Targeted maximum likelihood estimation of treatment effects in randomized controlled trials and drug safety analysis

by

Kelly L. Moore

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Biostatistics

in the

GRADUATE DIVISION
of the
UNIVERSITY OF CALIFORNIA, BERKELEY

Committee in charge:
Professor Mark J. van der Laan, Chair
Professor Ira B. Tager
Professor Alan E. Hubbard

Fall 2009

The dissertation of Kelly L. Moore, titled Targeted maximum likelihood estimation of treatment effects in randomized controlled trials and drug safety analysis, is approved:

Chair Date

Date

Date

University of California, Berkeley

Fall 2009

Targeted maximum likelihood estimation of treatment effects in randomized controlled trials and drug safety analysis

Abstract

Targeted maximum likelihood estimation of treatment effects in randomized controlled trials and drug safety analysis

by

Kelly L. Moore
Doctor of Philosophy in Biostatistics

University of California, Berkeley

Professor Mark J. van der Laan, Chair

In most randomized controlled trials (RCTs), investigators typically rely on estimators of causal effects that do not exploit the information in the many baseline covariates that are routinely collected in addition to treatment and the outcome. Ignoring these covariates can lead to a significant loss is estimation efficiency and thus power. Statisticians have underscored the gain in efficiency that can be achieved from covariate adjustment in RCTs with a focus on problems involving linear models. Despite recent theoretical advances, there has been a reluctance to adjust for covariates based on two primary reasons; 1) covariate-adjusted estimates based on non-linear regression models have been shown to be less precise than unadjusted methods, and, 2) concern over the opportunity to manipulate the model selection process for covariate adjustment in order to obtain favorable results. This dissertation describes statistical approaches for covariate adjustment in RCTs using targeted maximum likelihood methodology for estimation of causal effects with binary and right-censored survival outcomes.

Chapter 2 provides the targeted maximum likelihood approach to covariate adjustment in RCTs with binary outcomes, focusing on the estimation of the risk difference, relative risk and odds ratio. In such trials, investigators generally rely on the unadjusted estimate as the literature indicates that covariate-adjusted estimates based on logistic regression models are less efficient. The crucial step that has been missing when adjusting for covariates is that one must integrate/average the adjusted estimate over those covariates in order to obtain the population-level effect. Chapter 2 shows that covariate adjustment in RCTs using logistic regression models can be mapped, by averaging over the covariate(s), to obtain a fully robust and efficient estimator of the marginal effect, which equals a targeted maximum likelihood estimator. Simulation studies are provided that demonstrate that this targeted maximum likelihood method increases efficiency and power over the unadjusted method, particularly for smaller sample sizes, even when the regression model is misspecified.

Chapter 3 applies the methodology presented in Chapter 2 to a sampled RCT dataset with a binary outcome to further explore the origin of the gains in efficiency and provide a criterion for determining whether a gain in efficiency can be achieved with covariate adjustment over the unadjusted method. This chapter demonstrates through simulation studies and the data analysis that not only is the relation between R^2 and efficiency gain important, but also the presence of empirical confounding. Based on the results of these studies, a complete strategy for analyzing these type of data is formalized that provides a robust method for covariate adjustment while protecting investigators from misuse of these methods for obtaining favorable inference.

Chapters 4 and 5 focus on estimation of causal effects with right-censored survival outcomes. Time-to-event outcomes are naturally subject to right-censoring due to early patient withdrawals. In chapter 4, the targeted maximum likelihood methodology is applied to the estimation of treatment specific survival at a fixed end-point in time. In chapter 5, the same methodology is applied to provide a competitor to the logrank test. The proposed covariate adjusted estimators, under no or uninformative censoring, do not require any additional parametric modeling assumptions, and under informative censoring, are consistent under consistent estimation of the censoring mechanism or the conditional

hazard for survival. These targeted maximum likelihood estimators have two important advantages over the Kaplan-Meier and logrank approaches; 1) they exploit covariates to improve efficiency, and 2) they are consistent in the presence of informative censoring. These properties are demonstrated through simulation studies.

Chapter 6 concludes with a summary of the preceding chapters and a discussion of future research directions.

To Randy.

Contents

List of Figures

List of Tables

Acknowledgments

I would like to thank the people who helped make the doctoral program at Berkeley a wonderful experience. Their support has been instrumental in the completion of this dissertation.

I must first thank Mark van der Laan for being an incredible advisor. I cannot overstate his contribution to my personal and intellectual growth over the past four years. His generosity, encouragement, and mentorship have made my graduate studies both a challenging and rewarding experience.

Ira Tager has been an instrumental part of my doctoral education. Our collaborations have provided context and meaning for all of the interesting statistical methods I have learned. I am immensely grateful for his patience, guidance and support. I thank Alan Hubbard for serving on my qualifying and dissertation committees, and for being a great teacher with Mark in the causal inference classes I took during the Master's program.

I also thank Romain Neugebauer his generous support and insightful conversations. I will continue throughout my career to look to his conscientiousness, and excellent teaching and writing skills as examples.

I thank *Statistics in Medicine* for publishing a version of the material presented in Chapter 2, entitled "Covariate adjustment in randomized trials with binary outcomes: Targeted maximum likelihood estimation"; and Karl Peace for including a version of the material presented in Chapter 4 as a chapter in the book he edited, entitled "Design and Analysis of Clinical Trials with Time-to-Event Endpoints".

Many people helped make the everyday experience less painful, and even enjoyable. Sharon Norris has been invaluable in navigating through the administrative details often encountered by graduate students, particularly international students. I would like to thank my fellow doctoral students, including Raul Aguilar for the many laughs.

The support of my family has been invaluable in the completion of this dissertation. I thank my parents, Heather and Brian, for their endless love and support. I also thank my sister and friend Lesley for the endless hours on the phone and her beautiful family for giving a little perspective. My mother and father in law, Linda and Brian, have been like a second set of parents, and their optimism and positive attitude in difficult times are an inspiration. Finally, I thank my husband Randy for his supportive love, unwavering confidence in my abilities, and of course his never ending humor for keeping me sane.

Chapter 1

Introduction

In many randomized controlled trials (RCTs), the primary goal is to evaluate the causal effect of a treatment (e.g., new drug versus current standard of care) on an health outcome. For example, with binary outcomes, investigators are often interested in estimating the difference in the probability of an event, e.g. death, between the treated and untreated arms, referred to as the risk difference. This is typically estimated as the observed difference in the proportion of subjects with the event in the treated and untreated arms.

This estimator, which we refer to as the unadjusted estimator, is efficient when no other information is collected besides the data on the treatment, A, and the outcome, Y, of interest. Since the likelihood of the observed data is $P(A)P(Y|A)$ when the data only contain information on the treatment and outcome, the unadjusted estimator thus corresponds to the maximum likelihood estimator. From estimation theory, it is known that the nonparametric maximum likelihood estimator is the efficient estimator of the effect of interest [64].

In most RCTs, data are also collected on baseline (pre-treatment) covariates, W, in addition to the treatment and outcome of interest. In such cases, the unadjusted estimator of the risk difference is no longer equivalent to the maximum likelihood estimator since the likelihood of the observed data is now $P(W)P(A|W)P(Y|A,W)$, and the unadjusted estimator ignores the information from the covariates W. The unadjusted estimator can be viewed instead as a reduced data maximum likelihood estimator. It follows that ignoring covariate information by using the unadjusted estimator can lead to a loss in estimation efficiency (precision) in practice.

Causal effects are generally defined on the basis of counterfactual outcomes, i.e., outcomes that would have been observed on a subject had the treatment variable, possibly contrary to fact, been set at a particular level. However, in an actual study, we only observe a single counterfactual outcome for each subject, $Y = Y_A$ corresponding to the treatment that the subject actually received. For example, the outcome Y_1 is not observed for subjects that received treatment $A = 0$. Since the counterfactual outcome is missing for these subjects, we cannot directly estimate $E(Y_1)$. From the observed data, $E(Y_1 \mid A = 1)$ can be estimated, that is the mean of the counterfactual outcome for treated subjects $(a = 1)$ among those subjects who were actually observed at that treatment level $(A = 1)$. In an RCT, when treatment is randomized, $E(Y_1 \mid A = 1)$ is equal to $E(Y_1)$ since the group of subjects in the observed sample with $A = 1$ is indeed representative of the entire study population.

Now, suppose we the outcome Y is continuous. Let the parameter of interest be the marginal effect of A on Y, $\psi = E(Y_1) - E(Y_0)$. Since we know that $E(Y_1 \mid A = 1)$ is equal to $E(Y_1)$ in an RCT, investigators typically rely on the unadjusted estimate given by,

$$\hat{\psi}_1 = \hat{\mu}_1 - \hat{\mu}_0$$

where $\hat{\mu}_1 = \frac{1}{n_1}\sum_{i=1}^{n} I(A_i = 1)Y_i$ and $\hat{\mu}_0 = \frac{1}{n_0}\sum_{i=1}^{n} I(A_i = 0)Y_i$, where n_1 is the number of subjects in the treated arm, n_0 is the number of subjects in untreated arm, and $n = n_1 + n_0$.. Consider the conditional expectation of the outcome given treatment and covariates, denoted by $Q(A, W) = E(Y|A, W)$. This function can be estimated with a linear regression model such as,

$$\hat{Q}(A, W) = \hat{\beta}_0 + \hat{\beta}_1 A + \hat{\beta}_2 W.$$

In this setting, $\hat{\beta}_1$ coincides with and has been shown to be at least as precise as the unadjusted estimate $\hat{\psi}_1$ [43]. However, when $Q(A, W)$ is estimated as

$$\hat{Q}(A, W) = \hat{\beta}_0 + \hat{\beta}_1 A + \hat{\beta}_2 W + \hat{\beta}_3 AW,$$

then $\hat{\beta}_1$ no longer coincides with $\hat{\psi}_1$. In this case, to obtain the *marginal* effect, one must integrate out or average over the covariate(s) W. The G-computation estimator, introduced in Robins [46] and Robins [47] is an estimator that does indeed average over W and thus give a marginal effect,

$$\hat{\psi}_{Gcomp} = \frac{1}{n}\sum_{i=1}^{n}\left[\hat{Q}(1, W_i) - \hat{Q}(0, W_i)\right].$$

When $\hat{Q}(A, W)$ is estimated with a linear model, and it does not contain any interaction terms, then $\hat{\psi}_{Gcomp} = \hat{\beta}_1$. The G-computation estimator is not limited to a linear model for $Q(A, W)$ when estimating the treatment effect, for example, when the outcome is binary, one could use a logistic regression model to estimate $Q(A, W)$ and use the G-computation formula to obtain the estimated risk difference. With binary outcomes, the exponentiated coefficient for treatment from a logistic regression model provides an estimate of the odds ratio. The G-computation estimate in this case is given by,

$$\hat{\psi}_{OR-Gcomp} = \frac{\left[\frac{1}{n}\sum_{i=1}^{n}\hat{Q}(1, W_i)\right] / \left[1 - \frac{1}{n}\sum_{i=1}^{n}\hat{Q}(1, W_i)\right]}{\left[\frac{1}{n}\sum_{i=1}^{n}\hat{Q}(0, W_i)\right] / \left[1 - \frac{1}{n}\sum_{i=1}^{n}\hat{Q}(0, W_i)\right]}.$$

Even in the absence of interaction terms in the logistic regression model, $\hat{\psi}_{OR-Gcomp}$ is not necessarily equivalent to the estimate obtained from the logistic regression model. In fact, the logistic regression approach, which actually provides conditional effect estimates, has been shown to actually *reduce* the precision as compared to the unadjusted approach [2, 22, 53]. In these references, it has also been shown that the odds ratio point estimates (i.e., exponentiated coefficient for treatment), are typically further away from the null than the unadjusted estimates.

Because of this lack of correspondence between the linear and logistic settings with respect to consistency and precision, investigators are still resistant to adjusting for covariates and often rely on the unadjusted estimate. Furthermore, when the outcome is binary, investigators often wish to estimate not only the odds ratio, but the risk difference or relative risk.

Chapters 2 and 3 of this dissertation aim at clarifying the issue of efficiency gain with baseline covariates in RCTs with binary outcomes based on the recently developed framework of targeted maximum likelihood estimation, originally introduced in van der Laan and Rubin [65], and unify the different analytical protocols for covariate adjustment that have been proposed. We demonstrate the complete generality of the potential gain in efficiency that can be achieved from covariate adjustment in both the linear and logistic case and reconcile this result with the apparent contradictions in the literature. In short, reconciliation can be attained by noting that the typical adjusted estimator described above is not the maximum likelihood estimator of the marginal effect of interest in general but instead the maximum likelihood estimator of the conditional effect (given all covariates W). Only in the linear case is there correspondence between the estimands of the adjusted and unadjusted estimators. In other words, the adjusted estimator above does not correspond to the maximum likelihood estimator

in general and in particular not in the binary case. This explains the apparent loss in efficiency in comparison to the unadjusted estimator since one cannot in fact compare the efficiency performance of both approaches since they do not aim at evaluating the same effect in the logistic case. Rather, the maximum likelihood estimator of the marginal causal effect of A on Y is the weighted average of the conditional effect of A on Y given W, according to the distribution of W. For instance, when the marginal effect of interest is measured by the risk difference then its maximum likelihood estimator is the average of the stratum specific risk differences across all strata defined by categories of W, assuming W is discrete. Chapters 2 and 3 further explore this estimation issue and show that gains in efficiency can be achieved even when W is continuous and high dimensional. We hope that such clarifications will allow the broad application of these new techniques in RCTs.

In addition, we explore the origin of the gain in efficiency and criteria that can be used to anticipate whether the study design and the covariates collected can actually lead to increased estimation precision in practice. We illustrate how empirical confounding, i.e., an imbalance between the treated and untreated arms in the distribution of a covariate that also affects the outcome, explains the gain in efficiency that can be achieved from an adjusted analysis. Empirical confounding can occur due to bad luck in the randomization process, e.g., a higher percentage of older patients are assigned to the placebo arm and age has an effect on the outcome. Some RCT designs insure perfect covariate balance for some of the baseline covariates; however, there are typically many other covariates collected that do not have a perfect balance. Empirical confounding can introduce estimation error (sample bias) for which the unadjusted estimator cannot correct since it ignores covariate information. The method for covariate adjustment presented in this dissertation can account for such a covariate imbalance and thus improve over the poor finite sample performance of the unadjusted estimator due to empirical confounding.

In addition to the perceived discrepancy of the results for linear and non-linear settings, another issue that has obstructed the broad application of methodologies for covariate adjustment has been concerns about the selection of the parametric covariate adjustment. Incorrect covariate adjustment can indeed typically lead to estimation bias. However, we show, that in RCTs, the maximum likelihood estimator is doubly robust and describe its practical implication. Despite this double robust property, which ensures the estimator's consistency independent of the covariate adjustment, unease arises over the fact that investigators could still select the covariate adjustment that provides the most favorable inference without accounting for multiple testing. We provide a procedure, that would be outlined in an *a priori* analysis protocol, that protects investigators from guiding causal inferences based on selection of favorable covariates and their functional forms in a parametric model.

The reliance on these unadjusted estimators that ignore covariate information extend beyond RCTs with binary outcomes. In safety analysis in RCTs, patient reporting of adverse event (AE) occurrence usually occurs at many intervals throughout the study, often collected at follow-up interviews rather than only at a single fixed end-point. Similarly, efficacy studies are often constructed in such a manner. As such, time-to-event methods that exploit these data structures may provide further insight into the safety profile of the drug. In estimation of treatment specific survival at a fixed end point for right-censored survival outcomes, the standard approach is to apply the Kaplan-Meier (KM) estimator. Exploitation of covariate information can also improve the efficiency over the KM estimator. In addition, informative censoring is often present for such outcomes and in this case the KM will be biased and thus is not even reliable. Informative censoring occurs when the probability of censoring depends on the outcome the subject would have had in the absence of censoring. In particular, this arises if covariates affect both outcome and probability of being censored. Chapter 4 focuses on this common problem of estimation of treatment specific survival, using covariate adjustment to improve precision. The motivation for the methodology is based on safety analysis, and the specific issues therein, including multiple testing considerations due to the multivariate nature of AE, are discussed. Extensions to the methodology, including the adjustment for time-dependent covariates, and post-market safety analysis are discussed.

In both efficacy and safety analysis, even more common than estimation of the causal effect of treatment on survival at a fixed time, is the logrank test for testing for a treatment effect on survival, 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. Again, as in the previous examples, this unadjusted approach ignores the covariates and thus its efficiency can be improved upon. Such efficiency gains can directly translate in practice into smaller sample size requirements and even shorter trials [26].

The key principle in developing covariate adjusted estimators is to not require any additional assumptions beyond those required for the unadjusted method. This dissertation abides by this principle by proposing methods for covariate adjustment in RCTs using targeted maximum likelihood estimation [65]. This estimation procedure is a new approach to statistical learning that can be applied to any estimation problem. In short, 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 estimator (TMLE) is a familiar type of likelihood based estimator and it solves the efficient influence curve estimating equation. Due to this latter fact, it thereby inherits the properties of the solution of the efficient influence curve estimating equation, including asymptotic linearity and local efficiency [64]. These properties are of particular consequence in RCTs. Since the treatment mechanism is always known, when there is no censoring, our estimator is always consistent. The advantages of this methodology over other methodologies are provided for each of the specific problems discussed in each chapter.

In summary, this dissertation provides a new approach to covariate adjustment in RCTs through the application of targeted maximum likelihood estimation. Chapter 2 provides an approach to estimation of the risk difference, relative risk and odds ratio, in the context of RCTs with binary outcomes, both with and without censoring. Chapter 3 further elucidates the methodology through an application to a dataset, obtained by random sampling from a real RCT dataset. The origination of the efficiency gain and its relation to empirical confounding is explored. Chapter 4 applies the targeted maximum likelihood estimation approach to the estimation of treatment specific survival at a fixed end point for right-censored survival outcomes. Chapter 5 provides a covariate-adjusted competitor to the logrank test. Included in each chapter are extensive simulation studies to illustrate and explore the methodology presented. The dissertation concludes with a summary of the preceding chapters and a discussion of future research directions in chapter 6.

Chapter 2

Covariate adjustment in randomized controlled trials with binary outcomes

2.1 Introduction

Suppose we observe n independent and identically distributed observations of the random vector $O = (W, A, Y) \sim p_0$, where W is a vector of baseline covariates, A is the treatment of interest and $Y = \{0, 1\}$ is the binary outcome of interest, and p_0 denotes the density of O. Causal effects are based on a hypothetical full data structure $X = ((Y_a : a \in \mathcal{A}), W)$ containing the entire collection of counterfactual or potential outcomes Y_a for a ranging over the set of all possible treatments $\mathcal{A}$. The observed data structure O only contains a single counterfactual outcome $Y = Y_A$ corresponding to the treatment that the subject received. The observed data $O = (W, A, Y \equiv Y_A)$ is thus a missing data structure on X with missingness variable A. We denote the conditional probability distribution of treatment A by $g_0(a|X) \equiv P(A = a|X)$. The randomization assumption or coarsening at random assumption states that A is conditionally independent of the full data X given W, $g_0(A|X) = g_0(A|W)$. In a randomized controlled trial (RCT) in which treatment is assigned completely at random, we have $g_0(A|X) = g_0(A)$. For the sake of presentation, we assume the treatment A is binary and that A is completely randomized as in a typical RCT, but our methods are presented so that it is clear how our estimators generalize to observational studies or RCTs in which $g_0(A|W)$ is unknown. In the binary treatment case, $g_0(1) = p(A = 1) = \theta_0$ and $g_0(0) = p(A = 0) = 1 - \theta_0$ and n_1 the number of subjects in treatment group 1, n_0 the number of subjects in treatment group 0, and $n = n_1 + n_0$. The quantity of interest is causal effect of treatment A on Y, for example the risk difference $\psi = E(Y_1) - E(Y_0)$, where Y_1 and Y_0 are the counterfactual outcomes under treatments 1 and 0, respectively. We note that, as an alternative to the counterfactual presentation, we can write the parameter of interest as $\psi = E_W [E(Y|A = 1, W) - E(Y|A = 0, W)]$. This quantity is typically estimated in RCTs with the unadjusted estimate

$$\hat{\psi}_1 = \hat{\mu}_1 - \hat{\mu}_0$$

where $\hat{\mu}_1 = \frac{1}{n_1} \sum_{i=1}^{n} I(A_i = 1)Y_i$ and $\hat{\mu}_0 = \frac{1}{n_0} \sum_{i=1}^{n} I(A_i = 0)Y_i$. An adjusted effect is also sometimes obtained,

$$\hat{\psi}_W = \hat{P}(Y = 1|A = 1, W) - \hat{P}(Y = 1|A = 0, W).$$

Adjusting for baseline covariates and the issues involved has been discussed in Pocock et al. [43]. Although it has been recognized, at least for linear models, i.e., continuous outcomes, that adjusting for covariates increases the precision of the estimate of the marginal causal effect of treatment, investigators are still resistant to adjusting in logistic models and often rely on the unadjusted estimate. This generally appears to be due to confusion as to how to select the covariates and how to adjust for them

[43]. In addition, there is a concern that if data-adaptive procedures are used to select the model for $P(Y = 1|A, W)$ that investigators will be tempted to select the model that provides the most favorable results. However, we recommend that as long as the procedure is determined *a priori* then we can avoid this latter issue. Thus, a black box type data-adaptive procedure, e.g., forward selection, can still be applied as long as the algorithm and candidate covariates are specified *a priori*.

Adjusting for covariates with main terms in linear models, referred to as analysis of covariance (ANCOVA) in RCT literature, for the purpose of estimation of the marginal causal effect has been limited to no interaction terms with treatment. When there is such an interaction term, it is often not clear in the literature on analysis of RCT data how one uses this conditional model to obtain a marginal effect estimate. However, even in the absence of the interaction term, the increase in precision has not been observed for non-linear models such as the logistic model. In fact, it has actually been reported that the estimates are not more precise for logistic models [22, 53]. The crucial step that has been missing when the parameter of interest is the *marginal* causal effect of A on Y, is that when adjusting for covariates W, one must integrate/average the adjusted estimate over those W in order to obtain a marginal effect estimate that is comparable to the unadjusted effect estimate $\hat{\psi}_1$. This method of averaging over W has been referred to as the G-computation formula and is often applied in observational studies when the treatment or exposure has not been assigned randomly [46, 47]. We show that with this additional step of averaging over W, even when the outcome is binary, and even if the regression model is mis-specified, we obtain a more efficient estimate in the RCT setting. Such an approach allows for interactions between A and W in the model for $P(Y = 1|A, W)$ while still obtaining a marginal effect. We note that the conditional effect may be the parameter of interest in some studies, for example the effect of a drug conditional on age, and, thus, the investigator does not want to average over age.

In this chapter, we focus only on the marginal effect and using the covariates W to obtain the most efficient (precise) estimate of this marginal causal effect in a nonparametric model under the framework of targeted maximum likelihood estimation. This estimation procedure is a new approach to statistical learning introduced by van der Laan and Rubin [65]. This general targeted MLE methodology applies to any estimation problem, however, here it is applied to the estimation of the risk difference, relative risk and odds ratio, in the context of an RCT with and without censoring. This framework provides a new approach to covariate adjustment in RCTs. In a few special cases the targeted maximum likelihood estimator (TMLE) is equivalent to the double robust inverse probability of treatment weighted (DR-IPTW) based on plug-in maximum likelihood estimates of the nuisance parameters. The DR-IPTW estimator is defined as the solution to the efficient influence curve estimating equation [64, 38, 48, 50]. In Tsiatis et al. [59], the DR-IPTW estimator was applied to the estimation of the average difference in outcomes between treatment in RCTs with no censoring. This is an example where the efficient influence curve and TMLEs coincide. However, this is not generally true as demonstrated in examples provided in this chapter.

In summary, the goals of this chapter are threefold. First, we apply targeted maximum likelihood estimation as a method of covariate adjustment to the estimation of marginal effects of treatment in RCTs with binary outcomes. The second is to demonstrate the improved performance of this locally efficient estimator relative to the unadjusted method. The third goal is to compare different strategies of covariate adjustment, e.g., data-adaptive model selection algorithms, using simulation studies. The chapter is structured as follows. In section 2.2 we provide a brief overview of methods for covariate adjustment that have been proposed in literature. In section 2.4 we present the TMLEs for three marginal variable importance parameters: the risk difference, relative risk and odds ratio. We address missing data on the outcome and covariates, and estimation of the treatment mechanism. In section 2.6 we provide a formal relation between R^2 and efficiency gain. Section 2.5 provides testing and inference for the TMLE. In section 2.7 we present simulation studies that demonstrate the performance of the TMLE. Finally we conclude with a discussion in section 2.7.5.

2.2 Targeted maximum likelihood estimation

Traditional maximum likelihood estimation aims for a trade-off between bias and variance for the whole density of the observed data. Investigators however are typically not interested in the whole density of the data O, but rather a specific parameter of it. Targeted maximum likelihood estimation was purposefully named in that it carries out a bias reduction specifically *tailored* for the parameter of interest. For technical details about this general estimation approach we refer the reader to its seminal article [65].

Consider a model $\mathcal{M}$ which is a collection of possible probability distributions of the data, where the true distribution of the data is p_0. Consider an initial estimator $\hat{p}$. We are interested in a particular feature of the data, $\psi_0 = \psi(p_0)$. The goals of targeted maximum likelihood estimation are twofold. First, it aims to find a density $\hat{p}^* \in \mathcal{M}$ that solves the efficient influence curve for estimating equation for the parameter of interest resulting in a bias reduction as compared to the maximum likelihood estimate $\psi(\hat{p})$. Second, the algorithm requires that $\hat{p}^*$ also achieves a small increase in the log-likelihood relative to $\hat{p}$. The algorithm achieves these goals by identifying a "stretching" of the initial $\hat{p}$ so that yields a maximal change in ψ. This is done by constructing a path denoted by $\hat{p}(\epsilon)$ through $\hat{p}$ where ϵ is a free parameter. The score of this path at $\epsilon = 0$ equals the efficient influence curve. The optimal amount of "stretch" is obtained by maximizing the likelihood of the data over ϵ. Applying this optimal stretch to $\hat{p}$ yields $\hat{p}^1$ which is the first step of the targeted maximum likelihood algorithm. This process is iterated until the "stretch" is essentially zero. The final step of the algorithm gives the targeted maximum likelihood estimate $\hat{p}^*$ which solves the efficient influence curve estimating equation thereby achieving the desired bias reduction with a small increase in the likelihood. The resulting substitution estimator $\psi(\hat{p}^*)$ 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 [64]. Thus, targeted maximum likelihood estimation provides a fusion between likelihood and estimating function based methodologies.

However, targeted maximum likelihood has various important advantages relative to estimating equation methodology. Firstly, by just solving the efficient influence curve equation in p itself, it does not rely on the assumption that the efficient influence curve can be represented as an estimating function and/or the particular representation of this estimating function. Secondly, estimating equations provide no criterion to select among multiple solutions in the parameter of interest for a given estimate of the nuisance parameters in the estimating equation, while targeted maximum likelihood can simply use the likelihood criterion to select among various targeted maximum likelihood estimates indexed by different initial density estimators. Thirdly, in estimating equation methodology the parameter estimator is typically not compatible with the nuisance parameter estimates, while in the targeted maximum likelihood procedure, the estimator of the parameter of interest and the nuisance parameters in the efficient influence curve are all compatible with a single density estimator.

The targeted maximum likelihood estimators of the parameters studied in this chapter are double robust (DR) under uninformative censoring (missing at random) in the sense that they rely on either a consistent estimator of the treatment mechanism g or a consistent estimator of $Q(A,W) = E(Y|A,W)$. When the treatment is assigned completely at random, the treatment mechanism, $P(A|W) = P(A)$, is always known and thus the targeted maximum likelihood estimator is always consistent whatever the estimator for Q on which it relies. That is, even when the estimator $\hat{Q}(A,W)$ of $Q(A,W)$ is inconsistent (e.g., if it relies on a mis-specified model), the TMLE remains consistent and one should hence not be concerned with estimation bias with this method in RCTs. More specifically, if $\hat{Q}(A,W)$ converges to $Q^*(A,W) \neq Q(A,W)$ then targeted maximum likelihood estimators remain asymptotically linear and consistent in RCTs. In practice, this means that the investigator is protected even when the *a priori* specified model selection algorithm selects a mis-specified model for $Q(A,W)$. Note that if

$\hat{Q}(A,W)$ is a consistent estimator of $Q(A,W)$, then the TMLE is consistent but also efficient. In the special case that we use the true $P(A|W) = P(A)$ or a marginal estimate $\hat{\theta}$ in the targeting step of the algorithm, convergence is achieved in zero steps. Thus, in this case the TMLE coincides with the standard G-computation maximum likelihood estimator, thereby demonstrating that this latter estimator is already locally efficient. In appendix section A.3, in one of our settings, we provide a relation between the DR-IPTW, TMLE and G-computation estimator and the circumstances in which they coincide.

When censoring depends on baseline covariates, consistency of the TMLE relies on consistent estimation of the censoring mechanism or $Q(A,W)$. Even in this setting, for many causal parameters such as the causal risk difference, the targeted maximum likelihood algorithm converges in a single step.

The TMLE is a very practically attractive estimator since it can be achieved by simply adding a covariate to an initial estimate of the regression $Q(A,W)$. The corresponding coefficient ϵ for this new covariate can be estimated with standard software and thus has a straightforward implementation.

It was shown in Scharfstein et al. [54] (p. $1140 - 1141$) that to obtain a DR estimate of the difference in two mean outcomes, one can extend a parametric model for $Q(A,W)$ by adding the 2-dimensional covariate $(\frac{I(A=1)}{g(1|W)}, \frac{I(A=0)}{g(0|W)})$ where in the RCT setting, $g(1|W) = \theta$ and $g(0|W) = 1 - \theta$, and estimate the combined parameter by solving the maximum likelihood estimating equation. In section 2.4.1 we show that for this same additive effect the targeted maximum likelihood algorithm that targets both parameters $(P(Y_0 = 1), P(Y_1 = 1))$ also adds these two covariates, the first for $P(Y_1 = 1)$ and one for $P(Y_0 = 1)$, so that any function of these two parameters is estimated in a targeted manner. This TMLE still differs from the proposal in Scharfstein et al. [54] by fixing the initial regression, which can thus also represent a data adaptive machine learning fit, and simply estimating the coefficients for the additional covariates. The proposed estimator of [54] does not fix the initial regression but fits all coefficients for the parametric regression and the additional covariates simultaneously. This distinction in fixing the initial regression is important in that it allows one to apply data adaptive algorithms for the initial estimate and simply update the estimate with the targeting step. This is in contrast to the procedure proposed in Bang and Robins [3] and Scharfstein et al. [54] which appears to rely on a parametric estimate for the regression. In Bang and Robins [3] it is stated that when the initial model for $Q(A,W)$ is correct, one can obtain a more efficient DR estimate by adding the 1-dimensional (rather than 2-dimensional) covariate $\frac{I(A=1)}{g(1|W)} - \frac{I(A=0)}{g(0|W)}$. The covariate is equivalent to the TMLE covariate $\frac{I(A=1)}{g(1|W)} - \frac{I(A=0)}{g(0|W)}$, targeting the risk difference effect $P(Y_1 = 1) - P(Y_0 = 1)$. This covariate satisfies the condition of the targeting fluctuation that the score of the initial density $\hat{p}^0$ at $\epsilon = 0$ must include the efficient influence curve at $\hat{p}^0$. Again, the targeted maximum likelihood procedure fixes the initial regression and then estimates the coefficient for the additional covariate as opposed to the proposal in Bang and Robins [3] where all coefficients for the parametric regression and the additional covariate are fit simultaneously. We note that the covariate that is added in the targeted maximum likelihood algorithm is specific to the parameter one is estimating and thus differs when the parameter of interest is the relative risk or odds ratio as shown in section 2.4.

2.3 Current methods for obtaining covariate adjusted estimates

Suppose we observe $O = (W, A, Y)$ as above except the outcome Y is now continuous. Let the parameter of interest be the marginal effect of A on Y, $\psi = E(Y_1) - E(Y_0)$. For a continuous outcome Y,

$Q(A,W) = E(Y|A,W)$ is typically obtained using a linear regression model such as,

$$\hat{Q}(A,W) = \hat{\beta}_0 + \hat{\beta}_1 A + \hat{\beta}_2 W.$$

In this setting, $\hat{\beta}_1$ coincides with and has been shown to be at least as precise as the unadjusted estimate $\hat{\psi}_1$. In particular, the increase in precision occurs when the correlation between the covariate(s) and outcome is strong [2]. However, when $Q(A,W)$ is estimated as

$$\hat{Q}(A,W) = \hat{\beta}_0 + \hat{\beta}_1 A + \hat{\beta}_2 W + \hat{\beta}_3 AW,$$

then $\hat{\beta}_1$ no longer coincides with $\hat{\psi}_1$. In this case, to obtain the *marginal* effect, one must integrate out or average over the covariate(s) W. The G-computation estimator introduced in Robins [46] and Robins [47] is an estimator that does indeed average over W and thus give a marginal effect,

$$\hat{\psi}_{Gcomp} = \frac{1}{n}\sum_{i=1}^{n}\left[\hat{Q}(1,W_i) - \hat{Q}(0,W_i)\right].$$

When $\hat{Q}(A,W)$ is estimated with a linear model, and it does not contain any interaction terms, then $\hat{\psi}_{Gcomp} = \hat{\beta}_1$. The G-computation estimator is not limited to a linear model for $Q(A,W)$ when estimating the treatment effect, for example, when the outcome is binary, one could use a logistic regression model to estimate $Q(A,W)$ and use the G-computation formula to obtain the estimated risk difference. However, even in the absence of interaction terms, $\hat{\psi}_{Gcomp}$ is not necessarily equivalent to the estimate obtained from the logistic regression model.

In Tsiatis et al. [59], the DR estimator is applied to estimate the marginal effect where the authors recommend estimating two regression models separately: $Q_1(1,W) = E(Y|A=1,W)$ is obtained using only the subpopulation of individuals for whom $A=1$ and $Q_2(0,W) = E(Y|A=0,W)$ is obtained using only the subpopulation of individuals for whom $A=0$. This was proposed so that two different analysts could independently select these models to prevent the analysts from selecting the model providing the most favorable results. Another possibility is to select one model $Q(A,W) = E(Y|A,W)$ using the whole sample pooled together. When the procedure for selecting $Q(A,W)$ is specified a priori this additional step of estimating $Q_1(1,W)$ and $Q_2(0,W)$ is not necessary.

The method for estimating the marginal difference $E(Y_0) - E(Y_1)$ is provided in Tsiatis et al. [59]. However, when the outcome is binary, investigators are often also interested in not only the risk difference $E(Y_0) - E(Y_1) = P(Y_1 = 1) - P(Y_0 = 1)$, but the relative risk and odds ratios. In Zhang et al. [73], the approach in Tsiatis et al. [59] is expanded upon by applying the estimating function approach to the estimation of general parameters in RCTs. The corresponding covariate adjusted estimators are shown to provide an increase in precision over the unadjusted method. This general approach for constructing locally efficient double robust estimators that are guaranteed to improve on the unadjusted estimator can be found in van der Laan and Robins [64]. The approach provided in this chapter does not deviate from the line of research in Tsiatis et al. [59] and Zhang et al. [73], but instead applies the relatively new targeted maximum likelihood methodology to RCTs. A particular advantage of targeted maximum likelihood estimation over the estimating function approach is that for various effect parameters the efficient influence curve cannot even be represented as an estimating function in a parameter of interest and nuisance parameters, while the TMLE does not require such a representation. This is illustrated with the causal relative risk in a nonparametric model in section 2.4.

Covariate adjustment in logistic regression models for binary outcomes in RCTs has been studied in literature. These models provide conditional effect estimates and have been shown to actually *reduce* the precision as compared to unadjusted methods. In Robinson and Jewell [53], it was observed that adjusting for covariates in logistic regression models leads to an increase in power due to the fact that estimates of the treatment effect in the conditional logistic models are further away from the null even

though standard errors are larger for the adjusted effects. Hernández et al. [22] also demonstrated this fact using simulation studies and observed that the increase in power was related to the correlation between the covariate and the outcome. The simulations included only a single covariate and no interactions between the covariate and treatment. Similar results were indicated in Assmann et al. [2] with logistic regression models in that odds ratios were generally further away from the null but the standard errors were larger than the unadjusted estimates. It appears that in general, when adjusting for covariates in a logistic regression model, the standard error provided by the software, i.e., standard maximum likelihood procedures, is the standard error used by the investigator although it is often not explicitly stated [5, 15, 60, 44]. When adjusting for covariates in RCTs using logistic regression, often the investigator is interested in a conditional effect identified by continuous covariates in which case this may be an appropriate approach. We focus on the targeted maximum likelihood method for covariate adjustment that provides inference for the marginal (unconditional) effect. However, note that this method can be applied to different subgroups defined by categorical or discrete valued covariates by simple stratification.

2.4 Targeted maximum likelihood estimation of the risk difference, relative risk and odds ratio

In this section we present the targeted maximum likelihood method for adjusting for covariates in estimating the marginal effect of a binary treatment on a binary outcome with the following three parameters: risk difference, relative risk and odds ratio. We provide an overview of the derivation of the covariate that is added to an initial regression estimate. The covariate is derived in such a way that the update of the regression targets the specific parameter one is estimating and thus differs for each of the three we focus on in this chapter. For technical details we refer the reader to appendix sections A.1 through A.7.

Let $O = (W, A, Y) \sim p_0$ and $\mathcal{M}$ be the class of all densities of O with respect to an appropriate dominating measure: so $\mathcal{M}$ is nonparametric up to possible smoothness conditions. Let the parameter of interest be represented by $\Psi(p_0)$. The first step of the algorithm involves finding an initial density estimator $\hat{p}^0$ of the density p_0 of O, identified by $\hat{Q}^0(A, W)$, marginal distribution of A identified by $\hat{\theta} = \frac{1}{n}\sum_{i=1}^{n} A_i$, the marginal distribution of W being the empirical probability distribution of $W_1, ..., W_n$, and A being independent of W.

An initial fit $\hat{Q}^0(A, W)$ may be obtained in a number of ways. For example, we may fit a data-adaptive logistic regression model for the outcome Y fixing treatment A in the model and including covariates W as candidates. Since Y is binary, the density is given by,

$$\hat{p}^0(Y|A, W) = \left[\hat{Q}^0(A, W)\right]^Y \left[1 - \hat{Q}^0(A, W)\right]^{1-Y}.$$

We could choose a logistic regression model for $\hat{Q}^0(A, W)$,

$$\hat{Q}^0(A, W) = \frac{1}{1 + \exp -\hat{m}^0(A, W)}$$

for some function $\hat{m}^0$.

The targeted maximum likelihood estimation procedure updates the initial density by creating a parametric submodel through $\hat{p}^0$ indexed by parameter ϵ,

$$\hat{p}^0(\epsilon)(Y|A, W) = \left[\hat{Q}^0(\epsilon)(A, W)\right]^Y \left[1 - \hat{Q}^0(\epsilon)(A, W)\right]^{1-Y}.$$

In the case that the initial choice $\hat{Q}^0(A,W)$ is given by a logistic regression fit, then $\hat{Q}^0(\epsilon)(A,W)$ is given by the logistic regression model,

$$\hat{Q}^0(\epsilon)(A,W) = \frac{1}{1+\exp-\left[\hat{m}^0(A,W)+\epsilon h(A,W)\right]}.$$

The targeted maximum likelihood algorithm finds this covariate $h(A,W)$ by requiring that the score of $\hat{p}^0$ at $\epsilon = 0$ is equal to the efficient influence curve at $\hat{p}^0$. In van der Laan and Rubin [65], it was shown that the efficient influence curve $D(p_0)$ can be decomposed into three components corresponding with scores for $p(Y|A,W)$, $g_0(A|W)$ and the marginal probability distribution $p(W)$ of W which we refer to as $D_1(p_0)$, $D_2(p_0)$ and $D_3(p_0)$ respectively (see appendix A.1). Since the score for $\hat{p}^0$ at $\epsilon = 0$ corresponds with a zero score for $\hat{g}^0$, and the empirical distribution of W is a nonparametric maximum likelihood estimator, we only need to choose $h(A,W)$ so that the score of $\hat{p}^0(Y|A,W)$ at $\epsilon = 0$ includes the efficient influence curve component for $p(Y|A,W)$, i.e., $D_1(p_0)$. The next step of the algorithm involves estimating ϵ with maximum likelihood. The initial $\hat{Q}^0(A,W)$ is thus updated to obtain $\hat{Q}^1(A,W)$ and the algorithm is iterated by replacing $\hat{Q}^0(A,W)$ with $\hat{Q}^1(A,W)$.

It is a fortunate result that in RCTs, the covariate $h(A,W)$ is none other than a linear combination of A and an intercept only. It follows that if $\hat{m}^0(A,W)$ includes the main term A and the intercept, then $\hat{\epsilon} = 0$, and the TMLE for $Q_0(A,W)$ is given by $\hat{Q}^0(A,W)$ itself. Specifically, consider the following risk difference (RD), relative risk (RR) and odds ratio (OR) parameters,

$$P_0 \to \Psi_{RD}(p_0) = E_{p_0}\left[E(Y|A=1,W) - E(Y|A=0,W)\right], \tag{2.1}$$

$$P_0 \to \Psi_{RR}(p_0) = \frac{E_{p_0}\left[E(Y|A=1,W)\right]}{E_{p_0}\left[E(Y|A=0,W)\right])} = \frac{\mu_1}{\mu_0}, \tag{2.2}$$

and

$$P_0 \to \Psi(p_0) = \frac{E_{p_0}\left[E(Y|A=1,W)\right]/\left[1-E_{p_0}\left\{E(Y|A=1,W)\right\}\right]}{E_{p_0}\left[E(Y|A=0,W)\right]/\left[1-E_{p_0}\left\{E(Y|A=0,W)\right\}\right]} = \frac{\mu_1/(1-\mu_1)}{\mu_0/(1-\mu_0)}. \tag{2.3}$$

Now consider the initial logistic regression fit $\hat{Q}^0(A,W)$. It is straightforward to demonstrate (see appendix A.2, A.4 and A.6) that the corresponding covariates that update the initial fit for each of the above parameters are given by,

$$h_{RD}(A,W) = \frac{I(A=1)}{\hat{\theta}} - \frac{I(A=0)}{(1-\hat{\theta})},$$

$$h_{RR}(A,W) = \frac{1}{\mu_1}\frac{I(A=1)}{\hat{\theta}} - \frac{1}{\mu_0}\frac{I(A=0)}{(1-\hat{\theta})},$$

and

$$h_{OR}(A,W) = \left(\frac{1}{\mu_1}+\frac{1}{1-\mu_1}\right)\frac{I(A=1)}{\hat{\theta}} - \left(\frac{1}{\mu_0}+\frac{1}{1-\mu_0}\right)\frac{I(A=0)}{(1-\hat{\theta})}$$

Each of these covariates is simply a linear combination of A and an intercept. Thus if $\hat{m}^0(A,W)$ includes the main term A and the intercept, then $\hat{\epsilon} = 0$, and the TMLE for $Q_0(A,W)$ is given by $\hat{Q}^0(A,W)$ itself. In other words, the TMLE for ψ_{RD}, ψ_{RR} and ψ_{OR} are given by the standard maximum likelihood estimators,

$$\hat{\psi}_{RD-TMLE} = \frac{1}{n}\sum_{i=1}^{n}\left[\hat{Q}^0(1,W_i) - \hat{Q}^0(0,W_i)\right],$$

$$\hat{\psi}_{RR-TMLE} = \frac{\frac{1}{n}\sum_{i=1}^{n}\hat{Q}^0(1,W_i)}{\frac{1}{n}\sum_{i=1}^{n}\hat{Q}^0(0,W_i)},$$

and

$$\hat{\psi}_{OR-TMLE} = \frac{\left[\frac{1}{n}\sum_{i=1}^{n}\hat{Q}^0(1,W_i)\right] / \left[1-\frac{1}{n}\sum_{i=1}^{n}\hat{Q}^0(1,W_i)\right]}{\left[\frac{1}{n}\sum_{i=1}^{n}\hat{Q}^0(0,W_i)\right] / \left[1-\frac{1}{n}\sum_{i=1}^{n}\hat{Q}^0(0,W_i)\right]}.$$

It is interesting to note that in estimating the risk difference, in addition to the equivalence between the TMLE and G-computation (MLE), the TMLE solves the efficient influence curve estimating equation by definition and thus the DR-IPTW, MLE and TMLE all reduce to the same estimator in this general setting, for details see appendix A.3.

As an alternative to using a logistic fit for the initial $Q^0(A,W)$, we can instead choose a relative risk regression fit,

$$\log(\hat{Q}^0)(\epsilon)(A,W) = \hat{m}^0(A,W) + \epsilon h(A,W).$$

In estimation of the relative risk parameter, the corresponding covariate added to the initial regression model to obtain the TMLE is given by (see appendix A.5 for details),

$$h(A,W) = \left\{\frac{1}{\mu_1}\frac{I(A=1)}{\hat{\theta}} - \frac{1}{\mu_0}\frac{I(A=0)}{1-\hat{\theta}}\right\}\left[1-\hat{Q}^0(A,W)\right].$$

The maximum likelihood estimate,

$$\hat{\epsilon} = \arg\max_{\epsilon}\sum_{i=1}^{n}\log\hat{Q}^0(\epsilon)(A_i,W_i),$$

can be estimated in practice by fitting a relative risk regression in $\hat{m}^0(A,W)$ and $h(A,W)$, fixing the coefficient in front of $\hat{m}^0(A,W)$ to 1 and the intercept to 0. The resulting coefficient for $h(A,W)$ is $\hat{\epsilon}$. In this case, the covariate is no longer simply a function of A and thus $\hat{\epsilon}$ does not necessarily equal 0 and the convergence is no longer achieved in zero steps but rather iteratively. Now $\hat{Q}^k(A,W)$ is updated as,

$$\log\left[\hat{Q}^{k+1}(A,W)\right] = \hat{m}^k(A,W) + \hat{\epsilon}h^k(A,W),$$

setting $k = k+1$ and one iterates this updating step. One may also derive the updating covariate for targeting estimation of the risk difference or odds ratio as well using this initial regression in addition to the logistic initial choice.

2.4.1 Targeted maximum likelihood estimation of the two treatment specific means, and thereby for all parameters

Consider the odds ratio, as an example. An alternative for targeting the odds ratio is to simultaneously target both μ_1 and μ_0 and simply evaluate the odds ratio from the TMLEs of μ_1 and μ_0. This is a straightforward approach where two covariate extensions are added to the logistic fit $\hat{Q}^0$,

$$h_1(A,W) = \epsilon_1\frac{I(A=1)}{\hat{\theta}},$$

and,

$$h_2(A,W) = \epsilon_2\frac{I(A=0)}{(1-\hat{\theta})}.$$

Again, if the initial logistic regression fit already includes an intercept and main term A, then $\hat{\epsilon} = (\hat{\epsilon_1},\hat{\epsilon_2}) = 0$ so that this TMLE $\hat{Q} = \hat{Q}^0(\hat{\epsilon}) = \hat{Q}^0$ is not updated. This TMLE can now be used to map into a locally efficient estimator of any parameter of μ_0 and μ_1 such as the risk difference $\mu_1-\mu_0$, the relative risk μ_1/μ_0 and the odds ratio $\mu_1(1-\mu_0)/[(1-\mu_1)\mu_0]$.

2.4.2 Estimating the treatment mechanism as well

Even when the treatment mechanism (the way treatment was assigned) is known as it is in an RCT, it has been shown that efficiency is increased when estimating it from the data if $Q(A,W)$ is not correctly specified [64]. Estimating the treatment mechanism does not add any benefit to the G-computation estimator since it does not use this information. The TMLE can however leverage this information to obtain a more precise estimate of the treatment effect. This can be particularly beneficial when the model for $Q(A,W)$ is mis-specified. The TMLE is still consistent when $Q(A,W)$ is mis-specified, however, we can gain efficiency when estimating the treatment mechanism in such a case. The treatment mechanism can be estimated from the data using a logistic regression model, for example, $\hat{g}^0(1|W) = \frac{1}{1+\exp[-(\alpha_1 W_1+\alpha_2 W_2)]}$, but one can also augment an initial fit $\hat{g}^0$ with a targeted direction aiming for a maximal gain in efficiency [65]. We present the targeted maximum likelihood estimation procedure for the risk difference, however, this can be immediately extended to the relative risk and odds ratio as well.

The covariate that is added to the logistic regression $\hat{Q}^0(A,W)$ is given by,

$$h(A,W) = \frac{I(A=1)}{\hat{g}^0(1|W)} - \frac{I(A=0)}{\hat{g}^0(0|W)},$$

where, $\hat{\epsilon} = \arg\max_\epsilon \sum_{i=1}^n \log \hat{Q}^0(\epsilon)(A_i,W_i)$ can be estimated in practice by fitting a logistic regression in $\hat{m}^0(A,W)$ and $h(A,W)$, fixing the coefficient in front of $\hat{m}^0(A,W)$ to 1 and the intercept to 0. The resulting coefficient $\hat{\epsilon}$ for $h(A,W)$ is no longer necessarily (and not typically) equal to 0. Let the TMLE for $Q_0(A,W)$ be given by $\hat{Q}^*(A,W) = \hat{Q}^0(\hat{\epsilon})(A,W)$. The TMLE for ψ_0 is then,

$$\hat{\psi}_{RD-TMLE2} = \frac{1}{n}\sum_{i=1}^n \left[\hat{Q}^*(1,W_i) - \hat{Q}^*(0,W_i)\right].$$

Note that $\hat{Q}^0(A,W)$ is now updated, contrary to the case when we were not estimating the treatment mechanism as in previous subsections.

2.4.3 Missing data

Here we provide the TMLE for the case that the outcome Y is subject to missingness that can be informed by the baseline covariates W. In such a case the missingness cannot be ignored as it can lead to biased estimates since treatment groups are no longer balanced with respect to the covariates. Let Δ represent the indicator whether or not the outcome was observed. The observed data can be represented as $O = (W, A, \Delta, \Delta Y) \sim p_0$ and the full data is given by $X = ((Y_a : a \in \mathcal{A}), W)$. We assume that the conditional distribution of the joint censoring variable (A,Δ) given X satisfies coarsening at random (CAR), i.e., $g_0(A,\Delta|X) = g_0(A,\Delta|W)$. Let

$$P_0 \to \Psi(p_0) = E_{p_0}\left[E(Y|A=1,W) - E(Y|A=0,W)\right]$$

be the parameter of interest. We wish to estimate the risk difference with targeted maximum likelihood. In choosing an initial logistic regression fit $\hat{Q}^0(A,W)$, it can be shown (see appendix A.7) that the updating covariate is given by,

$$h(A,\Delta=1,W) = \frac{I(A=1)}{\hat{g}(1,1|W)} - \frac{I(A=0)}{\hat{g}(0,1|W)}.$$

The estimate of ϵ given by $\hat{\epsilon} = \arg\max_\epsilon \sum_{i=1}^n I(\Delta_i=1)\log \hat{Q}^0(\epsilon)(A_i,W_i)$. Now the logistic regression fit $\hat{Q}^0(Y|A,\Delta=1,W)$ can be updated by adding as covariate $h(A,\Delta=1,W)$ to obtain the TMLE,

$\hat{Q}^*(Y|A, \Delta = 1, W)$ for $Q_0(A, \Delta = 1, W)$ based on all observations with $\Delta_i = 1$. The estimate for $P(\Delta = 1|A = 0, W)$ as required to calculate the extra covariate $h(A, W)$ can be obtained by using a logistic regression model selected either data-adaptively or using a fixed pre-specified model for Δ conditional on $W, A = 0$. The TMLE for ψ_0 is given by,

$$\hat{\psi}_{RD-TMLE} = \frac{1}{n}\sum_{i=1}^{n}\left[\hat{Q}^*(1, 1, W_i) - \hat{Q}^*(0, 1, W_i)\right].$$

We note that the TMLE for missing covariate values is derived in exactly the same manner.

2.5 Testing and inference

Let $\hat{p}^*$ represent the TMLE of p_0. One can construct a Wald-type 0.95-confidence interval based on the estimate of the efficient influence curve, $\hat{IC}(O) = D(\hat{p}^*)$. The influence curves for the risk difference, relative risk, and odds ratio are provided in appendix A.2, A.4 and A.6 respectively. An estimate of the asymptotic variance of $\sqrt{n}(\hat{\psi} - \psi_0)$ can be estimated with,

$$\hat{\sigma}^2 = \frac{1}{n}\sum_{i=1}^{n}\hat{IC}^2(O_i).$$

The corresponding asymptotically conservative Wald-type 0.95-confidence interval is defined as $\hat{\psi} \pm 1.96\frac{\hat{\sigma}}{\sqrt{n}}$. The null hypothesis $H_0 : \psi_0 = 0$ can be tested with the test statistic,

$$\hat{T} = \frac{\hat{\psi}}{\frac{\hat{\sigma}}{\sqrt{n}}},$$

whose asymptotic distribution is $N(0, 1)$ under the null hypothesis. To establish the asymptotics of the TMLE, we apply Theorem 2.4 as provided in van der Laan and Robins [64] and also noted in van der Laan and Rubin [65]. If $\hat{Q}^0(A, W)$ converges to a mis-specified $\hat{Q}^*(A, W)$, then the TMLE $\hat{\psi}$ is aymptotically linear and consistent. Furthermore, if $\hat{Q}^0(A, W)$ is consistent, then TMLE ψ_n is aymptotically linear, consistent and efficient. For further details, see [65].

We note that if the true treatment mechanism is used, inference based on the corresponding influence curve is not conservative whereas with an estimated treatment mechanism the corresponding influence curve is conservative. One can improve the variance estimator by either applying the bootstrap procedure or by deriving the true analytical form of the influence curve. The latter can be achieved by projecting the influence curve (based on the estimated treatment mechanism) on the tangent space of the model for the treatment mechanism.

First order efficiency improves by using a larger model for the regression $Q(A, W)$ which can be achieved by applying machine learning algorithms (e.g., D/S/A). However, the asymptotic results on which the IC-based inference rely can break down in that the second order terms generated by data adaptive approaches can be an issue. Fortunately in the case that the treatment mechanism is known we have found that both the estimator and inference are not affected by the second order terms suggesting that fairly data adaptive regression approaches can be used. Based on our experience model selection algorithms based on cross-validation (e.g., D/S/A) generally result in similar estimates of the standard error using the influence curve or bootstrap procedure which does take into account model selection. Thus, if such a methodology is used in the model selection algorithm or if the model for $Q(A, W)$ is specified *a priori*, then one can rely on the IC-based standard error estimates.

2.6 Relation between R^2 and efficiency gain with targeted maximum likelihood estimation

An analytical relationship exists between the relative efficiency (RE) of the targeted maximum likelihood and unadjusted estimators and the predictive power of the baseline covariates W as expressed in the following formula,

$$\begin{aligned} RE &= \frac{\sigma^2(\text{TMLE}(Q(A)))}{\sigma^2(\text{TMLE}(Q(A,W)))} \\ &= 1 - \frac{\sigma^2(\text{TMLE}(Q(A)))}{\sigma^2(\text{TMLE}(Q(A))) - 4[E(Y-EY)^2 - E(Y-Q(W))^2]}, \end{aligned}$$

where $Q(A) = E(Y \mid A)$, $Q(W) = E(Y|W)$, $\sigma^2(TMLE(Q(A,W)))$ is the variance of the influence curve at $Q(A,W)$ and $\sigma^2(TMLE(Q(A)))$ is the variance of the influence curve at $Q(A)$ (i.e., the variance associated with the unadjusted estimator). The RE can also be expressed with respect to $R^2_{Q(W)} = 1 - \frac{E(Y-Q(W))^2}{E(Y-E(Y))^2}$ as,

$$RE = 1 - \frac{\sigma^2(\text{TMLE}(Q(A)))}{\sigma^2(\text{TMLE}(Q(A))) - 4E(Y-E(Y))^2 R^2_{Q(W)}}. \tag{2.4}$$

Thus, as the $R^2_{Q(W)}$ increases, so does the gain in efficiency, i.e., the ratio of the variances of the influence curves increases. This formula clearly shows that whenever $R^2_{Q(W)} > 0$, i.e., when outcome prediction with W through the model $Q(W)$ outperforms outcome prediction through the simple intercept model ($E(Y)$) then one achieves a gain efficiency by adjusting for the covariates W with the targeted maximum likelihood estimation approach relative to the unadjusted estimation approach. Note that this formula is based on the assumption that $P(A=1) = P(A=0) = 0.5$ and thus is consistent with the result in [16], which states that "adjustment is either neutral or helps" when $P(A=1) = P(A=0) = 0.5$. This is also in agreement with published work that demonstrated an increase in estimation precision with the maximum likelihood estimator when the correlation between the covariate(s) and the outcome (e.g., as measured by $R^2_{Q(W)}$) is strong [43, 2].

2.7 Simulation studies

2.7.1 Simulation 1: Strong and weak treatment effects

In this simulation, the treatment A and outcome Y are binary and W is a 2-dimensional covariate, $W = (W_1, W_2)$. The simulated data were generated according to the following laws:

1. $W_1 \sim N(2,2)$
2. $W_2 \sim U(3,8)$
3. $P(A=1) = \theta_0 = 0.5$
4. $Q_0(A,W) = P(Y=1|A,W) = \frac{1}{1+\exp\left[-(kA-5W_1^2+2W_2)\right]}$

We simulated the data for 2 scenarios based on the value for k in $P(Y=1|A,W)$. In the first scenario, $k = 1.2$ and there is a small treatment effect and in the second $k = 20$, and there is a larger treatment

effect. The risk difference, relative risk and odds ratio were estimated. The true values were given by $P(Y_1 = 1) = 0.372$, $P(Y_0 = 1) = 0.352$ and $(RD, RR, OR) = (0.019, 1.055, 1.087)$ for $k = 1.2$, $P(Y_1 = 1) = 0.583$, $P(Y_0 = 1) = 0.352$ and $(RD, RR, OR) = (0.231, 1.654, 2.570)$ for $k = 20$. The parameters were estimated using 4 methods. The first method "Unadjusted" is the unadjusted method of regressing Y on A using a logistic regression model. The second method "Correct" is the targeted maximum likelihood method which is equivalent to the standard G-computation (maximum likelihood) estimator with,

$$\hat{Q}(A, W) = 1/\left\{1 + \exp\left[-(\hat{\alpha}_0 + \hat{\alpha}_1 A + \hat{\alpha}_2 W_1^2 + \hat{\alpha}_3 W_2)\right]\right\}.$$

The third method "Mis-spec" used a mis-specified fit given by,

$$\hat{Q}(A, W) = 1/\left\{1 + \exp\left[-(\hat{\alpha}_0 + \hat{\alpha}_1 A + \hat{\alpha}_2 W_1)\right]\right\}.$$

For the fourth method ,"DSA", the estimate $\hat{Q}(A, W)$ was obtained using the Deletion/ Substitution/Addition algorithm (D/S/A). The D/S/A algorithm is a data-adaptive model selection procedure based on cross-validation that relies on deletion, substitution, and addition moves to search through a large space of possible functional forms, and is publicly available at `http://www.stat.berkeley.edu/~laan/Software/` [56]. The variable A was forced into the model and the D/S/A then selected from the remaining covariates. The maximum power for any term in the D/S/A algorithm was set to 2, meaning square terms and 2-way interactions were allowed. Standard errors for the TMLE were estimated using the estimated influence curve. For the odds ratio simulations, the estimator obtained by extracting the coefficient for A and the corresponding standard error from the logistic regression model fit is labelled "Adjusted". The simulation was run 5000 times for each sample size: $n = 250, 500, 1000$. For $k = 1.2$, W strongly predicts Y and thus the TMLE, which adjusts for W results in a large increase in efficiency over the unadjusted method as observed by the RE provided in Tables 2.1 for each of the risk difference, relative risk and odds ratio parameters. The largest gain in efficiency occurs as expected when $\hat{Q}(A, W)$ is correctly specified followed closely by the D/S/A method, which in general shows a slightly higher variability than the correctly specified model due to possible overfitting of $\hat{Q}(A, W)$. In the scenario where $k = 20$ (Tables 2.2), A is more strongly predictive of Y as compared to W and thus the increase in efficiency is not as marked as when $k = 1.2$. The largest increase in efficiency for both values of k occurs for the estimates of the odds ratio (Table 2.2). When $\hat{Q}(A, W)$ is mis-specified, there is still a noticeable increase in efficiency showing that it is advised to always adjust for covariates. This is a result of the double robustness of the estimator as discussed in section 2.4 and appendix A.3. A significant result is the increase in power of the TMLE as evidenced by the proportion of rejected tests. In particular when $k = 1.2$, that is when the effect of A is weaker and more difficult to detect, the increase in power is quite significant. When $k = 20$, the performance of the unadjusted estimator is similar to the targeted maximum likelihood estimator with respect to power. In the strong treatment effect case, the conditional logistic regression method for estimating the odds ratio (see "Adjusted" in Table 2.2) demonstrates the issue that the point estimates are further from the null value which is consistent with the findings in literature as discussed in section 2.3. This is reflected in the coverage probabilities for the 95% confidence intervals, ranging from 0% to 22%.

Another notable result is that the TMLE circumvents the issue of singularity, i.e., Y is perfectly predicted by A and W, that occurs when using the adjusted estimate. In this situation the adjusted estimate is drastically inflated and for this reason, the adjusted results were not included in the consistency plots. However, this is not an issue for the TMLE. The efficiency gain of the TMLE increases as the covariate becomes more predictive. This becomes even more drastic when the covariate is perfectly predictive, whereas the adjusted estimate completely breaks down. For example, in a single run of the simulation for the odds ratio with $k = 1.2$, with $n = 250$, the coefficient for the treatment term in the conditional logistic fit was 25.4 and thus an estimate odds ratio of approximately 10^{11}. The corresponding TMLE using this same model gives an estimate of 1.083, noting that the true value is 1.087. This is of particular importance for small sample sizes but still occurs even for large sample sizes as shown in the RE estimates for the "Adjusted" estimate in Tables 2.1 and 2.2.

We also note that the consistency plots provided in Figure 2.1 show a small positive bias for all methods for the odds ratio and relative risk for smaller sample sizes. The targeted maximum likelihood methods however are less biased than the unadjusted method for all sample sizes.

The relative risk regression initial model $Q^0(A, W)$ was also applied. The G-computation estimate based on $Q^0(A, W)$ was computed in addition to the targeted maximum likelihood estimate for which the update covariate was required as discussed in section 2.4. The G-computation estimate based on this relative risk regression model resulted in a small gain in efficiency of approximately 3% with an additional 1% gain achieved with the TMLE with the appropriate update.

Table 2.1: Simulation 1: k=1.2, Mean squared error (MSE) and power comparison by sample size. *This table summarizes the MSE for the unadjusted estimator, and compares it to the TMLE as an RE measure. The TMLE are based on three methods for selecting Q(A,W), correctly specified (Correct), mis-specified (Mis-spec), and D/S/A algorithm (DSA). For the odds ratio parameter, the exponentiated coefficient from the conditional logistic regression model is also provided (Adjusted). The power for each of the estimators is provided next to which in parentheses is the coverage probability of the 95 percent confidence interval.*

	$n = 250$	$n = 500$	$n = 1000$
Risk Difference			
Unadjusted MSE	3.8e-03	1.9e-03	9.5e-04
TMLE Correct RE	10.46	13.70	13.67
TMLE Mis-spec RE	2.14	2.19	2.18
TMLE DSA RE	11.72	13.31	13.49
Unadjusted Power	0.07 (0.94)	0.08 (0.95)	0.10 (0.95)
TMLE Correct Power	0.26 (0.94)	0.42 (0.94)	0.67 (0.95)
TMLE Mis-spec Power	0.08 (0.94)	0.10 (0.95)	0.16 (0.95)
TMLE DSA Power	0.26 (0.94)	0.43 (0.93)	0.67 (0.94)
Relative Risk			
Unadjusted MSE	3.6e-02	1.7e-02	8.2e-03
TMLE Correct RE	9.70	13.97	13.70
TMLE Mis-spec RE	2.22	2.27	2.25
TMLE DSA RE	12.50	13.59	13.53
Unadjusted Power	0.05 (0.95)	0.07 (0.95)	0.10 (0.95)
TMLE Correct Power	0.25 (0.94)	0.41 (0.94)	0.67 (0.95)
TMLE Mis-spec Power	0.05 (0.95)	0.05 (0.96)	0.10 (0.96)
TMLE DSA Power	0.19 (0.94)	0.37 (0.94)	0.64 (0.96)
Odds Ratio			
Unadjusted MSE	1.0e-01	4.6e-02	2.2e-02
Adjusted RE	0.46	0.51	0.48
TMLE Correct RE	2.83	14.60	14.04
TMLE Mis-spec RE	2.24	2.28	2.21
TMLE DSA RE	13.46	14.19	13.86
Unadjusted Power	0.06 (0.95)	0.08 (0.95)	0.10 (0.95)
Adjusted Power	0.05 (0.97)	0.07 (0.96)	0.11 (0.96)
TMLE Correct Power	0.26 (0.94)	0.42 (0.94)	0.67 (0.95)
TMLE Mis-spec Power	0.08 (0.94)	0.10 (0.95)	0.15 (0.95)
TMLE DSA Power	0.26 (0.94)	0.43 (0.93)	0.67 (0.95)

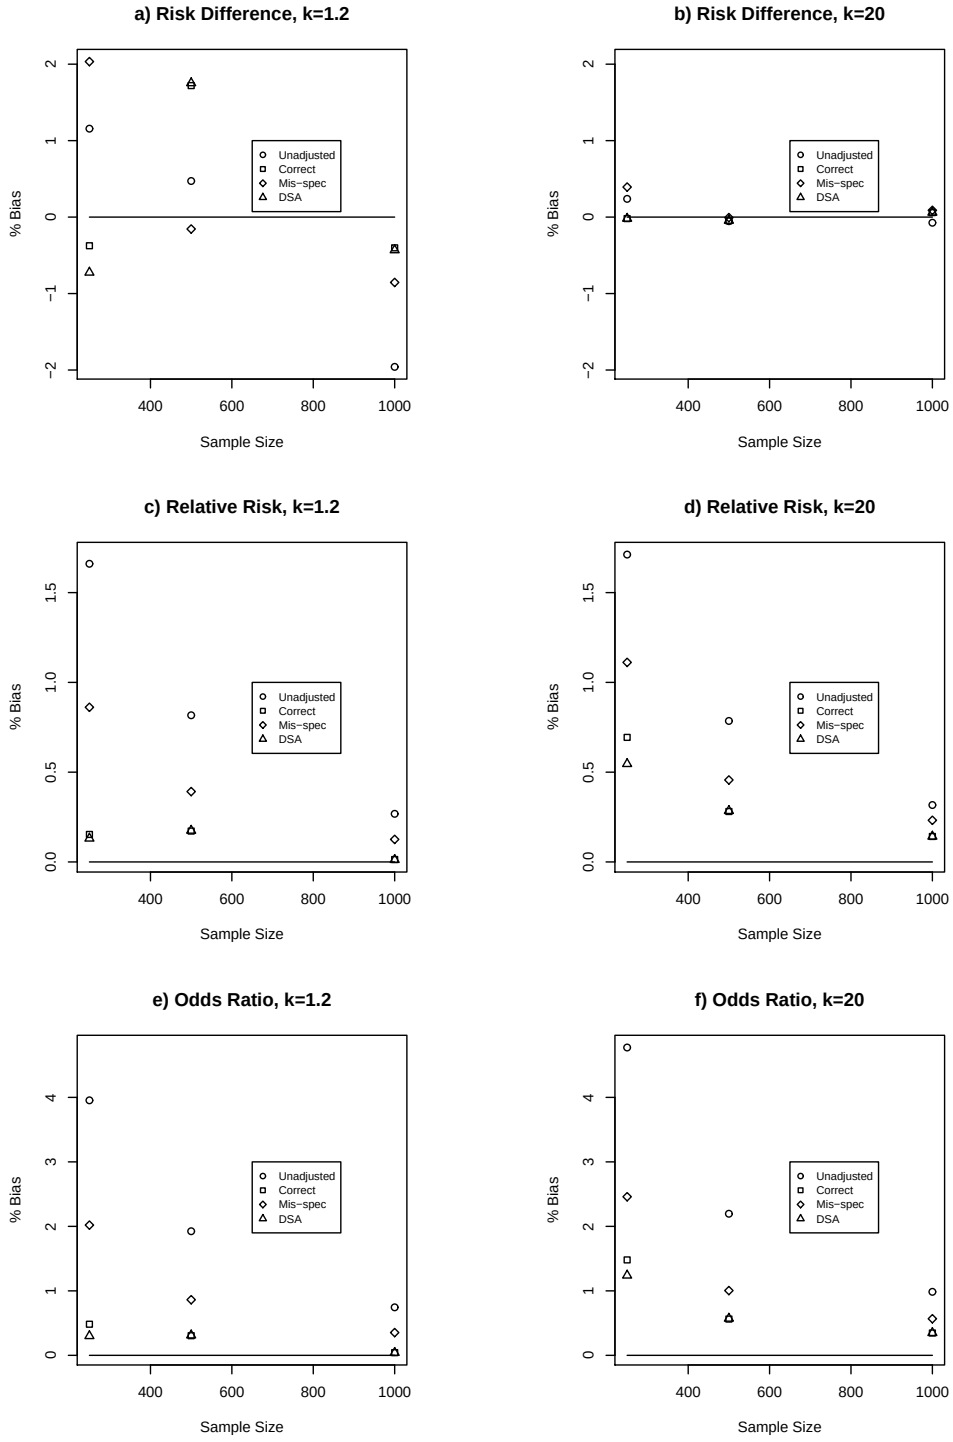

Figure 2.1: Simulation 1: Consistency graphs. a) risk difference, k=1.2; (b) risk difference, k=20; (c) relative risk, k=1.2; (d) relative risk, k=20; (e) odds ratio, k=1.2; and (f) odds ratio, k=20. *These plots show the bias as a percent of the true value for the average of the 5000 estimates for the unadjusted estimator and the TMLE with $Q(A,W)$ correctly specified (Correct), mis-specified (Mis-spec), and selected by the D/S/A algorithm (DSA).*

Table 2.2: Simulation 1: k=20, MSE and power comparison by sample size. *This table summarizes the MSE for the unadjusted estimator, and compares it to the TMLE as an RE measure. The TMLE are based on three methods for selecting Q(A,W), correctly specified (Correct), mis-specified (Mis-spec), and D/S/A algorithm (DSA). For the odds ratio parameter, the exponentiated coefficient from the conditional logistic regression model is also provided (Adjusted). The power for each of the estimators is provided next to which in parentheses is the coverage probability of the 95 percent confidence interval.*

	$n = 250$	$n = 500$	$n = 1000$
Risk Difference			
Unadjusted MSE	3.9e-03	2.0e-03	9.5e-04
TMLE Correct RE	3.58	4.70	4.60
TMLE Mis-spec RE	2.51	2.55	2.51
TMLE DSA RE	4.40	4.65	4.59
Unadjusted Power	0.96 (0.94)	1.00 (0.94)	1.00 (0.95)
TMLE Correct Power	1.00 (0.94)	1.00 (0.94)	1.00 (0.95)
TMLE Mis-spec Power	1.00 (0.95)	1.00 (0.94)	1.00 (0.95)
TMLE DSA Power	1.00 (0.94)	1.00 (0.94)	1.00 (0.94)
Relative Risk			
Unadjusted MSE	6.5e-02	3.1e-02	1.4e-02
TMLE Correct RE	2.14	4.13	3.97
TMLE Mis-spec RE	2.23	2.30	2.28
TMLE DSA RE	4.01	4.06	3.96
Unadjusted Power	0.95 (0.95)	1.00 (0.95)	1.00 (0.95)
TMLE Correct Power	1.00 (0.94)	1.00 (0.94)	1.00 (0.95)
TMLE Mis-spec Power	1.00 (0.99)	1.00 (0.99)	1.00 (1.00)
TMLE DSA Power	1.00 (0.99)	1.00 (0.99)	1.00 (1.00)
Odds Ratio			
Unadjusted MSE	5.7e-01	2.6e-01	1.2e-01
Adjusted RE	0.00	0.01	0.00
TMLE Correct RE	2.89	5.05	4.67
TMLE Mis-spec RE	2.63	2.65	2.52
TMLE DSA RE	5.01	4.97	4.67
Unadjusted Power	0.96 (0.95)	1.00 (0.94)	1.00 (0.95)
Adjusted Power	1.00 (0.22)	1.00 (0.03)	1.00 (0.00)
TMLE Correct Power	1.00 (0.94)	1.00 (0.94)	1.00 (0.95)
TMLE Mis-spec Power	1.00 (0.95)	1.00 (0.95)	1.00 (0.95)
TMLE DSA Power	1.00 (0.94)	1.00 (0.94)	1.00 (0.95)

2.7.2 Simulation 2: Odds ratio with interaction term

In this simulation, the treatment A and outcome Y are binary and W is a 2-dimensional covariate, $W = (W_1, W_2)$. The simulated data were generated according to the following laws:

1. $W_1 \sim N(2, 2)$
2. $W_2 \sim U(3, 8)$
3. $P(A = 1) = \theta_0 = 0.5$
4. $Q_0(A, W) = P(Y = 1|A, W) = \frac{1}{1+\exp\left[-(1.2A-5W_1^2+2W_2-5AW_1)\right]}$

The true values were given by $P(Y_1 = 1) = 0.312$, $P(Y_0 = 1) = 0.352$ and $OR = 0.833$. The same methods used in simulation 1 were used here to estimate the odds ratio. The simulation was run 5000 times for each sample size: $n = 250, 500, 1000$. For the "Mis-spec" TMLE, the mis-specified fit was given by,

$$\hat{Q}(A, W) = 1/\left\{1 + \exp\left[-(\hat{\alpha}_0 + \hat{\alpha}_1 A + \hat{\alpha}_2 W_1)\right]\right\}.$$

Figure 2.2 provides the consistency plot for each of the estimators. The results are similar to odds ratio for simulation 1 in that there is a small positive bias for all methods. The TMLE methods are again less biased than the unadjusted method for all sample sizes. Even when $\hat{Q}(A, W)$ is mis-specified the MSE is reduced as compared to the unadjusted estimate (Table 2.3). The D/S/A algorithm, which allows for interactions, shows a significant improvement in terms of the MSE. A notable increase in power is again observed for the TMLE over the unadjusted method.

Table 2.3: Simulation 2: MSE and power comparison by sample size where the true $Q(A, W)$ contains an interaction term between treatment and a covariate. *This table summarizes the MSE for the unadjusted estimator, and compares it to the TMLE as an RE measure. The TMLE are based on three methods for selecting Q(A,W), correctly specified (Correct), mis-specified (Mis-spec), and D/S/A algorithm (DSA). The exponentiated coefficient from the conditional logistic regression model is also provided (Adjusted). The power for each of the estimators is provided next to which in parentheses is the coverage probability of the 95 percent confidence interval.*

	$n = 250$	$n = 500$	$n = 1000$
Unadjusted MSE	6.1e-02	2.7e-02	1.3e-02
Adjusted RE	0.68	0.55	0.39
TMLE Correct RE	6.59	6.12	5.89
TMLE Mis-spec RE	2.58	2.52	2.46
TMLE DSA RE	7.47	7.87	7.40
Unadjusted Power	0.10 (0.95)	0.16 (0.95)	0.27 (0.96)
Adjusted Power	0.13 (0.94)	0.26 (0.92)	0.48 (0.87)
TMLE Correct Power	0.40 (0.94)	0.66 (0.94)	0.91 (0.95)
TMLE Mis-spec Power	0.20 (0.95)	0.34 (0.95)	0.57 (0.95)
TMLE DSA Power	0.50 (0.92)	0.78 (0.94)	0.97 (0.94)

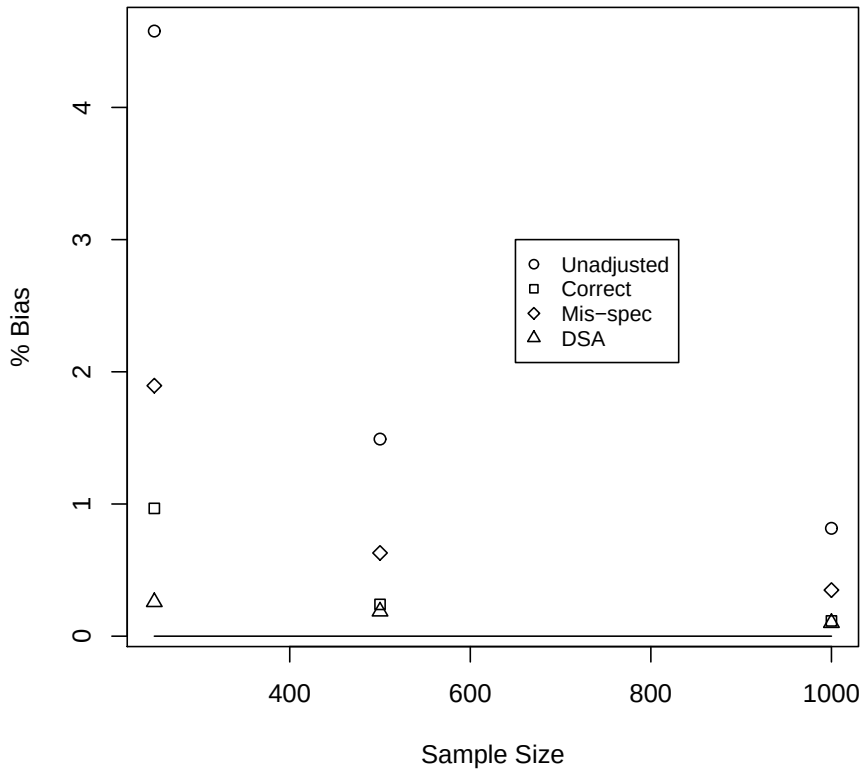

Figure 2.2: Simulation 2: Consistency graph where the true $Q(A, W)$ contains an interaction term between treatment and a covariate. *This plot shows the bias as a percent of the true value for the average of the 5000 estimates for the unadjusted estimator and the TMLE with $Q(A, W)$ correctly specified (Correct), mis-specified (Mis-spec), and selected by the D/S/A algorithm (DSA).*

2.7.3 Simulation 3: Estimating the treatment mechanism as well

In this simulation, the treatment mechanism, $\hat{P}(A|W)$ is estimated from the data using a logistic regression model with covariates that are predictive of the outcome Y. The simulated data were generated according to the following laws:

1. $W_1 \sim N(1, 2)$
2. $W_2 \sim U(1, 4)$
3. $W_3 \sim U(0, 20)$
4. $P(A = 1) = \theta_0 = 0.5$
5. $Q_0(A, W) = P(Y = 1|A, W) = \frac{1}{1+\exp\left[-(3A-2W_1^2-\log(W_2)+0.5W_3)\right]}$

The true values were given by $P(Y_1 = 1) = 0.569$, $P(Y_0 = 1) = 0.419$ and $RD = 0.150$. The treatment mechanism was estimated with the logistic regression model given by,

$$g(A|W) = 1/\left\{1 + \exp\left[-(\gamma_0 + \gamma_1 W_1 + \gamma_2 W_2 + \gamma_3 W_3)\right]\right\}.$$

The TMLE, represented as "Est tx" in Table 2.4 and Figure 2.3, with the estimated treatment mechanism is no longer equivalent to the G-computation estimator. The mis-specified fit for,

$$Q(A, W) = 1/\left\{1 + \exp\left[-(\alpha_0 + \alpha_1 A + \alpha_2 W_1)\right]\right\}$$

is used as the initial fit and the covariate $h(A, W)$ provided in section 2.4.2 is then added to this logistic regression. The TMLE is then estimated as usual. Thus, we are interested in comparing the mis-specified TMLE to the estimated treatment mechanism TMLE. Table 2.4 the efficiency is increased when estimating the treatment mechanism, from approximately 1.0 to 1.5. The power was approximately equal for the mis-specified and estimated treatment mechanism TMLE. The TMLE method using the D/S/A again shows a large improvement in efficiency and power over the unadjusted method.

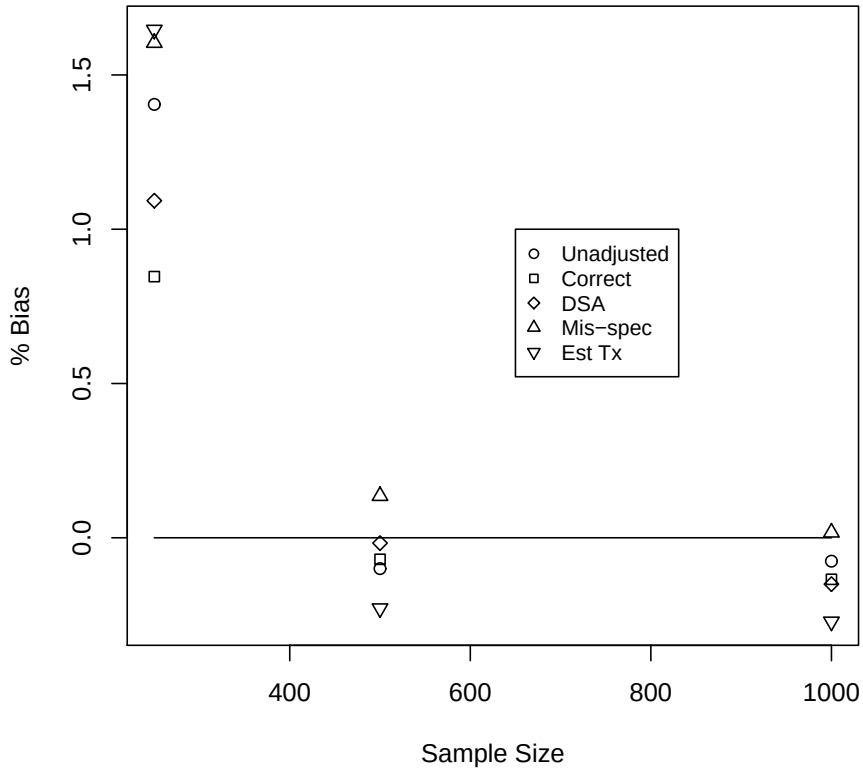

Figure 2.3: Simulation 3: Consistency graph where the treatment mechanism is also estimated. *This plot shows the bias as a percent of the true value for the average of the 5000 estimates for the unadjusted estimator and the TMLE with $Q(A, W)$ correctly specified (Correct), mis-specified (Mis-spec), and selected by the D/S/A algorithm (DSA). In addition the TMLE was also estimated with the correctly specified $Q(A, W)$ and the estimated treatment mechanism (Est Tx).*

Table 2.4: Simulation 3: MSE and power comparison by sample size where the treatment mechanism is also estimated *This table summarizes the MSE for the unadjusted estimator, and compares it to the TMLE as an RE measure. The TMLE are based on three methods for selecting $Q(A,W)$, correctly specified (Correct), mis-specified (Mis-spec), and D/S/A algorithm (DSA). In addition the TMLE was also estimated with the correctly specified $Q(A,W)$ and the estimated treatment mechanism (Est Tx). The power for each of the estimators is provided next to which in parentheses is the coverage probability of the 95 percent confidence interval.*

	$n = 250$	$n = 500$	$n = 1000$
Unadjusted MSE	2.6e-03	1.3e-03	6.5e-04
TMLE Correct RE	4.34	4.41	4.54
TMLE Mis-spec RE	1.01	1.03	1.01
TMLE DSA RE	4.17	4.35	4.51
TMLE Est tx RE	1.42	1.47	1.46
Unadjusted Power	0.25 (0.94)	0.42 (0.95)	0.69 (0.94)
TMLE Correct Power	0.79 (0.94)	0.97 (0.94)	1.00 (0.94)
TMLE Mis-spec Power	0.25 (0.94)	0.42 (0.95)	0.70 (0.95)
TMLE DSA Power	0.80 (0.94)	0.97 (0.93)	1.00 (0.94)
TMLE Est tx Power	0.21 (0.98)	0.40 (0.98)	0.73 (0.98)

2.7.4 Efficiency gain and R^2

To demonstrate the relation between efficiency gain and R^2, a simulation was run according to the following laws:

1. $\sqrt{W} \sim N(2, 2)$
2. $P(A = 1) = \theta_0 = 0.5$
3. $Q_0(A, W) = P(Y = 1|A, W) = \frac{1}{1+\exp[-(1.2A - cW)]}$

The data were sampled 5000 times with a sample size $n = 1000$ for each $c = \{0, 0.25, 2, 10\}$. That is covariate W is increasingly predictive. The R^2 was estimated in the ordinary least squares sense,

$$R^2 = 1 - \frac{\sum_{i=1}^{n} \left[Y_i - \hat{Q}(A, W)\right]^2}{\sum_{i=1}^{n} \left[Y_i - \bar{Y}\right]^2}.$$

A gain in R^2 was computed as the difference between R^2 in the covariate adjusted model and the covariate unadjusted model. Figure 2.4 depicts the RE to the unadjusted model for the TMLE for the risk difference and odds ratio against the gain in R^2. Clearly a gain in RE comparing the TMLE with the unadjusted corresponds with a gain in R^2.

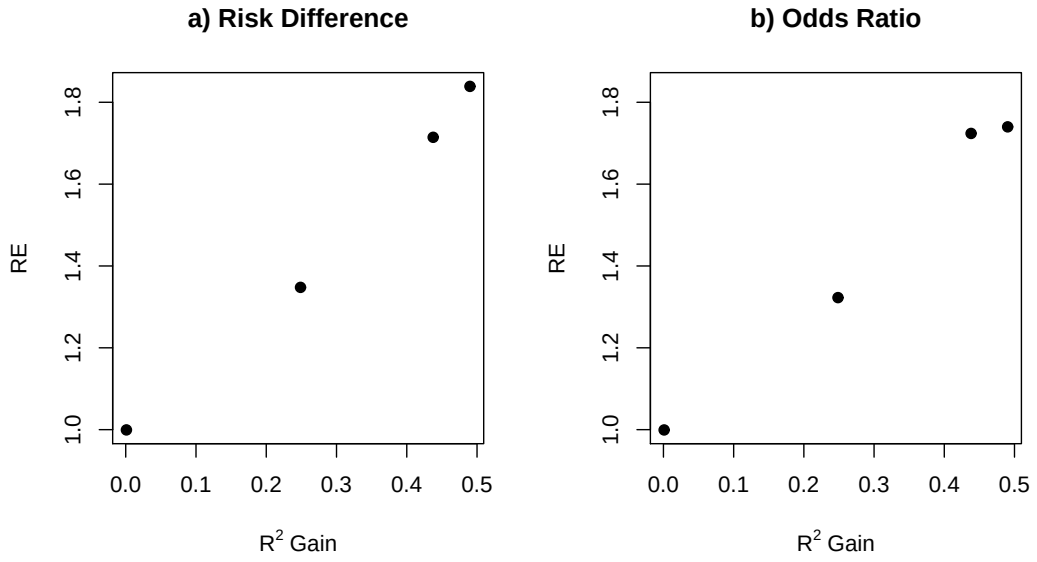

Figure 2.4: Efficiency Gain and R^2 for (a) risk difference; and (b) odds ratio. *These plots show the relationship between R^2 and the gains in relative MSE (RE) between the unadjusted and TMLE with correctly specified $Q(A, W)$*

2.7.5 Simulation discussion

The simulations were based on relatively simple data generating distributions but were useful in demonstrating the following results:

- The TMLE shows a clear increase in both efficiency and power over the unadjusted method, even when $Q(A, W)$ is not correctly specified.
- The D/S/A method for selecting $Q(A, W)$ provides a significant increase in efficiency and power over the mis-specified fixed $Q(A, W)$ method. The highest RE of approximately 13 was observed for the weak effect case with a sample size of $n = 1000$ in our simulations.
- The targeted maximum likelihood method circumvents the singularity issue that occurs when using the adjusted method of extracting the coefficient from the logistic regression model $Q(A, W)$.
- Interaction terms in the model for $Q(A, W)$ fit entirely into the framework of the targeted maximum likelihood.
- Estimating the treatment mechanism provides a further small increase in efficiency over targeting only $Q(A, W)$
- There is a clear relation between increasing R^2 and efficiency gain.
- The method of covariate adjustment that extracts the coefficient for treatment from the conditional logistic model demonstrated a loss in efficiency with a gain in power due to the inflated point estimates which corresponds with previous findings in the literature.

2.8 Discussion

Targeted maximum likelihood estimation provides a general framework that we applied to estimation of the marginal (unadjusted) effect of treatment in RCTs. We observed that the traditional method of covariate adjustment in RCTs using logistic regression models can be mapped, by averaging over the covariate(s), to obtain a fully robust and efficient estimator of the marginal effect, which equals the TMLE. We demonstrated that the TMLE does just this and results in an increase in efficiency and power over the unadjusted method, contrary to what has been reported in the literature for covariate adjustment for logistic regression. The simulation results demonstrated that data-adaptive model selection algorithms such as the D/S/A, which we used in this chapter, or forward selection, should be applied if the algorithm is specified *a priori*. However, we showed that even adjusting by a mis-specified regression model results in gain in efficiency and power. Thus, using an *a priori* specified model, even if it is mis-specified, can increase the power, and thus reduce the sample size requirements for the study. This is particularly important for trials with smaller sample sizes. The targeted maximum likelihood framework can also address missing data, either in the outcome as we demonstrated in section 2.4.3 for the risk difference, but also missingness in covariates and treatment as well for any of the parameters of interest. In these scenarios, derivation of the TMLE covariate may not be as straightforward as those that were presented in this chapter, but its derivation is analogous. We focused on logistic and relative risk regression, but the methodology can be extended to any other regression models for $Q(A, W)$. The targeted maximum likelihood framework can also be applied to other parameters of interest in RCTs such as an adjusted effect, for example by age or biomarker, and can also handle survival times as outcomes [65].

Chapter 3

Illustration of covariate adjustment with randomized controlled trial data and a recommended strategy for adjustment

3.1 Introduction

In an RCT, the effect of a treatment on a binary outcome is typically estimated as the observed difference in the proportion of subjects with the event in the treated and untreated arms. This unadjusted estimator is efficient when no other information is collected besides the data on the treatment, A, and outcome, Y, of interest. Indeed, from estimation theory [64], it is known that the nonparametric maximum likelihood estimator is the efficient estimator of the effect of interest. Since the likelihood of the observed data is $P(A)P(Y|A)$ when the data only contain information on the treatment and outcome, the unadjusted estimator does thus correspond to the maximum likelihood estimator.

When baseline covariates W are also collected, the likelihood of the observed data becomes

$$P(W)P(A|W)P(Y|A,W)$$

. Since the unadjusted estimator ignores the covariates, it can be viewed instead as a reduced data maximum likelihood estimator. It follows that ignoring covariate information by using the unadjusted estimator can lead to a loss in estimation efficiency (precision) in practice.

The targeted maximum likelihood estimation methodology presented in chapter 2 provided an estimation approach that, 1) incorporates the information of covariates W to improve estimation efficiency, and 2) does not require any additional parametric modeling assumptions to maintain their consistency property. This chapter aims at clarifying the issue of efficiency gain with baseline covariates in RCTs based on the framework of targeted maximum likelihood estimation and unify the different analytical protocols for covariate adjustment that have been proposed through an application to an RCT.

In addition, in this chapter, through application of this methodology to a sampled RCT dataset, we explore the origin of the gain in efficiency and criteria that can be used to anticipate whether the study design and the covariates collected can actually lead to increased estimation precision in practice. It is important to note that the criteria presented in this chapter rely on the assumption that the probability of receiving treatment is 0.5. Our recommendations for adjustment rely on this assumption and therefore do not deviate from the work of Freedman [16] where it is found that if the probability of receiving treatment is 0.5, that adjustment at worst simply does not result in a gain in efficiency. We illustrate how empirical confounding explains the gain in efficiency that can be achieved from an

adjusted analysis. Empirical confounding can occur due to bad luck in the randomization process, e.g. a higher percentage of older patients are assigned to the placebo arm and age has an effect on the outcome. Some RCT designs insure perfect covariate balance for some of the baseline covariates, however there are typically many other covariates collected that do not have a perfect balance. Empirical confounding, which occurs by chance for a given dataset, can introduce a large estimation error (sample bias), for which the unadjusted estimator cannot correct since it ignores covariate information. If the experiment were repeated many times, one would not expect this imbalance to occur in each sample and this estimation phenomenon can thus not be qualified as bias. We refer to it as sample bias or sample imbalance. The targeted maximum likelihood approach to covariate adjustment can account for such a covariate imbalance and thus improve over the poor finite sample performance of the unadjusted estimator due to empirical confounding.

Another issue that has obstructed the broad application of methodologies for covariate adjustment, in both the linear and logistic settings, has been concerns about the selection of the parametric covariate adjustment. Incorrect covariate adjustment can indeed typically lead to estimation bias. However, in chapter 2, the maximum likelihood estimator was shown to be doubly robust in RCTs. Despite this double robust property which ensures MLE consistency independent of the covariate adjustment, unease arises over the fact that investigators could still select the covariate adjustment that provides the most favorable inference without accounting for multiple testing. However, this is not an issue if one uses an *a priori* specified algorithm for model selection. When the model selection procedure is specified in the analysis protocol, the analysis is protected from investigators guiding causal inferences based on selection of favorable covariates and their functional forms in a parametric model.

It is important to note that the logistic models applied in this chapter are used for the purpose of improving efficiency in estimation of a marginal effect. They are not meant to infer information about subgroup (conditional) effects of treatment. The methodology can be extended to the estimation of subgroup specific effects but in this chapter we focus only on marginal treatment effects.

Throughout this chapter, we illustrate the proposed general methodology for covariate adjustment using a sampled dataset from an actual RCT. We outline and review all the aforementioned issues involved in covariate adjustment and suggest a concrete analytical protocol for covariate adjustment to improve estimation efficiency by accounting for empirical confounding. In section 3.2 we introduce the study and the data that are analyzed for illustration purposes throughout the chapter. We also outline the hypothesis tests of interest and their implementation. In section 3.3 using the data analysis, we provide an example of the apparent failure of conditional logistic models to improve estimation efficiency for the treatment effect of interest. In section 3.3.1, TMLE implementation and the selection of covariates are discussed based on the relation between R^2 and efficiency gain. In section 3.4 we illustrate how the performance and resulting inferences from the unadjusted and targeted maximum likelihood estimators compare based on the study and data presented in section 3.2. Section 3.4.1 explores the concept of empirical confounding and its relation to the origin of efficiency gain with the TMLE. Section 3.5 provides a recommended strategy for analyzing randomized trial data using covariates with the TMLE. Finally we conclude with a discussion in section 3.6.

3.2 Study and data example

The study of interest in this chapter is an international, multicenter, double-blind, parallel, placebo controlled randomized clinical trial which aims to evaluate safety based on mortality due to drug-to-drug interaction. The patients were randomized to receive either Drug1 or a placebo. All patients received Drug2 concomitantly as a background therapy. The primary objective was to determine whether the mortality rate between patients receiving Drug1 and placebo remained within a 1% margin or less.

We note that this is an exploratory safety analysis and an example to illustrate the methodology to evaluate whether there exists a drug-to-drug interaction that could result in higher mortality in the treated patients as compared to those that received the placebo.

The data consist of n independent and identically distributed observations of $O = (W, A, Y)$ where W is a vector of 40 baseline covariates, A is the treatment variable where $A = 1$ for the treated group and $A = 0$ for the placebo group, the outcome Y is all-cause mortality (0=Survived, 1=Died) at 28 days. The data available for this chapter were obtained from the original data by random sampling with replacement such that the distribution of the patient characteristics in the original study were maintained. The new subjects sampled were given a unique but different patient ID to protect the confidentiality of the original data. Therefore, the number of subjects, mortality and other crude rates may not be similar to that of the original clinical study. We treat each of the 2135 observations in the available data as independent realizations of O. Nine observations had missing outcomes and were deleted from the dataset leaving $n = 2126$ observations. Five of the 40 variables had proportions of missing values over 90%. The remaining 35 variables had proportions of missing values less than 1.3%. For the continuous variables, the missing values were imputed at the median for the given variable and the categorical variables were imputed at the category with the highest proportion of observed values. Corresponding indicator variables of whether or not the value was imputed were created, resulting in 80 baseline covariates. Dummy variables for each of the categorical outcomes were created resulting in 162 variables. The distribution of mortality by treatment is provided in Table 3.1. The proportion of subjects in the treated group is $\hat{\theta} = \frac{1}{n}\sum_{i=1}^{n} I(A = 1) = 0.504$ and the proportion in the placebo group is $1 - \hat{\theta} = 0.496$.

Table 3.1: Mortality by Treatment. *This table is the 2x2 contingency table for mortality and treatment.*

	Mortality	
	Survived (0)	Died (1)
Treatment		
PLACEBO	717	337
TEST DRUG	766	306

Since the primary research question of interest is whether the treated group has a mortality rate no worse than the placebo group within a 1% arbitrarily selected margin for equivalence, the hypothesis test for the risk difference is expressed as,

$$H_0 : P(Y = 1 \mid A = 1) - P(Y = 1 \mid A = 0) \geq 0.01$$
$$H_1 : P(Y = 1 \mid A = 1) - P(Y = 1 \mid A = 0) < 0.01.$$

In words,

$$H_0 : \text{Test inferior to placebo}$$
$$H_1 : \text{Test not much inferior to placebo.}$$

If the upper limit of the $100(1 - 2\alpha)\%$ confidence interval for $E(Y_1) - E(Y_0)$ is less than 0.01, then the null is rejected and conclusion is that the mortality rate of the test group is similar to that of placebo using a 1% equivalence margin. This is equivalent to testing the above hypothesis test at a level of α. The ICH guidelines state "The approach of setting Type I errors for one-sided tests at half the conventional Type I error used in two-sided tests is preferable in regulatory settings" [14]. Since the generally accepted level for a 2-sided hypothesis test is 0.05, we set the Type I error level for this 1-sided test to 0.025.

We note that a second test for superiority is expressed as,

$$H_0 : P(Y = 1 \mid A = 1) - P(Y = 1 \mid A = 0) \geq 0 \text{ (Test drug not superior)}$$
$$H_1 : P(Y = 1 \mid A = 1) - P(Y = 1 \mid A = 0) < 0 \text{ (Test drug superior).}$$

This test can be similarly performed at the 0.025 level as above by observing whether 95% upper confidence limit is less than 0.

To illustrate the proposed general methodology for extraction of covariate information to improve estimation efficiency over the standard unadjusted estimation approach used in RCTs, we compare results from both approaches for this latter test in addition to results from the noninferiority test described above.

3.3 Covariate adjustment with logistic models

To illustrate the decrease in estimation efficiency from conditional logistic models relative to unadjusted logistic models, we fit a conditional logistic model to the dataset and compare the corresponding estimate of the odds ratio to the unadjusted odds ratio. An estimate of the conditional odds ratio is easily obtained from a logistic model by simply exponentiating the coefficient for A. Note that this applies when there are no interaction terms between the treatment A and covariates W in the parametric model used. In this section, focus is placed on the odds ratio representation of the effect of interest versus the aforementioned risk difference measure because it is not clear how one typically derives a marginal risk difference from a model conditioning on W. We do provide a method in later sections based on averaging over the covariates W to obtain a marginal risk difference estimate from a conditional logistic model, however here we focus on the comparison of odds ratios to replicate and further illustrate published results [22, 53].

For clarity, covariate adjustment is based on a single covariate $BULTRA$, the indicator variable that a bilateral compression ultrasound was performed. This covariate was chosen since it is most correlated with the outcome. We later explain this selection criterion in section 3.3.1. The following logistic model was thus fit,

$$logit(P(Y = 1|A, W)) = \beta_0 + \beta_1 A + \beta_2 W,$$

where $W = BULTRA$.

The estimate $\hat{\beta}_1$ and the corresponding standard error for $\hat{\beta}_1$, in addition to the unadjusted estimates are provided in Table 3.2. Note that the odds ratio estimate is given by $\exp(\hat{\beta}_1)$ and that the standard error for $\hat{\beta}_1$ is indeed larger for the conditional estimate from the logistic model with $BULTRA$ than the unadjusted standard error. The upper limit of the 95% confidence interval includes 0 for the unadjusted estimate but does not for the conditional estimate. Thus, with the unadjusted method, one would not reject the null hypothesis at the 0.025 level for the following test,

$$H_0 : \log(OR) = 0 \;\; (OR \geq 1)$$

$$H_1 : \log(OR) < 0 \;\; (OR < 1).$$

However, one would in fact reject the null based on the conditional method. Even though the standard error is higher for the conditional estimate, the upper confidence limit is lower than the unadjusted due to the fact that the point estimate is further from the null as compared to the unadjusted, -0.217 and -0.163 respectively. The increased standard error is reflected in the width of the confidence interval which is wider for the conditional than the unadjusted. These results are consistent with those previously demonstrated in literature [22, 53].

Furthermore, note that these results are based on conditional models with no interaction terms between the treatment variable and covariates. With the presence of an interaction term, the logistic models becomes

$$\text{logit}(P(Y = 1|A, W)) = \beta_0 + \beta_1 A + \beta_2 W + \beta_3 AW,$$

and similar to the risk difference it is not clear what one should report as the estimate of the marginal causal odds ratio of interest. We note that this issue also applies in the linear model setting and thus no interaction terms are typically included in linear models.

Table 3.2: Conditional and unadjusted log odds ratio estimates. *This table provides the log odds ratio ($\log(OR)$), standard error (SE) and 95 percent confidence interval (95% CI) for the unadjusted estimate that is based on the extraction of the coefficient for treatment from the logistic regression fit that includes only treatment as a main term and the conditional estimate that is based on the extraction of the coefficient for treatment from the logistic regression fit that includes both treatment and the covariate BULTRA as main terms.*

	log(OR)	SE	95% CI
Unadjusted	-0.163	0.095	(-0.348,0.023)
Conditional	-0.217	0.108	(-0.428,-0.006)

3.3.1 Implementation of the TMLE and the selection of covariates for adjustment

The example in the previous section illustrated the results in RCT literature regarding the failure of covariate adjustment with logistic models. However, as demonstrated in chapter 2, proper covariate adjustment can improve estimation efficiency even with logistic models with targeted maximum likelihood estimation which provides an estimate of the marginal effect of interest rather than a conditional effect. We now continue with estimation of such a marginal effect in our data analysis using the methodology presented in chapter 2 and explore TMLE implementation issues and the selection of covariates for adjustment. The primary causal parameter of interest is the causal risk difference which is defined as $E_{p_0}(Y_1) - E_{p_0}(Y_0)$, which in words is the average difference over all subjects in the counterfactual outcomes corresponding with each treatment.

The implementation of the TMLE relies on estimating the nuisance parameter $Q(A, W)$ which will typically be based on a parametric model in practice (e.g. logistic model). Given the relation between R^2 and efficiency gain provided in section 2.6, one may be tempted to include as many covariates as possible in the model for $Q(A, W)$ since the corresponding $R^2_{Q(W)}$ increases as one increases the set of covariates W to predict Y. However, the correctness of inference based on the IC relies on the assumption that the model for Q used to derive $\hat{Q}(A, W)$ is not an overfit of the true nuisance parameter $Q(A, W)$. In an overfit situation, the asymptotic results on which the correctness of the IC-based inference rely do break down and one may run the risk of underestimating the variance with the IC, i.e. artificially attributing the gain in precision to covariate adjustments. This is due to the fact that IC-based inference is based on first order asymptotics whereas in finite samples second order terms can affect the inference (see [70], Chapter 3). Thus, we caution that the approach of including as many covariates as possible may result in overfitting $Q(W)$ which in turn can lead to incorrect (optimistic) inference from the influence curve associated with the TMLE, and more seriously, also result in loss in estimation efficiency.

To demonstrate this issue with the data analysis, we fit a model for $Q(A, W)$ including all 162 covariates as main terms. The standard error computed based on the influence curve was 0.016, whereas the

corresponding standard error based on the bootstrap procedure was 0.0229, a 43.4% increase. We note that as a comparison, the unadjusted influence curve and bootstrap standard errors were equivalent (within 0.3%). It thus appears that the bootstrap estimate of the TMLE variance is thus picking up second order terms contributing to the variability of the TMLE estimator that are ignored by the first order asymptotic approximation of the TMLE variance with the influence curve. In fact the bootstrap-based estimate of the TMLE standard error associated with the overfit model for $Q(A,W)$ is even higher than that of the unadjusted estimate and thus overfitting $Q(A,W)$ can result in a loss of efficiency.

The phenomenon illustrated above does not contradict the results from the relation between R^2 and relative efficiency provided in chapter 2, but instead can be explained by noting that the observed $R^2_{Q(W)}$ calculated on the same sample as the one on which $Q(W)$ is fit is not a good estimate of the true $R^2_{Q(W)}$. To obtain an appropriate estimate of the $R^2_{Q(W)}$, cross-validation is typically applied [67]. This cross-validated $R^2_{Q(W)}$ can subsequently be used as a model selection criterion for $Q(A,W)$ to avoid overfitting the nuisance parameter which results in incorrect inference from the influence curve and more importantly in a loss in estimation efficiency. As noted above, if the standard (non cross-validated) $R^2_{Q(W)}$ was used in practice to select a model for $Q(A,W)$, one would always use a model for $Q(A,W)$ that includes the maximum number of covariates since the associated $R^2_{Q(W)}$ would be largest. This is not necessarily the case if the cross-validated $R^2_{Q(W)}$ (cv-$R^2_{Q(W)}$) is applied instead.

In V-fold cross-validation, the data are divided into V subsets of equal size. In turn, each of the V data subsets are referred to as training set and the remaining subset of data is referred to as validation set. To compute cv-$R^2_{Q(W)}$ for a given model for $Q(W)$ with a V-fold cross-validation splitting scheme, one fits the model for $Q(W)$ on each of the V training sets and compute the associated $R^2_{Q(W)}$ with the observations from the corresponding validation sets. The mean of the V estimates for $R^2_{Q(W)}$ from each validation set is the cv-$R^2_{Q(W)}$.

Returning to our example with the overfitting of $Q(A,W)$, the corresponding observed $R^2_{Q(W)}$ was 0.36 as compared to the 5-fold ($V=5$) cv-$R^2_{Q(W)}$ of 0.23. When the model for $Q(A,W)$ only involves one single covariate, the observed $R^2_{Q(W)}$ is 0.23 and the cv-$R^2_{Q(W)}$ is 0.22. These results clearly demonstrate that the model for $Q(A,W)$ based on a single covariate does not overfit $Q(A,W)$ and thus the near equivalence between the cross-validated and standard $R^2_{Q(W)}$, whereas the very large model for $Q(A,W)$ shows a large difference in the two estimates signaling an overfitting problem.

Applying a cross-validated model selection criterion, one can avoid such overfitting issues when the algorithm is not overly agressive. There exist algorithms that use cross-validated criterion such as the Deletion/Substitution/Addition (DSA) algorithm which searches through a large space of possible polynomial models using the cross-validated risk as the selection criterion [56]. This algorithm can be computationally intensive. Other model selection algorithms that are based on a likelihood criterion, such as stepwise based on AIC, do not ensure an increase in cv-$R^2_{Q(W)}$ (decrease in risk) and thus do not guarantee a gain in efficiency.

We propose a more practical, simple and fast variant of the backward deletion algorithm for the selection of $Q(W)$ and thus $Q(A,W)$ (obtained by adding A to $Q(W)$) based on the maximum 5-fold cv-$R^2_{Q(W)}$. The algorithm is as follows:

1. Find all marginally associated covariates with False Discovery Rate (FDR) adjusted p-values less than 0.01. Let there be M such covariates.
2. Fit multivariate logistic regression including all M covariates and compute cv-$R^2_{Q(W)}$.
3. Delete covariate with largest p-value based on multivariate fit from previous step.

4. Fit new model with deleted covariate and compute new cv-$R^2_{Q(W)}$.
5. Repeat steps 3 and 4 until only 1 covariate remains in model.
6. Select the model among the M models with the largest cv-$R^2_{Q(W)}$ and add the treatment A to this model to obtain the model for $Q(A, W)$.

Note that the FDR procedure was applied in the first step due to the multiple tests performed to assess the univariate association of each covariate with the outcome [7].

We provide results from the application of the method described above to select the model for $Q(A, W)$ but also include results from the more simple approach that consists of selecting the model for $Q(A, W)$ based on the identification of the covariate most associated with the outcome. The single covariate selected and used to derive the model for $Q(A, W)$ in this second approach was the $BULTRA$ covariate discussed earlier. Confidence intervals corresponding with the tests, discussed in section 3.2, were based on the bootstrap procedure. For each bootstrap sample, the entire model selection process was run, including the ranking of the covariates by their FDR-adjusted p-values and the model selection procedure. Thus, the honest bootstrap procedure accounts for all sources of variability, including the second order terms discussed above. For each of these two methods for estimating $Q(A, W)$, the risk difference was estimated based on the targeted maximum likelihood method. Standard Errors were computed using 3 methods: the influence curve (IC) cross-validated IC (cv-IC) and the bootstrap procedure based on 20000 bootstrap samples. The cv-IC, as opposed to the standard IC was applied based on the same reasoning for using the cv-$R^2_{Q(W)}$ as opposed to the $R^2_{Q(W)}$. Using the cv-IC protects one from incorrectly claiming a significant result based on an overfit of $Q(A, W)$. However, if the model selection algorithm is based on cv-$R^2_{Q(W)}$, the standard IC can be used for inference since the use of cross-validation avoids the issue of overfit for $Q(A, W)$. We note that the cross-validated variance of the IC itself could also be used as a criterion for model selection, however, in this chapter we present results based on the cv-$R^2_{Q(W)}$ criterion only.

3.4 Results

The single most associated covariate had a strong univariate association with the outcome (p-value=$5.4e-78$). This covariate was used as a main term (with intercept) only model for the first method of selecting $Q(A, W)$. The backwards deletion method selected a model with 15 covariates, not including treatment. The plot of the cross-validated $R^2_{Q(W)}$ is provided in Figure 3.1 with the solid circle highlighting the maximum cross-validated $R^2_{Q(W)}$ corresponding with a model with 15 covariates. This plot shows there is little change in the cross-validated cv-$R^2_{Q(W)}$ for models with 15 covariates through 22 covariates. The largest gain results from the addition of a single covariate.

Table 3.3 provides all estimates including the unadjusted and targeted maximum likelihood methods. We first note that the upper confidence limit of the unadjusted test is greater than 0, and thus we would conclude that there is no evidence that the test drug is superior to the placebo at the 0.025 level. However, when we adjust for covariates using targeted maximum likelihood, using only the single most associated covariate, upper confidence limit is reduced to 0 with a relative efficiency of 1.14. The relative efficiency is calculated as $\frac{\hat{SE}_{un}}{\hat{SE}_{ajd}}$ where $\hat{SE}_{un}$ is the unadjusted standard error and $\hat{SE}_{adj}$ is the adjusted standard error. Furthermore, applying the backwards deletion algorithm, the upper confidence limit is reduced even further, and the relative efficiency increases to 1.211. Thus, applying the backwards deletion method we now have evidence to reject the null hypothesis and conclude that the mortality is lower is the test group as compared to the placebo. We note that the single most

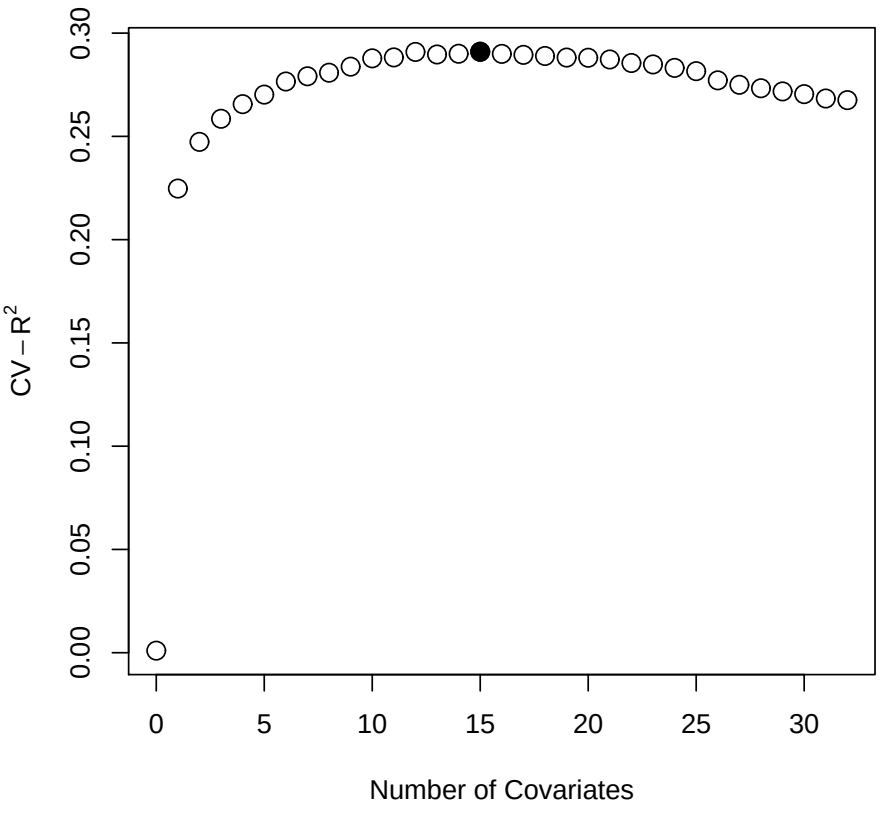

Figure 3.1: Cross-validated $R^2_{Q(W)}$ by number of covariates in model.

associated covariate method is right on the border of rejecting or accepting the null. However, we note that the research hypothesis was that the mortality for the test drug could be tolerated up to an increase of 1% in comparison to placebo. In this case, all of the upper confidence limits fall below 0.01 and the conclusion of noninferiority is the same using both the unadjusted or targeted maximum likelihood methods. It would still be of interest to the investigator that while the mortality falls within the pre-specified margin, that the observed reduction in mortality in the treated group is statistically significant.

Figure 3.2 shows that as the 5-fold cv-$R^2_{Q(W)}$ increases, so does the relative efficiency. Since this dataset contains highly predictive covariates of the outcome, and the study design did not balance treatment on covariates, the large gain in cv-$R^2_{Q(W)}$ also translates into a gain in estimation efficiency.

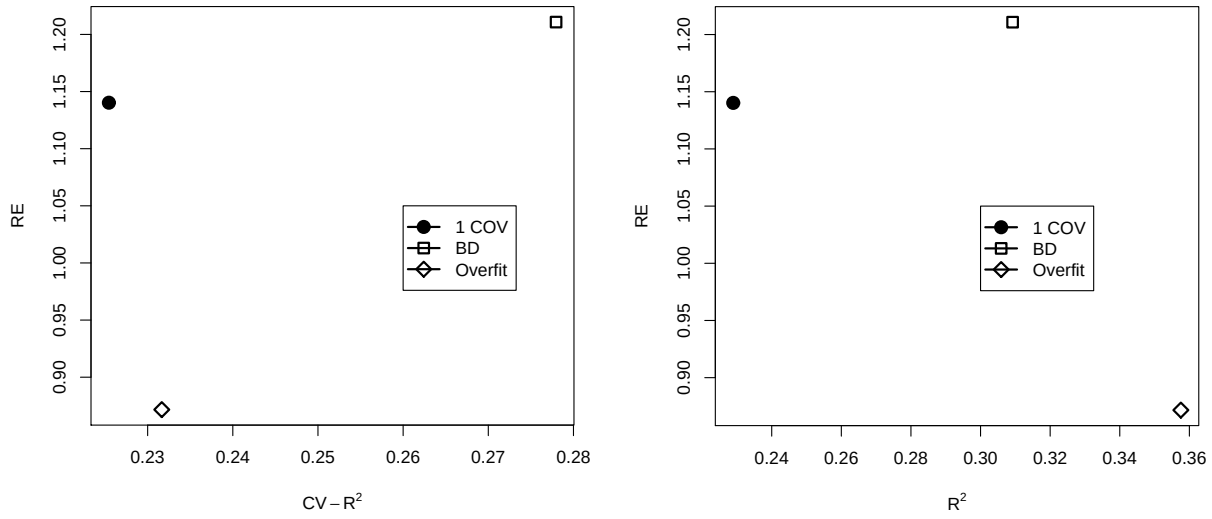

(a) Cross-validated $R^2_{Q(W)}$ (b) $R^2_{Q(W)}$

Figure 3.2: Relative Efficiency (RE) against cross-validated $R^2_{Q(W)}$ and $R^2_{Q(W)}$; a) cross-validated $R^2_{Q(W)}$, and b) $R^2_{Q(W)}$.

The results outlined above are based on standard errors computed based on the bootstrap procedure (20000 bootstrap samples). Table 3.4 provides a comparison of the IC, cv-IC and bootstrap standard errors. The standard errors for the unadjusted, single covariate and backwards deletion estimates are almost identical using the bootstrap as compared to the IC and cv-IC. The bootstrap standard error for the overfit method is 43.4% higher than the IC based standard error. However, the cv-IC method accounts for some of the overfit in that it is significantly higher than IC method. We again note that IC-based inference is valid in first order. However, with serious overfits, second order terms can play a role and the bootstrap procedure is indeed picking up these second order effects. We include this example as an extreme case to demonstrate when the methodology breaks down. It is our experience that using a cross-validated criterion in the model selection algorithm would avoid such a scenario. In our example based on the cv-$R^2_{Q(W)}$, the overfit would not have been selected.

It important to note that the error in the bootstrap based method is on the order of $1/\sqrt{(20000)} = 0.007$ where 20000 is the number of bootstrap samples. The difference in the standard errors are within this margin. Second, we note that the standard errors for the IC and the other 2 methods differ quite significantly (although still within this margin) for the overfit method. These results indicate that the cross-validated criterion is performing well with respect to not overfitting the data. They also indicate that when cross-validation is not used in the model fitting procedure, as in the overfit method, then

the cv-IC and IC based estimates differ. Thus one should always rely on the cv-IC or bootstrap based standard errors in such situations.

The results also show that the point estimates for the targeted maximum likelihood and the unadjusted methods differ. This difference can largely be attributed to the empirical confounder AGE. The probability of being treated among those patients over 75 years of age is 0.55 as compared to 0.49 for those under 75 years of age. The p-value for the univariate association between this discretized age covariate and treatment is 0.03. We note that this covariate is also associated with the outcome. The TMLE using the backward deletion method for selecting $Q(A, W)$ includes the covariate AGE. Thus, this estimate has adjusted for this small amount of empirical confounding and changes the point estimate from -0.034 to -0.042. The gains in efficiency are reflected in the standard error estimate as well as the point estimate. When the covariate AGE is removed from the backward deletion selected model and the TMLE is computed, the estimate becomes -0.038. The relative efficiency however is 1.196 and thus a gain is still achieved. The remaining difference in the point estimate and precision from the unadjusted after removing AGE indicates that there remains some empirical confounding due to other variables other than AGE.

Table 3.3: Comparison of unadjusted and targeted maximum likelihood estimates. *This table provides the unadjusted and targeted maximum likelihood estimates based on two methods for estimation of $Q(A, W)$; 1) The single most associated covariate (1 Cov), and 2) backwards deletion (BD). The 95 percent confidence interval (95 % CI) and Relative Efficiency (RE), computed as SE(TMLE)/SE(Unadjusted), are also provided.*

	Estimate	95% CI	RE
Unadjusted	-0.034	(-0.073,0.005)	1.000
TMLE 1 Cov	-0.035	(-0.07,-0.00)	1.140
TMLE BD	-0.041	(-0.074,-0.009)	1.210

Table 3.4: Comparison of influence curve (IC), cross-validated influence curve (cvIC) and bootstrap (Boot) based standard errors (SE). *This table summarizes the estimated standard errors for the TMLE based on 3 methods for fitting $Q(A, W)$: the single most associated covariate (1 Cov), backwards deletion (BD) and all covariates as main terms (Overfit). The comparison of SE IC and SE Boot is provided as the percent increase of SE Boot from SE IC (% Diff IC)*

	SE IC	SE cvIC	SE Boot	% Diff IC
Unadjusted	0.0199	0.0199	0.0200	0.3%
1 Cov	0.0175	0.0175	0.0175	0.1%
BD	0.0166	0.0167	0.0165	-0.4%
Overfit	0.0160	0.0175	0.0229	43.4%

3.4.1 Empirical confounding and the origin of efficiency gain with the TMLE

In this section, we conjecture that, for a given sample, the gain in efficiency from the TMLE is not only a function of the correlation between the covariates and the outcome ($R^2_{Q(W)} > 0$ as described above) but also a function of imbalances in these covariates with respect to treatment. More specifically, we conjecture that the origin of efficiency gain through covariate adjustment is empirical confounding,

i.e., if one had perfect balance in all covariates affecting the outcome, the adjusted estimate should not be more precise than the unadjusted estimate even though the covariates used for adjustment could be predictive of the outcome, i.e., $R^2_{Q(W)} > 0$. More specifically, let us hypothesize a RCT in which treatment is balanced perfectly in males and females and it is known that gender affects the outcome. In this case, $R^2_{Q(W)}$ (where W represents gender) is greater than 0, which would imply, based on the relation provided in chapter 2, that gender adjustment through TMLE would result in increased estimation efficiency. However, based on simulation results described below, we have found that the standard error of the adjusted estimate based on gender would be equivalent to the standard error of the unadjusted estimate. As is the case in most RCT, other covariates beyond gender would also be collected that have an effect on the outcome. Based on our simulation results again, we conjecture that adjusting for these other covariates in addition to gender lead however to an increase in precision and furthermore, adjusting jointly for gender and these other covariates appears to increase the precision even further, even though adjusting for gender only results in no gain. We note that ongoing work involves formalizing these hypotheses that relate efficiency gain in practice to empirical confounding.

To support our conjecture and illustrate the efficiency gains with the TMLE in the presence and absence of empirical confounding, 10,000 datasets of size $n = 1000$ were simulated with a binary treatment A, such that $P(A = 1) = P(A = 0) = 0.5$, two binary covariates, $W = (W_1, W_2)$, such that $P(W_1 = 1) = 0.4$ and $P(W_2 = 1) = 0.6$ and an outcome Y with $logit(P(Y = 1|A, W)) = 5A - 3W_1 - 3W_2$. The true risk difference is 0.616 with $E(Y_1) = 0.761$ and $E(Y_0) = 0.145$. In the first setting, the treatment arms were balanced (matched) perfectly on W_1 and W_2, in the second setting, the treatment arms were balanced on W_1 only and in the third setting the treatment arms were not balanced perfectly on either covariate. The unadjusted estimate and three adjusted estimates (TMLE) were computed under each of these three settings, where the three adjusted estimates correspond to adjusting for W_1 and W_2 ($\hat{Q}(A, W) = \hat{Q}(A, W_1, W_2)$), W_1 only ($\hat{Q}(A, W) = \hat{Q}(A, W_1)$) and W_2 only ($\hat{Q}(A, W) = \hat{Q}(A, W_2)$).

The mean squared error (MSE) results for each simulation setting and TMLE estimate are provided in Table 3.5. In the first setting where there is perfect balance of treatment on both covariates, no gain in MSE is achieved through covariate adjustment with any of the three TMLE estimates, even though both covariates have an effect on the outcome Y. In the second setting where treatment is perfectly balanced on W_1 only, a small amount of empirical confounding by W_2 is present and a precision gain is achieved by adjusting for W_2. Note that even though no gain is achieved by adjusting for W_1 only, adjusting for W_1 and W_2 results in a slightly lower MSE than adjusting for W_2 only. This may at first seem counterintuitive, however, although W_1 and W_2 are independent, for any given sample, a small amount of correlation exists between the two variables. Since W_1 is not perfectly balanced on W_2, adjusting for it as well as W_2 is akin to adjusting for empirical confounding. Therefore, even if there is a perfect balance of treatment on a covariate, adjusting for it when there exists another covariate on which there is not a perfect balance will provide a further gain in efficiency.

These results supports the conjecture that in the extreme scenario of a perfectly balanced trial in all covariates, one could not obtain a increase in precision through covariate adjustment, even if the covariates are strongly predictive of the outcome. The TMLE method of adjustment can then be viewed as an attempt to mimic this ideal setting in which perfect balance is present in all covariates. This is evidenced by comparing the MSE of the adjusted estimate in the third setting where treatment was not perfectly balanced on either covariate, and the MSE of the unadjusted estimate with perfect balancing on W in the first setting. The MSE are almost equivalent, indicating possibly an efficiency bound for the adjusted estimate.

These results support our claim that the gain in efficiency in RCT by adjusting for covariates is a result of adjusting for empirical confounding. We note that empirical confounding results in covariate imbalances that are typically quite small. In our dataset, only 4 covariates had significant associations with treatment (p-values < 0.05). For example, the probability of receiving treatment among those

patients over 75 years of age is 0.55 as compared to 0.49 for those under 75 years of age and adjusting for age results in a change in the point estimate. Note however that adjustment for variables not significantly associated with treatment but affecting the outcome, i.e., whose imbalances are very small with respect to the treatment, can still result in a gain in precision. Thus, testing for a significant imbalance is not a valid strategy since one could miss a covariate that is strongly associated with the outcome with only a small imbalance. Furthermore, it has been pointed out that tests for covariate imbalance do not make sense in RCT because by definition, all imbalances are due to chance. Thus, such a test is a test of a null hypothesis that is by definition true [55, 4, 42]. Since a perfect balance is very unlikely, it is a better strategy to included those covariates in the adjustment that are predictive of the outcome, based on the relation provided in chapter 2, as recommended in [43, 2, 55, 42].

Table 3.5: MSE comparison. *This table summarizes the mean MSE for the TMLE for the simulations scenarios that balance treatment on both covariates, one covariate and neither covariate.*

Balanced on	Adjusted For	MSE
W_1 and W_2	None	4.31e-04
W_1 and W_2	W_1 and W_2	4.31e-04
W_1 only	None	5.19e-04
W_1 only	W_1 and W_2	4.34e-04
W_1 only	W_1 only	5.19e-04
W_1 only	W_2 only	4.35e-04
None	None	6.12e-04
None	W_1 and W_2	4.36e-04

3.5 Recommended strategy for analyzing RCT Data

A clear strategy needs to be outlined in the study protocol detailing the analysis of clinical trials. We provide an approach based on our theoretical and simulation results presented in this chapter. The strategy is as follows.

1. As discussed throughout this chapter, gains in efficiency are related to gains in $R^2_{Q(W)}$. Thus, one should attempt to collect covariates known or speculated to be predictive of the outcome, which are not perfectly balanced on by design, and outline them in the study protocol.

2. Estimate the model for $Q(A, W)$.

   - Based on a model an *a priori* specified in the study protocol (for example include age only).

     OR

   - Apply a model selection algorithm, *a priori* specified in the study protocol and based on a cross-validated criterion (cv-$R^2_{Q(W)}$ or cross-validated variance of the IC) such as the backwards deletion algorithm provided in this chapter.

3. If cv-$R^2_{Q(W)}$ (or cross-validated variance of the IC) of the model for $Q(W)$ selected in the previous step is significantly different than 0 according to a test (work on such a test is in progress, see discussion) then proceed to step 4. Otherwise, no gain in efficiency can be achieved by covariate adjustment and the unadjusted estimate must be used (i.e., the next steps are skipped).

4. Apply TMLE based on the fitted model from the previous step to obtain an estimate of the parameter of interest.

5. Estimate standard error based on the IC or the bootstrap procedure. For honest bootstrap estimates, one must perform the entire model selection procedure (if used) on each bootstrap sample.

Note that steps 2 and 3 could involve a double layer or cross-validation. However, this can be replaced with a single layer of cross-validation if the algorithm is not overly aggressive. This can be achieved by specifying a maximum model size, such as 30 observations per term in the model.

3.6 Discussion

In this chapter, we have shown that covariate adjustment for binary outcomes using logistic models can indeed increase the estimation efficiency (precision) for the marginal effect of treatment. The difference from convenvional approaches for covariate adjustment using conditional logistic models lies in the fact that the method presented in this chapter averages over the covariates in the logistic model to obtain a marginal (unconditional) effect estimate that can be compared to the standard unadjusted effect estimate. The logistic models presented in this chapter are not meant to describe subgroup effects but rather have the purpose of increasing efficiency in the estimation of the marginal effect. We note that the method of targeted maximum likelihood estimation can also be applied to estimation of conditional or subgroup effects of treatment however we focused on marginal effects only in this chapter.

The gain in efficiency has real implications as was demonstrated with the fact that the test for equivalence would provide different conclusions using either the unadjusted or adjusted estimation approaches. However, we note that the test for equivalence within the pre-specified margin would result in the same conclusion using either method even though the confidence intervals were narrowed with the targeted maximum likelihood method.

Using an *a priori* specified algorithm for covariate adjustment protects the investigators from guiding their analyses in the direction that provides the most desirable results. The comparison of the bootstrap and analytic-based (IC and cv-IC) standard errors demonstrated the need for a cross-validated criterion for the selection of the covariate adjustment ($Q(A, W)$) to avoid the problem of overfitting which results in incorrect inference with the influence curve and a loss in the possible precision gain from covariate adjustment. We provided a fast and easy to implement algorithm based on the cross-validated $R^2_{Q(W)}$ that resulted in a relative efficiency of 1.211 as compared to the unadjusted method. Even adjusting for the single most associated covariate resulted in a significant gain. These results indicate that predictive covariates of the outcome, that do not have a perfect balance in treatment can significantly increase efficiency. This gain in efficiency is reflected in the reduction of the standard error and also possibly a change in the point estimate due to finite sample error. We conjecture that this gain in efficiency is the sole result of adjustment for empirical confounding.

Ongoing work includes studying and formalizing the relation between efficiency gain and empirical confounding. A model selection algorithm could be developed using this relation as a basis for the selection criterion. In addition, future work involves developing a formal test for the hypothesis $R^2_{Q(W)} > 0$, either a nonparametric permutation test of independence between W and Y or a model-based test for the fixed model approach similar to the likelihood ratio test comparing the model for $Q(W)$ to the intercept model.

Chapter 4

Assessment of safety in randomized controlled trials with time-to-event outcomes

4.1 Introduction

Safety analysis in RCTs involves estimation of the treatment effect on the numerous adverse events (AE) that are collected in the study. RCT are typically designed and powered for efficacy rather than safety. Even when assessment of AE is a major objective of study, the trial size is generally not increased to improve likelihood of detecting AE [17]. As a result, power is an important concern in the analysis of the effect of treatment on AE in RCT [41].

Typically in an RCT, crude incidences of each AE are reported at some fixed end point such as the end of study [18, 20, 32]. These crude estimates often ignore missing observations that frequently occur in RCT due to early patient withdrawals [35]. A review of published RCT in major medical journals found that that censored data are often inadequately accounted for in their statistical analyses [72]. A crude estimator that ignores censoring can be highly biased when the proportion of dropouts differs between treatment groups (see [18] for examples).

The crude incidence is an important consideration in the evaluation of safety for very rare, severe or unexpected AE. Such AE require clinical evaluation for each case and are not the focus of this chapter. Instead, we focus on those AE that are routinely collected in RCT and most often are not associated with a pre-specified hypothesis. These AE are typically reported as an observed rate with a confidence interval or p-value.

Patient reporting of AE occurrence usually occurs at many intervals throughout the study often collected at follow-up interviews rather than only at a single fixed end-point. As such, time-to-event methods that exploit these data structures may provide further insight into the safety profile of the drug. The importance of considering estimators of AE rates that account for time due to differential lengths of exposure and follow-up is discussed in O'Neill [40]. Furthermore, in most RCT in oncology, most if not all patients suffer from some AE [39], and thus investigators may be interested in the probability of the occurrence of a given AE by a certain time rather than simply the incidence. Time-to-event analysis techniques may be more sensitive than crude estimates in that they readily handle missing observations that frequently occur in RCT due to early patient withdrawals. For example, in Davis et al. [11], AE from the Beta-Blocker Heart Attack Trial were analyzed by comparing distributions of the time to the first AE in the two treatment arms. The results of this analysis were contrasted to the

cross-sectional crude percentage analysis and were found to be more sensitive in detecting a difference by taking into account the withdrawals. A vast amount of literature exists for time-to-event analysis but these methods are often not applied to the analysis of AE in RCT. A general review of survival analysis methods in RCT (without a particular focus on AE) is provided in Fleming and Lin [13].

In this chapter we focus on estimation of treatment specific survival at a fixed end point for right-censored survival outcomes using targeted maximum likelihood estimation [65]. Survival is estimated based on a hazard fit and thus the time-dependent nature of the data is exploited. There are two main goals of the methodology presented in this chapter over unadjusted crude proportions and Kaplan-Meier estimators. The first is to provide an estimator that exploits covariates to improve efficiency in the estimation of treatment-specific survival at fixed end points. The second is to provide a consistent estimator in the presence of informative censoring.

4.2 Motivation and outline

Consider the estimation of the effect of treatment on a particular AE at some fixed end point in the study. From estimation theory, it is known that the nonparametric maximum likelihood estimator (MLE) is the efficient estimator of the effect of interest [64]. In most RCT, data are collected on baseline (pre-treatment) covariates in addition to the treatment and the AE of interest. The unadjusted or crude estimator is defined as the difference in proportions of the AE between treatment groups. This 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.

Conflicting results in initial applications of covariate adjustment in RCT for estimating the treatment effect for fixed end-point efficacy studies were found. For continuous outcomes using linear models for adjustment demonstrated gains in precision over the unadjusted estimate [43]. However adjustment using logistic models for binary outcomes was shown to actually reduce precision and inflate point estimates [22, 53].

This apparent contradiction was resolved through the application of estimating function methodology [59, 73] and targeted maximum likelihood estimation [36]. In these references, consistent estimators that do not require parametric modeling assumptions were provided and shown to be more efficient than the unadjusted estimator, even with binary outcomes. It just so happens that the coefficient for the treatment variable in a linear regression that contains no interactions with treatment coincides with the efficient estimating function estimator and thus the targeted maximum likelihood estimator. This fortunate property does not hold for the logistic regression setting, i.e., the exponentiated coefficient for treatment from the logistic regression model does not equal the unadjusted odds ratio. This conditional estimator does not correspond to the marginal estimator in general and in particular not in the binary case. The efficient estimate of the marginal (i.e., unconditional) effect obtained from the conditional regression is the weighted average of the conditional effect of treatment on the outcome given covariates according to the distribution of the covariates.

With this principle of developing covariate adjusted estimators that do not require parametric modeling assumptions for consistency in mind, in this chapter we provide a method for covariate adjustment in RCT for the estimation of treatment specific survival at a fixed end point for right-censored survival outcomes. Thereby, we can estimate a comparison of survival between treatment groups at a fixed end point that is some function of the two treatment specific survival estimates. Examples of such parameters are provided in section 4.3 such as the marginal additive difference in survival at a fixed end point. Under no or uninformative censoring, the estimator provided in this chapter does not require any additional parametric modeling assumptions. Under informative censoring, the estimator

is consistent under consistent estimation of the censoring mechanism or the conditional hazard for survival.

It is important to note that the conditional hazard on which the estimate is based is not meant to infer information about subgroup (conditional) effects of treatment. By averaging over the covariates that have terms in the hazard model, we obtain a marginal or unconditional estimate. The methodology presented in this chapter can be extended to the estimation of subgroup specific effects however we focus only on marginal (unconditional) treatment effects on survival at fixed end point(s).

We also note that the methodology can be extended to provide a competitor test to the ubiquitous log-rank test. Methods have been proposed for covariate adjustment to improve power over the logrank test [23, 31, 33]. These are tests for an average effect of treatment over time. Our efficiency results are not in comparison to these methods but rather to the treatment-specific Kaplan-Meier estimate at that fixed end point.

In itself treatment specific survival at a fixed end point, and thereby the effect of treatment on survival at that end point can provide useful information about the given AE of interest. This is a very common measure to report (see [18, 20, 32, 35]), however most of the currently applied estimation approaches ignore covariates and censoring and do not usually exploit the time-dependent nature of the data.

We present our method of covariate adjustment under the framework of targeted maximum likelihood estimation originally introduced in van der Laan and Rubin [65]. Specifically, the chapter is outlined as follows. We first begin by outlining the data, model and parameter(s) of interest in section 4.3. The application of targeted maximum likelihood estimation to our parameter of interest with its statistical properties and inference are presented in section 4.4. In section 4.5 we present a simulation study to demonstrate the efficiency gains of the proposed method over the current methods in an RCT under no censoring and uninformative censoring. Furthermore, under informative censoring we demonstrate the bias that arises with the standard estimator in contrast to the consistency of our proposed estimator. The TMLE requires estimation of an initial conditional hazard. Methods for fitting this initial hazard as well as the censoring mechanism are provided in section 4.6. In section 4.7 we outline the inverse weighting assumption for the censoring mechanism. Alternative estimators and their properties are briefly outlined in section 4.8. AE data are multivariate in nature in that many AE are collected an analyzed in any given RCT. In section 4.9 we outline the multiple testing issues involved in the analysis of such data. Section 4.10 provides extensions to the methodology including time-dependent covariates, and post-market safety analysis. Finally, we conclude with a discussion in section 4.11.

4.3 Data, model and parameter of interest

We assume that in the study protocol, each patient is monitored at K clinical visits. At each visit, M AE are evaluated as having occurred or not occurred. We focus on the first occurrence of the AE and thus let T represent the first visit when the AE reported as occurring and thus can take values $\{1, ..., K\}$. 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 represents 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 the density of O. The conditional hazard is given by $\lambda_0(\cdot \mid A, W)$ and the corresponding conditional survival is given by $S_0(\cdot \mid A, W)$. The censoring mechanism is given by $\bar{G}(t_- \mid A, W) = P(C \geq t \mid A, W)$. We present the methodology for estimation of the treatment effect for a single AE out of the M total AE collected. This procedure would be repeated for each of the M AE. For multiplicity considerations see section 4.9.

Let T_1 represent a patient's time to the occurrence of an AE had she possibly contrary to fact been assigned to the treatment group and let T_0 likewise represent the time to the occurrence of the AE had the patient been assigned to the control group.

Let $\mathcal{M}$ be the class of all densities of O with respect to an appropriate dominating measure where $\mathcal{M}$ is nonparametric up to possible smoothness conditions. Let our parameter of interest be represented by $\Psi(p_0)$. Specifically, we aim to estimate the following treatment specific parameters,

$$P_0 \rightarrow \Psi_1(p_0)(t_0) = Pr(T_1 > t_0) = E_0(S_0(t_0|A = 1, W)), \tag{4.1}$$

and

$$P_0 \rightarrow \Psi_0(p_0)(t_0) = Pr(T_0 > t_0) = E_0(S_0(t_0|A = 0, W)), \tag{4.2}$$

where the subscript for Ψ denotes the treatment group, either 0 or 1. In order to estimate the effect of treatment A on survival T we can thereby estimate a parameter that is some combination of $Pr(T_1 > t_0)$ and $Pr(T_0 > t_0)$. Examples include the marginal log hazard of survival, the marginal additive difference in the probability of survival, and the marginal log relative risk of survival at a fixed time t_0 given respectively by,

$$P_0 \rightarrow \Psi_{HZ}(p_0)(t_0) = \log\left(\frac{\log(Pr(T_1 > t_0))}{\log(Pr(T_0 > t_0))}\right), \tag{4.3}$$

$$P_0 \rightarrow \Psi_{AD}(p_0)(t_0) = Pr(T_1 > t_0) - Pr(T_0 > t_0), \tag{4.4}$$

and

$$P_0 \rightarrow \Psi_{RR}(p_0)(t_0) = \log\left(\frac{Pr(T_1 > t_0)}{Pr(T_0 > t_0)}\right). \tag{4.5}$$

We note that if one averaged $\Psi_{HZ}(p_0)(t_0)$ over t, this would correspond with the Cox proportional hazards parameter and thus the parameter tested by the log rank test. However, we focus only on the t_0-specific parameter in this chapter.

4.4 Estimation of treatment specific survival at a fixed end point

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

Let the survival time be discrete and let the initial hazard fit $\hat{\lambda}(t \mid A, W)$ be given by a logistic regression model,

$$\text{logit}(\hat{\lambda}(t \mid A, W)) = \hat{\alpha}(t) + m(A, W \mid \hat{\beta}),$$

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

$$\text{logit}(\hat{\lambda}(\epsilon)(t \mid A, W)) = \hat{\alpha}(t) + m(A, W|\hat{\beta}) + \epsilon h(t, A, W). \tag{4.6}$$

The algorithm selects $h(t, A, W)$ 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 $\lambda(t \mid A, W))$ in the nonparametric model for the observed data assuming only coarsening at random (CAR).

The general formula for this covariate $h(t, A, W)$ for updating an initial hazard fit was provided in van der Laan and Rubin [66] and is given by,

$$h(t, A, W) = \frac{D^{FULL}(A, W, t \mid \hat{p}) - E_{\hat{p}}[D^{FULL}(A, W, T \mid \hat{p}) \mid A, W, T > t)]}{\bar{G}(t_- \mid A, W)}, \tag{4.7}$$

where D^{FULL} is the efficient influence curve of the parameter of interest in the model in which there is no right censoring. This is also the optimal estimating function in this model. This full data estimating function for $\Psi_1(p_0)(t_0)$ provided in (4.1) is given by,

$$\begin{aligned} &D_1^{FULL}(T, A, W \mid p)(t_0) = \\ &\quad [I(T > t_0) - S(t_0 \mid A, W)]\frac{I(A = 1)}{g(1|W)} + S(t_0 \mid 1, W) - \psi_1(p), \end{aligned} \tag{4.8}$$

and for $\Psi_0(p_0)(t_0)$ provided in (4.2) it is given by,

$$\begin{aligned} &D_0^{FULL}(T, A, W \mid p)(t_0) = \\ &\quad [I(T > t_0) - S(t_0 \mid A, W)]\frac{I(A = 0)}{g(0|W)} + S(t_0 \mid 0, W) - \psi_0(p), \end{aligned} \tag{4.9}$$

To obtain the specific covariates for targeting the parameters $\Psi_1(p_0)(t_0)$ and $\Psi_0(p_0)(t_0)$, the full data estimating functions provided in (4.8) and (4.9) at $t = t_0$ are substituted into (4.7). Evaluating these substitutions gives the covariates,

$$h_1(t, A, W) = -\frac{I(A = 1)}{g(1)\bar{G}(t_- \mid A, W)}\frac{S(t_0 \mid A, W)}{S(t \mid A, W)}I(t \leq t_0), \tag{4.10}$$

and

$$h_0(t, A, W) = -\frac{I(A = 0)}{g(0)\bar{G}(t_- \mid A, W)}\frac{S(t_0 \mid A, W)}{S(t \mid A, W)}I(t \leq t_0), \tag{4.11}$$

for the treatment specific parameters $\Psi_1(p_0)(t_0)$ and $\Psi_0(p_0)(t_0)$ respectively.

Finding $\hat{\epsilon}$ in the updated hazard provided in (4.6) to maximize the likelihood of the observed data can be done in practice by fitting a logistic regression in the covariates $m(A, W \mid \hat{\beta})$ and $h(t, A, W)$. The coefficient for $m(A, W \mid \hat{\beta})$ is fixed at one and the intercept is set to zero and thus the whole regression is not refit, rather only ϵ is estimated. These steps for evaluating $\hat{\epsilon}$ correspond with a single iteration of the targeted maximum likelihood algorithm. In the second iteration, the updated $\hat{\lambda}^1(t \mid A, W)$ now plays the role of the initial fit and the covariate $h(t, A, W)$ is then re-evaluated with the updated $\hat{S}^1(t \mid A, W)$ based on $\hat{\lambda}^1(t \mid A, W)$. In the third iteration $\hat{\lambda}^2(t|A, W)$ is fit and the procedure is iterated until $\hat{\epsilon}$ is essentially zero. The final hazard fit at the last iteration of the algorithm is denoted by $\hat{\lambda}^*(t \mid A, W)$ with the corresponding survival fit given by $\hat{S}^*(t \mid A, W)$.

As we are estimating two treatment specific parameters, we could either carry out the iterative updating procedure for each parameter separately or update the hazard fit simultaneously. To update the fit simultaneously, both covariates are added to the initial fit, i.e.,

$$\text{logit}(\hat{\lambda}(\epsilon)(t \mid A, W)) = \hat{\alpha}(t) + m(A, W|\hat{\beta}) + \epsilon_1 h_1(t, A, W) + \epsilon_2 h_0(t, A, W).$$

The iterative procedure is applied by now estimating two coefficients in each iteration as described above until both ϵ_1 and ϵ_2 are essentially zero.

Finally, the targeted maximum likelihood estimates of the probability of surviving past time t_0 for subjects in treatment arms 1 and 0 given by $\Psi_1(p_0)(t_0)$ and $\Psi_0(p_0)(t_0)$ are computed by,

$$\hat{\psi}_1^*(t_0) = \frac{1}{n}\sum_{i=1}^{n} \hat{S}^*(t_0 \mid 1, W_i).$$

and

$$\hat{\psi}_0^*(t_0) = \frac{1}{n}\sum_{i=1}^{n} \hat{S}^*(t_0 \mid 0, W_i).$$

4.4.1 Rationale for updating only initial hazard

The initial fit $\hat{p}^0$ of p_0 is identified by $\hat{\lambda}^0(t \mid A, W)$, $\hat{g}^0(A \mid W)$, $\hat{G}^0(t \mid A, W)$ and the marginal distribution of W. However the algorithm only updates $\hat{\lambda}^0(t \mid A, W)$. Assuming CAR the density of the observed data p factorizes in to the marginal distribution of W given by p_W, the treatment mechanism $g(A \mid W)$, the conditional probability of censoring up to time t given by $\bar{G}(t \mid A, W)$ and the product over time of the conditional hazard at $T = t$ given by $\lambda(t \mid A, W)$. This factorization implies the orthogonal decomposition of functions of O in the Hilbert space $L^2(p)$. We can thus apply this decomposition to the efficient influence curve $D(O \mid p)$. As shown in van der Laan and Robins [64], $D(O \mid p)$ is orthogonal to the tangent space $T_{CAR}(p)$ of the censoring and treatment mechanisms. Thus the components corresponding with $g(A \mid W)$ and $\bar{G}(t \mid A, W)$ are zero. This leaves the non zero components p_W and $\lambda(t \mid A, W)$. We choose the initial empirical distribution for W to estimate p_W which is the nonparametric maximum likelihood estimate for p_W and is therefore not updated. Thus the only element that does require updating is $\hat{\lambda}^0(t \mid A, W)$.

The efficient influence curve for $\Psi_1(p_0)(t_0)$ can be represented as,

$$\begin{aligned} D_1(p_0) = & \\ & \sum_{t<=t_0} h_1(g_0, G_0, S_0)(t, A, W)[I(\tilde{T} = t, \Delta = 1) - I(\tilde{T} >= t)\lambda_0(t \mid A = 1, W)] \\ & + S_0(t_0 \mid A = 1, W) - \Psi_1(p_0)(t_0), \end{aligned} \tag{4.12}$$

where $S_0(t_0 \mid A = 1, W)$ is a transformation of $\lambda_0(t \mid A = 1, W)$. This representation demonstrates the orthogonal decomposition described above. The empirical mean of the second component of $D_1(p_0)$ given by $S_0(t_0 \mid A = 1, W) - E_0 S_0(t_0 \mid A = 1, W)$ is always solved by using empirical distribution to estimate the marginal distribution of W. Thus, the TMLE solves this second component. The first component, the covariate times the residuals, is solved by performing the iterative targeted maximum likelihood algorithm with logistic regression fit of the discrete hazard $\lambda_0(t \mid A, W)$. We note that similarly, the efficient influence curve for $\Psi_0(p_0)(t_0)$ can be represented as,

$$\begin{aligned} D_0(p_0) = & \\ & \sum_{t<=t_0} h_0(g_0, G_0, S_0)(t \mid A, W)[I(\tilde{T} = t, \Delta = 1) - I(\tilde{T} >= t)\lambda_0(t \mid A = 0, W)] \\ & + S_0(t_0 \mid A = 0, W) - \Psi_0(p_0)(t_0). \end{aligned} \tag{4.13}$$

4.4.2 Double robustness properties

The targeted maximum likelihood estimate $\hat{p}^* \in \mathcal{M}$ of p_0 solves the efficient influence curve which is the optimal estimating equation for the parameter of interest. It can be shown that $E_0 D_1(p_0) = E_0 D_1(S, g, G) = 0$ if either $S = S_0(\cdot \mid A, W)$ (and thus $\lambda = \lambda_0(\cdot \mid A, W)$) or $g = g_0(A \mid W)$ and $G = G_0(\cdot \mid A, W)$. When the treatment is assigned completely at random as in an RCT, the treatment mechanism is known and $g(A \mid W) = g(A)$. Thus consistency of $\hat{\psi}_1^*(t_0)$ in an RCT relies on only consistent estimation of $\bar{G}_0(\cdot \mid A, W)$ or $S(\cdot \mid A, W)$. When there is no censoring or censoring is missing completely at random (MCAR), $\hat{\psi}_1^*(t_0)$ is consistent even when the estimator $\hat{S}(\cdot \mid A, W)$ of $S(\cdot \mid A, W)$ is inconsistent (e.g., if it relies on a mis-specified model). One is hence not concerned with estimation bias with this method in an RCT. Under informative or missing at random (MAR) censoring, if $\bar{G}_0(\cdot \mid A, W)$ is consistently estimated then $\hat{\psi}_1^*(t_0)$ is consistent even if $\hat{S}(\cdot \mid A, W)$ is mis-specified. If both are correctly specified then $\hat{\psi}_1^*(t_0)$ is efficient. These same statistical properties also hold for $\hat{\psi}_0^*(t_0)$.

4.4.3 Inference

Let $\hat{p}^*$ represent the targeted maximum likelihood estimate of p_0. One can construct a Wald-type 0.95-confidence interval for $\hat{\psi}_1^*(t_0)$ based on the estimate of the efficient influence curve $D_1(\hat{p}^*)(O)$ where $D_1(p)$ is given by (4.12). The asymptotic variance of $\sqrt{n}(\hat{\psi}_1^*(t_0) - \Psi_1(p_0)(t_0))$ can be estimated with

$$\hat{\sigma}^2 = \frac{1}{n} \sum_{i=1}^{n} D_1^2(\hat{p}^*)(O_i).$$

The corresponding asymptotically conservative Wald-type 0.95-confidence interval is defined as $\hat{\psi}_1^*(t_0) \pm 1.96 \frac{\hat{\sigma}}{\sqrt{n}}$. The null hypothesis $H_0 : \Psi_1(p_0)(t_0) = 0$ can be tested with the test statistic

$$T_n = \frac{\hat{\psi}_1^*(t_0)}{\frac{\hat{\sigma}}{\sqrt{n}}},$$

whose asymptotic distribution is $N(0, 1)$ under the null hypothesis. Similarly, confidence intervals and test statistics for $\Psi_0(p_0)(t_0)$ can be computed based on the estimate of the efficient influence curve $D_0(\hat{p}^*)(O)$ where $D_0(p)$ is given by (4.13).

If our parameter of interest is some function of the treatment specific survival estimates we can apply the δ-method to obtain the estimate of its influence curve. Specifically the estimated influence curve for the log hazard of survival, additive difference in survival, and relative risk of survival at t_0 provided in (4.3), (4.4), and (4.5) are respectively given by,

1. $\Psi_{HZ}(p_0)(t_0) : \frac{1}{\hat{\psi}_1^*(t_0) \log(\hat{\psi}_1^*(t_0))} D_1(\hat{p}^*)(O) - \frac{1}{\hat{\psi}_0^*(t_0) \log(\hat{\psi}_0^*(t_0))} D_0(\hat{p}^*)(O)$
2. $\Psi_{AD}(p_0)(t_0) : D_1(\hat{p}^*)(O) - D_0(\hat{p}^*)(O)$
3. $\Psi_{RR}(p_0)(t_0) : -\frac{1}{1-\hat{\psi}_1^*(t_0)} D_1(\hat{p}^*)(O) + \frac{1}{1-\hat{\psi}_0^*(t_0)} D_0(\hat{p}^*)(O)$

We can again compute confidence intervals and test statistics for these parameters using the estimated influence curve to estimate the asymptotic variance.

As an alternative to the influence curve based estimates of the asymptotic variance, one can obtain valid inference using the bootstrap procedure.

The inference provided in this section is for the estimates of the treatment effect for a single AE. For multiplicity adjustments for the analysis of a set of AE see section 4.9.

4.5 Simulation studies

The targeted maximum likelihood estimation procedure was applied to simulated data to illustrate the estimator's potential gains in efficiency. The conditions under which the greatest gains can be achieved over the standard unadjusted estimator were explored in addition to the estimators' performance in the presence of informative censoring.

4.5.1 Simulation protocol

Data were simulated to mimic an RCT in which the goal is to determine the safety of a new drug in comparison to the current standard of care on "survival" as measured by the occurrence of an adverse event at each time $t_0 \in \{2, ..., 10\}$. The probability of receiving the new treatment is 0.5. The covariate was negatively correlated with survival time, for example, this covariate might represent age in years (multiplied by 0.1). Specifically, 1000 replicates of sample size 300 were generated based on the following data generating distribution where time is discrete and takes values $t_0 \in \{1, ..., 10\}$:

- $Pr(A = 1) = Pr(A = 0) = 0.5$
- $W \sim U(0.2, 1.2)$
- $\lambda(t|A, W) = \frac{I(t_0<10)I(Y(t_0-1)=0)}{1+\exp(-(-3-A+\beta_W W^2))} + I(t_0 = 10)$
- $\lambda_C(t|A, W) = \frac{I(\Delta(t_0-1)=0)}{1+\exp(-(-\gamma_0-\gamma_1 A-\gamma_2 W))}$,

where $\lambda(t|A, W)$ is the hazard for survival and $\lambda_C(t|A, W)$ is the hazard for censoring. Two different data generating hazards for survival were applied corresponding with two values for β_W. These two values were set to $\beta_W \in \{1, 3\}$ corresponding with correlations between W and failure time of -0.22 and -0.63 respectively. We refer to the simulated data with $\beta_W = 1$ as the weak covariate setting and $\beta_W = 3$ as the strong covariate setting.

Three different types of censoring were simulated, no censoring, MCAR and MAR. Each type of censoring was applied to the weak and strong covariate settings for a total of six simulation scenarios. For both the weak and strong covariate settings, the MCAR an MAR censoring mechanisms were set such that approximately 33% of the observations were censored. The censoring was generated to ensure that $\bar{G}(t|A, W) > 0$ (see section 4.7 for details of this assumption). If censoring and failure time were tied, the subject was considered uncensored. For a summary of the simulation settings and the specific parameter values, see Table 4.1.

The difference in treatment-specific survival probabilities given by $\psi(t_0) = E_0(S_0(t_0|A = 1, W) - S_0(t_0|A = 0, W))$ was estimated at each time point $t_0 = 1$ through $t_0 = 9$. The unadjusted estimator is defined as the difference in the treatment specific Kaplan-Meier estimators at t_0. The TMLE was applied using two different initial hazard fits. The first initial hazard was correctly specified. The second initial hazard was mis-specified by including A and W as main terms and an interaction term between A and W. For both initial hazard fits, only time points 1 through 9 were included in the fit as the AE had occurred for all subjects by time point 10 and thus the hazard was one at $t_0 = 10$. In the MCAR censoring setting, the censoring mechanism was estimated using Kaplan-Meier. In the MAR censoring setting, the censoring mechanism was correctly specified. The update of the initial hazard was performed by adding to it the two covariates h_1 and h_0 provided in (4.10) and (4.11) respectively. The corresponding coefficients ϵ_1 and ϵ_2 were simultaneously estimated by fixing the offset from the

Table 4.1: Summary of simulation settings. *This table provides the parameter setting for the 6 simulation scenarios, including the correlation between W and failure time (Corr), the coefficients for the hazard for censoring ($\gamma = (\gamma_0, \gamma_1, \gamma_2)$) and the coefficient for W in the hazard for survival (β_W).*

Scenario	Censoring	γ	Corr W and T	β_W
1	No censoring	NA	-0.22 (Weak)	1
2	MCAR	(-2.7,0,0)	-0.22 (Weak)	1
3	MAR	(-1.75,1.35,-2.5)	-0.22 (Weak)	1
4	No censoring	NA	-0.65 (Strong)	3
5	MCAR	(-2,0,0)	-0.65 (Strong)	3
6	MAR	(-1.15,0.5,-2)	-0.65 (Strong)	3

initial fit and setting the intercept to 0. The procedure was iterated until ϵ_1 and ϵ_2 were sufficiently close to zero.

The estimators were compared using a relative efficiency 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 one indicates a gain in efficiency of the covariate adjusted TMLE over the unadjusted estimator.

In addition to these six simulation scenarios, to explore the relationship between relative efficiency and the correlation between the covariate and failure time, we generated data by varying β_W in the data generating distribution above for six values, $\beta_W \in \{0.5, 1, 1.5, 2, 2.5, 3\}$ corresponding with correlations between W and failure time of {-0.10,-0.22,-0.36,-0.46,-0.56,-0.63} under no censoring. The parameter $\psi(5)$ was estimated based on 1000 sampled datasets with sample size $n = 300$.

4.5.2 Simulation results and discussion

Strong covariate setting

In the no censoring and MCAR censoring scenarios, the bias should be approximately zero. Thus, the relative MSE is essentially comparing the variance of the unadjusted and targeted maximum likelihood estimates. Any gain in the MSE can therefore be attributed to a reduction in variance due to the covariate adjustment. 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 covariates h_1 and h_0 that are inverse weighted by the subjects' conditional probability of being observed at time t given their observed history.

Figure 4.1 provides the relative MSE results for $\hat{\psi}(t_0)$ for $t_0 \in \{1, ...9\}$ for the strong covariate setting with $\beta_W = 3$. Based on these results, we observe that indeed the expected gain in efficiency is achieved. The minimum observed relative MSE was 1.25 for $t_0 = 1$ in the MAR censoring setting with a misspecified initial hazard fit. A maximum relative MSE of 1.9 is observed under the no censoring setting with the correctly specified initial hazard at $t_0 = 3$. The approximate overall average relative MSE was 1.6 for the no censoring scenario. Consistently across all time points and censoring scenarios, the TMLE is outperforming the unadjusted estimator.

Figure 4.2 provides the bias as a percent of the truth for the two estimators under the MAR censoring setting with the correctly specified initial hazard. Clearly as t_0 increases, the bias of the unadjusted

estimates increases whereas the targeted maximum likelihood estimates is relatively close to zero in comparison. Thus the targeted maximum likelihood approach can not only provide gains in efficiency through covariate adjustment, but can also account for informative censoring as well.

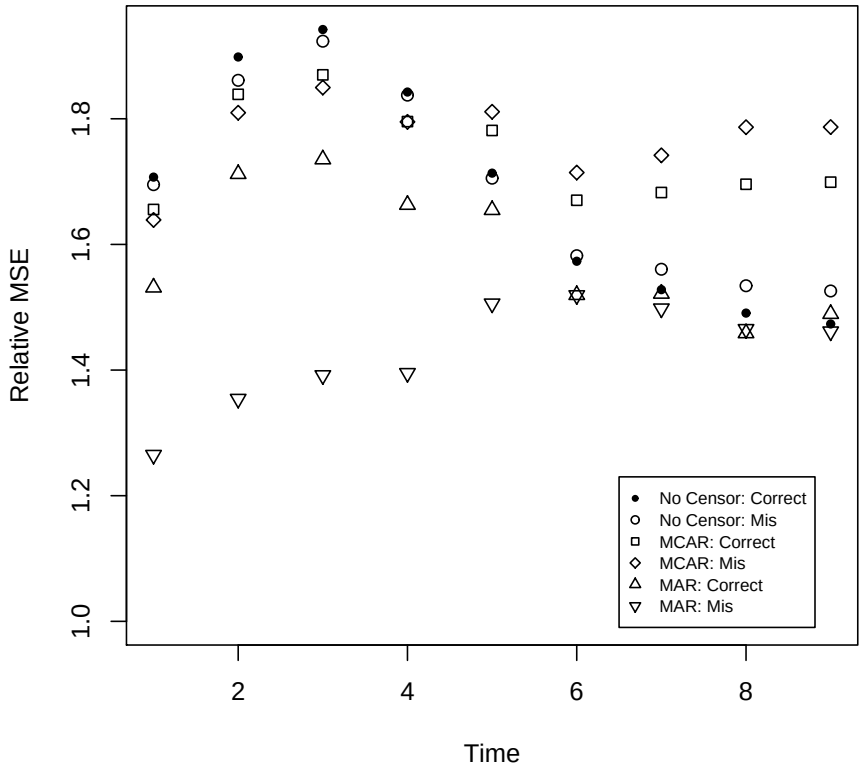

Figure 4.1: Comparison of relative MSE for the strong covariate setting ($\beta_W = 3$). *This table provides the ratio of the MSE of the Kaplan-Meier (KM) to the MSE of the TMLE for the three censoring scenarios: no censoring, MCAR and MAR.*

Weak covariate setting

In this weak covariate setting, again in the no censoring and MCAR censoring scenarios, the bias should essentially be zero. However, we expect a lesser gain in efficiency if any as compared to the strong covariate setting since the covariate in this setting is not as useful for hazard prediction. We do again expect a bias reduction in the MAR censoring setting for the TMLE over the unadjusted estimator.

Figure 4.3 provides the relative MSE results for the weak correlation simulation with $\beta_W = 1$. As expected, the relative MSE are all close to one indicating that only small efficiency gains are achieved when only weak covariates are present in the data. However, as small the gains are they are also achieved across all time points as in the strong covariate setting. Regardless of the correlation between the covariate and failure time, in the informative censoring scenario the targeted maximum likelihood estimate is consistent under consistent estimation of the censoring mechanism as evidenced in the plot of the % bias in Figure 4.4.

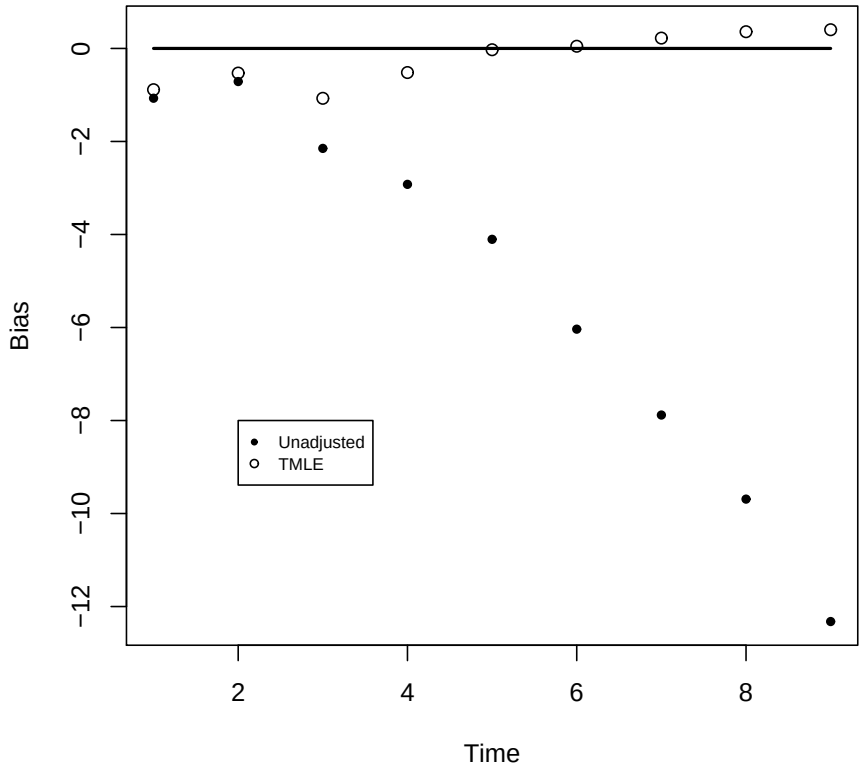

Figure 4.2: Comparison of bias for the strong covariate setting ($\beta_W = 3$) under the MAR censoring scenario. *This plot shows the bias of the Kaplan-Meier estimate (Unadjusted) and the TMLE.*

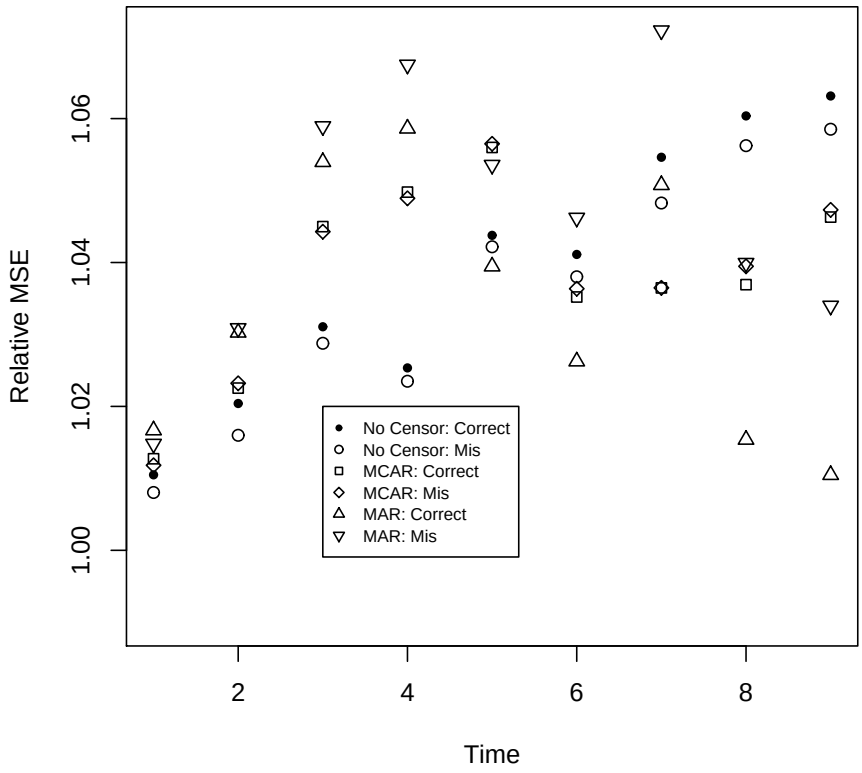

Figure 4.3: Comparison of relative MSE for the weak covariate setting ($\beta_W = 1$). *This table provides the ratio of the MSE of the Kaplan-Meier (KM) to the MSE of the TMLE for the three censoring scenarios: no censoring, MCAR and MAR.*

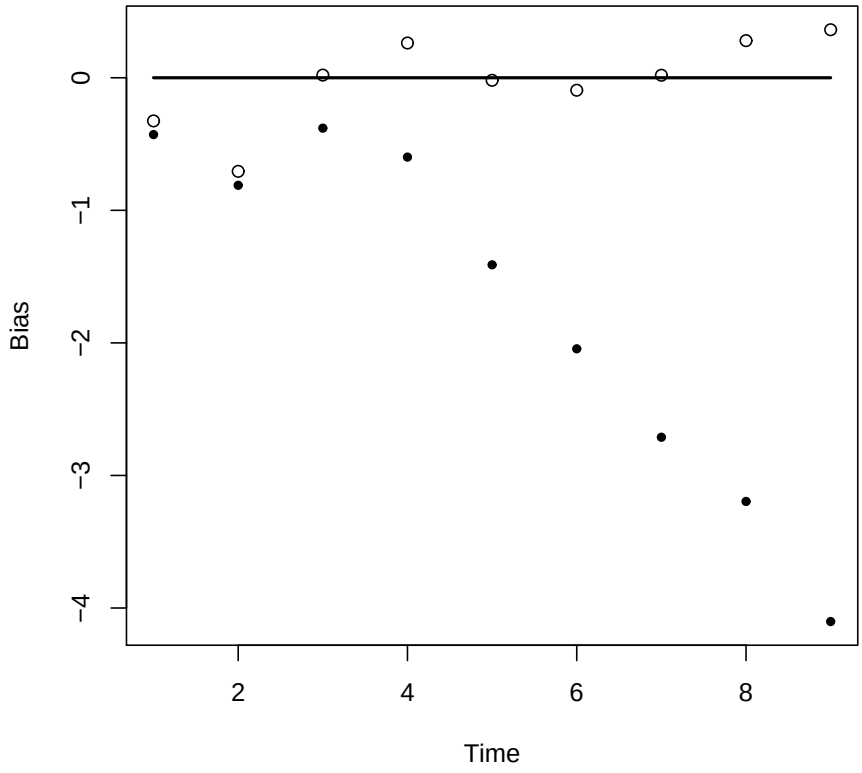

Figure 4.4: Comparison of bias for the weak covariate setting ($\beta_W = 1$) under the MAR censoring scenario. *This plot shows the bias of the Kaplan-Meier estimate (Unadjusted) and the TMLE.*

4.5.3 Relationship between correlation of covariate(s) and failure time with efficiency gain

As the correlation between W and failure time increases we expect to observe increasing gains in efficiency. Selecting an arbitrarily selected time point $t_0 = 5$ for ease of presentation, Figure 4.5 clearly demonstrates that as the correlation between W and failure time increases so does the relative MSE. In fact in for this particular data generating distribution, at time $t_0 = 5$ the relationship is nearly linear. These results reflect similar findings in RCT with fixed end point studies where relations between R^2 and efficiency gain have been demonstrated [36, 43]. This relationship indicates that if indeed the particular dataset contains covariates that are predictive of the failure time of the AE of interest, one can achieve gains in precision and thus power by using the TMLE.

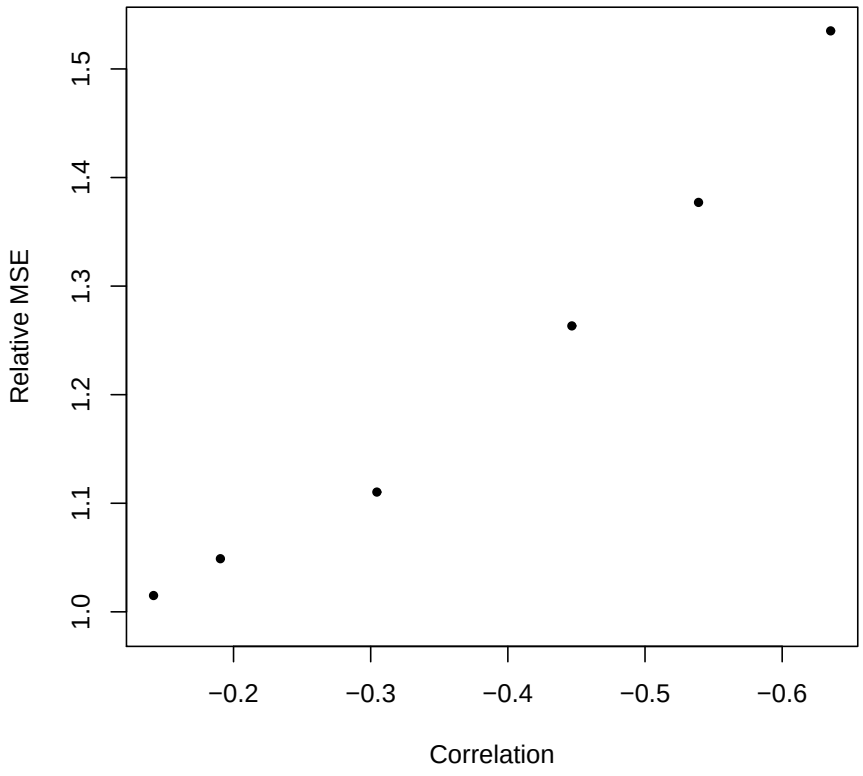

Figure 4.5: Relation between efficiency gain and correlation between covariate and failure time.

4.6 Fitting initial hazard and censoring mechanism

Despite these potential gains in efficiency as demonstrated by theory and simulation results, there has been concern with covariate adjustment in RCT with respect to investigators selecting covariates to obtain favorable inference. We conjecture that such cheating can be avoided if one uses an *a priori* specified algorithm for model selection. When the model selection procedure is specified in an analysis protocol, the analysis is protected from investigators guiding causal inferences based on selection of favorable covariates and their functional forms in a parametric model. In safety analysis, if investigator (sponsor) bias does indeed exist, it would be reasonable to assume that it would lean towards the treatment having no effect on the AE and thus the concerns are the reverse from efficacy analysis. The investigator bias would tend towards the less efficient unadjusted estimator. The analysis of AEs is often exploratory in nature and the results are meant to flag potential AE of concern which may reduce the motivation for dishonest inference using covariate adjustment. Regardless of the covariate selection strategy, it should be explicitly outlined to avoid any such concerns.

There are a number of model selection algorithms that can be applied to data-adaptively select the initial hazard fit. One such approach is the D/S/A algorithm that searches through a large space of functional forms using deletion, substitution and addition moves. One can apply this algorithm to the pooled data (over time) to fit the initial hazard [56]. One can also fit hazards using the hazard regression (HARE) algorithm developed by Kooperberg et al. [30], which uses piecewise linear regression splines and adaptively selects the covariates an knots. As another alternative, one could also include all covariates that have a strong univariate association with failure time in a hazard fit as main terms in addition to the treatment variable. Since one is often investigating many AE, a fast algorithm such as the latter may be an appropriate alternative for computational efficiency.

We also note that if weights are required as they are for the inverse probability of censoring weighted (IPCW) reduced data TMLEs as outlined in section 4.10.1, the D/S/A algorithm can be run with the corresponding weights.

In addition to the hazard for survival, the hazard for censoring must also be estimated. One of the algorithms discussed above can also be applied to estimate the censoring mechanism. We note that the application of the TMLE to a set of M AE requires M hazard fits whereas only one fit for censoring is required. Thus, the censoring mechanism is estimated once and for all and is used in the analysis of each of the M AE.

4.7 Inverse weighting assumption

The TMLE, as well as other inverse weighted estimators (see section 4.8) for the parameters presented in this chapter rely on the assumption that each subject has a positive probability of being observed (i.e., not censored) at time t, which can be expressed by,

$$\bar{G}(t_- \mid A, W) > 0, t = t_0.$$

This identifiability assumption has been addressed as an important assumption for right-censored data [51]. In Neugebauer and van der Laan [38] it was demonstrated that practical violations of this assumption can result in severely variable and biased estimates.

One is alerted of 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.

4.8 Alternative estimators

Prior to the introduction of targeted maximum likelihood estimation, there were two main approaches to estimating the treatment specific survival at a fixed end point t_0: maximum likelihood estimation and estimating function estimation. In the maximum likelihood approach, one obtains an estimate $\hat{p}$ for p identified by perhaps a Cox proportional hazards model for continuous survival or logistic regression for discrete survival. The parameter of interest is then evaluated via substitution, i.e., $\hat{\psi} = \psi(\hat{p})$. These maximum likelihood substitution estimators involve estimating some hazard fit using an *a priori* specified model or a model selection algorithm that is concerned with performing well with respect to the whole density rather than the actual parameter of interest, e.g., the difference in treatment specific survival at a specific time t_0. These type of estimators often have poor performance and can be heavily biased whenever the estimated hazard is inconsistent [49]. Furthermore, inference for such maximum likelihood estimators that rely on parametric models are overly optimistic and thus their corresponding p-values are particularly unreliable. This is in contrast to the inference for the TMLEs which respects that no *a priori* models are required.

An alternative to the likelihood based approach is the extensively studied estimating function based approach. Recall that the full data estimating functions provided in (4.8) and (4.9) are estimating functions that could be applied to estimate the treatment specific survival at time t_0 if we had access to the full data, i.e., the uncensored survival time. The full data estimating function can be mapped into a an estimating function based on the observed data using the IPCW method. The IPCW estimators based on the IPCW estimating function denoted by $D^{IPCW}(T, A, W \mid \psi_1, g, G)$ have been shown to be consistent and asymptotically linear if the censoring mechanism G can be well approximated [52, 64]. While the IPCW estimators have advantages such as simple implementation, they are not optimal in terms of robustness and efficiency. Their consistency relies on correct estimation of the censoring mechanism whereas maximum likelihood estimators rely on correct estimation of the full likelihood of the data.

The efficient influence curve can be obtained by subtracting from the IPCW estimation function the IPCW projection onto the tangent space T_{CAR} of scores of the nuisance parameter G [64]. The efficient influence curve is the optimal estimating function in terms of efficiency and robustness and the corresponding solution to this equation is the so-called double robust IPCW (DR-IPCW) estimator. The "double" robust properties of this estimator are equivalent to those of the TMLE as the TMLE solves the efficient influence curve estimating equation, see section 4.4.2. Despite the advantageous properties of such efficient estimating function based estimators, maximum likelihood based estimators are much more common in practice.

The more recently introduced targeted maximum likelihood estimation methodology that was applied in this chapter can be viewed as a fusion between the likelihood and estimating function based methods. A notable advantage of the TMLEs is their relative ease of implementation in comparison to estimating equations which are often difficult to solve.

4.9 Multiple testing considerations

An important consideration in safety analysis is multiple testing in that often as many as hundreds of AE are collected. The ICH guidelines indicate that it is recommended to adjust for multiplicity when hypothesis tests are applied [24]. However, the ICH guidelines do not provide any specific methods for adjustment. The need for adjustment is demonstrated by the following example outlined in Kaplan et al. [28]. In this study, out of 92 safety comparisons the investigators found a single significant

result according to unadjusted p-values. A larger hypothesis driven study for this AE that had no known clinical explanation was carried out and did not result in any significant findings. Such false positive results for testing the effect of treatment on a series of AE based on unadjusted p-values can cause undue concern for approval/labeling and can affect post-marketing commitments. On the other hand, over adjusting could also result in missing potentially relevant AE. Thus appropriate adjustment requires some balance between no adjustment and a highly stringent procedure such as Bonferroni.

Many advances have been made in the area of multiple testing over the Bonferroni-type methods including resampling based methods to control the familywise error rate (FWER), for example see van der Laan et al. [68], and the Benjamini-Hochberg method for controlling the false discovery rate (FDR) [7]. With FWER approaches, one is concerned with controlling the probability of erroneously rejecting one or more of the true null hypotheses, whereas the FDR approach controls the expected proportion of erroneous rejections among all rejections. The resampling based FWER method makes use of the correlation of test statistics which can provide a gain in power over assuming independence. However, the Benjamini-Hochberg FDR approach has been shown to be perform well with correlated test statistics as well [8]. The selection of the appropriate adjustment depends on whether or not a more conservative approach is reasonable. In safety analysis, one certainly does not want to miss flagging an important AE and thus might lean towards an FDR approach.

FDR methods have been proposed specifically in the analysis of AE in Mehrotra and Heyse [34]. Their method involves a two-step procedure that groups AE by body system and performs an FDR adjustment both within and across the body system. Presumably this method attempts to account for the dependency of the AE by grouping in this manner. Thus the multiple testing considerations and the dependency of the test statistics in safety analysis has indeed received some attention in literature.

The multiple testing adjustment procedure to be applied in the safety analysis should be provided in the study protocol to avoid potential for dishonest inference. In addition, the unadjusted p-values should continue to be reported with the adjusted p-values so all AE can be evaluated to assess their potential clinical relevance.

4.10 Extensions

4.10.1 Time-dependent covariates

It is not unlikely that many time-dependent measurements are collected at each follow-up visit in addition to the many AE and efficacy outcome measurements. Such time-dependent covariates are often predictive of censoring. The efficiency and robustness results presented in this chapter have been based on data structures with baseline covariates only. The targeted maximum likelihood estimation procedure for data structures with time-dependent covariates is more complex as demonstrated in van der Laan [61]. To overcome this issue and avoid modeling the full likelihood, van der Laan [61] introduced IPCW reduced data TMLEs. We provide only an informal description of this procedure here, for details we refer readers to the formal presentation provided in van der Laan [61].

In this framework, the targeted maximum likelihood estimation procedure is carried out for a reduced data structure X^r, which in this case is the data structure that only includes baseline covariates. The IPCW reduced data procedure differs from the procedure where X^r is the full data in that the log-likelihoods are weighted by a time-dependent stabilizing weight given by,

$$sw(t) = \frac{I(C > t)\bar{G}^r(t \mid X^r)}{\bar{G}(t \mid X)}.$$

This stabilizing weight is based on $\bar{G}^r(t \mid X^r)$ which is the censoring mechanism based on the reduced data structure that includes baseline covariates only and $\bar{G}(t \mid X)$ which is the censoring mechanism based on the complete data structure that includes time-dependent covariates.

In practice in estimation of the parameter $\psi(t_0) = E_0(S_0(t_0|A=1,W) - S_0(t_0|A=0,W))$, one must apply these weights anytime maximum likelihood estimation is performed. Thus, the IPCW reduced data targeted maximum likelihood estimation procedure differs from the standard targeted maximum likelihood procedure provided in section 4.4 in that each time the conditional hazard is fit it is weighted by $sw(t)$. These weights are time-specific and thus each subject receives a different weight at each point in time. The initial hazard estimate $\hat{\lambda}^0(t \mid A, W)$ is weighted by $sw(t)$. The algorithm then updates $\hat{\lambda}^0(t \mid A, W)$ by adding the time-dependent covariates $h_1(t, A, W$ and $h_0(t, A, W)$ and estimating their corresponding coefficients ϵ_1 and ϵ_2. In the IPCW reduced data targeted maximum likelihood estimation procedure one includes the weights $sw(t)$ in estimation of ϵ_1 and ϵ_2. These weights are applied in each iteration of the algorithm to obtain the final fit $\hat{\lambda}^*(t \mid A, W)$ that is achieved when $\hat{\epsilon}_1$ and $\hat{\epsilon}_2$ are sufficiently close to zero. Thus estimation can again be achieved using standard software with the only additional requirement of weighting each of the regressions by these time-dependent weights.

Estimation of these time-dependent weights requires estimation of $\bar{G}^r(t \mid X)$ and $\bar{G}(t \mid X)$. Model selection algorithms that can be applied to estimate $\bar{G}^r(t \mid X)$ were described in section 4.6. Similarly the censoring mechanism $\bar{G}(t \mid X)$ can be estimated using a Cox proportional hazards model with time-dependent covariates for continuous censoring times or logistic regression model with time dependent covariates for discrete censoring times. Model selection algorithms such as those described in section 4.6 can also be applied by including these time-dependent covariates as candidates.

Let $\hat{\psi}^r(t_0)$ represent the IPCW reduced data TMLE of $\psi(t_0)$. By applying this IPCW weighting in the reduced data targeted maximum likelihood estimation procedure a particular type of double robustness is obtained. If there are no time-dependent covariates that are predictive of censoring time, then the ratio of estimated survival probabilities of censoring in the above weight $sw(t)$ is one. In this case, if $\bar{G}(t \mid X)$ is consistently estimated or $\lambda(\cdot \mid A, W)$ is consistently estimated then $\hat{\psi}^r(t_0)$ is consistent; if both are consistent then it is even more efficient than the estimator that was based on the reduced data structure. If there are indeed time-dependent covariates that are predictive of censoring time, and $\bar{G}(t \mid A, W)$ is well approximated then $\hat{\psi}^r(t_0)$ is consistent and the desired bias reduction is achieved.

4.10.2 Post market data

As RCT are powered for efficacy, it is often the case that many AE are either not observed at all during the pre-market phase or so few are observed that statistically conclusive results are often exceptions [41]. In an RCT of a rotavirus vaccine in which the AE of intussusception among vaccine recipients compared to controls was not found to be statistically significant. After the vaccine was approved and had been widely used, an association between this AE and the vaccine was found and it was pulled off the market. A subsequent analysis demonstrated that to obtain power of 50% to detect a difference as small as the actual observed Phase III incidence of the AE, a sample size of approximately 90,000 would be required (6 times the actual sample size) [25]. Due to the high cost and complications involved in running an RCT, such large sample sizes are not feasible.

It is not only the rarity of many AE that causes issues in detection during RCT, but also the fact that RCT may have restrictive inclusion criteria whereas the drug is likely applied to a less restrictive population in post-market. Furthermore, the follow-up time in the pre-market phase may not be long enough to detect delayed AE. For a discussion regarding the difficulties in "proving" safety of a compound in general see Bross [10]. Post-market monitoring is therefore an important aspect of safety analysis.

There are a number of types of post-market data (for a thorough description of the various types of post-market data see Glasser et al. [19]) including spontaneous adverse event reporting systems (e.g., "MedWatch"). These data can be useful for detecting potentially new or unexpected adverse drug reactions that require further analysis however they often suffer from under-reporting by as much as a factor of 20 [12].

In this section, we focus on observational post-market studies or pharmacoepidemiological studies. Since patients in these type of studies are not randomized to a drug versus placebo (or competitor), confounding is typically present. Of particular concern is the fact that sicker patients are often selected to receive one particular drug versus another. There exists a vast amount of literature for controlling for confounding in epidemiological studies. Popular methods in pharmacoepidemiology include propensity score (PS) methods and regression based approaches. However, consistency with these methods rely on correct specification of the PS or the regression model used. Furthermore, it is not clear how informative censoring is accounted for with these methods. The TMLEs are double robust and are thus more advantageous than these commonly applied alternative approaches.

Before we proceed with discussion of estimation of causal effects with observational data, we first outline the data and assumptions. Suppose we observe n independent and identically distributed copies of $O = (\tilde{T}, \Delta, A, W) \sim p_0$ as defined in section 4.3. Causal effects are based on a hypothetical full data structure $X = (T_{1,1}, T_{1,0}, T_{0,1}, T_{0,0}, W)$ which is a collection of action specific survival times where this action is comprised of treatment and censoring. Note that we are only interested in the counterfactuals under this joint action-mechanism that consists of both censoring and treatment mechanisms where censoring equals zero, i.e., $T_{1,0}$ and $T_{0,0}$. In other words, we aim to investigate what would have happened under each treatment had censoring not occurred.

The consistency assumption states that the observed data consist of the counterfactual outcome corresponding with the joint action actually observed. The coarsening at random (CAR) assumption implies that the joint action is conditionally independent of the full data X given the observed data. We denote the conditional probability distribution of treatment A by $g_0(a \mid X) \equiv P(A = a \mid X)$. In observational studies, CAR implies $g_0(A \mid X) = g_0(A \mid W)$, in contrast to RCT in which treatment is assigned completely at random and $g_0(A \mid X) = g_0(A)$.

We aim to estimate $\psi(t_0) = E_0(S_0(t_0|A = 1, W) - S_0(t_0|A = 0, W)) = Pr(T_{1,0} > t_0) - Pr(T_{0,0} > t_0)$. Even under no censoring or MCAR, we are can no longer rely on the unadjusted treatment specific Kaplan-Meier estimates being unbiased due to confounding of treatment.

Under the assumptions above, the TMLE for $\psi(t_0)$ is double robust and locally efficient. Thus the targeted maximum likelihood estimation procedure described in this chapter is theoretically optimal in terms of robustness and efficiency. In our presentation, we assumed that treatment was assigned at random. In observational studies, in addition to estimating $\lambda(\cdot \mid A, W)$ and possibly $\bar{G}(\cdot \mid A, W)$ (when censoring is present), observational studies require estimation of the treatment mechanism $g(A \mid W)$ as well. It has been demonstrated that when censoring is MCAR in an RCT, the targeted maximum likelihood estimate $\hat{\psi}^*(t_0)$ is consistent under mis-specification of $\lambda(\cdot \mid A, W)$ since $g(A \mid W)$ is always correctly specified. However, even under MCAR, in observational studies, consistency of $\hat{\psi}^*(t_0)$ relies on consistent estimation of $\lambda(\cdot \mid A, W)$ or $g(A \mid W)$ and is efficient if both are consistently estimated [65]. When censoring is MAR, then consistency of $\hat{\psi}^*(t_0)$ also relies on consistent estimation of the joint missingness $g(A \mid W)$ and $\bar{G}(\cdot \mid A, W)$ or $\lambda(\cdot \mid A, W)$.

We also note that the TMLEs as well as the commonly applied PS methods rely on the experimental treatment assignment (ETA) assumption. Under this assumption, each patient must have a positive probability of receiving each treatment. The inverse weighted PS estimator is known to suffer severely from violations of this assumption in practice [38, 51, 71]. This poor performance is evident with inverse weighting, however we note that all other PS methods rely on this assumption as well, but are

not as sensitive to practical violations. This assumption is essentially about information in the data and violations of it indicate that for certain strata of the data, a given treatment level is never or rarely experienced. When the ETA is violated estimation methods rely on extrapolation.

If it is the case that a given treatment level is very rare or non-existent for given strata of the population, an investigator may want to re-consider the original research question of interest. To this end, van der Laan and Petersen [63] developed causal effect models for realistic intervention rules. These models allow estimation of the effect of realistic interventions, that is only intervening on patients for whom the intervention is reasonably "possible" where "possible" is defined by $g(A \mid W)$ greater than some value, e.g., 0.05. We note that targeted maximum likelihood estimation can be applied to estimate parameters from such models. For applications of such models see Bembom and van der Laan [6].

The ETA assumption and development of realistic causal models are simply examples of some of the many considerations that arise with observational data as compared to RCT data. However despite the many issues the rich field of causal inference provides promising methods for safety analysis in post-market data. As it is not possible to observe all AE in the pre-market phase, post-market safety analysis is an important and emerging area of research.

4.11 Discussion

Safety analysis is an important aspect in new drug approvals and has become increasingly evident with the recent cases of drugs withdrawn from the market (e.g., Vioxx). Increasing estimation efficiency is one area that can help overcome the issue that RCT are not powered for safety. Using covariate information is a promising approach to help detect AE that may have remained undetected with the standard crude analysis. Furthermore, time-to-event methods for AE analysis may be more appropriate particularly in studies where the AE often occur for all patients, such as oncology studies. Exploiting the time-dependent nature can further provide more efficient estimates for the effect of treatment on AE occurrence.

In this chapter we provided a method for covariate adjustment in RCT for estimating the effect of treatment on the AE failing to occur by a fixed end point. The method does not require any parametric modeling assumptions under MCAR censoring and thus is robust to mis-specification of the hazard fit. The methods advantages were twofold. The first is the potential efficiency gains over the unadjusted estimator. The second is that the TMLE accounts for informative censoring through inverse weighting of the covariate(s) that is added to an initial hazard fit. The standard unadjusted estimator is biased in the informative censoring setting.

The estimator has a relatively straightforward implementation. Given an initial hazard fit either logistic for discrete failure times or Cox proportional hazards for continuous survival times, one updates this fit by iteratively adding a time dependent covariate(s).

The simulation study demonstrated the potential gains in efficiency that can be achieved in addition to the relation of the correlation between the covariate(s) and failure time and efficiency gains. When no predictive covariates were present the relative efficiency was approximately one indicating that one is protected from actually losing precision from applying this method even when the covariates provide little information about failure time. The simulations also demonstrated the reduction in bias in the informative censoring setting.

Considerations for balancing the potential for false positives and the danger of missing possibly significant AE are an important aspect of safety analysis. The strategies from the rich field of multiple testing briefly discussed in this chapter can exploit the correlation of the AE outcomes and thus provide the

most powerful tests.

Chapter 5

Covariate adjusted analogue to logrank test

5.1 Introduction

Covariate adjustment in RCTs has been demonstrated to improve estimation efficiency over standard unadjusted methods in studies with continuous or binary outcomes at fixed end-points [29, 36, 43, 59, 73]. 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 [64], 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 RCT, data are additionally collected on baseline (pre-treatment) 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.

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 [33] 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 [57] and explored via simulation studies in Jiang et al. [27]. 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. [1]. 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 paper we develop covariate adjusted analogues to the logrank test in RCT. 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_0, relative to the cumulative hazard of controls at time t_0 , averaged over many time points t_0. However, we note that if the time scale is sufficiently fine, that this methodology is compatible with continuous survival outcomes.

We present our methods for covariate adjustment under the framework of targeted maximum likelihood estimation, originally introduced in van der Laan and Rubin [65]. 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 $\hat{p}^*$ 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 $\psi(\hat{p}^*)$ 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 [64].

There are several advantages to this methodology over estimating equation methodology as discussed in chapters 2 and 4. 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. Thus, estimating equation methodology fails for this particular parameter. As a result, the proof of the double robustness consistency properties does not follow in the usual obvious manner. Therefore, in this paper, 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. Although we present the conjecture that the first estimator is less efficient than the second and provide evidence of this fact through simulation studies, for the first estimator we can prove these double robustness properties based on theory, contrary to the latter estimator. Therefore, we include both of these covariate adjusted TMLEs of the logrank parameter.

The likelihood of the observed data (provided in section 5.5) 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 paper 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. Thus, the conditional hazard on which the TMLE is based is not meant to infer information about subgroup (conditional) effects of treatment. The methodology presented in this chapter can be extended to the

estimation of subgroup specific logrank analogues, however, in this paper we focus only on marginal (unconditional) treatment effects.

Specifically, the paper is outlined as follows. We first begin with a brief overview of the data, model and parameter(s) of interest in section 5.2. We then provide the discrete analogue to the logrank test in section 5.3, to which we compare the targeted maximum likelihood estimation approaches. We then review the methodology for estimating the treatment specific estimates of survival at a fixed endpoint t_0 as presented in chapter 4. This t_0 approach is then extended to provide our first of two analogues to the logrank test, the substitution based targeted maximum likelihood method (section 5.5). In the second analogue, we provide the direct targeted maximum likelihood approach which does not require estimation of the t_0-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 5.8. In section 5.9 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. Finally, we conclude with a discussion.

5.2 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 can take values $\{1, ..., K\}$. 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 the density of O. The conditional hazard is given by $\lambda_0(\cdot \mid A, W)$ and the corresponding conditional survival is given by $S_0(\cdot \mid A, W)$. The censoring mechanism is given by $\bar{G}(t_- \mid A, W) = P(C \geq t \mid A, W)$.

Let T_1 represent a patient's time to the occurrence of an event had she, possibly contrary to fact, been assigned to the treatment group and let T_0 likewise represent the time to the occurrence of the event had the patient been assigned to the control group.

Chapter 4 presented the targeted maximum likelihood estimation method for the estimation of the t_0 and treatment specific parameters,

$$P_0 \to \Psi_1(p_0)(t_0) = Pr(T_1 > t_0) = E_0(S_0(t_0 \mid A = 1, W)) = S_1(t_0), \tag{5.1}$$

and

$$P_0 \to \Psi_0(p_0)(t_0) = Pr(T_0 > t_0) = E_0(S_0(t_0 \mid A = 0, W)) = S_0(t_0), . \tag{5.2}$$

where the subscript for Ψ denotes the treatment group, either 0 or 1. 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,

$$P_0 \to \Psi_{t_0}(p_0) = \log\left(\frac{\log(Pr(T_1 > t_0))}{\log(Pr(T_0 > t_0))}\right) = \log\left(\frac{\log(S_1(t_0))}{\log(S_0(t_0))}\right). \tag{5.3}$$

In this paper, we are interested not in a test for the effect of treatment at a fixed end point t_0, but rather the average effect over time. Note that, in the continuous survival case, if one averaged $\Psi_{t_0}(p_0)$ 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,

$$\lambda(t \mid A) = \lambda(t) \exp(\psi_C A).$$

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

$$P_0 \to \Psi(p_0) = \sum_{t_0} w(t_0) \log \left(\frac{\log(S_1(t_0))}{\log(S_0(t_0))} \right), \tag{5.4}$$

for some weight function $w(t_0)$ which we discuss in section 5.7. 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 $\psi_0 = \Psi(p_0)$.

We note that we could have chosen a number of other parameters to evaluate the average effect of treatment on survival over time, such as the difference given by,

$$\Psi_D(p_0) = \sum_{t_0} w(t_0) \left(S_1(t_0) - S_0(t_0) \right),$$

or the log of the relative risk of survival, given by

$$\Psi_{RR}(p_0) = \sum_{t_0} w(t_0) \log \left(\frac{S_1(t_0)}{S_0(t_0)} \right).$$

However, we focus on $\Psi(p_0)$ as defined in (5.4) so that we can compare the power of the targeted maximum likelihood test to that of the discrete analogue to the logrank test as outlined in the next section.

5.3 Unadjusted estimation of $\Psi(p_0)$

The discrete extension to the Cox proportional hazards model is a model for the odds of dying at t, given survival up to time t, given by,

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

where

$$\log \left(\frac{\lambda(t \mid A)}{1 - \lambda(t \mid A)} \right) = \beta_1 I(t = 1) + \beta_2 I(t = 2) + ... + \beta_K (t = K) + \beta_A A. \tag{5.5}$$

Thus, $\beta_1, ..., \beta_K$ 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 nonparameteric 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,

$$L(O_r) = \prod_{i=1}^{n} Pr(T_i = \tilde{t}_i)^{\Delta_i} Pr(T_i > \tilde{t}_i)^{(1-\Delta_i)} \tag{5.6}$$

$$= \prod_{i=1}^{n} \left(\lambda_i(\tilde{t}_i) \prod_{t=1}^{\tilde{t}_i - 1} (1 - \lambda_i(t)) \right)^{\Delta_i} \left(\prod_{t=1}^{\tilde{t}_i} (1 - \lambda_i(t)) \right)^{1-\Delta_i}, \tag{5.7}$$

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 $\bar{y}_i = (y_{i1}, ..., y_{i\tilde{t}_i})$ denote the event history for individual i where $(y_{i1}, ..., y_{i\tilde{t}_i - 1}) = (0, ..., 0)$, 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,

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

Note that this likelihood is equivalent to that of a sequence of independent Bernouilli trials and thus we can use standard logistic regression software to obtain the maximum likelihood estimates for the coefficients β in (5.5).

In practice, the logistic regression model is fit with the dataset that includes repeated measures for each subject up until the time that 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 (5.4). Thus, we use the logistic regression fit for the hazard, denoted by $\hat{\lambda}(t \mid A)$, to obtain estimates for $\hat{\lambda}_1(t) = \hat{\lambda}(t \mid A = 1)$ and $\hat{\lambda}_0(t) = \hat{\lambda}(t \mid A = 0)$. Based on these estimates, we use the relation,

$$S(t_0) = \prod_{j \leq t_0} (1 - \lambda(j)),$$

to obtain estimates $\hat{S}_1(t_0)$ and $\hat{S}_0(t_0)$. 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_0-specific estimates. We note that, alternatively, one could estimate $S_1(t_0)$ and $S_0(t_0)$ for $t_0 \in 1, ..., K$ using Kaplan-Meier and use these estimates as plug-ins into (5.4). The use of Kaplan-Meier would be more nonparametric since the proportional odds model assumes proportionality of the hazards. 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 paper, we use the proportional odds approach only.

5.4 General approach to targeted maximum likelihood estimation of $\Psi(p_0)$

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

$$P(W)g(A \mid W)\left[\lambda(\tilde{t} \mid A,W)\prod_{t=1}^{\tilde{t}-1}(1-\lambda(t \mid A,W))\right]^{\delta}\left[\prod_{t=1}^{\tilde{t}}(1-\lambda(t \mid A,W))\right]^{1-\delta} \times \left[\bar{G}(\tilde{t}_{-} \mid A,W)^{\delta}P(C=\tilde{t} \mid A,W)^{1-\delta}\right], \quad (5.8)$$

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

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

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

$$\text{logit}\left(\hat{\lambda}^0(t \mid A,W)\right) = m(t,A,W),$$

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,

$$\text{logit}\left(\hat{\lambda}^0(t \mid A,W)\right) = \hat{\alpha}(t) + k(A,W \mid \hat{\beta}),$$

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.,

$$\text{logit}(\hat{\lambda}^0(\epsilon)(t \mid A,W)) = m(t,A,W) + \epsilon h(t,A,W). \quad (5.9)$$

Now, ϵ is estimated by fitting (5.9) 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 $\hat{\epsilon}h(t,A,W)$. The covariate $h(t,A,W)$ is then re-evaluated 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., $\hat{\epsilon}$ 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 ($S(t \mid A,W)$) and the censoring mechanism ($\bar{G}(t_{-} \mid A,W)$). The covariate is re-evaluated at each step of the algorithm based on the updated hazard estimate for $\lambda(t \mid A,W)$ (and thus $S(t \mid A,W)$). The censoring mechanism is not updated in the algorithm. The rationale for updating only the hazard was provided in chapter 4.

The covariate $h(t,A,W)$ is selected such 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 $\lambda(t \mid A,W)$ in the nonparametric model for the observed data, assuming only coarsening at random (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_0 specific parameter, $S_1(t_0)$ and $S_0(t_0)$, at each $t_0 \in 1, ..., K$ 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.

5.5 Method 1: Substitution TMLE

In this first method, we describe the procedure for targeted maximum likelihood estimation of the t_0 and treatment specific parameters $S_1(t_0)$ and $S_0(t_0)$ Thereby, we can estimate any parameter that is a combination of them, such as the average parameter $\Psi(p_0)$ defined in (5.4). We note 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_0 and treatment specific parameters $S_1(t_0)$ and $S_0(t_0)$ as defined in (5.1) and (5.2) respectively, were provided in chapter 4. In short, the algorithm selects the covariates, $h_{1t_0}(t, A, W)$ and $h_{0t_0}(t, A, W)$, corresponding with targeting parameters $S_1(t_0)$ and $S_0(t_0)$ respectively. For the parameter $S_1(t_0)$, $h_{1t_0}(t, A, W)$ is defined such that the score for the hazard model at $\epsilon_1 = 0$ is equal to the projection of the efficient influence curve of $S_1(t_0)$ on scores generated by the parameter $\lambda(t \mid A, W)$ in the nonparametric model for the observed data, assuming only CAR. Similarly, for $S_0(t_0)$, $h_{0t_0}(t, A, W)$ is defined such that the score for the hazard model at $\epsilon_2 = 0$ is equal to the projection of the efficient influence curve of $S_0(t_0)$ on scores generated by the parameter $\lambda(t \mid A, W)$ in the nonparametric model for the observed data, assuming only CAR. These covariates corresponding with parameters $S_1(t_0)$ and $S_0(t_0)$ are respectively given by,

$$h_{1t_0}(t, A, W) = -\frac{I(A = 1)}{g(1)\bar{G}(t_- \mid A, W)} \frac{S(t_0 \mid A, W)}{S(t \mid A, W)} I(t \leq t_0), \tag{5.10}$$

and

$$h_{0t_0}(t, A, W) = -\frac{I(A = 0)}{g(0)\bar{G}(t_- \mid A, W)} \frac{S(t_0 \mid A, W)}{S(t \mid A, W)} I(t \leq t_0). \tag{5.11}$$

The parameters $S_1(t_0)$ and $S_0(t_0)$ can be targeted simultaneously by addition of both covariates $h_1(t, A, W)$ and $h_0(t, A, W)$ as in (5.9), and finding the two-dimensional updating coefficients, $\hat{\epsilon} = \{\hat{\epsilon}_1, \hat{\epsilon}_2\}$, i.e.,

$$\text{logit}(\hat{\lambda}^0(\epsilon)(t \mid A, W)) = m(t, A, W) + \epsilon_1 h_{1t_0}(t, A, W) + \epsilon_2 h_{0t_0}(t, A, W). \tag{5.12}$$

Finding $\hat{\epsilon} = \{\hat{\epsilon}_1, \hat{\epsilon}_2\}$ in the updated hazard provided in (5.12) 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 $h_{1t_0}(t, A, W)$ and $h_{0t_0}(t, A, W)$. 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, rather only ϵ is estimated. These steps for evaluating $\hat{\epsilon}$, and thus obtaining the updated hazard fit $\hat{\lambda}^1(t \mid A, W)$ correspond with a single iteration of the targeted maximum likelihood algorithm. In the second iteration, the updated $\hat{\lambda}^1(t \mid A, W)$ now plays the role of the initial fit and the covariates $h_{1t_0}(t, A, W)$ and $h_{0t_0}(t, A, W)$ are then re-evaluated with the updated $\hat{S}^1(t \mid A, W)$ based on $\hat{\lambda}^1(t \mid A, W)$ and $\hat{\epsilon}$ is estimated again. Based on this update, $\hat{\lambda}^2(t \mid A, W)$ is obtained. In the third iteration, $\hat{\lambda}^3(t \mid A, W)$ is fit and the procedure is iterated until $\hat{\epsilon}$ is essentially zero. The final hazard fit at the last iteration of the algorithm is denoted by $\hat{\lambda}^*(t \mid A, W)$ with the corresponding survival fit given by $\hat{S}^*(t \mid A, W)$.

The t_0-specific parameter $\Psi_{t_0}(p_0)$, defined in (5.3), is estimated by,

$$\hat{\psi}^*_{t_0} = \log\left(\frac{\log(\frac{1}{n}\sum_{i=1}^n \hat{S}^*(t_0 \mid 1, W_i))}{\log(\frac{1}{n}\sum_{i=1}^n \hat{S}^*(t_0 \mid 0, W_i))}\right).$$

Finally, the parameter of interest $\Psi(p_0)$ can be estimated by plugging in each of the t_0-specific estimates $\hat{\psi}_{t_0}$ for $t_0 \in 1, ..., K$. That is,

$$\hat{\psi}^* = \sum_{t_0} w(t_0)\hat{\psi}_{t_0}.$$

5.5.1 Efficient influence curves

The efficient influence curves for the parameters $S_1(t_0)$ and $S_0(t_0)$, denoted by IC_{1t_0} and IC_{0t_0}, were provided in chapter 4, and are respectively given by,

$$\begin{aligned} &IC_{1t_0}(g_0, G_0, S_0) = \\ &\quad \sum_{t \le t_0} h_1(g_0, G_0, S_0)(t \mid A, W)[I(\tilde{T} = t, \Delta = 1) - I(\tilde{T} \ge t)\lambda_0(t \mid A = 1, W)] \\ &\qquad\qquad + S_0(t_0 \mid A = 1, W) - \Psi_1(p_0)(t_0), \end{aligned} \tag{5.13}$$

and

$$\begin{aligned} &IC_{0t_0}(g_0, G_0, S_0) = \\ &\quad \sum_{t \le t_0} h_0(g_0, G_0, S_0)(t \mid A, W)[I(\tilde{T} = t, \Delta = 1) - I(\tilde{T} \ge t)\lambda_0(t \mid A = 0, W)] \\ &\qquad\qquad + S_0(t_0 \mid A = 0, W) - \Psi_0(p_0)(t_0). \end{aligned} \tag{5.14}$$

The efficient influence curve, denoted by IC_{t_0}, for the parameter $\Psi_{t_0}(p_0)$, can be obtained by application of the δ-method to the influence curves IC_{1t_0} and IC_{0t_0}. We have,

$$IC_{t_0} = a(t_0)IC_{1t_0} + b(t_0)IC_{0t_0}, \tag{5.15}$$

which is a linear combination of IC_{1t_0} and IC_{0t_0}, with coefficients only a function of t_0. With some algebra, one can easily show that the coefficients are given by $a(t_0) = \frac{1}{S_1(t_0)\log(S_1(t_0))}$ and $b(t_0) = \frac{-1}{S_0(t_0)\log(S_0(t_0))}$.

Our parameter of interest $\Psi(p_0)$ is the average (possibly weighted) of the t_0-specific log ratios of the logs of survival, i.e., average over t_0 of $\Psi_{t_0}(p_0)$. Therefore, its efficient influence curve is given by,

$$\bar{IC} = \sum_{t_0} w(t_0)IC_{t_0} \tag{5.16}$$

5.5.2 Double robustness consistency properties of the S-TMLE

In chapter 4, the statistical properties of the treatment and t_0-specific estimators were provided. Consider the parameter $S_1(t_0)$. The targeted maximum likelihood estimate $\hat{p}^* \in \mathcal{M}$ of p_0 solves the efficient influence curve estimating equation, given by $\sum_{i=1}^n IC_{1t_0}(g_0, \hat{G}, \hat{S}^*)(O_i) = 0$, which is the optimal estimating equation for the parameter of interest. It can be shown that $E_0 IC_{1t_0}(S, g, G) = 0$ if either:

1) $S = S_0(\cdot \mid A, W)$ (and thus $\lambda = \lambda_0(\cdot \mid A, W)$) or, 2) $g = g_0(A \mid W)$ and $G = G_0(\cdot \mid A, W)$ (see Appendix B.1). In an RCT, the treatment mechanism is known and $g_0(A \mid W) = g_0(A)$. Therefore, the consistency of the estimator $\hat{\psi}_1^*(t_0)$ of $S_1(t_0)$ in an RCT relies only on consistent estimation of $G_0(\cdot \mid A, W)$ or $S_0(\cdot \mid A, W)$. When there is no censoring or censoring is missing completely at random (MCAR), $\hat{\psi}_1^*(t_0)$ is consistent even when the estimator $\hat{S}(\cdot \mid A, W)$ of $S_0(\cdot \mid A, W)$ is inconsistent (e.g., if it relies on a mis-specified model). Hence, in an RCT, one is not concerned with estimation bias due to mis-specification of the hazard model. Under informative or missing at random (MAR) censoring, if $G_0(\cdot \mid A, W)$ is consistently estimated then $\hat{\psi}_1^*(t_0)$ is consistent even if $S_0(\cdot \mid A, W)$ is not consistently estimated. If both are correctly specified then $\hat{\psi}_1^*(t_0)$ is efficient. These same properties hold for the estimator $\hat{\psi}_0^*(t_0)$ for $S_0(t_0)$.

Since the S-TMLE $\hat{\psi}_{t_0}^*$ for the t_0-specific parameter given by (5.3) 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_0-specific log ratios of the logs of survival, clearly the properties of the t_0-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 mis-specified. 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 S-TMLEs for the parameters $\Psi_{t_0}(p_0)$ and $\Psi(p_0)$ are inherited from the properties of $\hat{\psi}_1^*(t_0)$ and $\hat{\psi}_0^*(t_0)$, that $\hat{\psi}_{t_0}^*$ and $\hat{\psi}^*$ also solve their corresponding efficient influence curve estimating equations based on (5.15) and (5.16) respectively.

5.5.3 Inference for the S-TMLE

We first consider the parameter $S_1(t_0)$. Since the TMLE is a solution to the efficient influence curve estimating equation, then from estimating equation theory (see van der Laan and Robins [64]) if g_0 and G_0 are known, then the estimator is asymptotically linear with influence curve $IC_{1t_0}(g_0, G_0, S)$. However, even though g_0 is typically known in an RCT, G_0 is not. In this case, the influence curve is given by,

$$IC_{1t_0}(g_0, G, S) - \Pi(IC_{1t_0} \mid T_G),$$

that is one must subtract from IC_{1t_0} 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 $\hat{\psi}_1^*(t_0)$ based on the estimate of the efficient influence curve for $S_1(t_0)$ ignoring this projection and using (5.13), i.e., using $IC_{1t_0}(g_0, \hat{G}, \hat{S})$, where this confidence interval is given by $\hat{\psi}_1^*(t_0) \pm 1.96\frac{\hat{\sigma}_1}{\sqrt{n}}$, and,

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

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

$$\hat{T} = \frac{\hat{\psi}^*}{\frac{\hat{\sigma}}{\sqrt{n}}},$$

whose asymptotic distribution is $N(0, 1)$ under the null hypothesis. Inference is derived in the same manner for the TMLE $\hat{\psi}_0^*(t_0)$ of $S_0(t_0)$.

As with the estimators of treatment specific survival, if g_0 and G_0 are known, then the S-TMLE of $\Psi_{t_0}(p_0)$ is asymptotically linear with influence curve $IC_{t_0}(g_0, G_0, S)$. Since G_0 is not known, the

influence curve is given by,

$$\begin{aligned} IC_{t_0} &= a(t_0)(IC_{1t_0} - \Pi(IC_{1t_0} \mid T_G)) + b(t_0)(IC_{0t_0} - \Pi(IC_{0t_0} \mid T_g)) \\ &= a(t_0)(IC_{1t_0}) + b(t_0)IC_{0t_0} - \Pi\left((a(t_0)IC_{1t_0} - b(t_0)IC_{0t_0}) \mid T_G\right). \end{aligned}$$

Therefore, one can construct an asymptotically conservative Wald-type 0.95-confidence interval for $\hat{\psi}^*_{t_0}$ based on the estimate of the efficient influence curve for $\Psi_{t_0}(p_0)$ using (5.15), i.e, using $IC_{t_0}(g_0, \hat{G}, \hat{S})$, as above. Similarly, the test statistic can be constructed as above.

Finally, for our parameter of interest, the average (possibly weighted) of the t_0-specific log of the ratio of logs of survival, the influence curve when G_0 is estimated is given by,

$$\begin{aligned} I\bar{C} &= \sum_{t_0} w(t_0)\left[a(t_0)(IC_{1t_0} - \Pi(IC_{1t_0} \mid T_G)) + b(t_0)(IC_{0t_0} - \Pi(IC_{0t_0} \mid T_g))\right] \\ &= \sum_{t_0} w(t_0)\left[a(t_0)(IC_{1t_0}) + b(t_0)IC_{0t_0}\right] - \Pi\left(w(t_0)(a(t_0)IC_{1t_0} - b(t_0)IC_{0t_0}) \mid T_G\right). \end{aligned}$$

Therefore, one can construct an asymptotically conservative Wald-type 0.95-confidence interval for $\hat{\psi}^*$ based on the estimate of the efficient influence curve for $\Psi_(p_0)$ using (5.16), i.e., using $I\bar{C}(g_0, \hat{G}, \hat{S})$, as above. Similarly, the test statistic can be constructed as above.

5.6 Method 2: Directly targeted method

In the previous method, we provided a substitution based procedure based on the t_0-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 (5.4). 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 $t_0 \in \{1, ..., K\}$, 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 (5.16) in section 5.5.1. We provided the time-dependent covariates $h_{1t_0}(t, A, W)$ and $h_{0t_0}(t, A, W)$ for targeting the t_0-specific survival $S_1(t_0)$ and $S_0(t_0)$, required to generate the corresponding components of IC_{1t_0} and IC_{0t_0} of the form

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

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

$$\bar{h}(t, A, W) = \sum_{t_0} w(t_0)a(t_0)h_{1t_0} + \sum_{t_0} w(t_0)b(t_0)h_{0t_0}.$$

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

$$\text{logit}(\hat{\lambda}^0(\epsilon)(t \mid A, W)) = m(t, A, W) + \epsilon\bar{h}(t, A, W). \tag{5.17}$$

Again, as in the substitution based method, we now find $\hat{\epsilon}$ by fitting equation (5.17) 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 $\hat{\lambda}^1(t \mid A, W)$ is obtained. This represents the first step of the algorithm. The hazard fit $\hat{\lambda}^1(t \mid A, W)$ now plays the role of the initial fit and the covariate $\bar{h}(t, A, W)$ is then re-evaluated based on $\hat{\lambda}^1(t \mid A, W)$ (and thus $\hat{S}^1(t \mid A, W)$). Based on $\hat{\epsilon}$, estimated as described above, $\hat{\lambda}^2(t \mid A, W)$ is obtained. On the third iteration, $\hat{\lambda}^3(t \mid A, W)$ 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 $\hat{\lambda}^*(t \mid A, W)$ with the corresponding survival fit given by $\hat{S}^*(t \mid A, W)$.

The procedure for estimation of the parameter of interest, $\Psi(p_0)$, now follows exactly that of the S-TMLE in section 5.5. That is, estimates of the t_0-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 $\sum_{i=1}^n IC_{1t_0}(\hat{\lambda}^*, g_0, \hat{G})(O_i)$ and $\sum_{i=1}^n IC_{0t_0}(\hat{\lambda}^*, g_0, \hat{G})(O_i)$ as does the S-TMLE.

5.6.1 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 do 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 mis-specified such that it solves $P_0(\bar{IC}(\lambda', g_0, G_0)) = 0$, however, $\psi(\lambda')$ 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 5.9, and an extensive study of over 100 data generating distributions (results provided in Appendix C.1), 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. This was tested for a number of mis-specified 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 $S_1(t_0)$ and $S_0(t_0)$ based on λ'^* are not necessarily (and are typically not) consistent. Thus, the following can, and often does occur: $\sum_{t_0} \log\left(\frac{\log(S_1(t_0)(\lambda'^*))}{\log(S_0(t_0)(\lambda'^*))}\right) = \sum_{t_0} \log\left(\frac{\log(S_1(t_0)(\lambda_0))}{\log(S_0(t_0)(\lambda_0))}\right)$, but $S_1(t_0)(\lambda'^*) \neq S_1(t_0)(\lambda_0)$ and $S_0(t_0)(\lambda'^*) \neq S_0(t_0)(\lambda_0)$. Thus, the overall logrank parameter is consistently estimated, however its components (i.e., $S_1(t_0)$ and $S_0(t_0)$) are not.

5.6.2 Inference for D-TMLE

The theorem provided in Appendix C.2 can be used to derive the influence curve when $\hat{\lambda}^*$ converges to some mis-specified 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 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 (5.16). For correct inference in this situation, one can apply this theorem to derive the formal influence

curve and use this for inference.

Since deriving 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 5.5.3. 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 5.5.3.

5.7 Weight function

The parameter provided in (5.4) 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_0-specific log ratios. A weight function that down weights the time points $t_0 \in 1, ..., K$ in the tail in 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 $\log\left(\frac{\log(S_1(t_0))}{\log(S_0(t_0))}\right)$, that is $w(t_0) = \frac{1}{var(IC_{t_0})}$, where IC_{t_0} is the efficient influence curve of $\log\left(\frac{\log(S_1(t_0))}{\log(S_0(t_0))}\right)$. This weight function puts more emphasis on those points $t_0 \in 1, ..., K$ for which there is less censoring, and thus more information, providing a more stable estimate of the parameter and the variance. In this paper, we apply only unit weights.

5.8 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 fit 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/subtitution/addition (D/S/A) algorithm [56]. 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,

$$\lambda_0(t \mid, A, W) = \arg\min_\lambda E_0 L(\lambda),$$

where the actual repeated measures loss function is given by,

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

where $w(t, A, W)$ is an arbitrary weight function. Alternatively, we can use the log-likelihood loss function for $\lambda(t \mid A, W)$, 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 [69]. For example, one such learner could be the D/S/A algorithm. The super-learner algorithm then selects α that minimizes,

$$E_{B_n} P^1_{n,B_n} L(\sum_j \alpha(j)\hat{\lambda}_j(P^0_{n,B_n})),$$

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^1_{n,B_n} and P^0_{n,B_n}, 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 above 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, the one can use Kaplan-Meier to estimate the censoring mechanism.

5.9 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 $t_k \in \{1, ..., 9\}$:

- $Pr(A = 1) = Pr(A = 0) = 0.5$
- $W_1 \sim U(2, 6)$
- $W_2 \sim N(10, 10)$
- $\lambda(t|A, W) = \frac{I(t_0<9)I(Y(t_0-1)=0)}{1+\exp(-(-8-0.75A+0.3W_1^2+0.25W_2))} + I(t_0 = 9)$

where $\lambda(t|A, W)$ is the hazard for survival and $Y(t_0)$ is the indicator that the event has occurred at or before time t_0. 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 MAR censoring mechanism is as follows: For $t_0 \in 2, ..., 9$

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

For $t_0 = 1$, $\lambda_C(t \mid A, W_1) = 0$.

The unadjusted estimator was applied as defined in section 5.3. The two targeted maximum likelihood methods provided in sections 5.5 and 5.6 were applied using three different initial hazard fits. The first initial hazard was correctly specified. The second initial hazard was mis-specified by including only a main term for A and W_1. The third initial hazard was mis-specified 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 one 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.

Lastly, in section 5.9.3, 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.

5.9.1 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 5.1, 5.2 and 5.3 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 5.1). Although the gains are more modest when the initial hazard is mis-specified, 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 mis-specification were somewhat smaller, with an increases in power between 0.06 and 0.51 from mis-specified to correctly specified initial hazards (Table 5.2). 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 ($\approx 21\%$) 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 RCT [72].

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 at this small sample size, there happens to be a tiny amount of negative finite sample bias for the D-TMLE (average of the 2500 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.

Table 5.1: No censoring: power and efficiency comparison. *This table compares the two targeted maximum likelihood approaches to the unadjusted logrank under the no censoring setting. Correctly specified initial $\lambda(t \mid A, W)$ (COR), mis-specified initial $\lambda(t \mid A, W)$ includes only a main term for treatment and W_1 (MIS1), and mis-specified initial $\lambda(t \mid A, W)$ includes only a main term for treatment and W_2 (MIS2).*

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

5.9.2 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 5.1 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 RCT with fixed-end point studies where relations between R^2 and efficiency gain have been demonstrated [36, 43, 59].

Table 5.2: MCAR censoring: power and efficiency comparison. *This table compares the two targeted maximum likelihood approaches to the unadjusted logrank, under MCAR censoring. Correctly specified initial $\lambda(t \mid A, W)$ (COR), mis-specified initial $\lambda(t \mid A, W)$ includes only a main term for treatment and W_1 (MIS1), and mis-specified initial $\lambda(t \mid A, W)$ includes only a main term for treatment and W_2 (MIS2).*

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

Table 5.3: MAR censoring: power and efficiency comparison. *This table compares the two targeted maximum likelihood approaches to the unadjusted logrank, under MAR censoring. Correctly specified initial $\lambda(t \mid A, W)$ (COR), mis-specified initial $\lambda(t \mid A, W)$ includes only a main term for treatment and W_1 (MIS1), and mis-specified initial $\lambda(t \mid A, W)$ includes only a main term for treatment and W_2 (MIS2).*

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

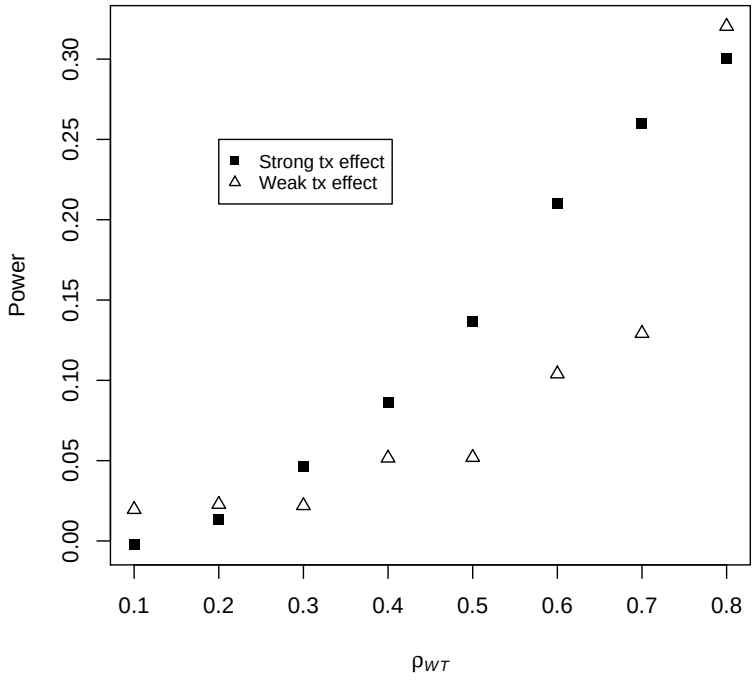

Figure 5.1: Power by increasing correlation between covariate and survival time. *This plots show the relationship between the correlation between covariate and survival time (ρ_{WT}) and the gains in power between the unadjusted and S-TMLE for both strong and weak treatment (tx) effects on survival.*

5.9.3 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 [69].

We simulated 500 replicates of sample size 500 from the following data generating distribution where time is discrete and takes values $t_0 \in \{1, ..., 8\}$:

- $Pr(A = 1) = Pr(A = 0) = 0.5$
- $W_1 \sim U(2, 5)$
- $W_2 \sim N(-2, 2)$
- $Pr(W_3 = 1) = 0.3 = 1 - Pr(W_3 = 0)$
- $W_4 \sim N(1, 1)$
- $W_5 \sim U(-2, 4)$
- $\lambda(t|A, W) = \frac{I(t_0<8)I(Y(t_0-1)=0)}{1+\exp(-(-3-0.5A+0.2W_1W_2+0.2W_2W_4-0.4W_4W_5+0.5W_5W_2))} + I(t_0 = 8)$

The data were generated with no censoring.

In the first method, the initial hazard was estimated using the super-learner algorithm (see section 5.8) which included as candidate learners:

- Simple linear regression with all 5 covariates as main terms.
- Lasso logistic regression [58].
- Random forest [9].
- Generalized additive models [21].
- K-nearest neighbor classification [45].

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 2 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 5.4. 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.

Table 5.4: Power and initial hazard estimation. *This table provides 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$_{SL}$).*

	% 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

5.10 Discussion

The simulation studies provided in this paper 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 paper 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 paper differ from adjusting through Cox-proportional hazards models as was done in Hernández et al. [23], which requires additional assumptions about the proportionality of hazards. Furthermore, the method presented in this paper 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 paper.

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 $\bar{G}(t_- \mid A, W) > 0, t_0 \in 1, ..., K$. This identifiability assumption has been addressed as an important assumption for right-censored data [51]. One is alerted of 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 $\bar{G}(t_- \mid A, W) > 0.1$ in practice. When violations of this assumption are present, an new advance to the targeted maximum likelihood approach

presented in this paper, namely collaborative targeted maximum likelihood, can be applied [62]. In this approach, a sequence of TMLEs are generated with increasing likelihood that correspond with increasingly nonparametric 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 paper 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 paper, as well as more complex parameters as defined by marginal structural models. A further extension, in both RCT and observational studies, is the inclusion of time-dependent covariates which are often predictive of censoring and/or survival [61, 37]. Future work includes the application of these methods in observational studies.

Chapter 6

Conclusions

6.1 Summary

The preceding chapters presented a statistical approach for the robust extraction of covariate information in RCTs for gains in efficiency, and thus power, over the standard unadjusted approaches. The methodology of targeted maximum likelihood was originally proposed in van der Laan and Rubin [65]. Here we have applied this general statistical approach to the estimation of parameters in the most common types of RCTs, including those with fixed end point binary outcomes, and those with right-censored survival outcomes. These applications represent novel contributions to the literature for analyzing data from RCTs.

Chapter 2 provided the targeted maximum likelihood approach to the estimation of marginal causal effects of treatment in RCTs. The approach for the estimation of three parameters representing such effects was provided: the risk difference, relative risk, and odds ratio. It was observed that covariate adjustment in RCT using logistic regression models through extraction of the exponentiated coefficient for treatment represented the conditional treatment effect. The missing step in the estimation of a *marginal* causal effect, is that when adjusting for covariates, these covariates must be integrated/averaged over in order to obtain a marginal effect estimate that is comparable to the unadjusted effect estimate. It was demonstrated that the TMLE does include this averaging step and is a fully robust and efficient estimator of the marginal causal effect. An analytical relationship between the relative efficiency of the targeted maximum likelihood and unadjusted estimators and the R^2 was also provided. This relation showed that when outcome prediction with covariate(s) W through the model $Q(W)$ outperforms outcome prediction through the simple intercept model, then a gain in efficiency by adjusting for the covariates W with the targeted maximum likelihood estimation approach is achieved relative to the unadjusted estimation approach. Missing data were are shown to be easily addressed by the targeted maximum likelihood framework.

The simulation studies presented in section 2.7 demonstrated that the TMLE is more efficient than the unadjusted estimator, with relative efficiencies as high as 13 for the TMLE with the D/S/A algorithm over the unadjusted estimator. These results suggest that data-adaptive model selection algorithms should be applied if the algorithm is specified *a priori*. However, we showed that even adjusting with a mis-specified regression model results in gain in efficiency and power. Thus, using an *a priori* specified model, even if it is mis-specified, can increase the power, and thus reduce the sample size requirements for the study. In addition, the simulation studies demonstrated the clear relation between increasing R^2 and efficiency gain.

Chapter 3, through application of the methodology presented in Chapter 2 to a sampled RCT dataset

with a binary outcome, the issue of efficiency gain with baseline covariates in RCT based on the framework of targeted maximum likelihood estimation was further clarified. The origin of the gains in efficiency and criteria that can be used to anticipate whether the study design and the covariates collected can actually lead to increased estimation precision in practice was explored. It was found that not only is the relation between R^2 and efficiency gain important, but also that empirical confounding can help explain the gain in efficiency that can be achieved by adjusting for covariates. Simulation studies supported the conjecture that in order to achieve a gain in efficiency through covariate adjustment, some imbalance between the treated and untreated arms in the distribution of a covariate must be present. The data analysis elucidated these points and demonstrated that with this near real-life dataset, gains in power are clearly achievable. The 95% confidence interval for the targeted maximum likelihood estimate of the risk difference was $(-0.074, -0.009)$, which is clearly notably narrower than the unadjusted estimate's confidence interval of $(-0.073, 0.005)$, with a relative efficiency of 1.21. Based on the findings of the data analysis and simulation studies, a complete strategy for analyzing these type of data was provided that will protect investigators from misuse of these methods for obtaining favorable inference.

Chapter 4 applied the targeted maximum likelihood methodology to the estimation of treatment specific survival at a fixed end point in time for right-censored survival outcomes. Particular emphasis was placed on safety analysis in RCT due to the fact that power is particularly important in such studies since they are typically powered for efficacy rather than safety. It was demonstrated that the targeted maximum likelihood method does not require any parametric modeling assumptions under MCAR censoring and thus is robust to mis-specification of the hazard model. It was shown that informative censoring can be directly handled by the proposed method, in contrast to the standard unadjusted Kaplan-Meier estimator which is biased in this censoring setting. Extensions to the methodology for inclusion of time-dependent covariates and the application to observational data were also provided.

The simulation studies presented in section 4.5 demonstrated the potential gains in efficiency that can be achieved, with a relative efficiency as high as 1.9, even when the initial hazard was mis-specified. The relation of the correlation between the covariate(s) and failure time and efficiency gain was also demonstrated. In the simulation study, as the correlation increased, so did the relative efficiency in an almost linear pattern. The results also provided evidence refuting an important concern, which is whether or not one can actually lose in efficiency by applying the covariate adjusted strategy. The simulation results showed that when no predictive covariates were present, the relative efficiency was approximately one indicating that investigators are protected from actually losing precision from applying this method even when the covariates provide little information about failure time. The simulations also demonstrated the reduction in bias of the TMLE as compared to the Kaplan-Meier estimator in the informative censoring setting.

Chapter 5 applied the targeted maximum likelihood methodology to provide a covariate adjusted competitor to the ubiquitous logrank test. Two methods were provided. In the first, the methodology for estimating treatment specific survival from Chapter 4 was applied to obtain a substitution based estimator of the logrank parameter. This robust approach requires targeting of each of the time and treatment specific survival parameters, and thus includes many targeting steps. As an alternative, in the second approach, the logrank parameter was targeted directly. In contrast to the substitution estimator, the latter estimator only requires a single targeting step. It was shown through estimating equation theory that the substitution estimator is doubly robust. Estimating equation theory could not be applied to the directly targeted estimator. However, empirical evidence was provided to support the conjecture that these properties also hold for the directly targeted estimator. An important step in both approaches is the estimation of the initial hazard. A data-adaptive approach for its estimation was discussed using the super learner methodology [69].

The simulation studies provided in section 5.9 demonstrated that significant gains in efficiency and

thus power can be achieved over the unadjusted logrank method through covariate adjustment using targeted maximum likelihood. Relative efficiencies as high as 4 and 1.6 were observed for the targeted maximum likelihood estimator over the unadjusted estimator with correctly and mis-specified initial hazards, respectively. With informative censoring, the significant bias (21%) of the unadjusted estimator was eliminated with the TMLE. As with covariate adjustment of the other parameters discussed in this dissertation, a similar relation between R^2 and efficiency gain with the TMLE over the unadjusted estimator was demonstrated with the logrank parameter. The simulation study results also demonstrated the importance of the estimation of the initial hazard in optimizing gains in power. The gain in power for the TMLE with super-learner was 0.23 as compared 0.12 for the targeted maximum likelihood estimator with the main-term D/S/A over the unadjusted estimator. These results provide evidence that the approach for achieving maximum gains in power is a two step procedure, 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, targeted maximum likelihood is applied as a bias-reduction step for the parameter of interest.

6.2 Directions for future research

The methods discussed in this dissertation can be extended in a number of useful ways. In many RCT, time-dependent measurements are collected at each follow-up visit in addition to the many outcome measurements. Such time-dependent covariates are often predictive of censoring. This dissertation focused on data structures with baseline covariates only, however, chapter 2 briefly discussed data structures with time-dependent covariates. Since these latter data structures are not uncommon in RCT, further investigation through simulation studies into targeted maximum likelihood estimation with the inclusion of time-dependent covariates would be desirable.

The identifiability assumption, required for the estimators presented in chapters 4 and 5, is that each subject has a positive probability of being observed at each point in time. When violations of this assumption are present, an new advance to the targeted maximum likelihood approach presented in this dissertation, namely collaborative targeted maximum likelihood estimation, can be applied [62]. Essentially, in this latter 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. Future work is needed to examine the application of the collaborative targeted maximum likelihood methodology in RCT with right-censored survival outcomes.

As post-market studies are becoming a more common tool in assessing and monitoring the overall safety of drugs. Robust methods for estimating causal effects are required since in such studies unadjusted methods are no longer consistent due to confounding. The focus is no longer only on efficiency and power gains, but rather in bias reduction. The targeted maximum likelihood methodology applied in this dissertation can be extended for the estimation of the causal parameters presented here, as well as more complex parameters as defined by marginal structural models.

...

In summary, this dissertation presents a novel contribution to the literature for the estimation of causal effects in RCT by providing methods for the robust extraction of covariate information using targeted maximum likelihood. The simulation studies presented here suggest that these methods represent useful additions to the current set of tools available for covariate adjustment.

Bibliography

[1] K. Akazawa, T. Nakamura, and Y. Palesch. Power of logrank test and cox regression model in clinical trials with heterogenous samples. *Statistics in Medicine*, 16(5):583–597, 1997.

[2] S.F. Assmann, S.J. Pocock, L.E. Enos, and L.E. Kasten. Subgroup analysis and other (mis)uses of baseline data in clinical trials. *Lancet*, 355(9209):1064–1069, March 2000.

[3] H. Bang and J.M. Robins. Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61(4):962–973, 2005.

[4] C. Begg. Significance tests of covariate imbalance in clinical trials. *Controlled Clinical Trials*, 11: 223–225, 1990.

[5] F.J. Belda, L. Aguilera, J. Garcia de la Asuncion, J. Alberti, R. Vicente, L. Ferrandiz, R. Rodriguez, Roque Company, D.I. Sessler, G. Aguilar, S.G. Botello, R. Orti, and for the Spanish Reduccion de la Tasa de Infeccion Quirurgica Group. Supplemental Perioperative Oxygen and the Risk of Surgical Wound Infection: A Randomized Controlled Trial. *JAMA*, 294(16):2035–2042, 2005.

[6] O. Bembom and M.J. van der Laan. Analyzing sequentially randomized trials based on causal effect models for realistic individualized treatment rules. Technical Report 216, Division of Biostatistics, University of California, Berkeley, 2007.

[7] Y. Benjamini and Y. Hochberg. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society, Series B*, 57(1):289–300, 1995.

[8] Y. Benjamini, Y. Hochberg, and Y. Kling. False discovery rate control in multiple hypotheses testing using dependent test statistics. Technical Report 97-1, Tel Aviv University, 1997.

[9] L. Breiman. Random forests. *Machine Learning*, V45(1):5–32, October 2001. doi: 10.1023/A:1010933404324.

[10] I.D. Bross. Why proof of safety is much more difficult than proof of hazard. *Biometrics*, 31(3): 785–793, 1985.

[11] B.R. Davis, C.D. Furberg, and C.B. Williams. Survival analysis of adverse effects data in the beta-blocker heart attack trial. *Clin Pharmacol Ther*, 41:611–615, December 1987.

[12] S.A. Edlavitch. Adverse drug event reporting. Improving the low us reporting rates. *Arch Intern Med*, 148:14991503, 1988.

[13] T.R. Fleming and D.Y. Lin. Survival analysis in clinical trials: Past developments and future directions. *Biometrics*, 56(4):971–983, 2000.

[14] US Food and Drug Administration. International conference on harmonisation: Guidance on statistical principles for clinical trials. *Federal Register*, 63(179):49583–49598, 1998.

[15] N. Frasure-Smith, F. Lespérance, R.H. Prince, P. Verrier, R. Garber, M. Juneau, C. Wolfson, and M. Bourassa. Randomised trial of home-based psychological nursing intervention for patients recovering from myocardial infarction. *Lancet*, 350:473479, 1997.

[16] David A. Freedman. On regression adjustments to experimental data. *Advances in Applied Mathematics*, 40(2):180 – 193, 2008.

[17] L.M. Friedman, C.D. Furberg, and D.L. DeMets. *Fundamentals of Clincal Trials*. Springer-Verlag, New York, 3rd edition, 1998.

[18] J.E. Gait, S. Smith, and S.L. Brown. Evaluation of safety data from controlled clinical trials: the clinical principles explained. *Drug Information Journal*, 34:273–287, 2000.

[19] S.P. Glasser, M. Salas, and E. Delzell. Importance and challenges of studying marketed drugs: What is a phase IV study? Common clinical research designs, registries, and self-reporting systems. *J. Clin. Pharmacol.*, 47(9):1074–1086, 2007.

[20] A. Güttner, J. Kübler, and I. Pigeot. Multivariate time-to-event analysis of multiple adverse events of drugs in integrated analyses. *Statistics in Medicine*, 26(7):1518–1531, March 2007.

[21] T. Hastie and R. Tibshirani. *Generalized Additive Models*. Chapman & Hall/CRC, December 1990. ISBN 0412343908.

[22] A.V. Hernández, E.W. Steyerberg, and J.D. Habbema. Covariate adjustment in randomized controlled trials with dichotomous outcomes increases statistical power and reduces sample size requirements. *J Clin Epidemiol*, 57(5):454–460, 2004.

[23] A.V. Hernández, M.J. Eijkemans, and E.W. Steyerberg. Randomized controlled trials with time-to-event outcomes: How much does prespecified covariate adjustment increase power? *Annals of Epidemiology*, 16(1):41 – 48, 2006.

[24] ICH. Harmonised tripartite guideline structure and content of clinical study reports e3. *61 Federal Register*, 37320, 1996.

[25] R.M. Jacobson, A. Adedunni, V.S. Pankratzc, and G.A. Poland. Adverse events and vaccination-the lack of power and predictability of infrequent events in pre-licensure study. *Vaccine*, 19: 2428–2433, 2001.

[26] C. Jennison and B.W. Turnbull. *Group Sequential Methods: Applications to Clinical Trials*. Chapman & Hall, Boca Raton, Fla, 2000.

[27] H. Jiang, J. Symanowski, S. Paul, Y. Qu, A. Zagar, and S. Hong. The type i error and power of non-parametric logrank and wilcoxon tests with adjustment for covariates-a simulation study. *Statistics in Medicine*, 27(28):5850–5860, 2008.

[28] K.M. Kaplan, S.A. Rusche, H.D. Lakkis, G. Bottenfield, F.A. Guerra, J. Guerrero, H. Keyserling, E. Felicione, T.M. Hesley, and J.W. Boslego. Post-licensure comparative study of unusual high-pitched crying and prolonged crying following comvax(tm) and placebo versus pedvaxhib(tm) and recombivax hb(tm) in healthy infants. *Vaccine*, 21(3-4):181 – 187, 2002.

[29] G. Koch, C.M. Tangen, J. Jung, and I.A. Amara. Issues for covariance analysis of dichotomous and ordered categorical data from randomized clinical trials and non-parametric strategies for addressing them. 17(15-16):1863–1892, 1998.

[30] C. Kooperberg, C.J. Stone, and Y.K. Truong. Hazard regression. *Journal of the American Statistical Association*, 90:78–94, 1995.

[31] Z. Li. Covariate adjustment for non-parametric tests for censored survival data. *Statistics in Medicine*, 20(12):1843–1853, 2001.

[32] F.G. Liu, J. Wang, K. Liu, and D.B. Snavely. Confidence intervals for an exposure adjusted incidence rate difference with applications to clinical trials. *Statistics in Medicine*, 25(8):1275–1286, 2006.

[33] X Lu and A.A Tsiatis. Improving the efficiency of the log-rank test using auxiliary covariates. *Biometrika*, 95(3):679–694, 2008.

[34] D.V. Mehrotra and J.F. Heyse. Use of the false discovery rate for evaluating clinical safety data. *Statistical Methods in Medical Research*, 13(3):227–238, 2004.

[35] S.S. Menjoge. On estimation of frequency data with censored observations. *Pharmaceutical Statistics*, 2(3):191–197, 2003.

[36] K.L. Moore and M.J. van der Laan. Covariate adjustment in randomized trials with binary outcomes: Targeted maximum likelihood estimation. *Statistics in Medicine*, 28(1):39–64, 2009.

[37] K.L. Moore and M.J. van der Laan. Application of time-to-event methods in the assessment of safety in clinical trials. In K Peace, editor, *Design and Analysis of Clinical Trials with Time-to-Event Endpoints*. Taylor & Francis, 2009.

[38] R. Neugebauer and M.J. van der Laan. Why prefer double robust estimators in causal inference? *Journal of the American Statistical Association*, 129:405–426, 2005.

[39] M. Nishikawa, T. Tango, and M. Ogawa. Non-parametric inference of adverse events under informative censoring. *Statistics in Medicine*, 25(23):3981–4003, December 2006.

[40] R.T. O'Neill. The assessment of safety. In K Peace, editor, *Biopharmaceutical Statistics for Drug Development*. Marcel Dekker, New York, 1988.

[41] K. Peace. Design, monitoring and analysis issues relative to adverse events. *Drug Information Journal*, 21(1):21–28, 1987.

[42] Thomas Permutt. Testing for imbalance of covariates in controlled experiments. *Statistics in medicine*, 9(12):1455–1462, 1990.

[43] S.J. Pocock, S.E. Assmann, L.E. Enos, and L.E. Kasten. Subgroup analysis, covariate adjustment and baseline comparisons in clinical trial reporting: current practice and problems. *Statistics in Medicine*, 21:29172930, 2002.

[44] A.G. Randolph, D. Wypij, S.T. Venkataraman, J.H. Hanson, R.G. Gedeit, K.L. Meert, P.M. Luckett, P. Forbes, M. Lilley, J. Thompson, I.M. Cheifetz, P. Hibberd, R. Wetzel, P.N. Cox, J.H. Arnold, for the Pediatric Acute Lung Injury, and Sepsis Investigators Network. Effect of Mechanical Ventilator Weaning Protocols on Respiratory Outcomes in Infants and Children: A Randomized Controlled Trial. *JAMA*, 288(20):2561–2568, 2002.

[45] B.D. Ripley. *Pattern Recognition and Neural Networks*. Cambridge University Press, January 1996. ISBN 0521460867.

[46] J.M. Robins. A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. *Math. Modelling*, 7(9-12): 1393–1512, 1986. ISSN 0270-0255. Mathematical models in medicine: diseases and epidemics, Part 2.

[47] J.M. Robins. A graphical approach to the identification and estimation of causal parameters in mortality studies with sustained exposure periods. *Journal of Chronic Disease*, 40:139S–161S, 1987.

[48] J.M. Robins. Robust estimation in sequentially ignorable missing data and causal inference models. In *Proceedings of the American Statistical Association*. American Statistical Association, Alexandria, VA, 2000.

[49] J.M. Robins and Y. Ritov. Toward a curse of dimensionality appropriate (coda) asymptotic theory for semi-parametric models. *Statistics in Medicine*, 16:285–319, 1997.

[50] J.M. Robins and A. Rotnitzky. Comment on the Bickel and Kwon article, "Inference for semiparametric models: Some questions and an answer". *Statistica Sinica*, 11(4):920–936, 2001.

[51] J.M. Robins and A. Rotnitzky. Recovery of information and adjustment for dependent censoring using surrogate markers. In *AIDS Epidemiology*, Methodological issues. Bikhäuser, 1992.

[52] J.M. Robins and A. Rotnitzky. Inverse probability weighted estimation in survival analysis. In P Armitage and T Colton, editors, *Encyclopedia of Biostatistics*. Wiley & Sons, New York, second edition, 2005.

[53] L.D. Robinson and N.P. Jewell. Some surprising results about covariate adjustment in logistic regression models. *International Statistical Review*, 59:227240, 1991.

[54] D.O. Scharfstein, A. Rotnitzky, and J.M. Robins. Adjusting for nonignorable drop-out using semiparametric nonresponse models. *Journal of the American Statistical Association*, 94(448): 1096–1120, 1999.

[55] S. Senn. Testing for baseline balance in clinical trials. *Statistics in Medicine*, 13(17):1715–1726, 1994.

[56] S. Sinisi and M.J. van der Laan. The deletion/substitution/addition algorithm in loss function based estimation: Applications in genomics. *Statistical Applications in Genetics and Molecular Biology*, 3(1), 2004.

[57] C.M. Tangen and G.G. Koch. on-parametric analysis of covariance for hypothesis testing with logrank and wilcoxon scores and survival-rate estimation in a randomized clinical trial. *Journal of Biopharmaceutical Statistics*, 9(2):307–338, 1999.

[58] R. Tibshirani. Regression shrinkage and selection via the lasso. *J. Roy. Statist. Soc. Ser. B*, 58 (1):267–288, 1996.

[59] A.A. Tsiatis, M. Davidian, M. Zhang, and X. Lu. Covariate adjustment for two-sample treatment comparisons in randomized clinical trials: A principled yet flexible approach. *Statistics in Medicine*, 27(23):4658–4677, 2008.

[60] C.M. van der Horst, M.S. Saag, G.A. Cloud, R.J. Hamill, J.R. Graybill, J.D. Sobel, P.C. Johnson, C.U. Tuazon, T. Kerkering, B.L. Moskovitz, W.G. Powderly, and W.E. Dismukes. Treatment of cryptococcal meningitis associated with the acquired immunodeficiency syndrome. national institute of allergy and infectious diseases mycoses study group and aids clinical trials group. *The New England Journal of Medicine*, 337(1):15–21, 1997.

[61] M.J. van der Laan. The construction and analysis of adaptive group sequential designs. Technical Report 232, Division of Biostatistics, University of California, Berkeley, 2008.

[62] M.J. van der Laan and S. Gruber. Collaborative double robust targeted penalized maximum likelihood estimation. Technical Report 246, Division of Biostatistics, University of California, Berkeley, 2009.

[63] M.J. van der Laan and M.L. Petersen. Causal effect models for realistic individualized treatment and intention to treat rules. *The International Journal of Biostatistics*, 3(1), 2007.

[64] M.J. van der Laan and J.M. Robins. *Unified methods for censored longitudinal data and causality*. Springer, New York, 2003.

[65] M.J. van der Laan and D. Rubin. Targeted maximum likelihood learning. *The International Journal of Biostatistics*, 2(1):Article 11, 2006.

[66] M.J. van der Laan and D. Rubin. A note on targeted maximum likelihood and right censored data. Technical Report 226, Division of Biostatistics, University of California, Berkeley, 2007.

[67] M.J. van der Laan, S. Dudoit, and S. Keles. Asymptotic optimality of likelihood based cross-validation. Technical Report 125, Division of Biostatistics, University of California, Berkeley, 2003.

[68] M.J. van der Laan, S. Dudoit, and K.S. Pollard. ugmentation procedures for control of the generalized family-wise error rate and tail probabilities for the proportion of false positives. *Statistical Applications in Genetics and Molecular Biology*, 3(1), 2004.

[69] M.J. van der Laan, E.C. Polley, and A.E. Hubbard. Super learner. Technical Report 222, Division of Biostatistics, University of California, Berkeley, 2007.

[70] A.W. van der Vaart and J.A. Wellner. *Weak Convergence and Empirical Processes: With Applications to Statistics*. Springer, November 2000. ISBN 0387946403.

[71] Y. Wang, M.L. Petersen, D. Bangsberg, and M.J. van der Laan. Diagnosing bias in the inverse-probability-of-treatment-weighted estimator resulting from violation of experimental treatment assignment. Technical Report 211, Division of Biostatistics, University of California, Berkeley, 2006.

[72] A.M. Wood, I.R. White, and S.G. Thompson. Are missing outcome data adequately handled? A review of published randomized controlled trials in major medical journals. *Clinical Trials*, 1(4): 368–376, 2004.

[73] M. Zhang, A.A. Tsiatis, and M. Davidian. Improving efficiency of inferences in randomized clinical trials using auxiliary covariates. *Biometrics*, 64(3):707 – 715, 2008.

Appendix A

Finding the targeted maximum likelihood covariate and the relation between TMLE, DR-IPTW and G-computation estimators of the risk difference

A.1 Decomposition of efficient influence curve

Following the strategy of [65], the efficient influence curve $D(p_0)$ can be decomposed as,

$$\begin{aligned} D(p_0) &= D(p_0) - E(D(p_0)|A,W) + E(D(p_0)|A,W) - E(D(p_0)|W) + \\ &\quad + E(D(p_0)|W) - E(D(p_0)) \end{aligned}$$

Let, $D_1(p_0) = D(p_0) - E(D(p_0)|A,W)$, $D_2(p_0) = E(D(p_0)|A,W) - E(D(p_0)|W)$ and $D_3(p_0) = E(D(p_0)|W) - E(D(p_0))$. Then, $D_1(p_0)$ is a score for $p(Y|A,W)$, $D_2(p_0)$ is a score for $g_0(A|W)$ and $D_3(p_0)$ is a score for the marginal probability distribution $p(W)$ of W. Note that in this RCT setting, $g_0(A|W) = g_0(A) = \theta_0^A(1-\theta_0)^{(1-A)}$.

A.2 Finding covariate for risk difference TMLE based on logistic regression submodel

The efficient influence curve for the risk difference is given by,

$$\begin{aligned} D(p_0) &= \frac{I(A=1)}{\theta_0}[Y - Q_0(1,W)] \\ &\quad - \frac{I(A=0)}{1-\theta_0}[Y - Q_0(0,W)] \\ &\quad + Q_0(1,W) - Q_0(0,W) - \Psi(p_0), \end{aligned}$$

Consider an initial density estimator $\hat{p}^0$ of the density p_0 of O identified by a regression fit $\hat{Q}^0(A,W)$, marginal distribution of A identified by $\hat{\theta} = \frac{1}{n}\sum_{i=1}^n A_i$, the marginal distribution of W being the

empirical probability distribution of $W_1, ..., W_n$, and A being independent of W. Since Y is binary, we have the following density,

$$\hat{p}^0(Y|A,W) = (\hat{Q}^0(A,W))^Y(1-\hat{Q}^0(A,W))^{1-Y}$$

where,

$$\hat{Q}^0(A,W) = \frac{1}{1+\exp -\hat{m}^0(A,W)}$$

for some function $\hat{m}^0$. Now, consider the parametric submodel through $\hat{p}^0$ indexed by parameter ϵ,

$$\hat{p}^0(\epsilon)(Y|A,W) = (\hat{Q}^0(\epsilon)(A,W))^Y(1-\hat{Q}^0(\epsilon)(A,W))^{1-Y}$$

where $\hat{Q}^0(\epsilon)(A,W)$ is given by the logistic regression model,

$$\hat{Q}^0(\epsilon)(A,W) = \frac{1}{1+\exp -(\hat{m}^0(A,W)+\epsilon h(A,W))}$$

with an extra covariate $h(A,W)$, which needs to be chosen so that the score of ϵ at $\epsilon = 0$ includes the efficient influence curve component $D_1(p^0)$ [65]. The required choice h will be specified below. We estimate ϵ with the maximum likelihood estimator $\hat{\epsilon} = \arg\max_\epsilon \sum_{i=1}^n \log \hat{Q}^0(\epsilon)(A_i,W_i)$. The score for this logistic regression model at $\epsilon = 0$ is given by,

$$\frac{d}{d\epsilon_1}\log p^0(\epsilon)(A,W)\Big|_{\epsilon=0} = h(A,W)(Y-\hat{Q}^0(A,W))$$

We now set the score equal to the part of the efficient IC for $p(Y|A,W)$, that is D_1, at $\hat{p}^0$ to obtain,

$$h(A,W)(Y-\hat{Q}^0(A,W)) = (Y-\hat{Q}^0(A,W))\left(\frac{I(A=1)}{\hat{\theta}} - \frac{I(A=0)}{(1-\hat{\theta})}\right).$$

This equality in $h(A,W)$ is solved by

$$h(A,W) = \frac{I(A=1)}{\hat{\theta}} - \frac{I(A=0)}{(1-\hat{\theta})}.$$

A.3 Relation between TMLE, DR-IPTW and G-computation estimators

The efficient influence curve $D(p_0)$ can be represented as an estimating function in ψ indexed by Q and g, $D(p_0) = D(Q_0, g_0, \Psi(p_0))$. In this RCT setting, $g_0 = \theta_0^A(1-\theta)^{1-A}$. The DR-IPTW estimate is the solution to the corresponding estimating equation in ψ, $\frac{1}{n}\sum_{i=1}^n D(\hat{Q}^0(A_i,W_i),\hat{\theta},\psi) = 0$ and is given by,

$$\begin{aligned}\hat{\psi}_{DR-IPTW} &= \frac{1}{n}\sum_{i=1}^n \frac{I(A_i=1)}{\hat{\theta}}\left[Y_i - \hat{Q}^0(1,W_i)\right] - \frac{1}{n}\sum_{i=1}^n \frac{I(A_i=0)}{1-\hat{\theta}}\left[Y_i - \hat{Q}^0(0,W_i)\right] \\ &\quad + \frac{1}{n}\sum_{i=1}^n \hat{Q}^0(1,W_i) - \frac{1}{n}\sum_{i=1}^n \hat{Q}^0(0,W_i),\end{aligned}$$

where $\hat{\theta} = \frac{1}{n}\sum_{i=1}^{n} A_i$. In the logistic regression fit, $\log\left[\frac{\hat{Q}(A,W)}{1-\hat{Q}(A,W)}\right] = \hat{\alpha}X$, where $X = (1, A, W)$, the maximum likelihood estimate $\hat{\alpha}$ solves the score equations given by,

$$0 = \sum_{i=1}^{n} X_{ij}\left[Y_i - \hat{Q}(A_i, W_i)\right],$$

for $j = 1, ..., p$. The linear span of scores includes the covariate,

$$x_j = \frac{I(A=1)}{\hat{\theta}} - \frac{I(A=0)}{1-\hat{\theta}},$$

when A and an intercept are included in X. Thus, it follows that

$$0 = \frac{1}{n}\sum_{i=1}^{n} \frac{I(A_i=1)}{\hat{\theta}}\left[Y_i - \hat{Q}^0(1, W_i)\right] - \frac{1}{n}\sum_{i=1}^{n} \frac{I(A_i=0)}{1-\hat{\theta}}\left[Y_i - \hat{Q}^0(0, W_i)\right].$$

Hence,

$$\hat{\psi}_{DR} = \frac{1}{n}\sum_{i=1}^{n}\left[\hat{Q}(1, W_i) - \frac{1}{n}\sum_{i=1}^{n}\hat{Q}(0, W_i)\right] = \hat{\psi}_{Gcomp} = \hat{\psi}_{RD-TMLE}$$

Thus in this quite general scenario, we have that the DR-IPTW, the G-computation estimator, and the TMLE, all reduce to the same estimator.

A.4 Finding covariate for relative risk TMLE based on logistic submodel

We apply the delta method to obtain the efficient influence curve of the log of the relative risk parameter, i.e. $\log(\frac{\mu_1}{\mu_0}) = \log(\mu_1) - \log(\mu_0)$. The efficient influence curve is given by,

$$\begin{aligned} D(p_0) &= \frac{1}{\mu_1}\left(\frac{I(A=1)}{\theta_0}(Y - Q_0(1,W)) + Q_0(1,W) - \mu_1\right) - \\ &\quad -\frac{1}{\mu_0}\left(\frac{I(A=0)}{(1-\theta_0)}(Y - Q_0(0,W)) + Q_0(0,W) - \mu_0\right) \\ &= \frac{1}{\mu_1}\left(\frac{I(A=1)}{\theta_0}(Y - Q_0(1,W)) + Q_0(1,W)\right) - \\ &\quad -\frac{1}{\mu_0}\left(\frac{I(A=0)}{(1-\theta_0)}(Y - Q_0(0,W)) + Q_0(0,W)\right) \end{aligned}$$

In order to find the covariate $h(A, W)$ that is added to the regression model, we note the following equality given in [64],

$$V(Y, A, W) = (V(1, A, W) - V(0, A, W))(Y - Q(A, W)), \tag{A.1}$$

if V is a function with conditional mean 0 given A and W. We apply this equality to $D(p_0) = V(Y, A, W)$ to obtain $h(A, W)$.

Let $\hat{p}^0(\epsilon_1)$ be the logistic regression fit with an extra covariate extension $\epsilon_1 h(A,W)$. Based on A.1 we can immediately observe that the covariate $h(A,W)$ added to the logistic regression is $V(1,A,W) - V(0,A,W)$ since,

$$\begin{aligned}\frac{d}{d\epsilon}\log \hat{p}^0(\epsilon)(A,W)\Big|_{\epsilon=0} &= h(A,W)(Y - \hat{Q}^0(A,W)) \\ &= (V(1,A,W) - V(0,A,W))(Y - \hat{Q}^0(A,W))\end{aligned}$$

Thus, evaluating $D(\hat{p}_0)$ at $Y=1$ and $Y=0$ gives,

$$h(A,W) = \frac{1}{\mu_1}\frac{I(A=1)}{\hat{\theta}} - \frac{1}{\mu_0}\frac{I(A=0)}{(1-\hat{\theta})}.$$

A.5 Finding covariate for relative risk TMLE based on relative risk regression submodel

Consider now the parametric submodel $\hat{p}^0$ indexed by parameter ϵ,

$$\hat{p}^0(\epsilon)(Y|A,W) = (\hat{Q}^0(\epsilon)(A,W))^Y (1 - \hat{Q}^0(\epsilon)(A,W))^{1-Y}$$

where $\hat{Q}^0(\epsilon)(A,W)$ is given by the relative risk regression model,

$$\log(\hat{Q}^0)(\epsilon)(A,W) = \hat{m}^0(A,W) + \epsilon h(A,W).$$

The score for this model evaluated at $\epsilon = 0$ is given by,

$$\frac{d}{d\epsilon}\log \hat{p}^0(\epsilon)(A,W)\Big|_{\epsilon=0} = \frac{h(A,W)}{1-\hat{Q}^0(A,W)}(Y - \hat{Q}^0(A,W)),$$

and it follows that the covariate added to logistic regression model to obtain the TMLE is given by,

$$h(A,W) = \left(\frac{1}{\mu_1}\frac{I(A=1)}{\hat{\theta}} - \frac{1}{\mu_0}\frac{I(A=0)}{(1-\hat{\theta})}\right)(1 - \hat{Q}^0(A,W)).$$

A.6 Finding covariate for odds ratio based on logistic regression submodel

We apply the delta method to obtain the efficient influence curve of the log of this parameter, i.e. $\log(\frac{\mu_1/(1-\mu_1)}{\mu_0/(1-\mu_0)}) = \log(\frac{\mu_1}{(1-\mu_1)}) - \log(\frac{\mu_0}{(1-\mu_0)})$. We have, $\frac{d}{d\mu_1}\log(\frac{\mu_1}{1-\mu_1}) = \frac{1}{\mu_1} + \frac{1}{1-\mu_1}$ and thus the efficient influence curve is given by,

$$\begin{aligned}D(p_0) &= \left(\frac{1}{\mu_1} + \frac{1}{1-\mu_1}\right)\left(\frac{I(A=1)}{\theta_0}(Y - Q_0(1,W)) + Q_0(1,W) - \mu_1\right) - \\ &\quad - \left(\frac{1}{\mu_0} + \frac{1}{1-\mu_0}\right)\left(\frac{I(A=0)}{(1-\theta_0)}(Y - Q_0(0,W)) + Q_0(0,W) - \mu_0\right)\end{aligned}$$

Applying equality A.1 to $D(\hat{p}^0)$, we obtain,

$$h(A,W) = \left(\frac{1}{\mu_1} + \frac{1}{1-\mu_1}\right)\frac{I(A=1)}{\hat{\theta}} - \left(\frac{1}{\mu_0} + \frac{1}{1-\mu_0}\right)\frac{I(A=0)}{(1-\hat{\theta})}$$

A.7 Finding covariate for risk difference with missing data based on logistic regression submodel

The efficient influence curve is given by,

$$\begin{aligned} D(p_0) &= \frac{I(A=1)}{g_0(1,1|W)}\left[Y - Q_0(1,1,W)\right] \\ &\quad -\frac{I(A=0)}{(g_0(0,1|W))}\left[Y - Q_0(0,1,W)\right] \\ &\quad +Q_0(1,1,W) - Q_0(0,1,W) - \Psi(p_0), \end{aligned}$$

where $g_0(A=1,\delta|W) = \theta_0 g(\delta|A=1,W)$ and $g_0(A=0,\delta|W) = (1-\theta_0)g(\delta|A=0,W)$. We now present the analogue to the derivation of the TMLE for ψ_0. Consider the parametric submodel through $\hat{p}^0$ indexed by parameter ϵ,

$$\hat{p}^0(\epsilon)(Y|A,\Delta=1,W) = \left(\hat{Q}^0\left[\epsilon\right)(A,\Delta=1,W)\right]^Y \left[1-\hat{Q}^0(\epsilon)(A,\Delta=1,W)\right]^{1-Y}$$

where $\hat{Q}^0(\epsilon)(A,\Delta=1,W)$ is given by the logistic regression model,

$$\hat{Q}^0(\epsilon)(A,\Delta=1,W) = \frac{1}{1+\exp-\left[\hat{m}^0(A,\Delta=1,W)+\epsilon h(A,\Delta=1,W)\right]}.$$

At $\Delta=0$, the likelihood of $P(Y \mid A,\Delta,W)$ provides as contribution a factor 1, which can thus be ignored. The score for this logistic regression model at $\epsilon=0$ is given by,

$$\left.\frac{d}{d\epsilon}\log p^0(\epsilon)(A,\Delta,W)\right|_{\epsilon=0} = I(C=1)h(A,\Delta=1,W)(Y-\hat{Q}^0(A,\Delta=1,W))$$

We now set this score equal to the component of the efficient influence curve which equals a score for $P(Y|A,\Delta=1,W)$, at $\hat{p}^0$, to obtain the equality

$$\begin{aligned} &h(A,\Delta=1,W)(Y-\hat{Q}^0(A,\Delta=1,W)) \\ &\quad = (Y-\hat{Q}^0(A,\Delta=1,W))\left(\frac{I(A=1)}{\hat{g}(1,1|W)}-\frac{I(A=0)}{\hat{g}(0,1|W)}\right). \end{aligned}$$

Solving for $h(A,\Delta=1,W)$ we obtain,

$$h(A,\Delta=1,W) = \frac{I(A=1)}{\hat{g}(1,1|W)}-\frac{I(A=0)}{\hat{g}(0,1|W)}.$$

Appendix B

Double robust consistency properties for treatment specific survival at a fixed end point

B.1 TMLE for t_0 specific parameter is doubly robust

We now show that,

$$E_0 D^*(S, g_0, G_0, S_{0,1}(t_0)) = 0,$$

for any $S = S(\cdot \mid A, W)$, where $D^*(S, g, G, \Psi(S))$ is the efficient influence curve of $\Psi(S) = S_1(t_0)$ at the data generating distribution identified by S, g, G and the marginal distribution of W.

We have,

$$\begin{array}{l} D^*(S, g_0, G_0, S_{0,1}(t_0)) = \\ -\sum_{t \le t_0} \frac{I(A=1)}{\bar{G}_0(t_- \mid A=1,W) g_0(1 \mid W)} \frac{S(t_0 \mid A=1,W)}{S(t \mid A=1,W)} (dN(t) - E(dN(t) \mid \bar{N}(t_-), A=1, W)) \\ +S(t_0 \mid A=1, W) - S_{0,1}(t_0). \end{array}$$

Rewrite this as,

$$\begin{array}{l} D^*(S, g_0, G_0, S_{0,1}(t_0)) = \\ -\sum_{t \le t_0} \frac{I(C \ge t)}{\bar{G}_0(t_- \mid A=1,W)} \frac{I(A=1)}{g_0(1 \mid W)} \frac{S(t_0 \mid A=1,W)}{S(t \mid A=1,W)} (I(T=t) - I(T \ge t) \lambda(t \mid A=1, W)) \\ +S(t_0 \mid A=1, W) - S_{0,1}(t_0). \end{array}$$

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

$$\begin{array}{l} -P_0 D^*(S, g_0, G_0, S_{0,1}(t_0)) = \\ P_0 \sum_{t \le t_0} \frac{I(C \ge t)}{\bar{G}_0(t_- \mid A=1,W)} \frac{I(A=1)}{g_0(1 \mid W)} \frac{S(t_0 \mid A=1,W)}{S(t \mid A=1,W)} I(T \ge t) (\lambda_0(t \mid A=1, W) - \lambda(t \mid A=1, W)) \\ +P_0 S(t_0 \mid A=1, W) - S_{0,1}(t_0) \end{array}$$

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

$$\begin{array}{l} -P_0 D^*(S, g_0, G_0, S_{0,1}(t_0)) = \\ P_0 \sum_{t \le t_0} \left[\frac{I(C \ge t)}{\bar{G}_0(t_- \mid A=1,W)} \frac{I(A=1)}{g_0(1 \mid W)} \frac{S(t_0 \mid A=1,W)}{S(t \mid A=1,W)} S_0(t- \mid A, W) \times \right. \\ \left. (\lambda_0(t \mid A=1, W) - \lambda(t \mid A=1, W)) \right] \\ +P_0 S(t_0 \mid A=1, W) - S_{0,1}(t_0) \end{array}$$

Define now the term

$$W^*(t) \equiv \frac{S(t_0 \mid A=1,W)}{S(t \mid A=1,W)} S_0(t- \mid A,W)(\lambda_0(t \mid A=1,W) - \lambda(t \mid A=1,W)).$$

If

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

then we obtain

$$\begin{array}{l} -P_0 D^*(S, g_0, G_0, S_{0,1}(t_0)) = \\ P_0 \sum_{t \leq t_0} \frac{S(t_0 \mid A=1,W)}{S(t \mid A=1,W)} S_0(t- \mid A,W)(\lambda_0(t \mid A=1,W) - \lambda(t \mid A=1,W)) \\ +P_0 S(t_0 \mid A=1,W) - S_{0,1}(t_0) \end{array}$$

In particular, this identity applies if $\bar{G}_0(t- \mid A=1,W)$, and $g_0(1 \mid W)$ are the true conditional distributions, given the whole W, but the above requirement is weaker by only requiring that the censoring mechanism G_0 and treatment mechanism g_0 are conditioning on 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. Firstly, note that it can be rewritten as:

$$\begin{array}{l} -P_0 D^*(S, g_0, G_0, \psi_0) = \\ -P_0 \sum_{t \leq t_0} S(t_0 \mid A=1,W) \left[\frac{S(t_- \mid A=1,W) f_0(t \mid A=1,W) - S_0(t_- \mid A=1,W) f(t \mid A=1,W)}{S(t \mid A=1,W)) S(t_- \mid A=1,W)} \right] \\ +P_0 S(t_0 \mid A=1,W) - S_{0,1}(t_0) \end{array}$$

Here we used ψ_0 to indicate the true target parameter $S_{0,1}(t_0)$. Now we can use the algebraic trick, $ab - cd = (a-c)d + a(b-d)$ to obtain the expression,

$$\begin{array}{l} E_0 D^*(S, g_0, G_0, \psi_0) = \\ -E_0 \sum_{t \leq t_0} S(t_0 \mid A=1,W) [\frac{(S(t_- \mid A=1,W) - S_0(t \mid A=1,W)) f(t \mid A=1,W)}{S(t \mid A=1,W)) S(t_- \mid A=1,W)} \\ + \frac{S(t_- \mid A=1,W)(f_0(t \mid A=1,W) - f(t \mid A=1,W))}{S(t \mid A=1,W)) S(t_- \mid A=1,W)}] + P_0 S(t_0 \mid A=1,W) - S_{0,1}(t_0) \end{array}$$

For the continuous survival case, we note that,

$$\begin{array}{rl} & \int_{[0,t_0]} \left(\frac{S(t_- \mid A=1,W) - S_0(t_- \mid A=1,W)}{S(t_- \mid A=1,W)} \right) \\ = & \int_{[0,t_0]} - \frac{(S(t_- \mid A=1,W) - S_0(t_- \mid A=1,W)) f(t_- \mid A=1,W)}{S^2(t_- \mid A=1,W)} \\ & \qquad + \frac{f(t_- \mid A=1,W)}{S(t_- \mid A=1,W)} - \frac{f_0(t_- \mid A=1,W)}{S(t_- \mid A=1,W)} \\ = & \frac{(S(t_0 \mid A=1,W) - S_0(t_0 \mid A=1,W)) f(t_- \mid A=1,W)}{S^2(t_0 \mid A=1,W)} + \\ & \qquad \frac{(f_0(t_0 \mid A=1,W) - f(t_0 \mid A=1,W)) S(t_0 \mid A=1,W)}{S^2(t_0 \mid A=1,W)}. \end{array}$$

Therefore, we have,

$$\begin{aligned}
&E_0(D^*(S, g_0, G_0, \psi_0 \mid X)) \\
&= -E_0 \frac{S(t_0 \mid A=1,W)(S(t_- \mid A=1,W) - S_0(t_- \mid A=1,W))}{S(t_0 \mid A=1,W)} \\
&\quad + E_0 S(t_0 \mid A=1,W) - E_0 S_{0,1}(t_0) \\
&= -E_0 S(t_0 \mid A=1,W) + S_{0,1}(t_0) + E_0 S(t_0 \mid A=1,W) - S_{0,1}(t_0) \\
&= 0.
\end{aligned}$$

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

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

Furthermore,

$$\begin{aligned}
&\sum_{t_0 \le t_0} \frac{(S(t_k \mid A=1,W) - S_0(t_k \mid A=1,W))}{S(t_k \mid A=1,W)} - \\
&\qquad \sum_{t_0 \le t_0} \frac{(S(t_{k-1} \mid A=1,W) - S_0(t_{k-1} \mid A=1,W))}{S(t_{k-1} \mid A=1,W)} \\
&= \frac{S(t_0 \mid A=1,W) - S_0(t_0 \mid A=1,W))}{S(t_0 \mid A=1,W}.
\end{aligned}$$

Therefore for discrete survival we can write,

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

Appendix C

Properties of the D-TMLE

C.1 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 minimal parameter directly. We consider the scenario where $\hat{\lambda^*}$ 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 4 data generating mechanisms, the two covariates 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, the treatment mechanism was given by $P(A=1 \mid W) = \frac{1}{1+\exp(-(-0.75+.3*W_2))}$ and $P(A=0 \mid W) = 1 - P(A=1 \mid W)$. The settings are summarized as follows:

1. Constant hazard, treatment randomized:
$$\lambda(t \mid A, W) = \frac{I(t_0 < 10) I(Y(t_0 - 1) = 0)}{1 + \exp(-(\beta_0 - \beta_1 A + \beta_2 W_1^2 + \beta_3 W_2))} + I(t_0 = 10)$$
.

2. Hazard changed over time, treatment randomized:
$$\lambda(t \mid A, W) = \frac{I(t_0 < 10) I(Y(t_0 - 1) = 0)}{1 + \exp(-(\beta_0 - \beta_1 A + \beta_2 W_1^2 + \beta_3 W_2 + \beta_4 * t_0))} + I(t_0 = 10)$$
.

3. Hazard changed over time (interaction between time and covariate), treatment randomized:
$$\lambda(t \mid A, W) = \frac{I(t_0 < 10) I(Y(t_0 - 1) = 0)}{1 + \exp(-(\beta_0 - \beta_1 A + \beta_2 W_1 * t_0 + \beta_3 W_2))} + I(t_0 = 10)$$
.

4. Constant hazard, treatment not randomized:
$$\lambda(t \mid A, W) = \frac{I(t_0 < 10) I(Y(t_0 - 1) = 0)}{1 + \exp(-(\beta_0 - \beta_1 A + \beta_2 W_1^2 + \beta_3 W_2))} + I(t_0 = 10)$$
.

For each settings, 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 $n = 25000$ observations was generated for each of the 25 different randomly selected parameter values for each of the 4 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 non-randomized 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 C.1, C.2 and C.3, which correspond with different misspecified initial hazard estimates. Clearly, the D-TMLE, even for the intercept only initial hazard (Figure C.1) 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 C.2 and C.3). 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).

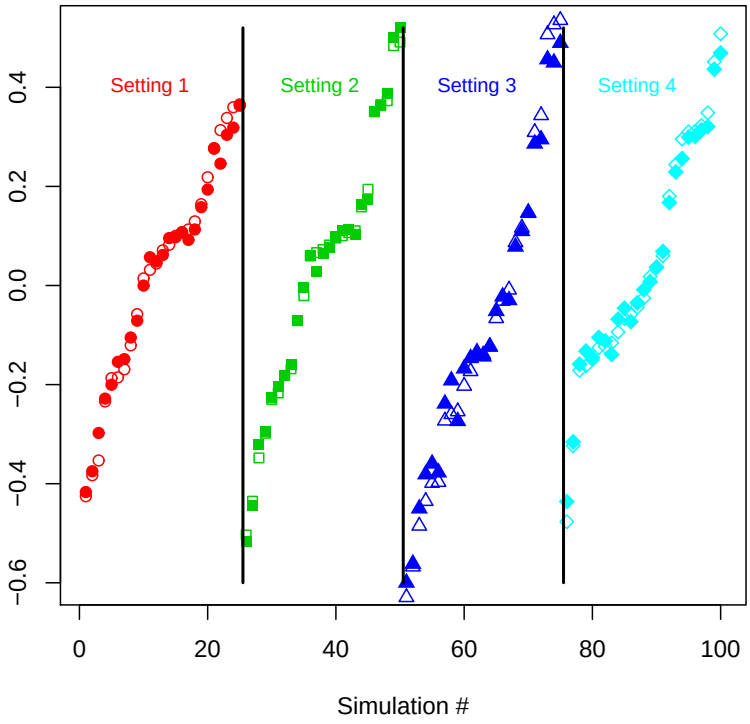

Figure C.1: D-TMLE results for 100 data generating distributions with misspecified initial hazard that includes intercept only. *Open points are the true values, filled in points are the directly targeted maximum likelihood estimates.*

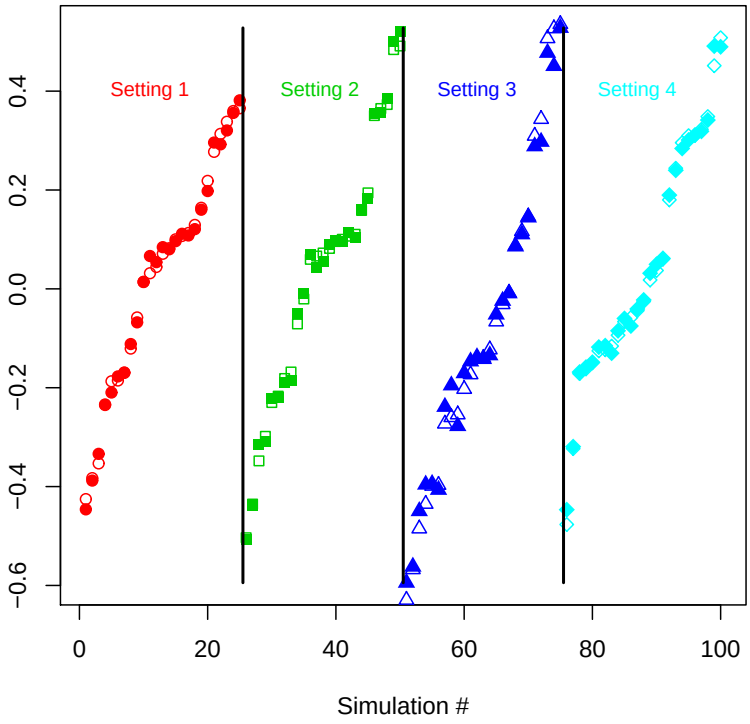

Figure C.2: D-TMLE results for 100 data generating distributions with misspecified initial hazard that incorrectly excludes W_2 term. *Open points are the true values, filled in points are the directly targeted maximum likelihood estimates.*

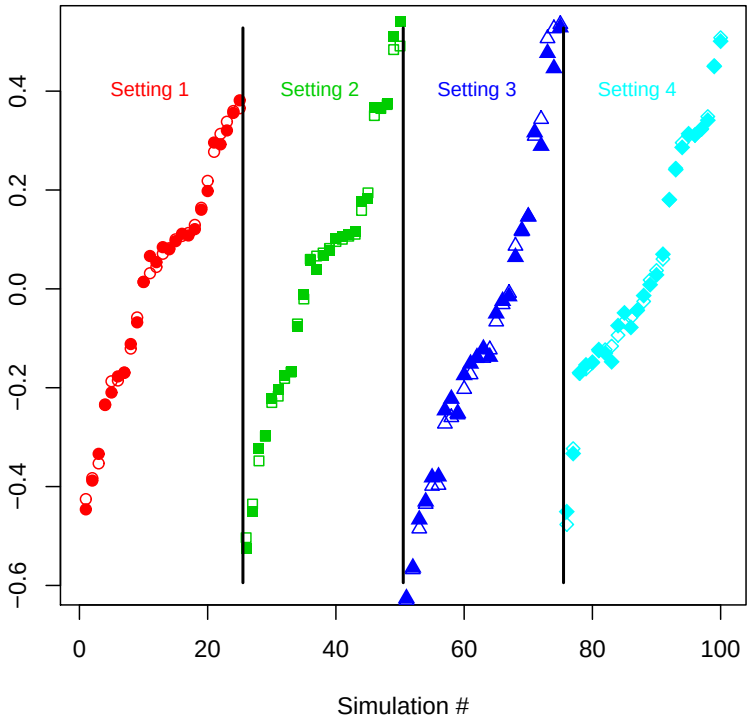

Figure C.3: D-TMLE results for 100 data generating distributions with misspecified initial hazard that incorrectly excludes W_1 term. *Open points are the true values, filled in points are the directly targeted maximum likelihood estimates.*

C.2 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 path-wise 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 $p_n^* = \hat{\lambda}^*\hat{G}^*$. We consider the case that $\hat{G}^*$ is assumed to be consistent for G_0. We proceed as follows:

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

where

$$R_{1n} = P_0\{D(\hat{\lambda}^*, \hat{G}^*) - D(\hat{\lambda}^*, G_0)\} - P_0\{D(Q_0^*, G_0) - D(Q_0^*, G_0)\}.$$

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

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

so that the term $P_0\{D(\lambda_0^*, \hat{G}^*) - D(\lambda_0^*, G_0)\}$ equals $\Phi(\hat{G}^*) - \Phi(G_0)$. We now assume that $\Phi(\hat{G}^*)$ 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 $T_g(P_0)$. It remains to consider $P_0 D(\hat{\lambda}^*, G_0)$. By the general representation Theorem 1.3 in [64] it follows that

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

where $D^F(\lambda)$ is a gradient in the full data model λ for the parameter $\lambda \to \Psi(\lambda)$, and P_{λ_0} denotes the full data distribution. Again, by path-wise differentiability of Ψ in the full data model, if $D^F(\hat{\lambda}^*)$ consistently estimates $D^F(\lambda_0)$, then one expects $P_{\lambda_0} D^F(\hat{\lambda}^*) = \psi_0 - \Psi(\hat{\lambda}^*) + o_P(1/\sqrt{n})$. In general, we note that, if $\hat{\lambda}^*$ converges to some possibly misspecified λ^* for which $\Psi(\lambda^*) = \Psi(\lambda_0)$ and $P_0 D^F(\lambda^*) = 0$, we have

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

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

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

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

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

To stay general, we assume the expansion:

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

for some $D_1(P_0)$. By Theorem 2.3 in van der Laan and Robins [64] the influence curve of $\Phi(\hat{G}^*)$ equals $-\Pi(D(\lambda_0^*, G_0) + D_1(P_0) \mid T_g(P_0)^\perp)$. This proves the following theorem which provides a template for establishing asymptotic linearity of the D-TMLE in CAR censored data models.

Theorem 1 *Let $O_1, \ldots, O_n \sim P_0$ 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 $\mathcal{M} = \{p_{\lambda G} = \lambda G : \lambda \in \lambda, G \in \mathcal{G}\}$, where $\mathcal{G}$ is a subset of all CAR distributions. Let $\Psi : \lambda \to \mathbb{R}^d$ 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 $p_n^* = \hat{\lambda}^* \hat{G}^*$ satisfying $P_n D(\hat{\lambda}^*, \hat{G}^*) = 0$.*

- *Define*
$$R_{1n} \equiv P_0\{D(\hat{\lambda}^*, \hat{G}^*) - D(\hat{\lambda}^*, G_0)\} - P_0\{D(\lambda_0^*, G_0) - D(\lambda_0^*, G_0)\}.$$
Assume $R_{1n} = o_P(1/\sqrt{n})$.

- *Define*
$$\Phi(g_n^*) \equiv P_0 D(\lambda_0^*, \hat{G}^*),$$
where P_0 and λ_0^* are treated as given. Assume that $\Phi(\hat{G}^*)$ 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:*
$$P_0 D(\hat{\lambda}^*, G_0) = \psi_0 - \Psi(\hat{\lambda}^*) + \frac{1}{n}\sum_{i=1}^{n} D_1(P_0) + o_P(1/\sqrt{n}),$$
for some $D_1(P_0)$.

- *Assume $D(\hat{\lambda}^*, \hat{G}^*)$ falls in a P_0-Donsker class. Then, $\Psi(P_n^*) - \psi_0 = O_P(1/\sqrt{n})$.*

- *In addition, assume $P_0\{D(\hat{\lambda}^*, \hat{G}^*) - D(\lambda_0^*, G_0)\}^2 \to 0$ in probability as $n \to \infty$ 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
$$IC(P_0) \equiv \Pi\left(D(\lambda_0^*, G_0) + D_1(P_0) \mid T_g(P_0)^{\perp}\right),$$
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.