis of the NNRTI treatment. The effect modifier variable, V , equals 1 for high baseline CD4+ cell count and 0 for low baseline CD4+ cell count. Thus, a positive marginal difference in the log relative hazard, equation (6), and a negative difference in the risk difference, equation (5), indicates that EFV has a larger beneficial causal effect in the high CD4+ cell count strata than in the low CD4+ cell count strata.

The causal effect modification by CD4 is significant for all the parameter estimates presented. The Cox model estimate is statistically significant at the $p=0.05$ level; however, as was seen above the p -value for the TMLE is more highly significant. These results are consistent with a decrease in bias or increase in efficiency from use of TMLE. Fig. 6 shows the survival curves for the FDA outcome setting individuals to CD4 level and treatment group. The figure depicts the significant effect modification seen in Table 2. Not only is the effect of treatment, EFV/NVP, for individuals with high CD4 different from the effect for those with low CD4, the effects are in opposite directions. Individuals with high CD4, assigned to NVP have a higher probability of experiencing viral failure, death or treatment modification over time compared to those assigned to EFV, whereas for those with low CD4, this effect is reversed, although the effect is smaller (Fig. 6). For the TLOVR outcome the average risk difference between the effect in the high CD4 group versus the low CD4 group is 12 percent (p -value = 0.023). One possible explanation for this is that NVP-treated adults have higher overall rates of treatment modifying toxicities (TMTs), that occur along a full continuum of CD4+ cell count ranges (particularly at high (> 250) CD4+ values), some of which may be life-threatening (i.e. hepatotoxicity and/or cutaneous hypersensitivity reactions). Such toxicities may be particularly problematic among persons with higher CD4 cell counts, resulting in higher rates of TMTs as such persons generally feel healthier and experience far fewer co-morbidities and are therefore more likely to report earlier grade toxicities and have lower thresholds to request a change in their treatment.

The survival curves presented on the left in Fig. 6 present the targeted maximum likelihood estimates where the target is the survival probability at each time point and the initial estimate of the event hazard is based on the super learner as described in section 4. Thus, great care has been taken to estimate the event hazard as well as possible. The survival curves on the right of Fig. 6 shows when the initial hazard is intentionally mis-specified. In fact, the initial estimates for each of the four groups of differing CD4 level and treatment level are not at all different and a main terms logistic model that only accounts for t and t^2 was used. The red line in Fig. 6 shows the initial estimate of all four survival curves when mis-specified. Super learner was then used to estimate the treatment distribution as well as possible, and since the only censoring event is the end of study, so that censoring is known to be independent of the baseline covariates, Kaplan- Meier estimates were used to estimate the censoring process. It is seen that by using targeted maximum likelihood with a completely mis-specified initial hazard, the effect modification is recovered, depicted by the separation of the four survival curves. This exemplifies the value of the double robustness of targeted maximum likelihood.

Figure 6: Time until virologic failure, Death or Treatment Modification By NNRTI and CD4 Strata. The survival curves on the left were estimated with the initial hazard correctly specified. The survival curves on the right were estimated with a mis-specified initial hazard.

Table 2: Parameter Estimates To Asses Causal Effect Modification Due To Baseline CD4 Level.

	Parametric	Targeted Maximum Likelihood		
	Cox-PH	Mean RH	Mean RD	RD at t = 34
Estimate	0.675	0.829	−0.115	−0.144
SE	0.317	0.356	0.051	0.071
p-value	0.033	0.020	0.023	0.043

7.3 Causal Effect Modification Due to Gender

Table 3 presents the estimates that address whether or not effect modification due to gender on the effect of cART treatment (NVP/EFV). This is a situation where it is known that the effect modifier comes before the baseline covariates. The first column presents estimates from the Cox model. Main terms were included for A,

EFV/NVP, V , gender; cross term between A and V are also included. The estimate presented is the β in front of the cross term, $A * V$ in the Cox model. The targeted maximum likelihood estimates presented are the stratified effect modification parameters that quantify the difference in the risk difference and difference in the marginal log hazard of survival, equations (7) and (8), respectively. The mean is taken over the first 34 months after randomization, and weighting is based on the estimated variance of the influence curve. The difference in the risk difference at the final time point, 34 months, is also presented. The treatment variable, A , is set as for analyzing the effect of the NNRTI treatment. The effect modifier variable, V , equals 1 for males and 0 for females. Thus, a negative marginal difference in the log relative hazard, equation (8), and a positive difference in the risk difference, equation (7), indicates that EFV is having a larger beneficial causal effect for females.

Figure 7: Time to virologic failure, Death or Treatment Modification Setting NNRTI and Gender.

The treatment effect modification by gender on the TLOVR outcome is statistically significant for all three TMLE estimates and the Cox proportional hazards estimate. In fact, it is highly significant for the mean difference in the marginal log hazard of survival, and the difference in the marginal log hazard of survival at $t_k = 34$. Fig. 7 shows the survival curves for the TLOVR outcome. Thus, we conclude that gender does modify the effect of EFV versus NVP on the time until TLOVR. Women tend to have more favorable outcomes using EFV while males tend to have similar outcomes with both drug therapies. For TLOVR the average causal risk difference between the effect in the males versus the females is 12 percent (p.value = 0.001). A possible reason for this result is that female NVP users tend to modify their treatment at a higher rate than the other groups. Based on Fig. 7 the major difference in the modification rate tends to occur soon after starting the NVP therapy, which is consistent with the published literature showing that the greatest risk for moderate/severe NVP-associated toxicity is within the first 18 weeks following NVP initiation.

Table 3: Parameter Estimates To Assess Stratified Effect Modification Due To Gender.

Parametric		Targeted Maximum Likelihood		
Cox-PH		Mean RH	Mean RD	RD at t = 34
FDA				
Estimate	-0.952	-0.816	0.116	0.193
SE	0.353	0.329	0.035	0.054
p-value	0.007	0.013	0.001	0.000

8 Discussion

The above results illustrate the advantages of using TMLE over Cox proportional hazards regression in assessing causal effects and causal effect modification. It has been shown that the parameters estimated using TMLE are much easier to interpret than the parameter that is estimated using Cox regression. Furthermore, TMLE is double robust and locally efficient resulting in advantages over Cox regression in both reducing bias and gaining efficiency. The double robustness was illustrated by Fig. 6, where the effect modification due to CD4+ cell count strata was regained after starting with an initial estimate of the four CD4/treatment specific survival

curves that was the same for all four combinations of CD4/treatment. Though a single data analysis should not be used to validate gains in efficiency and reductions in bias due to a particular method the results presented here are consistent with previous theoretical results and simulations that have exhibited these advantages. Since TMLE targets the parameter of interest and uses data adaptive loss based estimation rather than relying on an *a priori* specified model to estimate the conditional hazard, it produces consistent estimates of a specified parameter of interest under less restrictive model assumptions. By contrast, the Cox proportional hazard estimate is only consistent under its strict model assumptions. For these reasons the parameter estimates and significance levels produced by the TMLE are more reliable than those produced using the currently available methods for survival analysis.

The methods presented here assessed effect modification for binary effect modifiers. It was discussed above how these methods may be extended to effect modification parameters with multiple levels. Extending these methods to estimate parameters for possible continuous effect modifiers may be done through the use of a marginal structural model as is recommended in Neugebauer and van der Laan (2006). This approach for estimating parameters for continuous treatments based on a working marginal structural model using TMLE may be seen in Rosenblum and van der Laan (2010).

The next step in developing these methods is to extend them to include adjustments for time-dependent covariates. In a forthcoming article, Stitelman and van der Laan (2011), we will present a TMLE of the causal effect of treatment on the survival function utilizing both the baseline covariates and time-dependent covariates. This approach will both improve efficiency of the estimates as well as remove bias due to informative dropout. Furthermore, we will present a method that can be easily adapted to handle general longitudinal data structures and target parameters, such as causal parameters for continuous outcomes and causal effect of dynamic treatment regimens.

A Inference For Targeted Maximum Likelihood Estimation

In this appendix we will show how 95-percent confidence intervals may be constructed for the TMLE estimates discussed here. In Stitelman and van der Laan (2010) we show that quantile based bootstrap estimates of the confidence intervals are the preferred method for inference for their robustness, especially in situations where positivity violations are an issue. Moreover, such an approach takes into account the variance associated with data adaptively estimating Q_0 . Alternatively, confidence intervals may be constructed by relying on the fact that the TMLE solves the efficient influence curve estimating equation. This approach will be described here.

The following is the the efficient influence curve estimating equation:

$$0 = \sum_{i=1}^n D^*(Q_n^*, g_n)(O_i),$$

where $D^*(Q_0, g_0) = D^*(Q_0, g_0, \Psi(Q_0))$ is the efficient influence curve for a particular parameter of interest (e.g., $\Psi_{RD}(P_0)(t_k)$). Efficient influence curves, D^* , for several different parameters of interest, $\Psi(P_0)$, are displayed below.

One can also state that $\Psi(Q_n^*)$ solves the estimating equation in ψ_0 :

$$0 = \sum_i D^*(Q_n^*, g_n, \Psi(Q_n^*))(O_i),$$

as defined by this efficient influence curve equation. Under regularity conditions, it can be shown that $\Psi(Q_n^*)$ is asymptotically linear with an influence curve $D^*(Q, g_0, \psi_0) + D_1$ for the case that Q_n^* possibly converges to a mis-specified Q , and g_n converges to the true g_0 (van der Laan and Robins (2003), Section 2.3.7). If $Q = Q_0$, or $g_n = g_0$ (i.e., g_0 is known), then $D_1 = 0$. In addition, if g_n is a MLE-based estimator of g_0 (such that a particular smooth functional of g_n is an efficient estimator of the corresponding function of g_0), then ignoring the contribution D_1 results in an asymptotically conservative influence curve and variance estimator. Likewise, if Q_n converges to the true Q_0 and g_n converges to a possibly misspecified g , then asymptotic linearity may be established as well, but the influence curve will now depend on the particular estimator Q_n^* (not just on its limit as in the case that g_n is consistent for g_0). Thus, if either Q_n converges to Q_0 or g_n converges to g_0 , then,

under appropriate regularity conditions, we have that ψ_n^* is asymptotically linear with an influence curve $D(P_0)$:

$$n^{1/2}(\psi_n^* - \Psi(P_0)) = n^{-1/2} \sum_{i=1}^n D(P_0)(O_i) + o_p(1).$$

If g_n consistently estimates g_0 , one can conservatively replace $D(P_0)$ by $D^*(Q, g_0)$, thereby allowing for convenient influence curve based confidence intervals and variance estimators. Furthermore, in van der Laan and Gruber (2010) a weaker set of conditions under which the TMLE $\Psi(Q_n^*)$ is asymptotically linear is presented. In that article the asymptotic linearity result is generalized to hold for targeted MLE when Q_n converges to a possibly misspecified Q , and g_n converges to a true conditional censoring/treatment mechanism that adjusts for covariates that predict the residual bias between Q and Q_0 . Using the asymptotic linearity results above, by the central limit theorem,

$$n^{1/2}(\psi_n^* - \Psi(P_0)) \xrightarrow{D} N(0, E(D^2(P_0)(O))),$$

as sample size n converges to infinity.

So, if g_n is consistent for g_0 , the asymptotic variance of $n^{1/2}(\psi_{n,a}^* - \Psi_a(P_0))$ may be estimated by:

$$\sigma_n^2 = \frac{1}{n} \sum_{i=1}^n D_a^2(P_n^*)(O_i),$$

where P_n^* is the targeted maximum likelihood estimate of P_0 and $D_a^2(P_n^*)(O_i)$ is the efficient influence curves for the treatment specific survival curves above evaluated at P_n^* . Now 95-percent confidence intervals for the treatment specific survival curve at a particular time point may be constructed in the following way $\psi_{n,a}^* \pm 1.96 \frac{\sigma_n}{\sqrt{n}}$. Similarly, hypothesis tests can be constructed using the estimate of the asymptotic variance, σ_n^2 . Specifically, the following test statistic:

$$T_n = \frac{\psi_{n,a}^*}{\frac{\sigma_n}{\sqrt{n}}}$$

may be used to test the null hypothesis $H_0 : \Psi_a(P_0)(t_k) = 0$.

Variance estimates for the parameters of interests above which combine $\Psi_1(P_0)(t_k)$ and $\Psi_0(P_0)(t_k)$ may be estimated through the use of the delta-method.

The resulting efficient influence curves for the parameters of interest presented in section 4 are:

$$D_{RD}^*(P_n^*)(t_k) = D_1^*(P_n^*)(t_k) - D_0^*(P_n^*)(t_k), \quad (10)$$

$$D_{RR}^*(P_n^*)(t_k) = -\frac{1}{1 - \psi_{n,1}^*(t_k)} D_1^*(P_n^*)(t_k) + \frac{1}{1 - \psi_{n,0}^*(t_k)} D_0^*(P_n^*)(t_k), \quad (11)$$

$$D_{RH}^*(P_n^*)(t_k) = -\frac{1}{\psi_{n,1}^*(t_k) \log(\hat{\psi}_1^*)} D_1^*(P_n^*)(t_k) + \frac{1}{\psi_{n,0}^*(t_k) \log(\psi_{n,0}^*)} D_0^*(P_n^*)(t_k). \quad (12)$$

Confidence intervals may now be constructed for these parameters at a particular time point using the above estimates of the corresponding efficient influence curve. Furthermore, the estimated influence curve for estimates which are means of these parameters may be constructed by taking means of the estimated efficient influence curves (equations (10) - (12)) over the desired time points.

For statistical inference only relying on either g_n or Q_n^* being consistent we recommend the bootstrap, since the calculation of the influence curve requires involved analytical calculations.

References

- Altman, D. and P. Andersen (1999): "Calculating the number needed to treat for trials where the outcome is time to an event," *British Medical Journal*, 319, 1492–1495.
- Andersen, P., O. Borgan, R. Gill, and N. Keiding (1993): *Statistical Models Based on Counting Processes*, Springer-Verlag New York.
- Bang, H. and J. Robins (2005): "Doubly robust estimation in missing data and causal inference models," *Biometrics*, 61, 962–972.
- Box, G., J. Hunter, and W. Hunter (2005): *Statistics For Experimenters*, Hoboken, NJ: John Wiley and Sons, Hoboken.
- Collett, D. (1991): *Modeling Survival Data in Medical Research*, London: Chapman and Hall.
- Cox, D. (1972): "Regression models and life-tables," *Journal of the Royal Statistical Society. Series B (Methodological)*, 34, 187–220.
- Gruber, S. and M. van der Laan (2010): "A targeted maximum likelihood estimator of a causal effect on a bounded continuous outcome," *The International Journal of Biostatistics*, 6.
- Hernan, M. and J. Robins (2006): "Estimating causal effects from epidemiological data," *J Epidemiol Community Health*, 60, 578–586.

- Kang, J. and J. Schafer (2007): “Demystifying double robustness: A comparison of alternative strategies for estimating a population mean from incomplete data (with discussion),” *Statistical Science*, 22, 523–39.
- Laubender, R. P. and R. Bender (2010): “Estimating adjusted risk difference (rd) and number needed to treat (nnt) measures in the cox regression model,” *Statistics in Medicine*, 29, 851–859.
- Moore, K. and M. van der Laan (2009): “Application of time-to-event methods in the assessment of safety in clinical trials,” in K. E. Peace, ed., *in Design, Summarization, Analysis & Interpretation of Clinical Trials with Time-to-Event Endpoints*, Chapman and Hall.
- Neugebauer, R. and M. van der Laan (2005): “Why prefer double robust estimates,” *Journal of Statistical Planning and Inference*, 129, 405–26.
- Neugebauer, R. and M. van der Laan (2006): “Nonparametric causal effects based on marginal structural model,” *Journal of Statistical Planning and Inference*, 137, 419–434.
- Pearl, J. (2008): *Causality: Models, Reasoning, and Inference*, Cambridge University Press, Cambridge.
- Porter, K., S. Gruber, M. van der Laan, and J. Sekhon (forthcoming 2011): “The relative performance of targeted maximum likelihood estimators,” *The International Journal of Biostatistics*.
- Robins, J. (1986): “A new approach to causal inference in mortality studies with sustained exposure periods - application to control of the healthy worker survivor effect,” *Mathematical Modelling*, 7, 1393–1512.
- Robins, J., M. Sued, Q. Lei-Gomez, and A. Rotnitzky (2007): “Comment: Performance of double-robust estimators when “inverse probability” weights are highly variable,” *Statistical Science*, 22, 544–559.
- Rosenblum, M. and M. van der Laan (2010): “Nonparametric causal effects based on marginal structural model,” *The International Journal of Biostatistics*, 6.
- Rudser, K., M. LeBlanc, and S. Emerson (2008): “Estimation for arbitrary functionals of survival,” *UW Biostatistics Working Paper Series*, Working Paper 335.
- Stitelman, O. and M. van der Laan (2010): “Collaborative targeted maximum likelihood for time to event data,” *The International Journal of Biostatistics*, 6.
- Stitelman, O. and M. van der Laan (2011): “Targeted maximum likelihood estimation of time-to-event parameters with time-dependent covariates,” *Forthcoming*.
- Stone, C. (1996): *A Course in Probability and Statistics*, Duxbury Press, Belmont.
- Tchetgen Tchetgen, E. (2010): “On the interpretation, robustness, and power of varieties of case-only tests of gene-environment interaction,” *American Journal of Epidemiology*, 172, 1335–1338.

- van der Laan, M. and S. Dudoit (2003): “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, Berkeley.
- van der Laan, M., S. Dudoit, and S. Keles (2004): “Asymptotic optimality of likelihood-based cross-validation,” *Statistical Applications in Genetics and Molecular Biology*, 3.
- van der Laan, M. and S. Gruber (2010): “Collaborative double robust targeted maximum likelihood estimation,” *The International Journal of Biostatistics*, 6.
- van der Laan, M., E. Polley, and A. Hubbard (2007): “Super learner,” *Statistical Applications in Genetics and Molecular Biology*, 6.
- van der Laan, M. and J. Robins (2003): *Unified methods for censored longitudinal data and causality*, Springer, New York.
- van der Laan, M. and D. Rubin (2006): “Targeted maximum likelihood learning,” *The International Journal of Biostatistics*, 2.
- Vansteelandt, S. et al. (2008): “Multiply robust inference for statistical interactions,” *Journal of the American Statistical Association*, 103, 1693–1704.
- Wester, C. et al. (2010): “Non-nucleoside reverse transcriptase inhibitor outcomes among combination antiretroviral therapy-treated adults in botswana,” *AIDS*, 24, S27–S36.