Published in final edited form as:

J Biopharm Stat. 2009 November ; 19(6): 1099–1131. doi:10.1080/10543400903243017.

INCREASING POWER IN RANDOMIZED TRIALS WITH RIGHT CENSORED OUTCOMES THROUGH COVARIATE ADJUSTMENT

K. L. Moore and M. J. van der Laan

Division of Biostatistics, University of California, Berkeley, Berkeley, California, USA

Abstract

Targeted maximum likelihood methodology is applied to provide a test that makes use of the covariate data that are commonly collected in randomized trials, and does not require assumptions beyond those of the logrank test when censoring is uninformative. Under informative censoring, the logrank test is biased, whereas the test provided in this article is consistent under consistent estimation of the censoring mechanism or the conditional hazard for survival. Two approaches based on this methodology are provided: (1) a substitution-based approach that targets treatment and time-specific survival from which the logrank parameter is estimated, and (2) directly targeting the logrank parameter.

Keywords

Clinical trials; Covariate adjustment; Logrank test; Targeted maximum likelihood estimation

1. INTRODUCTION

Covariate adjustment in randomized controlled trials (RCTs) has been demonstrated to improve estimation efficiency over standard unadjusted methods in studies with continuous or binary outcomes at fixed endpoints (Koch et al., 1998; Moore and van der Laan, 2009b; Pocock et al., 2002; Tsiatis et al., 2008; Zhang et al., 2008). It is often the case in RCTs that the outcome is time-to-event in nature and subject to right censoring. The standard approach for testing for a treatment effect on survival is the logrank test, or asymptotically equivalently the test, $H_0 : \psi = 0$, where ψ is the coefficient for treatment in the Cox proportional hazards model that includes only a main term for treatment. From estimation theory (van der Laan and Robins, 2003), it is known under the proportional hazards assumption that this maximum likelihood estimator (MLE) is the efficient estimator of the effect of interest, given that the data include only treatment and survival times. In most RCTs, data are additionally collected on baseline (pretreatment) covariates. This unadjusted estimator ignores the covariates and is thus not equivalent to the full MLE. It follows that application of the unadjusted estimator can lead to a loss in estimation efficiency (precision) in practice.

Copyright © Taylor & Francis Group, LLC

Address correspondence to Kelly L. Moore, Division of Biostatistics, University of California, Berkeley, 101 Haviland Hall, Berkeley, CA 94720, USA; klmoores@stat.berkeley.edu.

Publisher's Disclaimer: Full terms and conditions of use: <http://www.informaworld.com/terms-and-conditions-of-access.pdf> This article may be used for research, teaching and private study purposes. Any substantial or systematic reproduction, re-distribution, re-selling, loan or sub-licensing, systematic supply or distribution in any form to anyone is expressly forbidden. The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

The key principle in developing covariate adjusted estimators is to not require any additional assumptions beyond those required for the unadjusted method. For example, a Cox proportional hazards model that includes covariates in addition to treatment requires heavy parametric modeling assumptions and thus is not a suitable method of covariate adjustment. Lu and Tsiatis (2008) demonstrated how the efficiency of the logrank test can be improved with covariate adjustment based on estimating equation methodology. Their method, which does not make assumptions beyond those of the logrank test, is more efficient and was shown to increase power over the logrank test. A nonparametric method for a covariate adjusted method that uses logrank or Wilcoxon scores was proposed in Tangen and Koch (1999) and explored via simulation studies in Jiang et al. (2008). This method attempts to adjust for the random imbalances that occur in covariate distributions between treatment groups. The authors use a linear regression model to estimate the difference in vectors between treatment groups, where these vectors include the average rank score and average value for each covariate within the given treatment group. This latter method limits the flexibility in adjusting for covariates by only allowing the comparison of their mean differences between treatment groups. Furthermore, it does not allow for adjustment for informative censoring. The importance of adjusting for covariates to gain power over the logrank test is also discussed in Akazawa et al. (1997). However, their stratified approach does not ensure a gain in power, and can actually lose power over the logrank test for certain stratification strategies.

With the principle of not making any assumptions beyond those required for the unadjusted test, in this article we develop covariate adjusted analogues to the logrank test in RCTs. We present our methodology for discrete survival outcomes where the logrank parameter represents an effect of treatment by comparing the cumulative hazard of treated subjects at time t_k , relative to the cumulative hazard of controls at time t_k , averaged over many time points t_k . However, we note that if the time scale is sufficiently fine, this methodology is compatible with continuous survival outcomes.

We present our methods for covariate adjustment under the framework of targeted maximum likelihood estimation, which is an approach to statistical learning originally introduced in van der Laan and Rubin (2006). Targeted maximum likelihood estimation is an estimation procedure that carries out a bias reduction specifically targeted for the parameter of interest. This is in contrast to traditional maximum likelihood estimation, which aims for a bias–variance trade-off for the whole density of the observed data, rather than a specific parameter of it. The targeted maximum likelihood methodology aims to find a density p^* that solves the efficient influence curve estimating equation for the parameter of interest that results in a bias reduction and also achieves a small increase in the log-likelihood as compared to the maximum likelihood estimate. The resulting substitution estimator $\psi(p^*)$ is a familiar type of likelihood-based estimator and due to the fact that it solves the efficient influence curve estimating equation it thereby inherits its properties, including asymptotic linearity and local efficiency (van der Laan and Robins, 2003).

In particular, for the parameter estimated in this article, the marginal log hazard of survival, the targeted maximum likelihood algorithm involves estimation of an initial conditional hazard. This hazard is then updated iteratively by adding to it a particular targeting covariate that is derived based on the specific parameter of interest. That is, two different parameters will have two different updating covariates. The algorithm is iterated until convergence, i.e., the efficient influence estimating equation for the parameter of interest is solved.

There are several advantages to this methodology over estimating equation methodology as discussed in Moore and van der Laan (2009a,b). One important advantage is that the methodology does not rely on the assumption that the efficient influence curve can be represented as an estimating equation in the parameter of interest. This is of particular

consequence for the logrank parameter since the efficient influence curve cannot be represented as an estimating equation in this parameter. As a result, the proof of the double robustness consistency properties does not follow in the usual obvious manner. Therefore, in this article, we provide two methods for covariate adjustment using the targeted maximum likelihood methodology. The first is a substitution-based approach that targets the time- and treatment-specific survival parameters for the treated and untreated arms. The corresponding estimates are used as plug-ins to evaluate the logrank parameter. Here we can prove the double robustness properties for the time and treatment specific estimators using the usual estimating equation approach. We then show how we can extend these properties to the logrank parameter. In the second approach, we target the logrank parameter directly. However, since we cannot use the estimating equation approach for proof of the double robustness properties due to the fact that the efficient influence curve cannot be represented as an estimating equation in this parameter, we rely on empirical validation of these properties. Therefore, we include both of these covariate adjusted targeted maximum likelihood estimators (TMLEs) of the logrank parameter.

The likelihood of the observed data (provided in Section 4.1) can be expressed in terms of the hazard of survival, conditional on treatment and covariates. It is important to note that the TMLEs rely on estimation of the conditional hazard. This initial hazard estimate can be maximum likelihood based, but can involve sieve-based estimation and selection of fine-tuning parameters/algorithms/models using likelihood-based cross-validation, since nonparametric maximum likelihood estimation is not possible. Machine learning algorithms can be applied to obtain an initial hazard estimate, after which the targeting step is applied as a means of bias reduction for the parameter of interest. Part of the targeting step involves averaging over the covariates that have terms in the hazard model to obtain a marginal or unconditional estimate. In summary, the methods presented in this article involve two steps. First, an initial hazard of survival, conditional on treatment and covariates, must be estimated. Second, the targeting step is applied as a bias reduction step for the parameter of interest.

Specifically, the article is outlined as follows. We begin with a brief overview of the data, model and parameter of interest in Section 2. We then provide an unadjusted method for testing for a treatment effect in Section 3, to which we compare the targeted maximum likelihood estimation approaches. In the continuous case, this unadjusted test is the logrank test, and in the discrete case, we provide a method based on the proportional odds model. We then review the methodology for estimating the treatment specific estimates of survival at a fixed endpoint t_k as presented in Moore and van der Laan (2009a). This t_k approach is then extended to provide our first of two analogues to the logrank test, the substitution-based targeted maximum likelihood method (Section 4.1). In the second analogue, we provide the direct targeted maximum likelihood approach, which does not require estimation of the t_k -specific survival estimates, but rather directly targets the average (over time) effect of treatment on survival. Since the TMLE requires estimation of an initial conditional hazard, methods for fitting it as well as the censoring mechanism are provided in Section 5. In Section 6 we present simulation studies to demonstrate the efficiency gains of the proposed methods over the logrank test in an RCT under no censoring and uninformative censoring. Furthermore, under informative censoring we demonstrate the bias that arises with the standard approach in contrast to the consistency of our proposed estimator. A second simulation study demonstrates the importance of data-adaptive model selection algorithms in the estimation of the initial hazard used by the targeted maximum likelihood algorithm in order to obtain maximal power. A third simulation study demonstrates the methodology as it would be applied to single dataset as in a typical data analysis. The data were simulated based on a real RCT study. Finally, we conclude with a discussion.

2. DATA, MODEL, AND PARAMETER OF INTEREST

We assume that in the study protocol, each patient is monitored at K equally spaced clinical visits. At each visit, an outcome is evaluated as having occurred or not occurred. Let T represent the first visit at which the event was reported, and thus T can take values $\{1, \dots, J\}$. The censoring time C is the first visit when the subject is no longer enrolled in the study. Let $A \in \{0, 1\}$ represent the treatment assignment at baseline and W represent a vector of baseline covariates. The observed data are given by $O = (\tilde{T}, \Delta, A, W) \sim p_0$ where $\tilde{T} = \min(T, C)$, $\Delta = I(T \leq C)$ is the indicator that that subject was not censored and p_0 denotes density of O . The conditional hazard is given by $\lambda_0(\cdot | A, W)$ and the corresponding conditional survival is given by $S_0(\cdot | A, W)$. The censoring mechanism is given by $\bar{G}(t_- | A, W) = P(C \geq t | A, W)$.

Consider the data structure $(W, A, T = T_A)$. Within the counterfactual framework for causal inference, since this data structure only contains a single counterfactual outcome $T = T_A$ corresponding with the treatment the patient actually received, this data structure is a missing data structure on the full data $X = (T_0, T_1, W)$ with missingness variable A , where T_1 represents a patient's time to the occurrence of an event had she, possibly contrary to fact, been assigned to the treatment group, and T_0 likewise represents the time to the occurrence of the event had the patient been assigned to the control group. Our target parameter, which we define later; is a parameter of this full data structure. The randomization assumption states that A is conditionally independent of the full data X , given W . In an RCT, this assumption clearly holds. For the sake of presentation, we assume that the treatment A is completely randomized; however, we present our methods under the counterfactual framework so that it is clear how our estimators generalize to observational studies or randomized trials in which the treatment assignment mechanism is unknown. The actual observed data O are a right-censored version of $(W, A, T = T_A)$ and can thus be viewed as a censored data structure of the full data X with censoring variables (A, C) . We assume coarsening at random (CAR), which means that the coarsening process only depends on the observed data. Under these counterfactual assumptions, our observed data representations of the parameters discussed below are also full data parameters.

In Moore and van der Laan (2009a), the authors presented the targeted maximum likelihood estimation method for the estimation of the t_k and treatment-specific parameters,

$$p_0 \rightarrow \Gamma_1(p_0)(t_k) = Pr(T_1 > t_k) = E_0(S_0(t_k | A=1, W)) = S_{0,1}(t_k) \quad (1)$$

and

$$p_0 \rightarrow \Gamma_0(p_0)(t_k) = Pr(T_0 > t_k) = E_0(S_0(t_k | A=0, W)) = S_{0,0}(t_k) \quad (2)$$

where the subscripts for Γ and T denote the treatment group, either 0 or 1, elsewhere the subscript 0 denotes the truth, and $S_{0,1}$ is the true survival function for treatment group 1 and $S_{0,0}$ is the true survival function for treatment group 0. Thereby, any linear combination of these parameters can be estimated to evaluate the effect of treatment A on survival T , e.g., the marginal log hazard of survival,

$$p_0 \rightarrow \Gamma(p_0)(t_k) = \log\left(\frac{\log(Pr(T_1 > t_k))}{\log(Pr(T_0 > t_k))}\right) = \log\left(\frac{\log(S_{0,1}(t_k))}{\log(S_{0,0}(t_k))}\right) \quad (3)$$

In this article, we are interested not in a test for the effect of treatment at a fixed end point t_k , but rather in the average effect over time. Note that in the continuous survival case, if one averaged $\Gamma(p_0)(t_k)$ over all t , this parameter would correspond with the Cox proportional hazards parameter (i.e., coefficient for treatment in Cox proportional hazards model) and thus the parameter tested by the ubiquitous logrank test, given by

$$\lambda(t|A) = \lambda(t) \exp(\psi_C A) \quad (4)$$

More formally, let $\mathcal{M}$ be the class of all densities of O with respect to an appropriate dominating measure and be nonparametric up to possible smoothness conditions. Let our parameter of interest be represented by $\Psi(p_0)$, where

$$p_0 \rightarrow \Psi(p_0) = \sum_{t_k} w(t_k) f(S_1(t_k), S_0(t_k)) \quad (5)$$

for some weight function $w(t_k)$ that we discuss in Section 4.3 and some function f . For example, g may be defined as in Eq. (3), and if T is continuous (or close to continuous), then this parameter is commonly tested with testing that ψ_C in Eq. (4) is equal to 0, or equivalently the logrank test. If T is discrete, one could instead define f as the marginal log odds of failure at t_k . In the latter case, this parameter is commonly tested using logistic regression (see Section 3). However, even in the discrete case, one can still choose f as in Eq. (3) as we do in this article, but can no longer use the Cox proportional hazards model (or equivalently the logrank test). Note that although we focus on the marginal log hazard of survival as our parameter of interest, the derivation of the TMLE for the average of the log odds parameter is very similar.

Thus, the targeted maximum likelihood test for the effect of treatment on survival is a test for $H_0: \psi_0 = 0$ against $H_A: \psi_0 \neq 0$, where $\psi_0 = \Psi(p_0)$.

3. UNADJUSTED ESTIMATION OF $\Psi(p_0)$

In the case that T is continuous, the unadjusted estimation of $\Psi(p_0)$ could be carried out by fitting a Cox proportional hazards model with a main term for treatment only, or equivalently carry out the logrank test. The proportional odds model, introduced in Cox (1972) for discrete survival times, is a model for the odds of dying at t , given survival up to time t , given by

$$\frac{\lambda(t|A)}{1 - \lambda(t|A)} = \frac{\lambda(t)}{1 - \lambda(t)} \exp(\beta_A A)$$

where

$$\log\left(\frac{\lambda(t|A)}{1 - \lambda(t|A)}\right) = \beta_1 I(t=1) + \beta_2 I(t=2) + \cdots + \beta_J I(t=J) + \beta_A A \quad (6)$$

Thus, $\beta_1, \dots, \beta_J$ capture the logit of the baseline hazard function, and β_A is the effect of treatment on the logit of the hazard. Such a parameterization leaves the baseline hazard unspecified and since A is a binary variable, β_A is the nonparametric formulation of the effect of treatment on the logit of the hazard.

The likelihood function for the discrete hazard process, where $O_r = (\tilde{T}, \Delta, A)$, can be expressed as

$$L(O_r) = \prod_{i=1}^n \Pr(T_i = \tilde{t}_i)^{\Delta_i} \Pr(T_i > \tilde{t}_i)^{(1-\Delta_i)} \quad (7)$$

$$= \prod_{i=1}^n \left(\lambda_i(\tilde{t}_i) \prod_{t=1}^{\tilde{t}_i-1} (1 - \lambda_i(t)) \right)^{\Delta_i} \left(\prod_{t=1}^{\tilde{t}_i} (1 - \lambda_i(t)) \right)^{1-\Delta_i} \quad (8)$$

where $\tilde{t}_i = \min(t_i, c_i)$ is the last time point at which individual i was observed (i.e., either censored or the event occurred). Let $\tilde{y}_i = (y_{i1}, \dots, y_{i\tilde{t}_i})$ denote the event history for individual i where $(y_{i1}, \dots, y_{i\tilde{t}_i-1}) = (0, \dots, 0)$, and $y_{i\tilde{t}_i} = 1$ if $\Delta_i = 1$ and $y_{i\tilde{t}_i} = 0$ if $\Delta_i = 0$. It can be shown that,

$$L(O_r) = \prod_{i=1}^n \prod_{t=1}^{\tilde{t}_i} \lambda(t | A_i)^{y_{it}} (1 - \lambda(t | A_i))^{(1-y_{it})}$$

Note that this likelihood is equivalent to that of a sequence of independent Bernoulli trials and thus we can use standard logistic regression software to obtain the maximum likelihood estimates for the coefficients β in Eq. (6).

In practice, the logistic regression model is fitted with the dataset that includes repeated measures for each subject until the time when the subject either dies or is censored—e.g., if a given subject dies or is censored at time point 5, this subject would contribute 5 rows of data to the new dataset. The outcome variable is zero up until the event occurs, where it is set to 1. If the subject is censored, then the outcome remains 0, even at the last time point.

An estimate of the effect of treatment on the logit of the hazard can be obtained by extracting the coefficient for A ; however, our parameter of interest is the average of the log of the ratio of log of survival under the two treatment regimens, as given by $\Psi(p_0)$ as defined in Eq. (5). Thus, we use the logistic regression fit for the hazard, denoted by $\hat{\lambda}(t | A)$, to obtain estimates for $\hat{\lambda}_1(t) = \hat{\lambda}(t | A=1)$ and $\hat{\lambda}_0(t) = \hat{\lambda}(t | A=0)$. Based on these estimates, we use the relation

$$S(t_k) = \prod_{j \leq t_k} (1 - \lambda(j))$$

to obtain estimates $\hat{S}_1(t_k)$ and $\hat{S}_0(t_k)$. The unadjusted estimate of $\Psi(p_0)$ is then computed as the crude average over time of the log of the ratio of the logs of these t_k -specific estimates. We note that, alternatively, one could estimate $S_1(t_k)$ and $S_0(t_k)$ for $t_k \in 1, \dots, J$ using Kaplan–Meier and use these estimates as plug-ins into Eq. (5). The use of Kaplan–Meier would be more nonparametric since the proportional odds model assumes proportionality of the odds of failure. However, both provide valid tests of $H_0 : \psi_0 = 0$ in the nonparametric model since under H_0 , both methods provide consistent estimators of ψ_0 . In this article, we use the proportional odds approach only.

4. TARGETED MAXIMUM LIKELIHOOD ESTIMATION OF $\Psi(p_0)$

The log likelihood for the observed data O , which includes covariates W , for a single observation is given by

$$P(W) g(A | W) \left[\lambda(\tilde{t} | A, W) \prod_{t=1}^{\tilde{t}-1} (1 - \lambda(t | A, W)) \right]^{\delta} \left[\prod_{t=1}^{\tilde{t}} (1 - \lambda(t | A, W)) \right]^{1-\delta} \times \left[\bar{G}(\tilde{t}_- | A, W)^\delta P(C = \tilde{t} | A, W)^{1-\delta} \right] \quad (9)$$

where $\delta = 1$ if $T = \tilde{t}$ and $C \geq \tilde{t}$.

Consider an initial fit p^0 of the density of the observed data O , identified by a hazard fit $\bar{\lambda}^0(t | A, W)$, the distribution of A identified by $\bar{g}^0(1 | W)$ and $\bar{g}^0(0 | W) = 1 - \bar{g}^0(1 | W)$, the censoring mechanism $\bar{G}^0(t | A, W)$, and the marginal distribution of W being the empirical probability distribution of $W_1, \dots, W_n$. In an RCT, treatment is randomized and

$$\bar{g}^0(1 | W) = \frac{1}{n} \sum_{i=1}^n A_i.$$

Let the initial hazard fit be denoted by $\bar{\lambda}^0(t | A, W)$. This initial hazard can be represented as

$$\text{logit}(\bar{\lambda}^0(t | A, W)) = m(t, A, W)$$

where m is any function of t , A , and W . We show that representing the initial hazard in this manner allows us to obtain its update (fluctuation) using standard software (e.g., glm in R). For example, we could consider the initial hazard

$$\text{logit}(\bar{\lambda}^0(t | A, W)) = \bar{\alpha}(t) + k(A, W | \bar{\beta})$$

where k is some function of A and W . The targeted maximum likelihood algorithm updates this initial fit by adding to it the term $\epsilon h(t, A, W)$, i.e.,

$$\text{logit}(\bar{\lambda}^0(\epsilon)(t | A, W)) = m(t, A, W) + \epsilon h(t, A, W) \quad (10)$$

Now, ϵ is estimated by fitting Eq. (10) with standard logistic regression software and fixing the coefficient for $m(t, A, W)$ at one and setting the intercept to zero. The initial hazard fit is then updated by adding to it $\bar{\epsilon} h(t, A, W)$. The covariate $h(t, A, W)$ is then reevaluated based on this updated hazard (and thus survival) fit. The newly updated hazard now plays the role of the initial hazard and ϵ is then again estimated based on the newly updated covariate $h(t, A, W)$. The procedure is iterated until convergence, i.e., $\bar{\epsilon}$ is essentially zero. The targeted maximum likelihood estimate is based on the hazard obtained in the final step of the algorithm. The covariate $h(t, A, W)$, which is defined in the following sections, is a function of the conditional survival function ($S(t | A, W)$) and the censoring mechanism ($\bar{G}(t_- | A, W)$). The covariate is reevaluated at each step of the algorithm based on the updated hazard estimate for $\lambda(t | A, W)$ (and thus $S(t | A, W)$). The censoring mechanism is not updated in the algorithm. For the rationale for updating only the hazard, see Moore and van der Laan (2009a).

The covariate $h(t, A, W)$ is selected such that the score for this hazard model at $\epsilon = 0$ is equal to the projection of the efficient influence curve on scores generated by the parameter $\lambda(t | A,$

W) in the nonparametric model for the observed data, assuming only CAR. Thus, the TMLE, that is, the estimator based on this iteratively updated hazard fit, solves the efficient influence curve estimating equation.

We provide two different approaches to the estimation of $\Psi(p_0)$ based on the targeted maximum likelihood methodology. This specific covariate $h(t, A, W)$ that updates the hazard is dependent on the approach. In the first approach, the targeted maximum likelihood estimates for each t_k specific parameter, $S_1(t_k)$ and $S_0(t_k)$, at each $t_k \in 1, \dots, J$ are obtained. The substitution estimate for $\Psi(p_0)$ is then based on these targeted maximum likelihood estimates. In the second method, the estimate for $\Psi(p_0)$ is obtained by directly targeting that parameter.

4.1. Method 1: Substitution TMLE

In this first method, we describe the procedure for targeted maximum likelihood estimation of the t_k and treatment-specific parameters $S_1(t_k)$ and $S_0(t_k)$. Thereby, we can estimate any parameter that is a combination of them, such as the average parameter $\Psi(p_0)$ defined in Eq. (5). We note that with this procedure one could choose a number of other parameters such as the average difference in survival or the average ratio of survival. However, here we focus on the substitution targeted maximum likelihood estimator (S-TMLE) of $\Psi(p_0)$ only.

The covariates for targeting the t_k and treatment specific parameters $S_1(t_k)$ and $S_0(t_k)$ as defined in Eqs. (1) and (2), respectively, were provided in Moore and van der Laan (2009a). In short, the algorithm selects the covariates, $h_{1t_k}(t, A, W)$ and $h_{0t_k}(t, A, W)$, corresponding with targeting the parameters $S_1(t_k)$ and $S_0(t_k)$, respectively. For the parameter $S_1(t_k)$, $h_{1t_k}(t, A, W)$ is defined such that the score for the hazard model at $\varepsilon_1 = 0$ is equal to the projection of the efficient influence curve of $S_1(t_k)$ on scores generated by the parameter $\lambda(t | A, W)$ in the nonparametric model for the observed data, assuming only CAR. Similarly, for $S_0(t_k)$, $h_{0t_k}(t, A, W)$ is defined such that the score for the hazard model at $\varepsilon_2 = 0$ is equal to the projection of the efficient influence curve of $S_0(t_k)$ on scores generated by the parameter $\lambda(t | A, W)$ in the nonparametric model for the observed data, assuming only CAR. These covariates corresponding with parameters $S_1(t_k)$ and $S_0(t_k)$ are respectively given by

$$h_{1t_k}(t, A, W) = - \frac{I(A=1)}{g(1) \bar{G}(t_- | A, W)} \frac{S(t_k | A, W)}{S(t | A, W)} I(t \leq t_k) \quad (11)$$

and

$$h_{0t_k}(t, A, W) = - \frac{I(A=0)}{g(0) \bar{G}(t_- | A, W)} \frac{S(t_k | A, W)}{S(t | A, W)} I(t \leq t_k) \quad (12)$$

The parameters $S_1(t_k)$ and $S_0(t_k)$ can be targeted simultaneously by addition of both covariates $h_1(t, A, W)$ and $h_0(t, A, W)$ as in Eq. (10), and finding the two-dimensional updating coefficients, $\widehat{\varepsilon} = \{\widehat{\varepsilon}_1, \widehat{\varepsilon}_2\}$, i.e.,

$$\text{logit}(\widehat{\lambda}^0(\varepsilon)(t | A, W)) = m(t, A, W) + \varepsilon_1 h_{1t_k}(t, A, W) + \varepsilon_2 h_{0t_k}(t, A, W) \quad (13)$$

Finding $\widehat{\varepsilon} = \{\widehat{\varepsilon}_1, \widehat{\varepsilon}_2\}$ in the updated hazard provided in Eq. (13) that maximizes the likelihood of the observed data can be done in practice by fitting a logistic regression in the covariates $m(t, A, W)$ and $h_{1t_k}(t, A, W)$ and $h_{0t_k}(t, A, W)$. The coefficient for $m(t, A, W)$ is fixed at one and the intercept is set to zero and thus the whole regression is not refit, but rather only ε is estimated.

These steps for evaluating $\widehat{\epsilon}$, and thus obtaining the updated hazard fit $\widehat{\lambda}^1(t|A, W)$ correspond with a single iteration of the targeted maximum likelihood algorithm. In the second iteration, the updated $\widehat{\lambda}^1(t|A, W)$ now plays the role of the initial fit, and the covariates $h_{1t_k}(t, A, W)$ and $h_{0t_k}(t, A, W)$ are then reevaluated with the updated $\widehat{S}^1(t|A, W)$ based on $\widehat{\lambda}^1(t|A, W)$ and $\widehat{\epsilon}$ is estimated again. Based on this update, $\widehat{\lambda}^2(t|A, W)$ is obtained. In the third iteration, $\widehat{\lambda}^3(t|A, W)$ is fitted and the procedure is iterated until $\widehat{\epsilon}$ is essentially zero. The final hazard fit at the last iteration of the algorithm is denoted by $\widehat{\lambda}^*(t|A, W)$ with the corresponding survival fit given by $\widehat{S}^*(t|A, W)$.

The t_k -specific parameter $\Gamma(p_0)(t_k)$, defined in Eq. (3), is estimated by

$$\widehat{\gamma}(t_k)^* = \log \left(\frac{\log \left(\frac{1}{n} \sum_{i=1}^n \widehat{S}^*(t_k | 1, W_i) \right)}{\log \left(\frac{1}{n} \sum_{i=1}^n \widehat{S}^*(t_k | 0, W_i) \right)} \right) \quad (14)$$

Finally, the parameter of interest $\Psi(p_0)$ can be estimated by plugging in each of the t_k -specific estimates $\widehat{\gamma}(t_k)^*$ for $t_k \in 1, \dots, J$. That is,

$$\widehat{\psi}^* = \sum_{t_k} w(t_k) \widehat{\gamma}(t_k)^*$$

4.1.1. Efficient Influence Curves—The efficient influence curves for the parameters $S_1(t_k)$ and $S_0(t_k)$, denoted by IC_{1t_k} and IC_{0t_k} , were provided in Moore and van der Laan (2009a), and are respectively given by

$$IC_{1t_k}(g_0, G_0, S_0) = \sum_{t \leq t_k} h_{1t_k}(g_0, G_0, S_0)(t|A, W) \left[I(\widetilde{T}=t, \Delta=1) - I(\widetilde{T} \geq t) \times \lambda_0(t|A=1, W) \right] + S_0(t_k|A=1, W) - \Gamma_1(p_0)(t_k) \quad (15)$$

and

$$IC_{0t_k}(g_0, G_0, S_0) = \sum_{t \leq t_k} h_{0t_k}(g_0, G_0, S_0)(t|A, W) \left[I(\widetilde{T}=t, \Delta=1) - I(\widetilde{T} \geq t) \times \lambda_0(t|A=0, W) \right] + S_0(t_k|A=0, W) - \Gamma_0(p_0)(t_k) \quad (16)$$

The efficient influence curve, denoted by IC_{t_k} , for the parameter $\Gamma_0(p_0)(t_k)$, can be obtained by application of the δ -method to the influence curves IC_{1t_k} and IC_{0t_k} . We have

$$IC_{t_k} = a(t_k) IC_{1t_k} + b(t_k) IC_{0t_k} \quad (17)$$

which is a linear combination of IC_{1t_k} and IC_{0t_k} , with coefficients only a function of t_k . With some algebra, one can easily show that the coefficients are given by

$$a(t_k) = \frac{1}{S_1(t_k) \log(S_1(t_k))} \text{ and } b(t_k) = \frac{-1}{S_0(t_k) \log(S_0(t_k))}.$$

Our parameter of interest $\Psi(p_0)$ is the average (possibly weighted) of the t_k -specific log ratios of the logs of survival, i.e., average over t_k of $\Gamma(p_0)(t_k)$. Therefore, its efficient influence curve is given by

$$\bar{IC} = \sum_{t_k} w(t_k) IC_{t_k} \quad (18)$$

4.1.2. Double Robustness Consistency Properties of the S-TMLE—In Moore and van der Laan (2009a), the statistical properties of the treatment and t_k -specific estimators were provided. Consider the parameter $S_1(t_k)$. The targeted maximum likelihood estimate $\hat{p}^* \in \mathcal{M}$ of p_0 solves the efficient influence curve estimating equation, given by

$\sum_{i=1}^n IC_{1t_k}(g_0, \hat{G}, \hat{S}^*)(O_i) = 0$, which is the optimal estimating equation for the parameter of interest. It can be shown that $E_0 IC_{1t_k}(S, g, G) = 0$ if either (1) $S = S_0(\cdot | A, W)$ (and thus $\lambda = \lambda_0(\cdot | A, W)$) or (2) $g = g_0(A | W)$ and $G = G_0(\cdot | A, W)$ (see Appendix A). In an RCT, the treatment mechanism is known and $g_0(A | W) = g_0(A)$. Therefore, the consistency of the estimator $\hat{\gamma}_1^*(t_k)$ of $S_1(t_k)$ in an RCT relies only on consistent estimation of $G_0(\cdot | A, W)$ or $S_0(\cdot | A, W)$. When there is no censoring or censoring is missing completely at random (MCAR), $\hat{\gamma}_1^*(t_k)$ is consistent even when the estimator $\hat{S}(\cdot | A, W)$ is inconsistent (e.g., if it relies on a misspecified model). Hence, in an RCT, one is not concerned with estimation bias due to misspecification of the hazard model. Under informative or missing at random (MAR) censoring, if $G_0(\cdot | A, W)$ is consistently estimated then $\hat{\gamma}_1^*(t_k)$ is consistent even if $S_0(\cdot | A, W)$ is not consistently estimated. If both are correctly specified then $\hat{\gamma}_1^*(t_k)$ is efficient. These same properties hold for the estimator $\hat{\gamma}_0^*(t_k)$ for $S_0(t_k)$.

Since the substitution based estimator $\hat{\gamma}(t_k)^*$ for the t_k -specific parameter given by Eq. (3) is simply a function of these two treatment-specific estimators, it inherits these same double robustness properties. Similarly, since our parameter of interest, $\Psi(p_0)$ is the average (possibly weighted) of the t_k -specific log ratios of the logs of survival, clearly the properties of the t_k -specific log ratio estimator directly extend to the estimator of the average (over time) parameter (i.e., the logrank analogue parameter). That is, if there is no censoring or censoring is MCAR, this method provides a covariate adjusted estimator that is consistent even when the hazard is misspecified. Thus, if one captures only part of the relevant covariate information, one can still gain in efficiency over the unadjusted method, without the risk of introducing bias.

We note that although the double robustness properties of the substitution estimators for the parameters $\Gamma(p_0)(t_k)$ and $\Psi(p_0)$ are inherited from the properties of $\hat{\gamma}_1^*(t_k)$ and $\hat{\gamma}_0^*(t_k)$, the $\hat{\gamma}(t_k)^*$ and $\hat{\psi}^*$ also solve their corresponding efficient influence curve estimating equations based on Eqs. (17) and (18) respectively.

4.1.3. Inference for S-TMLE—We first consider the parameter $S_1(t_k)$. Since the TMLE is a solution to the efficient influence curve estimating equation, then from estimating equation theory (see van der Laan and Robins, 2003) if g_0 and G_0 are known, then the estimator is asymptotically linear with influence curve $IC_{1t_k}(g_0, G_0, S)$. However, even though g_0 is typically known in an RCT, G_0 is not. In this case, the influence curve is given by

$$IC_{1t_k}(g_0, G, S) - \Pi(IC_{1t_k} | T_G)$$

That is, one must subtract from IC_{1t_k} its projection on the tangent space of the model for the censoring mechanism. Therefore, one can construct an asymptotically conservative Wald-type 0.95-confidence interval for $\widehat{\gamma}_1^*(t_k)$ based on the estimate of the efficient influence curve for $S_1(t_k)$ ignoring this projection and using Eq. (15), i.e., using $IC_{1t_k}(g_0, \widehat{G}, \widehat{S})$, where this

confidence interval is given by $\widehat{\gamma}_1^*(t_k) \pm 1.96 \frac{\widehat{\sigma}_1}{\sqrt{n}}$, and

$$\widehat{\sigma}^2 = \frac{1}{n} \sum_{i=1}^n IC^2(g_0, \widehat{G}, \widehat{S})$$

The null hypothesis $H_0 : \gamma_1 = 0$ can be tested with the test statistic

$$\widehat{T} = \frac{\widehat{\gamma}_1^*}{\frac{\widehat{\sigma}}{\sqrt{n}}},$$

whose asymptotic distribution is $N(0, 1)$ under the null hypothesis. Inference is derived in the same manner for the TMLE $\widehat{\gamma}_0^*(t_k)$ of $S_0(t_k)$.

As with the estimators of treatment specific survival, if g_0 and G_0 are known, then the S-TMLE of $\Gamma(p_0)(t_k)$ is asymptotically linear with influence curve $IC_{t_k}(g_0, G_0, S)$. Since G_0 is not known, the influence curve is given by

$$\begin{aligned} IC_{t_k} &= a(t_k)(IC_{1t_k} - \Pi(IC_{1t_k} | T_G)) + b(t_k)(IC_{0t_k} - \Pi(IC_{0t_k} | T_G)) \\ &= a(t_k)(IC_{1t_k}) + b(t_k)IC_{0t_k} - \Pi((a(t_k)IC_{1t_k} - b(t_k)IC_{0t_k}) | T_G) \end{aligned}$$

Therefore, one can construct an asymptotically conservative Wald-type 0.95-confidence interval for $\widehat{\gamma}$ based on the estimate of the efficient influence curve for $\Gamma(p_0)(t_k)$ using Eq. (17), i.e., using $IC_{t_k}(g_0, \widehat{G}, \widehat{S})$, as given earlier. Similarly, the test statistic can be constructed as already shown.

Finally, for our parameter of interest, the average (possibly weighted) of the t_k -specific log of the ratio of logs of survival, the influence curve when G_0 is estimated is given by

$$\begin{aligned} \bar{IC} &= \sum_{t_k} w(t_k) \left[a(t_k)(IC_{1t_k} - \Pi(IC_{1t_k} | T_G)) + b(t_k)(IC_{0t_k} - \Pi(IC_{0t_k} | T_G)) \right] \\ &= \sum_{t_k} w(t_k) \left[a(t_k)(IC_{1t_k}) + b(t_k)IC_{0t_k} \right] - \Pi(w(t_k)(a(t_k)IC_{1t_k} - b(t_k)IC_{0t_k}) | T_G) \end{aligned}$$

Therefore, one can construct an asymptotically conservative Wald-type 0.95-confidence interval for $\widehat{\psi}^*$ based on the estimate of the efficient influence curve for $\Psi(p_0)$ using Eq. (18), i.e., using $\bar{IC}(g_0, \widehat{G}, \widehat{S})$, as done earlier. Similarly, the test statistic can be constructed as already shown.

4.2. Method 2: Direct TMLE

In the previous method, we provided a substitution-based procedure based on the t_k -specific estimators of survival. However, we did not directly target the parameter of interest, which is the average of the log of the ratio of logs of survival in the treatment and control groups, given by $\Psi(p_0)$ as defined in Eq. (5). In this section we present the targeted maximum likelihood algorithm for targeting this single parameter. Contrary to the S-TMLE, which requires two covariates to update the hazard at each time $t_k \in \{1, \dots, J\}$, the algorithm targeting the single parameter only requires a single covariate. This suggests that application of this direct targeted maximum likelihood estimator (D-TMLE) results in finite sample improvements in efficiency over the S-TMLE.

The efficient influence curve of our parameter of interest $\Psi(p_0)$ is given by Eq. (18) in Section 4.1.1. We provided the time-dependent covariates $h_{1t_k}(t, A, W)$ and $h_{0t_k}(t, A, W)$ for targeting the t_k -specific survival $S_1(t_k)$ and $S_0(t_k)$, required to generate the corresponding components of IC_{1t_k} and IC_{0t_k} of the form

$$h_j(t, A, W)(dN(t) - E(dN(t) | \text{past}(t))), j=0, 1$$

where $N(t)$ is a counting process. Thus, to generate the component $\bar{IC}_t$ of $\bar{IC}$ for the t -factor of the likelihood, we now select the single time-dependent covariate,

$$\bar{h}(t, A, W) = \sum_{t_k} w(t_k) a(t_k) h_{1t_k} + \sum_{t_k} w(t_k) b(t_k) h_{0t_k}$$

Now, in the targeted maximum likelihood algorithm, the first updating step of the initial hazard, denoted by $\hat{\lambda}^0(t | A, W)$, is given by

$$\text{logit}(\hat{\lambda}^0(\epsilon)(t | A, W)) = m(t, A, W) + \epsilon \bar{h}(t, A, W) \quad (19)$$

Again, as in the substitution-based method, we now find $\hat{\epsilon}$ by fitting Eq. (19) using standard logistic regression software by setting the coefficient for $m(t, A, W)$ to one and setting the intercept to zero. The corresponding updated hazard $\hat{\lambda}^1(t | A, W)$ is obtained. This represents the first step of the algorithm. The hazard fit $\hat{\lambda}^1(t | A, W)$ now plays the role of the initial fit and the covariate $\bar{h}(t, A, W)$ is then reevaluated based on $\hat{\lambda}^1(t | A, W)$ (and thus $\hat{S}^1(t | A, W)$). Based on $\hat{\epsilon}$, estimated as described earlier, $\hat{\lambda}^2(t | A, W)$ is obtained. On the third iteration, $\hat{\lambda}^3(t | A, W)$ is obtained and the process is iterated until $\hat{\epsilon}$ is essentially zero. The final hazard fit at the last iteration of the algorithm is denoted by $\hat{\lambda}^*(t | A, W)$ with the corresponding survival fit given by $\hat{S}^*(t | A, W)$.

The procedure for estimation of the parameter of interest, $\Psi(p_0)$, now follows exactly that of the S-TMLE method described in Section 4.1. That is, estimates of the t_k -specific parameter are obtained and then averaged over time to obtain the targeted maximum likelihood estimate $\hat{\psi}^*$ of $\Psi(p_0)$. Note that $\hat{\psi}^*$ solves the efficient influence curve estimating equation $\sum_{i=1}^n \bar{IC}(\hat{\lambda}^*, g_0, \hat{G})(O_i)$; however, it does not solve $\sum_{i=1}^n IC_{1t_k}(\hat{\lambda}^*, g_0, \hat{G})(O_i)$ and $\sum_{i=1}^n IC_{0t_k}(\hat{\lambda}^*, g_0, \hat{G})(O_i)$ as does the S-TMLE.

4.2.1. Double Robustness Consistency Properties of the D-TMLE—It can be shown that the efficient influence curve $\bar{IC}$ cannot be written as an estimating equation in the parameter of interest $\Psi(p_0)$. Therefore, the formal proof of double robustness consistency properties does not follow in an obvious manner as with the usual estimating function approach.

Let $\psi(\lambda)$ be a solution to $P_0(\bar{IC}(\lambda, g_0, G_0)) = 0$. Empirically, we have found that solving $P_0(\bar{IC}(\lambda, g_0, G_0)) = 0$ does not imply $\psi(\lambda) = \psi_0$. In this case, the estimating function method breaks down since there are multiple solutions to $P_0(D(\lambda, g_0, G_0)) = 0$ and they do not all guarantee $\psi(\lambda) = \psi_0$. Furthermore, through simulation studies, we have found that if the initial

hazard (λ') is misspecified such that it solves $P_0(\bar{IC}(\lambda', g_0, G_0)) = 0$ but $\psi(\lambda')$ is inconsistent, then the targeted maximum likelihood algorithm does not update. That is, adding the ε -covariate results in a convex function in ε that should only have a single maximum in ε . In this very special case, the TMLE is inconsistent. However, through the simulation study presented in Section 6, and an extensive study of over 100 data-generating distributions (results provided in Appendix B), we have found that if the initial hazard does not solve

$P_0(\bar{IC}(\lambda', g_0, G_0)) = 0$, i.e., the algorithm has a chance to iterate, then the estimator $\psi(\lambda'^*)$ (where λ'^* is the updated hazard based on the targeting algorithm) is indeed consistent. That is, if the initial hazard solves the given equation, the algorithm ends in zero steps and no iterations can take place. If the initial hazard does not solve the given equation, the algorithm will iterate and, based on our results, the estimator based on the updated hazard from the last step of the algorithm is consistent. This was tested for a number of misspecified initial hazards including an intercept-only model. Thus, based on the empirical results, we conjecture that the double robustness properties of the S-TMLE also hold for the D-TMLE as well.

Another interesting note is that the estimates for $S_1(t_k)$ and $S_0(t_k)$ based on λ'^* are not necessarily (and are typically not) consistent. Thus, the following can, and often does, occur:

$$\sum_{t_k} \log \left(\frac{\log(S_1(t_k)(\lambda'^*))}{\log(S_0(t_k)(\lambda'^*))} \right) = \sum_{t_k} \log \left(\frac{\log(S_1(t_k)(\lambda_0))}{\log(S_0(t_k)(\lambda_0))} \right), \text{ but } S_1(t_k)(\lambda'^*) \neq S_1(t_k)(\lambda_0) \text{ and } S_0(t_k)(\lambda'^*) \neq S_0(t_k)(\lambda_0).$$

Thus, the overall logrank parameter is consistently estimated, but its components (i.e., $S_1(t_k)$ and $S_0(t_k)$) are not.

4.2.2. Inference for D-TMLE—The theorem provided in Appendix E can be used to derive the influence curve when $\hat{\lambda}^*$ converges to some misspecified λ^* , solving

$P_0 \bar{IC}(\lambda^*, g_0, G_0) = 0$ and satisfying $\psi(\lambda^*) = \psi_0$. When this is the case, from the preamble to the theorem, we can see that a contribution to the influence curve comes from $P_0 D^F(\hat{\lambda}^*)$, where D^F is the efficient influence curve for the full data, i.e., if there is no censoring. Thus, there is a contribution to the influence curve in this situation that would not be accounted for in Eq. (18). For correct inference in this situation, one can apply this theorem to derive the formal influence curve and use this for inference.

Since deriving the formal influence curve is not trivial in this case in which λ^* does not correctly specify all unknowns in the full data efficient influence curve, we can consider two straightforward alternatives. We first note that our empirical results suggest that the D-TMLE is more efficient than the S-TMLE. If indeed this property holds in general, then as the first alternative, we can use the influence curve of the S-TMLE to obtain a conservative estimate of the variance of the D-TMLE. The results then apply as in Section 4.1.3. A second alternative is the bootstrap procedure, although we note that this procedure can be computationally

intensive, particularly when K and n are large. The confidence intervals and t -statistics can be constructed as in Section 4.1.3.

4.3. Weight Function

The parameter provided in Eq. (5) is really a whole class of logrank parameters, indexed by a choice of weight function. Each can be used to provide a valid test for testing that the treatment effect is 0, i.e., $\psi = 0$. Therefore, it is of interest to choose a weight function that is most likely to have more power at the alternative. A perfectly reasonable choice of weight function is one that equally weights each of t_k -specific log ratios. A weight function that down weights the time points $t_k \in 1, \dots, J$, in the tail of which there is heavy censoring, may improve the power over the unit weight function. One such weight function is one that takes into account the

variance of $\log \left(\frac{\log(S_1(t_k))}{\log(S_0(t_k))} \right)$, that is, $w(t_k) = \frac{1}{\text{var}(IC_{t_k})}$, where IC_{t_k} is the efficient influence curve of $\log \left(\frac{\log(S_1(t_k))}{\log(S_0(t_k))} \right)$. This weight function puts more emphasis on those points $t_k \in 1, \dots, J$ for which there is less censoring, and thus more information, providing a more stable estimate of the parameter and the variance. In this article, we apply only unit weights.

5. INITIAL HAZARD ESTIMATION

To avoid the potential for selection of covariates to obtain favorable inference, it is imperative to use an a priori specified algorithm for the selection of the initial hazard that is specified in an analysis protocol. Since the hazard can be fitted with a logistic regression model, the initial hazard can be estimated with any model selection algorithm used to estimate logistic regression models with repeated measures. One such approach is the deletion/substitution/addition (D/S/A) algorithm (Sinisi and van der Laan, 2004). In this algorithm, the parameter of interest (in this case the conditional hazard for survival) is defined in terms of a loss function. Candidate estimators are then generated using deletion/substitution/addition moves that minimize, over subsets of variables (e.g., polynomial basis functions), the empirical risk of subset-specific estimators of the parameter of interest. Among these candidates, the estimator is selected using cross-validation. For purposes of this algorithm, the parameter of interest is the conditional hazard of survival, and we define it as

$$\lambda_0(t | A, W) = \arg \min_{\lambda} E_0 L(\lambda)$$

where the actual repeated measures loss function is given by

$$L(\lambda) = \sum_t w(t, A, W) \left[I(\tilde{T}=t, \Delta=1) - I(\tilde{T} \geq t) \lambda(t | A, W) \right]^2$$

where $w(t, A, W)$ is an arbitrary weight function. Alternatively, we can use the log-likelihood loss function for $\lambda(t | A, W)$, and apply the D/S/A algorithm as a standard logistic regression model selection with repeated measures for each subject, where candidate estimators are generated also based on the log-likelihood loss function.

An even more optimal algorithm for estimation of the initial conditional hazard for survival is the super-learner algorithm, which begins by selecting a set of candidate prediction algorithms ("learners") which ideally cover different basis functions (van der Laan et al., 2007). For example, one such learner could be the D/S/A algorithm. The super-learner algorithm then selects α that minimizes

$$E_{B_n} P_{n,B_n}^1 L \left(\sum_j \alpha(j) \widehat{\lambda}_j(P_{n,B_n}^0) \right)$$

where $B_n \in \{0, 1\}$ denotes a random binary vector whose realizations define a split of the learning sample into a training sample $\{i : B_n(i) = 0\}$ and validation sample $\{i : B_n(i) = 1\}$, and P_{n,B_n}^1 and P_{n,B_n}^0 denote the empirical probability distributions of the validation and training sample, respectively. This minimization problem can be solved by formulating it as a least-squares regression problem. Thus, the algorithm finds optimal weighted combinations of the candidate estimators with respect to the squared error (L_2) loss function, with weights defined by α . We note that candidate estimators can be based on the log-likelihood loss function.

In addition to the hazard for survival, the hazard for censoring must also be estimated. One of the algorithms discussed earlier can also be applied to estimate the censoring mechanism. In particular, the super-learner algorithm can be applied to obtain an estimate for the hazard for censoring in the same manner as for the hazard for survival. If censoring is uninformative, then one can use Kaplan–Meier to estimate the censoring mechanism.

6. SIMULATION STUDIES

Data were simulated to mimic an RCT in which the goal is to determine the effectiveness of a new drug in comparison to the current standard of care on “survival” as measured by a occurrence of an event (e.g., particular marker falling below a given level) by 9 months into treatment. The probability of receiving the new treatment is 0.5. Two covariates were negatively correlated with survival time; for example, these covariates might represent age in years (multiplied by 0.1) and weight gain in the year prior to baseline. Specifically, the 2500 replicates of sample size 500 were generated based on the following data-generating distribution where time is discrete and takes values $t_k \in \{1, \dots, 9\}$:

- $Pr(A = 1) = Pr(A = 0) = 0.5$
- $W_1 \sim U(2, 6)$
- $W_2 \sim N(10, 10)$
- $$\lambda(t | A, W) = \frac{I(t_k < 9) I(Y(t_k - 1) = 0)}{1 + \exp\left(-\left(-8 - 0.75A + 0.3W_1^2 + 0.25W_2\right)\right)} + I(t_k = 9)$$

where $\lambda(t | A, W)$ is the hazard for survival and $Y(t_k)$ is the indicator that the event has occurred at or before time t_k . The linear correlations between $\{W_1, W_2\}$ and failure time were approximately $\{-0.62, -0.52\}$. Three different types of censoring were simulated: no censoring, MCAR, and MAR. The MCAR and MAR censoring mechanisms were set such that approximately 27% and 20% of the observations were censored, respectively. The censoring was generated to ensure that $\bar{G}(t | A, W) > 0$ (see the Discussion section for details of this assumption). If censoring and failure time were tied, the subject was considered uncensored. Under MCAR, the hazard for censoring was $\lambda_C(t) = 0.15$. Under MAR censoring, the hazard for censoring depends on A and W_1 where the treated subjects ($A = 1$) have a much higher hazard for censoring for high levels of W_1 than the untreated subjects, whereas the untreated subjects have a much higher hazard for censoring than the treated subjects for low levels of W_1 . The specific censoring mechanism is provided in Appendix C.

The unadjusted estimator was applied as defined in Section 3. The two targeted maximum likelihood methods provided in sections 4.1 and 4.2 were applied using three different initial

hazard fits. The first initial hazard was correctly specified. The second initial hazard was misspecified by including only a main term for A and W_1 . The third initial hazard was misspecified by including only a main term for A and W_2 . In the MCAR censoring setting, the censoring mechanism was correctly estimated using Kaplan–Meier. In the MAR censoring setting, the censoring mechanism was correctly specified. Inference for the D-TMLE was based on the variance of the S-TMLE.

The estimators were compared using a relative efficiency (RE) measure based on the mean squared error (MSE), computed as the MSE of the unadjusted estimates divided by the MSE of the targeted maximum likelihood estimates. Thus, a value greater than 1 indicates a gain in efficiency of the covariate adjusted TMLE over the unadjusted estimator.

In addition to these three simulation scenarios, to explore the relationship between RE and the correlation between the covariate and failure time, we generated data with a hazard that is based on A and a single covariate W . The data were simulated such that the correlation between W and failure time ranged from -0.1 through -0.8 while the effect of A on survival remained constant. The data were simulated with these increasing correlations between W and T with both weak effect and strong effects of treatment. Note that in each of these simulation scenarios, there was in fact a treatment effect. We also simulated data according to the same distributions with the exception that the coefficient for treatment was zero to check the Type I error. A sample of the results is provided in Appendix D, where we found that indeed the 0.05 level was correctly maintained.

To demonstrate that the updating of the initial hazard in the TMLE algorithm does indeed result in a bias reduction for the parameter of interest, we also estimated this parameter directly based on the initial hazard. That is, in Eq. (14), instead of $S^*(t | A, W)$ based on $\lambda^*(t | A, W)$, we use $S^0(t | A, W)$ based on the initial hazard $\lambda^0(t | A, W)$. Such a substitution estimator relies on correct specification of the regression model, just as the Cox proportional hazards regression model would result in a model-dependent substitution estimator. Thus, we only include these results to demonstrate the bias reduction of the TMLE when the initial hazard is misspecified.

Lastly, in Section 6.3, we provide a simulation study to demonstrate the importance of the use of data-adaptive algorithms in the estimation of the initial hazard with respect to maximal gains in power.

6.1. Simulation Results and Discussion for Various Censoring Scenarios

In the no censoring and MCAR censoring scenarios, the bias should be approximately zero. In this strong covariate setting, exploiting this covariate by applying the TMLE should provide a gain precision due to a reduction in the residuals. In the informative censoring setting (MAR), in addition to the expected gain in efficiency we expect a reduction in bias of the TMLE with the correctly specified treatment mechanism over the unadjusted estimator. The informative censoring is accounted for through the covariate h that is inverse weighted by the subjects' conditional probability of being observed at time t given their observed history.

Tables 1–3 provide the bias, relative MSE, and power based for the unadjusted and two targeted maximum likelihood approaches for the no censoring, MCAR censoring, and MAR censoring settings respectively. The results show that indeed the expected gain in efficiency is achieved in the no censoring and MCAR censoring scenarios. When the initial hazard was correctly specified, the gain in power for the TMLE was as high as 0.57 in the no censoring scenario over the unadjusted estimator (Table 1). Although the gains are more modest when the initial hazard is misspecified, the gain in power was as high as 0.22 for the TMLE over the unadjusted estimator. Under the MCAR censoring scenario, the gains under misspecification were somewhat smaller, with an increases in power between 0.06 and 0.51 from misspecified to

correctly specified initial hazards (Table 2). These results demonstrate that when the initial hazard fit can be closely approximated, the potential reduction in standard error and thus increase in power is substantial.

Under the MAR setting, the unadjusted estimate is severely biased (~21%), whereas both the TMLEs remain consistent. In such a setting, one must account for the informative censoring as the results from the unadjusted method are completely unreliable. This is a strong advantage of this methodology as it accounts for this bias-inducing censoring, which is often ignored or not correctly handled in RCTs (Wood et al., 2004).

We also note that for all censoring scenarios, the REs are all greater for the D-TMLE. However, the actual power is slightly lower than the S-TMLE. This is due to the fact that at this small sample size, there happens to be a tiny amount of negative finite sample bias for the D-TMLE (average of the 2,500 point estimates is slightly smaller in absolute value than the truth), whereas the S-TMLE is slightly positively biased. Thus, even though the D-TMLE is more efficient than the S-TMLE, the absolute values of the point estimates are slightly smaller, causing the t -statistics to be smaller as well. Thus, the power is lower as well. For larger sample sizes, as the finite sample is eliminated, the power for the D-TMLE will be at least as large as the power for the S-TMLE. Also, as expected, the inference was slightly conservative as compared to the S-TMLE, although one cannot observe this from the presented results that are rounded to 10^{-2} . The bootstrap procedure would provide less conservative inference.

The results for the substitution estimator based on the initial hazard in the no censoring scenario are provided in Table 4. As expected, when the initial hazard is correctly specified, the performance of the substitution estimator is similar to both TMLEs. However, when the initial hazard is misspecified, it becomes clear that the targeting algorithm is indeed updating the hazard in such a way that the bias is essentially removed. If the updating were not done, from Table 4, we can see that the bias for the first misspecified hazard was 6% and for the second misspecified hazard the bias was 21%.

6.2. Relationship Between Correlation of Covariate(s) and Failure Time with Efficiency Gain

As the correlation between covariates and failure time increases we expect to observe increasing gains in efficiency. In this simulation study, we include only a single covariate W with no censoring. For simplicity, we include the results for the S-TMLE only. Figure 1 clearly demonstrates that as the correlation between W and failure time increases, so does the gain in power of the S-TMLE over the unadjusted. The gain in power for the strong treatment effect setting has a nearly linear relationship with increasing correlation between W and failure. The weak treatment effect setting has moderate gains until the covariate effect is very strong, when the corresponding gain in power of the S-TMLE is very high. These results reflect similar findings in RCTs with fixed-end point studies where relations between R^2 and efficiency gain have been demonstrated (Moore and van der Laan, 2009b; Pocock et al., 2002; Tsiatis et al., 2008).

6.3. Importance of the Use of Data-Adaptive Algorithms for Initial Hazard Estimation

In this section, we provide a simulation example to demonstrate the power that one can obtain by using an aggressive algorithm for estimation of the initial hazard, in comparison to a simpler main term only model selection algorithm. These methods are compared to the standard unadjusted approach. This section is not meant to examine the super-learner algorithm in detail, but rather to demonstrate that its use with targeted maximum likelihood estimation can result in significant gains in power over targeted maximum likelihood estimation with less aggressive algorithms. For details on the algorithm as well as some of the candidate learners, we refer the reader to the original paper (van der Laan et al., 2007).

We simulated 500 replicates of sample size 500 from the following data-generating distribution where time is discrete and takes values $t_k \in \{1, \dots, 8\}$:

- $\Pr(A = 1) \Pr(A = 0) = 0.5$
- $W_1 \sim U(2, 5)$
- $W_2 \sim N(-2, 2)$
- $\Pr(W_3 = 1) = 0.3 = 1 - \Pr(W_3 = 0)$
- $W_4 \sim N(1, 1)$
- $W_5 \sim U(-2, 4)$
- $$\lambda(t|A, W) = \frac{I(t_k < 8) I(Y(t_k - 1) = 0)}{1 + \exp(-(-3 - 0.5A + 0.2W_1W_2 + 0.2W_2W_4 - 0.4W_4W_5 + 0.5W_5W_2))} + I(t_k = 8)$$

The data were generated with no censoring.

In the first method, the initial hazard was estimated using the super-learner algorithm (see Section 5), which included as candidate learners:

- Simple linear regression with all five covariates as main terms.
- Lasso logistic regression (Tibshirani, 1996).
- Random forest (Breiman, 2001).
- Generalized additive models (Hastie and Tibshirani, 1990).
- K -nearest neighbor classification (Ripley, 1996).

In the second method, the D/S/A algorithm was applied, allowing only main terms (i.e., no interactions or terms with powers greater than 1 were considered). The targeted maximum likelihood method was then applied using these two different initial hazard estimates. For brevity, we include only the results for the S-TMLE. Lastly, for comparison, the unadjusted method was applied.

The percent bias, power, and RE results are provided in Table 5. The results demonstrate that both methods of adjustment result in a gain in power over the unadjusted method. The power for the S-TMLE with super-learner approach is nearly double that of the unadjusted method, and is 30% higher than the S-TMLE with the D/S/A using main terms only. Thus, a significant loss in power would result if an aggressive algorithm for the initial hazard estimation is not utilized. It is of note that a 50% gain in power over the unadjusted was achieved using S-TMLE with the simple main term only approach. It is clear from these results that the targeted maximum likelihood method of covariate adjustment provides gains in efficiency, even with suboptimal methods for initial hazard estimation. However, for even larger gains in efficiency and thus power, more aggressive algorithms such as the super-learner should be used in combination with the targeted maximum likelihood. We note that the method for selection of the initial hazard should be specified in the analysis protocol.

6.4. Analysis of Single Dataset

In this section, we analyze a single dataset that was simulated based on an actual RCT. The RCT aimed to evaluate the effect of a new treatment versus a standard on the time until viral load fell below 400 copies/ml. The $n = 750$ patients were followed monthly for a period of 12 months. Four covariates were associated with the viral load that was measured over time: baseline viral load, age, regional indicator Asia, and regional indicator Europe. The covariates were sampled with replacement from the original data. The median of the baseline viral load

was approximately 615 copies/ml, the median age was 33, and the probabilities that a patient was from Asia or Europe were 0.3 and 0.15, respectively. The probability of receiving the new treatment was 0.5. Viral loads at each subsequent month were simulated from these four variables in addition to treatment. The first month at which the viral load was below 400 copies/ml was the survival time T for a given patient. In the dataset analyzed, the correlations between baseline viral load, age, regional indicator Asia, and regional indicator Europe with survival time were 0.3, -0.08 , 0.12 , and -0.07 . Approximately 20% of patients were randomly censored. The S-TMLE was applied using the super-learner to estimate the initial hazard. Kaplan–Meier was used to estimate the censoring mechanism. The unadjusted estimator as outlined in Section 3 was also applied.

The point estimate, standard error and p values are provided in Table 6. The standard error was reduced by 11% by adjusting for the covariates using the S-TMLE estimator, and the corresponding p value was reduced from 0.07 to 0.04, which would change the conclusion of the test of the null hypothesis of no treatment effect. Thus, this dataset, which was simulated to mimic a real RCT, demonstrates that even moderate gains in efficiency can change the conclusions of the study.

7. DISCUSSION

The simulation studies provided in this article clearly demonstrate that significant gains in efficiency and thus power can be achieved over the unadjusted ubiquitous logrank method through covariate adjustment using the targeted maximum likelihood approach. Both the targeted maximum likelihood methods for covariate adjustment presented in this article do not require additional assumptions beyond those required for the logrank test. With the S-TMLE, we were able to show the double robustness consistency properties based on estimating function methodology. With the D-TMLE, the estimating function methodology could not be applied and therefore we provided extensive empirical evidence that these properties also held for this estimator.

We note that the methods presented in this article differ from adjusting through the Cox proportional hazards models as was done in Hernández et al. (2006), which requires additional assumptions about the proportionality of hazards. Furthermore, the method presented in this article provides a method for estimation of the marginal or population level effect of treatment, rather than a conditional effect from a Cox or covariate-adjusted logistic hazard model. We note that the method of targeted maximum likelihood estimation can also be applied to the estimation of conditional or subgroup effects of treatment; however, we focused on marginal effects only in this article.

The simulation study results also demonstrate the importance of the estimation of the initial hazard in optimizing gains in power. The ideal approach includes two steps, where in the first the initial hazard is estimated based on an aggressive data-adaptive approach such as the super-learner algorithm, and in the second the targeting maximum likelihood step is applied as a bias-reduction step for the parameter of interest. These two steps combined provide consistent estimates of the treatment effect with large gains in power over the procedure that ignores covariates.

It is also important to note that the TMLE, like other inverse weighted estimators, relies on the assumption that each subject has a positive probability of being observed (i.e., not censored) just before time t . More formally, this assumption is $\bar{G}(t_- | A, W) > 0$, $t_k \in 1, \dots, J$. This identifiability assumption has been addressed as an important assumption for right-censored data (Robins and Rotnitzky, 1992). One is alerted to such violations by observing very small probabilities of remaining uncensored based on the estimated censoring mechanism, i.e., there

are patients with a probability of censoring of almost one given their observed past. We recommend that one should check that $\bar{G}(t_- | A, W) > 0.1$ in practice. When violations of this assumption are present, an new advance to the targeted maximum likelihood approach presented in this article, namely, collaborative targeted maximum likelihood, can be applied (van der Laan and Gruber, 2009). In this approach, a sequence of TMLEs is generated with increasing likelihood indexed by increasingly non-parametric estimates of the censoring mechanism. The censoring mechanism estimator, for which the targeted maximum likelihood step results in the most effective bias reduction with respect to the parameter of interest, is selected using likelihood based cross-validation. Essentially, in this approach, covariates are only included in the censoring mechanism fit if they improve the targeting of the parameter of interest while not grossly affecting the MSE.

The methodology presented in this article can easily be extended to estimation of causal effects in observational studies, such as post-market safety studies. This includes estimation of the causal parameters presented in this article, as well as more complex parameters as defined by marginal structural models. A further extension, in both RCTs and observational studies, is the inclusion of time-dependent covariates, which are often predictive of censoring and/or survival (van der Laan, 2008; Moore and van der Laan, 2009a). Future work includes the application of these methods in observational studies.

APPENDIX A: TMLE FOR t_k -SPECIFIC PARAMETER IS (COLLABORATIVE) DOUBLY ROBUST

We now show that

$$E_0 D^*(S, g_0, G_0, S_{0,1}(t_k)) = 0$$

for any $S = S(\cdot | A, W)$, where $D^*(S, g, G, \Psi(S))$ is the efficient influence curve of $\Psi(S) = S_1(t_k)$ at the data-generating distribution identified by S, g, G and the marginal distribution of W .

We have

$$\begin{aligned} D^*(S, g_0, G_0, S_{0,1}(t_k)) &= - \sum_{t \leq t_k} \frac{I(A=1)}{\bar{G}_0(t_- | A=1, W) g_0(1 | W)} \frac{S(t_k | A=1, W)}{S(t | A=1, W)} \\ &\quad \times \left(dN(t) - E\left(dN(t) \mid \bar{N}(t_-), A=1, W\right) \right) \\ &\quad + S(t_k | A=1, W) \\ &\quad - S_{0,1}(t_k) \end{aligned}$$

Rewrite this as

$$\begin{aligned}
& D^*(S, g_0, G_0, S_{0,1}(t_k)) \\
&= - \sum_{t \leq t_k} \frac{I(C \geq t)}{\bar{G}_0(t- | A=1, W)} \frac{I(A=1)}{g_0(1 | W)} \frac{S(t_k | A=1, W)}{S(t | A=1, W)} \\
&\quad \times (I(T=t) - I(T \geq t) \lambda(t | A=1, W)) \\
&\quad + S(t_k | A=1, W) \\
&\quad - S_{0,1}(t_k)
\end{aligned}$$

We now take the conditional expectation given $I(\tilde{T} \geq t)$, A , W , to obtain

$$\begin{aligned}
& -P_0 D^*(S, g_0, G_0, S_{0,1}(t_k)) \\
&= P_0 \sum_{t \leq t_k} \frac{I(C \geq t)}{\bar{G}_0(t- | A=1, W)} \frac{I(A=1)}{g_0(1 | W)} \frac{S(t_k | A=1, W)}{S(t | A=1, W)} \\
&\quad \times I(T \geq t) (\lambda_0(t | A=1, W) - \lambda(t | A=1, W)) \\
&\quad + P_0 S(t_k | A=1, W) \\
&\quad - S_{0,1}(t_k)
\end{aligned}$$

We now take the conditional expectation, given $I(C \geq t)$, A , W , to obtain

$$\begin{aligned}
& -P_0 D^*(S, g_0, G_0, S_{0,1}(t_k)) \\
&= P_0 \sum_{t \leq t_k} \frac{I(C \geq t)}{\bar{G}_0(t- | A=1, W)} \frac{I(A=1)}{g_0(1 | W)} \frac{S(t_k | A=1, W)}{S(t | A=1, W)} \\
&\quad \times S_0(t- | A, W) (\lambda_0(t | A=1, W) - \lambda(t | A=1, W)) \\
&\quad + P_0 S(t_k | A=1, W) \\
&\quad - S_{0,1}(t_k)
\end{aligned}$$

Define now the term

$$W^*(t) \equiv \frac{S(t_k | A=1, W)}{S(t | A=1, W)} S_0(t- | A, W) (\lambda_0(t | A=1, W) - \lambda(t | A=1, W))$$

If

$$E\left(I(C \geq t, A=1) | W^*(t), \bar{G}_0(t- | A=1, W) g_0(1 | W)\right) = \bar{G}_0(t- | A=1, W) g_0(1 | W)$$

then we obtain

$$\begin{aligned}
& -P_0 D^*(S, g_0, G_0, S_{0,1}(t_k)) \\
&= P_0 \sum_{t \leq t_k} \frac{S(t_k | A=1, W)}{S(t | A=1, W)} S_0(t- | A, W) (\lambda_0(t | A=1, W) - \lambda(t | A=1, W)) \\
&\quad + P_0 S(t_k | A=1, W) \\
&\quad - S_{0,1}(t_k)
\end{aligned}$$

In particular, this identity applies if $\bar{G}_0(t- | A=1, W)$ and $g_0(1 | W)$ are the true conditional distributions, given the whole W , but the requirement just shown is weaker by only requiring that the censoring mechanism G_0 and treatment mechanism g_0 are conditioning on a function of W depending directly on $\lambda - \lambda_0$. In particular, if $\lambda = \lambda_0$, then treatment and censoring mechanism do not need to condition on any covariates at all.

We now take the right-hand side of the latter identity as starting point and prove that it equals zero for all S . Note that this term has nothing to do with the censoring and treatment mechanism anymore. First, note that it can be rewritten as:

$$\begin{aligned} & -P_0 D^*(S, g_0, G_0, \psi_0) \\ &= -P_0 \sum_{t \leq t_k} S(t_k | A=1, W) \\ & \quad \times \left[\frac{S(t_- | A=1, W) f_0(t | A=1, W) - S_0(t_- | A=1, W) f(t | A=1, W)}{S(t | A=1, W) S(t_- | A=1, W)} \right] \\ &+ P_0 S(t_k | A=1, W) \\ & - S_{0,1}(t_k) \end{aligned}$$

Here we used ψ_0 to indicate the true target parameter $S_{0,1}(t_k)$. Now we can use the algebraic trick, $ab - cd = (a - c)d + a(b - d)$ to obtain the expression

$$\begin{aligned} & E_0 D^*(S, g_0, G_0, \psi_0) \\ &= -E_0 \sum_{t \leq t_k} S(t_k | A=1, W) \left[\frac{(S(t_- | A=1, W) - S_0(t_- | A=1, W)) f(t | A=1, W)}{S(t | A=1, W) S(t_- | A=1, W)} + \frac{S(t_- | A=1, W) (f_0(t | A=1, W) - f(t | A=1, W))}{S(t | A=1, W) S(t_- | A=1, W)} \right] \\ &+ P_0 S(t_k | A=1, W) \\ & - S_{0,1}(t_k) \end{aligned}$$

For the continuous survival case, we note that

$$\begin{aligned} \int_{[0, t_k]} \left(\frac{S(t_- | A=1, W) - S_0(t_- | A=1, W)}{S(t_- | A=1, W)} \right) &= \int_{[0, t_k]} - \frac{(S(t_- | A=1, W) - S_0(t_- | A=1, W)) f(t_- | A=1, W)}{S^2(t_- | A=1, W)} + \frac{f(t_- | A=1, W)}{S(t_- | A=1, W)} - \frac{f_0(t_- | A=1, W)}{S(t_- | A=1, W)} \\ &= \frac{(S(t_k | A=1, W) - S_0(t_k | A=1, W)) f(t_- | A=1, W)}{S^2(t_k | A=1, W)} + \frac{(f_0(t_k | A=1, W) - f(t_k | A=1, W)) S(t_k | A=1, W)}{S^2(t_k | A=1, W)}. \end{aligned}$$

Therefore, we have

$$\begin{aligned} E_0(D^*(S, g_0, G_0, \psi_0 | X)) &= -E_0 \frac{S(t_k | A=1, W) (S(t_- | A=1, W) - S_0(t_- | A=1, W))}{S(t_k | A=1, W)} + E_0 S(t_k | A=1, W) - E_0 S_{0,1}(t_k) \\ &= -E_0 S(t_k | A=1, W) + S_{0,1}(t_k) + E_0 S(t_k | A=1, W) - S_{0,1}(t_k) \\ &= 0 \end{aligned}$$

Now for the discrete survival case, it can be shown with some algebra that

$$\begin{aligned}
& \frac{S(t_k | A=1, W) - S_0(t_0 | A=1, W)}{S(t_k | A=1, W)} \\
& - \frac{S(t_{k-1} | A=1, W) - S_0(t_{k-1} | A=1, W)}{S(t_{k-1} | A=1, W)} \\
& = \frac{S(t_k | A=1, W) - S_0(t_k | A=1, W) f(t_k)}{S(t_k | A=1, W) S(t_{k-1} | A=1, W)} \\
& - \frac{f(t_k | A=1, W) - f_0(t_k | A=1, W) S(t_k | A=1, W)}{S(t_k | A=1, W) S(t_{k-1} | A=1, W)}
\end{aligned}$$

Furthermore,

$$\sum_{t_k \leq t_k} \left[\frac{S(t_k | A=1, W) - S_0(t_k | A=1, W)}{S(t_k | A=1, W)} - \frac{S(t_{k-1} | A=1, W) - S_0(t_{k-1} | A=1, W)}{S(t_{k-1} | A=1, W)} \right] = \frac{S(t_k | A=1, W) - S_0(t_k | A=1, W)}{S(t_k | A=1, W)}$$

Therefore, for discrete survival we can write,

$$\begin{aligned}
E_0 D^*(S, g_0, G_0, \psi_0 | X) &= -E_0 \sum_{t \leq t_k} S(t_k | A=1, W) \left[\frac{S(t_0 | A=1, W) - S_0(t_0 | A=1, W)}{S(t_0 | A=1, W)} \right] + E_0 S(t_k | A=1, W) - S_{0,1}(t_k) \\
&= 0
\end{aligned}$$

APPENDIX B: EMPIRICAL VALIDATION FOR D-TMLE CONSISTENCY

In this appendix, results based on extensive simulations are included with the purpose of providing empirical validation of the consistency of the D-TMLE when targeting the single parameter directly. We consider the scenario where $\widehat{\lambda}^*$ converges to some misspecified λ^* but the efficient influence curve estimating equation is solved and parameter is consistently estimated. Remarkably, even though the efficient influence curve estimating equation is not an estimating equation with a variation-independent parameterization in the parameter of interest (ψ) and the nuisance parameters, we still obtain consistent estimates of ψ even with this misspecified λ .

Four different types data-generating distributions were used. For all four data-generating mechanisms, the two covariate distributions were given by $W_1 \sim U(2, 6)$ and $W_2 \sim N(1, 1)$. The data generating distributions differed by definition of the hazard and treatment mechanism. In the first three settings, treatment was randomized, $P(A=1) = P(A=0) = 0.5$. In the fourth

setting, treatment mechanism given by $P(A=1 | W) = \frac{1}{1 + \exp(-(-0.75 + 3 * W_2))}$ and $P(A=0 | W) = 1 - P(A=1 | W)$. The settings are summarized as follows:

1. Constant hazard, treatment randomized:

$$\lambda(t | A, W) = \frac{I(t_k < 10) I(Y(t_k - 1) = 0)}{1 + \exp(-(\beta_0 - \beta_1 A + \beta_2 W_1^2 + \beta_3 W_2))} + I(t_k = 10)$$

2. Hazard changed over time, treatment randomized:

$$\lambda(t|A, W) = \frac{I(t_k < 10) I(Y(t_k - 1) = 0)}{1 + \exp(-(\beta_0 - \beta_1 A + \beta_2 W_1^2 + \beta_3 W_2 + \beta_4 * t_k))} + I(t_k = 10)$$

3. Hazard changed over time (interaction between time and covariate), treatment randomized:

$$\lambda(t|A, W) = \frac{I(t_k < 10) I(Y(t_k - 1) = 0)}{1 + \exp(-(\beta_0 - \beta_1 A + \beta_2 W_1 * t_k + \beta_3 W_2))} + I(t_k = 10)$$

4. Constant hazard, treatment not randomized:

$$\lambda(t|A, W) = \frac{I(t_k < 10) I(Y(t_k - 1) = 0)}{1 + \exp(-(\beta_0 - \beta_1 A + \beta_2 W_1^2 + \beta_3 W_2))} + I(t_k = 10)$$

For each setting, 25 sets of parameter values that define the hazard were selected from the following distributions, respectively:

1. $\beta_0 \sim U(-8, -1)$, $\beta_1 \sim U(-1, 1)$, $\beta_2 \sim U(-0.5, 0.5)$, and $\beta_3 \sim U(-2.5, 2.5)$.
2. $\beta_0 \sim U(-3, 1)$, $\beta_1 \sim U(-1, 1)$, $\beta_2 \sim U(-0.4, 0.4)$, $\beta_3 \sim U(-2, 2)$, and $\beta_4 \sim U(-0.2, 0.4)$.
3. $\beta_0 \sim U(-5, -1)$, $\beta_1 \sim U(-1, 1)$, $\beta_2 \sim U(-0.1, 0.2)$, and $\beta_3 \sim U(0, 2)$.
4. $\beta_0 \sim U(-8, -1)$, $\beta_1 \sim U(-1, 1)$, $\beta_2 \sim U(-0.5, 0.5)$, and $\beta_3 \sim U(-2.5, 2.5)$.

For simplicity, there was no censoring. Generating the data in such a way provided different levels of correlation between the covariates and the outcome, as well as differing effects of treatment on survival, from negative to positive. A single large dataset of with $n = 25,000$ observations was generated for each of the 25 different randomly selected parameter values for each of the four settings (i.e., 100 simulation settings). The D-TMLE of the average minimal parameter was applied. The initial hazard was estimated based on (1) an intercept only model, (2) treatment and covariate W_1 only, and (3) treatment and covariate W_2 only. In the nonrandomized setting, the treatment mechanism was correctly specified.

The results for the 100 simulations, segmented by simulation setting and within each setting ordered by ψ_0 are provided in Figures B1–B3, which correspond with different misspecified initial hazard estimates. Clearly, the D-TMLE, even for the intercept only initial hazard (Figure B1) is consistently estimating ψ . There remains some noise due to the fact that the initial hazard is so grossly misspecified; however, the D-TMLE is still performing well. Once covariates are included, even though the initial hazard is still misspecified, the D-TMLE is performing even better (Figures B2 and B3). These results provide empirical evidence that indeed although the consistency properties of the D-TMLE cannot be derived based on the usual estimating function methodology, that the estimator is consistent as long as the initial estimate is not already a solution to the efficient estimating equation (i.e., the algorithm has a chance to iterate).

APPENDIX C: MAR CENSORING MECHANISM USED IN SIMULATION STUDY

For $t_k \in 2, \dots, 9$,

$$\lambda_c(t|A, W_1) = \begin{cases} 0.25 & \text{if } W_1 > 4.5 \text{ and } A=1 \\ 0.2 & \text{if } 4.5 \leq W_1 < 3.5 \text{ and } A=1 \\ 0.05 & \text{if } 3.5 \leq W_1 < 2.5 \text{ and } A=1 \\ 0 & \text{if } W_1 > 3.5 \text{ and } A=0 \\ 0.25 & \text{if } 3.5 \leq W_1 < 2.5 \text{ and } A=0 \\ 0.05 & \text{if } W_1 \leq 2.5 \end{cases}$$

For $t_k = 1$, $\lambda_c(t|A, W_1) = 0$.

APPENDIX D: CHECKING TYPE I ERROR CONTROL

The data were simulated using the same data generating distribution as in Section 6 with the exception that the coefficient for treatment was set to zero. The results for the no censoring scenario for the D-TMLE are provided in Table D1 and show that indeed the type I error control was maintained.

APPENDIX E: ASYMPTOTIC LINEARITY OF D-TMLE AND TEMPLATE FOR DERIVATION OF INFLUENCE CURVE

In this appendix, we establish the asymptotic linearity of the D-TMLE for the parameter of interest without having to use that the gradient or canonical gradient of the pathwise derivative can be represented as an estimating function for the parameter of interest. In addition, a template for the formal derivation of the influence curve is outlined.

Consider CAR-censored data models so that $D(P) = D(\lambda(P), G(P))$, $\Psi(P)$ depends on P through $\lambda(P)$ only, and the density factorizes as $p = \lambda(p)G(p)$. Let $p_n^* = \widehat{\lambda}^* \widehat{G}^*$. We consider the case that $\widehat{G}^*$ is assumed to be consistent for G_0 . We proceed as follows:

$$\begin{aligned} P_0 D(\widehat{\lambda}^*, \widehat{G}^*) &= P_0 D(\widehat{\lambda}^*, G_0) - \{P_0 D(\widehat{\lambda}^*, \widehat{G}^*) - D(\widehat{\lambda}^*, G_0)\} \\ &= P_0 D(\widehat{\lambda}^*, G_0) + P_0 \{D(\lambda_0^*, \widehat{G}^*) - D(\lambda_0^*, G_0)\} - R_{1n} \end{aligned}$$

where

$$R_{1n} = P_0 \{D(\widehat{\lambda}^*, \widehat{G}^*) - D(\widehat{\lambda}^*, G_0)\} - P_0 \{D(\lambda_0^*, G_0) - D(\lambda_0^*, G_0)\}$$

Here R_{1n} is a second-order term and therefore it is natural to make it an assumption that $R_{1n} = o_p(1/\sqrt{n})$. Second, we define

$$\Phi(\widehat{G}^*) \equiv P_0 D(\lambda_0^*, \widehat{G}^*),$$

so that the term $P_0 \{D(\lambda_0^*, \widehat{G}^*) - D(\lambda_0^*, G_0)\}$ equals $\Phi(\widehat{G}^*) - \Phi(G_0)$. We now assume that $\Phi(\widehat{G}^*)$ is an efficient estimator of the parameter $\Phi(G_0)$ in model $\mathcal{M}(\mathcal{G}) = \{p_{\lambda, G} = \lambda G : G \in \mathcal{G}\}$, where we denote the tangent space generated by model $\mathcal{G}$ for G_0 at $P_0 = \lambda_0 G_0$ with $T_{\mathcal{G}}(P_0)$. It remains to consider $P_0 D(\widehat{\lambda}^*, G_0)$. By the general representation Theorem 1.3 in van der Laan and Robins (2003) it follows that

$$P_0 D(\widehat{\lambda}^*, G_0) = P_{\lambda_0} D^F(\widehat{\lambda}^*)$$

where $D^F(\lambda)$ is a gradient in the full data model λ for the parameter $\lambda \rightarrow \Psi(\lambda)$, and P_{λ_0} denotes the full data distribution. Again, by pathwise differentiability of Ψ in the full data model, if $D^F(\widehat{\lambda}^*)$ consistently estimates $D^F(\lambda_0)$, then one expects

$P_{\lambda_0} D^F(\widehat{\lambda}^*) = \psi_0 - \Psi(\widehat{\lambda}^*) + o_p(1/\sqrt{n})$. In general, we note that, if $\widehat{\lambda}^*$ converges to some possibly misspecified λ^* for which $\Psi(\lambda^*) = \Psi(\lambda_0)$ and $P_0 D^F(\lambda^*) = 0$, we have

$$P_{\lambda_0} D^F(\widehat{\lambda}^*) = P_{\lambda^*} D^F(\widehat{\lambda}^*) + P_{\lambda_0 - \lambda^*} \{D^F(\widehat{\lambda}^*) - D^F(\lambda^*)\}$$

By pathwise differentiability, and the convergence of $\widehat{\lambda}^*$ to λ^* , the first-order Taylor expansion suggests

$$P_{\lambda^*} D^F(\widehat{\lambda}^*) = \psi_0 - \Psi(\widehat{\lambda}^*) + o_p(1/\sqrt{n})$$

A separate study of the other term (which can be represented as $\Phi(\widehat{\lambda}^*) - \Phi(\lambda^*)$ for some Φ) will result in an asymptotic linearity result:

$$P_{\lambda_0 - \lambda^*} \{D^F(\widehat{\lambda}^*) - D^F(\lambda^*)\} = (P_n - P_0) D_1(P_0) + o_p(1/\sqrt{n})$$

To stay general, we assume the expansion:

$$P_0 D(\widehat{\lambda}^*, G_0) = \psi_0 - \Psi(\widehat{\lambda}^*) + \frac{1}{n} \sum_{i=1}^n D_1(P_0) + o_p(1/\sqrt{n})$$

for some $D_1(P_0)$. By Theorem 2.3 in van der Laan and Robins (2003) the influence curve of $\Phi(\widehat{G}^*)$ equals $-\Pi(D(\lambda_0^*, G_0) + D_1(P_0) | T_g(P_0)^\perp)$. This proves the following which provides a template for establishing asymptotic linearity of the D-TMLE in CAR censored data models.

Theorem 1. Let $O_1, \dots, O_n \sim P_0$ be n i.i.d. copies of $O = \Phi(C, X)$ for some many to one mapping Φ of censoring variable C and full data structure X . Assume that the conditional distribution G_0 of C , given X , satisfies CAR so that $p_0 = \lambda_0 G_0$ w.r.t to appropriate dominating measure, G_0 is a density of G_0 and λ_0 a function of distribution of full data X . Let $\mathcal{M} = \{p_{\lambda G} : \lambda \in \Lambda, G \in \mathcal{G}\}$, where $\mathcal{G}$ is a subset of all CAR distributions. Let $\Psi : \lambda \rightarrow \mathbb{R}^d$ be the Euclidean target parameter of interest. Let $D(P) = D(\lambda(P), G(P))$ be a gradient of Ψ at $P \in \mathcal{M}$. Consider an estimator P_n^* with density $p_n^* = \widehat{\lambda}^* \widehat{G}^*$ satisfying $P_n D(\widehat{\lambda}^*, \widehat{G}^*) = 0$.

- Define

$$R_{1n} \equiv P_0 \{D(\widehat{\lambda}^*, \widehat{G}^*) - D(\widehat{\lambda}^*, G_0)\} - P_0 \{D(\lambda_0^*, G_0) - D(\lambda_0^*, G_0)\}$$

Assume $R_{1n} = o_p(1/\sqrt{n})$.

- Define

$$\Phi(g_n^*) \equiv P_0 D(\lambda_0^*, \widehat{G}^*)$$

where P_0 and λ_0^* are treated as given. Assume that $\Phi(\widehat{G}^*)$ is an efficient estimator of the parameter $\Phi(G_0)$ in model $\mathcal{M}(\mathcal{G}) = \{p_{\lambda G} = \lambda G : \lambda \in \Lambda, G \in \mathcal{G}\}$, and let $T_G(P_0)$ denote the tangent space generated by model $\mathcal{G}$ for G_0 at $P_0 = \lambda_0 G_0$.

- Assume the expansion:

$$P_0 D(\widehat{\lambda}^*, G_0) = \psi_0 - \Psi(\widehat{\lambda}^*) + \frac{1}{n} \sum_{i=1}^n D_1(P_0) + o_p(1/\sqrt{n})$$

for some $D_1(P_0)$.

- Assume $D(\widehat{\lambda}^*, \widehat{G}^*)$ falls in a P_0 -Donsker class. Then, $\Psi(P_n^*) - \psi_0 = o_p(1/\sqrt{n})$.
- In addition, assume $P_0 \{D(\widehat{\lambda}^*, \widehat{G}^*) - D(\lambda_0^*, G_0)\}^2 \rightarrow 0$ in probability as $n \rightarrow \infty$ for some λ_0^* and $D(\lambda_0^*, G_0)$ in the P_0 -Donsker class.

Then,

$$\Psi(P_n^*) - \psi_0 = (P_n - P_0) IC(P_0) + o_p(1/\sqrt{n})$$

where

$$IC(P_0) \equiv \Pi(D(\lambda_0^*, G_0) + D_1(P_0) \mid T_g(P_0)^\perp)$$

Π is the projection operator in $L_0^2(P_0)$ endowed with inner product $\langle f, g \rangle_{P_0} = E_{P_0} fg$ onto the orthogonal complement of $T_g(P_0)$. If $D_1(P_0) = 0$ and $D(\lambda_0^*, G_0) = D^*(\lambda_0^*, G_0)$ where D^* is the canonical gradient, then $\Psi(P_n^*)$ is asymptotically efficient.

REFERENCES

- Akazawa K, Nakamura T, Palesch Y. Power of logrank test and cox regression model in clinical trials with heterogenous samples. *Stat. Med* 1997;16(5):583–597. [PubMed: 9089965]
- Breiman L. Random forests. *Machine Learning* 2001;V45(1):5–32.
- Cox DR. Regression models and life-tables. *J. Roy. Stat. Society Ser. B (Methodological)* 1972;34(2):187–220.
- Hastie, T.; Tibshirani, R. Generalized Additive Models. Chapman & Hall/CRC; Boca Raton, FL: 1990.
- Hernández AV, Eijkemans MJ, Steyerberg EW. Randomized controlled trials with time-to-event outcomes: How much does prespecified covariate adjustment increase power? *Ann. Epidemiol* 2006;16(1):41–48. [PubMed: 16275011]
- Jiang H, Symanowski J, Paul S, Qu Y, Zagar A, Hong S. The type I error and power of non-parametric logrank and Wilcoxon tests with adjustment for covariates—A simulation study. *Stat. Med* 2008;27(28):5850–5860. [PubMed: 18759373]

- Koch G, Tangen CM, Jung J-W, Amara IA. Issues for covariance analysis of dichotomous and ordered categorical data from randomized clinical trials and non-parametric strategies for addressing them. *Statistics in Medicine* 1998;17(15–16):1863–1892. [PubMed: 9749453]
- Lu X, Tsiatis A. Improving the efficiency of the log-rank test using auxiliary covariates. *Biometrika* 2008;95(3):679–694.
- Moore, K.; van der Laan, M. Application of time-to-event methods in the assessment of safety in clinical trials.. In: Peace, K., editor. *Design and Analysis of Clinical Trials with Time-to-Event Endpoints*. Taylor & Francis; Philadelphia: 2009a. p. 455–482.
- Moore KL, van der Laan MJ. Covariate adjustment in randomized trials with binary outcomes: Targeted maximum likelihood estimation. *Stat. Med* 2009b;28(1):39–64. [PubMed: 18985634]
- Pocock S, Assmann S, Enos L, Kasten L. Subgroup analysis, covariate adjustment and baseline comparisons in clinical trial reporting: Current practice and problems. *Stat. Med* 2002;21:2917–2930. [PubMed: 12325108]
- Ripley, BD. *Pattern Recognition and Neural Networks*. Cambridge University Press; Cambridge: 1996.
- Robins, J.; Rotnitzky, A. Recovery of information and adjustment for dependent censoring using surrogate markers.. In: Jewell, N.; Dietz, K.; Farewell, V., editors. *AIDS Epidemiology, Methodological Issues*. Birkhäuser; Boston, MA: 1992. p. 297–331.
- Sinisi S, van der Laan MJ. The deletion/substitution/addition algorithm in loss function based estimation: Applications in genomics. *Stat. Appl. Genet. Mol. Biol* 2004;3(1) Article 18.
- Tangen CM, Koch GG. Non-parametric analysis of covariance for hypothesis testing with logrank and wilcoxon scores and survival-rate estimation in a randomized clinical trial. *J. Biopharm. Stat* 1999;9(2):307–338. [PubMed: 10379696]
- Tibshirani R. Regression shrinkage and selection via the lasso. *J. Roy. Stat. Soc. Ser. B* 1996;58(1):267–288.
- Tsiatis AA, Davidian M, Zhang M, Lu X. Covariate adjustment for two-sample treatment comparisons in randomized clinical trials: A principled yet flexible approach. *Stat. Med* 2008;27(23):4658–4677. [PubMed: 17960577]
- van der Laan, M.; Gruber, S. Technical Report 246, Division of Biostatistics. University of California; Berkeley: 2009. Collaborative Double Robust Targeted Penalized Maximum Likelihood Estimation..
- van der Laan, M.; Polley, E.; Hubbard, A. Technical Report 222, Division of Biostatistics. University of California; Berkeley: 2007. Super Learner..
- van der Laan, MJ. Technical Report 232, Division of Biostatistics. University of California; Berkeley: 2008. The Construction and Analysis of Adaptive Group Sequential Designs..
- van der Laan, MJ.; Robins, JM. *Unified Methods for Censored Longitudinal Data and Causality*. Springer; New York: 2003.
- van der Laan MJ, Rubin D. Targeted maximum likelihood learning. *Int. J. Biostatist* 2006;2(1) article 11.
- Wood A, White I, Thompson S. Are missing outcome data adequately handled? A review of published randomized controlled trials in major medical journals. *Clin. Trials* 2004;1:368–376. [PubMed: 16279275]
- Zhang M, Tsiatis AAA, Davidian M. Improving efficiency of inferences in randomized clinical trials using auxiliary covariates. *Biometrics* 2008;64(3):707–715. [PubMed: 18190618]

Figure 1.

Power by increasing correlation between covariate and survival time. Relationship between the correlation between covariate and survival time (ρ_{WT}) and the gains in power between the unadjusted and S-TMLE for both strong and weak treatment (tx) effects on survival.

Figure B1.

D-TMLE results for 100 data-generating distributions with misspecified initial hazard that includes intercept only. Open points are the true values; filled-in points are the directly targeted maximum likelihood estimates.

Figure B2.

D-TMLE results for 100 data-generating distributions with misspecified initial hazard that incorrectly excludes W_2 term. Open points are the true values; filled-in points are the directly targeted maximum likelihood estimates.

Figure B3.

D-TMLE results for 100 data-generating distributions with misspecified initial hazard that incorrectly excludes W_1 term. Open points are the true values; filled-in points are the directly targeted maximum likelihood estimates.

Table 1

No censoring: Power and efficiency comparison

Method	% Bias	Power	95% Coverage	RE
Unadjusted	-2	0.39	0.96	1.00
S-TMLE COR	0	0.96	0.94	3.99
D-TMLE COR	0	0.96	0.94	4.00
S-TMLE MIS1	2	0.61	0.95	1.50
D-TMLE MIS1	-1	0.60	0.95	1.59
S-TMLE MIS2	1	0.53	0.94	1.21
D-TMLE MIS2	-2	0.51	0.95	1.29

Note. Comparison of the two targeted maximum likelihood approaches the unadjusted logrank under the no censoring setting. Correctly specified initial $\lambda(t / A, W)$ (COR), misspecified initial $\lambda(t / A, W)$ includes only a main term for treatment and W_1 (MIS1), and misspecified initial $\lambda(t / A, W)$ includes only a main term for treatment and W_2 (MIS2).

Table 2

MCAR censoring: Power and efficiency comparison

Method	% Bias	Power	95% Coverage	RE
Unadjusted	1	0.43	0.94	1.00
S-TMLE COR	1	0.94	0.94	3.84
D-TMLE COR	2	0.95	0.95	4.12
S-TMLE MIS1	2	0.58	0.95	1.46
D-TMLE MIS1	0	0.56	0.95	1.60
S-TMLE MIS2	1	0.52	0.95	1.31
D-TMLE MIS2	-2	0.49	0.95	1.40

Note. Comparison of the two targeted maximum likelihood approaches to the unadjusted logrank, under MCAR censoring. Correctly specified initial $\lambda(t/A, W)$ (COR), misspecified initial $\lambda(t/A, W)$ includes only a main term for treatment and W_1 (MIS1), and misspecified initial $\lambda(t/A, W)$ includes only a main term for treatment and W_2 (MIS2).

Table 3

MAR censoring: Power and efficiency comparison

Method	% Bias	Power	95% Coverage	RE
Unadjusted	21	0.55	0.88	1.00
S-TMLE COR	2	0.94	0.95	3.89
D-TMLE COR	2	0.94	0.95	4.58
S-TMLE MIS1	1	0.57	0.95	1.50
D-TMLE MIS1	0	0.55	0.96	1.83
S-TMLE MIS2	1	0.51	0.94	1.25
D-TMLE MIS2	-2	0.48	0.95	1.53

Note. Comparison of the two targeted maximum likelihood approaches to the unadjusted logrank, under MAR censoring. Correctly specified initial $\lambda(t/A, W)$ (COR), misspecified initial $\lambda(t/A, W)$ includes only a main term for treatment and W_1 (MIS1), and misspecified initial $\lambda(t/A, W)$ includes only a main term for treatment and W_2 (MIS2).

Table 4

No censoring: Substitution estimator based on initial hazard (i.e., hazard not updated by TMLE algorithm)

Method	% Bias	Power	95% Coverage
COR	0	0.97	0.95
MIS1	6	0.60	0.95
MIS2	21	0.59	0.93

Note. Correctly specified initial $\lambda(t / A, W)$ (COR), misspecified initial $\lambda(t / A, W)$ includes only a main term for treatment and W_1 (MIS1), and misspecified initial $\lambda(t / A, W)$ includes only a main term for treatment and W_2 (MIS2).

Table 5

Power and initial hazard estimation

	% Bias	Power	RE
Unadjusted	-2	0.25	1.00
S-TMLE _{MT}	1	0.37	1.39
S-TMLE _{SL}	2	0.48	1.54

Note. Comparison of the S-TMLE using the main term only D/S/A algorithm for the initial hazard estimation (S-TMLE_{MT}), and S-TMLE using the super-learner algorithm for the initial hazard estimation (S-TMLE_{SL}).

Table 6

Single dataset analysis

	Point estimate	SE	<i>p</i> Value
Unadjusted	0.157	0.086	0.071
S-TMLE	0.162	0.078	0.038

Note. Comparison of the S-TMLE using the super-learner algorithm for the initial hazard estimation (S-TMLE) and the unadjusted estimator.

Table D1

Type I error control: No censoring, no treatment effect

Method	Rejections (%)
Unadjusted	0.048
D-TMLE COR	0.051
D-TMLE MIS1	0.048
D-TMLE MIS2	0.052