Data-adaptive Estimation in Longitudinal Data Structures with Applications in Vaccine Efficacy Trials

David Benkeser

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2015

Reading Committee:

Peter Gilbert, Chair

Marco Carone, Chair

Yu Dai

Program Authorized to Offer Degree:
Biostatistics

University of Washington

Abstract

Data-adaptive Estimation in Longitudinal Data Structures with Applications in Vaccine Efficacy Trials

David Benkeser

Co-Chairs of the Supervisory Committee:
Dr. Peter Gilbert
University of Washington Department of Biostatistics

Dr. Marco Carone
University of Washington Department of Biostatistics

This dissertation develops methodology for data-adaptive estimation of parameters defined on longitudinal data structures, while this abstract serves as an introduction to the material covered herein. The dissertation is organized into three related, but distinct chapters. Each chapter considers a similar data structure, wherein subjects are enrolled and followed over a period of time to obtain additional measurements, for example their failure status. During this followup period, subjects may drop out and therefore researchers are unable to observe the entire study population at all time points. Using the observed data, this dissertation develops asymptotically efficient estimators that may draw valid inferences on the original study population.

Data from preventive vaccine trials serve as the motivation for much of the work in this dissertation. In such trials, subjects are randomized to receive an active vaccine or placebo vaccine and are subsequently followed over some period of time to ascertain infection status. This infection data may be augmented with pathogen genetic data. Scientific interest may lie in assessing the vaccine's efficacy to prevent infections of a certain genotype; this problem is considered in Chapter 1. Researchers may possess additional information on the expected incidence of an infection in the population under study. For example, such information may

be ascertained from previous studies in the same population. In Chapter 2, we show how this information may be included in the estimation procedure to improve performance. The third and final chapter explores the construction of estimators that enjoy the unique property of being robust to model misspecification in terms of both estimation and inference drawn from the estimator.

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGMENTS

The author wishes to express sincere appreciation to his co-advisers Peter Gilbert and Marco Carone. The developments herein would not have been possible without their guidance.

DEDICATION

to my wife Katherine and my dear friend Tally.

Chapter 1

EFFICIENT AND ROBUST ESTIMATION OF CUMULATIVE INCIDENCE IN THE SETTING OF COMPETING RISKS

1.1 Introduction

In many studies where time to an event is of interest, subjects may fail due to one of multiple competing risks. For example, in preventive vaccine efficacy trials for the human immunodeficiency virus (HIV), trial participants become HIV infected with one of many genotypes of HIV [Rolland et al., 2012]. Competing risks are also prevalent in clinical trials with several outcomes of interest, such as in cancer prevention trials [Women's Health Initiative Investigators, 2002]. In these cases, interest lies in assessing the ability of a treatment to prevent events of a certain type in the presence of other competing risks. This is often done by comparing between treatment arms the probability that failure occurs by a fixed time and due to a particular cause, also referred to as the cumulative incidence.

As a motivating example, consider the RV144 ALVAC/AIDSVAX vaccine efficacy trial, a randomized trial of a vaccine to prevent HIV-1 infection conducted in Thailand [Rerks-Ngarm et al., 2009]. This was the first successful trial of a preventative HIV vaccine, though the estimated efficacy was insufficient to gain licensure. A primary aim of follow-up analyses was to determine whether the vaccine was more protective against viruses that were genetically similar to those included in the vaccine [Rolland et al., 2012]. This differential efficacy may be assessed by comparing the treament's effect on cumulative incidence for genetically similar and dissimilar infections.

The foremost estimator of cumulative incidence is the Aalen-Johansen estimator [Aalen, 1978]. This estimator requires few assumptions to achieve several desirable properties. Provided censoring is uninformative, this estimator is consistent. If additionally there are no measured prognostic covariates, it is also efficient. However, informative censoring is a common concern in prospective trials and for many endpoints prognostic covariates are available and routinely collected. For example, sexual risk behaviors are collected in HIV vaccine trials and have been shown to be predictive of infection [Hammer et al., 2013]. In such settings, it may be possible to improve on the Aalen-Johansen estimator.

More recent methods have been developed to allow for dependent censoring and incorporate prognostic covariate information. The primary methods for covariate adjustment are semiparametric models, including proportional hazards regression models for cause-specific hazards [Lunn and McNeil, 1995, Prentice et al., 1978] and subdistribution hazards [Fine and Gray, 1999]. These methods have recently been extended to also allow for dependent censoring [Chen et al., 2002, Scheike et al., 2008]. While this may represent an improvement relative to the Aalen-Johansen estimator, these gains come at the cost of more restrictive modeling assumptions. When such assumptions are violated, the estimands of these methods are difficult to interpret. For example, the estimand of a misspecified Cox model is known to generally involve the censoring distribution [Struthers and Kalbfleisch, 1986], an undesirable property for assessing treatment efficacy in practice [Stitelman et al., 2011].

To motivate the present work, consider the problem of non-proportional hazards and its oft-proposed solution: allowing cause-specific hazard ratios to vary in time. To retain the desired interpretation of the estimand, a correct specification of how the hazard ratios vary with time is required. Furthermore, for inference to be properly calibrated, this specification must be made a priori. However, in many applications background knowledge is simply insufficient to allow a correct pre-specification of the shape of a conditional hazard function. A more appealing approach might then be to use cross-validation to select from fits based on a number of different smooth functions of time and possibly employ state-of-the-art algorithms from machine learning to estimate conditional cause-specific hazards. Estimates from these models could furthermore be combined via ensemble learning, e.g., the Super Learner [van der Laan and Polley, 2007]. By capitalizing on these more data-adaptive approaches, interpretability of the estimand would be preserved over a much larger class of potential data-generating distributions than if a single proportional hazards model with time-varying hazard ratios had been used. However, when using data-adaptive approaches to generate candidate estimators, performing valid statistical inference is challenging. In this work, we utilize the framework of targeted minimum loss-based estimation [van der Laan and Rubin, 2006] to overcome these difficulties and propose two novel estimators of the cumulative incidence function.

The remainder of this article is organized as follows. In Section 1.2, we establish notation and define our statistical estimand of interest. We discuss our framework for estimation in Section 1.3. In Section 1.4, we present an introduction to targeted minimum loss-based estimation and propose two novel estimators of cumulative incidence. In Section 1.5, we discuss the statistical properties of our estimator, including how inference may be performed, and in Section 6, we evaluate the finite-sample performance of the proposed estimators in two simulation studies. We illustrate the methods with data from two HIV vaccine efficacy trials in Section 7. We conclude the chapter with a discussion. The methods herein were implemented using R code that we have made available online. In the supplementary material, we provide a short summary of how to use this code.

1.2 Notation and assumptions

Suppose n subjects are randomly assigned to an active treatment ($Z = 1$) or placebo ($Z = 0$). Baseline covariates (W) are collected prior to randomization and may be used in the randomization procedure, as in a stratified design. Subjects are then followed until the first occurrence of one of k different events or failures whose type we denote by $J \in \{1, ..., k\}$. Let T denote the time from baseline until the first event and assume T is integer-valued, e.g., corresponding to clinic visits until disease diagnosis, and denote the last clinic visit of follow-up by τ. If event ascertainment occurs in continuous time, we may discretize the follow-up period with an arbitrarily fine grid of time points so that no information is lost. Let U denote the censoring time. If a subject completes follow-up without experiencing any event, set $U = \tau$; if a subject drops out during follow-up, $U \in \{1, ..., \tau - 1\}$. The observed data are thus a collection of independent realizations of $O := (\tilde{T}, \Delta J, Z, W) \sim P_0$, where $\tilde{T} := \min\{T, U\}$, $\Delta := I(\tilde{T} = T)$, and P_0 denotes the true distribution of O.

We will often express O in terms of discrete counting processes. Specifically, we will write $O = (W, Z, \{dN(t), dC(t), Y(t) : t = 1, ..., \tau\})$, where $dN(t) := \{dN_1(t), ..., dN_k(t)\} := \{I(\tilde{T} = t, \Delta J = 1), ..., I(\tilde{T} = t, \Delta J = k)\}$, $dC(t) = I(\tilde{T} = t, \Delta J = 0)$, and $Y(t) = I(\tilde{T} \geq t)$. We use $\bar{A}(t) = \{A(1), ..., A(t)\}$ to denote the history of counting process A up until time t. We also let $Pa\{N(t)\} = \{Z, W, \bar{N}(t-1), \bar{C}(t-1)\}$ and $Pa\{C(t)\} = \{Z, W, \bar{N}(t), \bar{C}(t-1)\}$ denote the parents of the counting processes $N(t)$ and $C(t)$ in a graph, respectively. We will use $o(t) := \{w, z, \bar{n}(t), \bar{c}(t), \bar{y}(t)\}$ to denote a realization of the components of O available at time t. We also will utilize the notation Pf to denote $\int f(o)dP(o)$ and $P_n f = \int f(o)dP_n(o)$ where P_n is the empirical cumulative distribution function.

To simplify the exposition, we assume without loss of generality that there exist only two causes of failure, say $J \in \{1, 2\}$. Our goal is estimation of the cumulative incidence of events of type $J = 1$ at a fixed time $t_0 \in \{1, 2, \ldots, \tau\}$ under fixed treatment assignment $Z = z_0 \in \{0, 1\}$, defined as $F_1^{z_0}(t_0) := P(T \leq t_0, J = 1 \mid Z = z_0)$. For notational simplicity, we henceforth adopt subscripts to denote the type of failure, while the superscript z_0 indicates the treatment assignment. Because the labelling of these causes is arbitrary, our results immediately apply to estimation of the cumulative incidence of events of type $J = 2$. If more causes are active, we may select the failure type of interest as $J = 1$ and label all other types as $J = 2$. Since in our developments the choice of t_0 is arbitrary, our results also apply to pointwise estimation of the entire cumulative incidence function.

Our parameter of interest can be written as

$$F_1^{z_0}(t_0) \;=\; E_W\{F_1^{z_0}(t_0; W)\} = \int_w \sum_{t=1}^{t_0} \left[\bar{Q}_1^{z_0}(t; w) \prod_{s=1}^{t-1} \{1 - \bar{Q}_{\cdot}^{z_0}(s; w)\} \right] dQ_W(w) \;, \tag{1.1}$$

where $\bar{Q}_j^{z_0}(t; w) := E\{dN_j(t) \mid z = z_0, c(t-1) = \bar{0}(t-1), \bar{n}(t-1) = \bar{n}(t-1), W = w\}$ is the cause-specific hazard for type j events, $\bar{Q}_{\cdot}^{z_0}(t; w) := \sum_j \bar{Q}_j^{z_0}(t; w)$ is the combined hazard,

Q_W is the distribution function of baseline covariates, and $\bar{0}(t-1)$ denotes the zero vector of length $t-1$. For notational simplicity, we will continue to suppress the conditioning of these quantities on remaining at-risk.

We conclude this section by noting that the observed data parameter (1.1) corresponds to the counterfactual cumulative incidence under conditional independence of treatment assignment, positivity and coarsening at random. These assumptions are discussed in Appendix A.

1.3 Framework for estimation

Note that (1.1) can be viewed as a mapping $\Psi : \mathcal{M} \to [0,1]$, where $\mathcal{M}$ is a nonparametric model. The mapping could be applied to an estimate of P_0, say $\hat{P} \in \mathcal{M}$, to obtain an estimator $\psi_n := \Psi(\hat{P})$ of $\psi_0 := \Psi(P_0)$. We refer to such estimators as substitution estimators and note that these estimators are desirable in that they always obey bounds on the parameter space and will generally exhibit robust finite-sample behavior. We also see in (1.1) that the cumulative incidence function does not depend on all of P but rather only on a portion of P given by $Q = Q(P)$, where $Q(P)$ consists of the cause-specific hazard functions and the distribution of baseline covariates. Thus, we may write the parameter as $Q \mapsto \Psi(Q)$ to emphasize the dependence on Q alone. Practically, this means that in order to estimate cumulative incidence, we need only estimate cause-specific hazards and the distribution of baseline covariates. We need not necessarily estimate the portions of P that pertain to treatment or censoring probabilities, which we generically denote by $g = g(P)$. Nevertheless, we shall see that to construct an estimator Q_n of Q_0 such that $\psi_n = \Psi(Q_n)$ has desirable statistical properties, such as asymptotic linearity and efficiency, an estimate of $g_0 := g(P_0)$ is needed.

Efficiency theory stipulates that a regular estimator ψ_n of ψ_0 is asymptotically efficient if and only if it satisfies that $\psi_n - \psi_0 = P_n D^*(P_0) + o_P(n^{-1/2})$ for some function $D^*(P_0)$, the so-called efficient influence function, such that $P_0 D^*(P_0) = 0$ and $P_0 D^*(P_0)^2 < \infty$ [Bickel et al., 1997]. An application of the central limit theorem then readily implies that, upon appropriate centering and scaling, such an estimator converges in distribution to a normal variate with mean 0 and variance $\sigma_0^2 := P_0 D^*(P_0)^2$. Here, we use the notation $Pf = \int f(o)dP(o)$ for any function f and probability distribution P, and we denote the empirical distribution by P_n. Thus, finding the form of the efficient influence function is a crucial first step in pursuit of an asymptotically efficient estimator. We provide its explicit form below and include a proof in Appendix B.

Let $G_C^{z_0}(t;w) := E\{C(t-1) \mid Z = z_0, W = w\}$ be the conditional probability of remaining uncensored through time $t-1$; $gz(z_0;w) := P(Z = z_0 \mid W = w)$, the probability of receiving treatment z_0 given $W = w$; and $S^{z_0}(t;w) := P(T > t \mid Z = z_0, W = w)$, the conditional probability of not having an event of either type by time t.

Theorem 1. *The efficient influence function for the parameter defined in (1.1) in a non-parametric model is given by*

$$D^{z_0,*}(P_0;o) = \sum_{t=1}^{\tau} D_1^{z_0}(t;o) + \sum_{t=1}^{\tau} D_2^{z_0}(t;o) + D_W(o) \ , \tag{1.2}$$

where

$$\begin{aligned}
D_1^{z_0}(t;o) &:= I(t \le t_0) H_g^{z_0}(t;o)\big\{1 - R_1^{z_0}(t;w)\big\}\big\{dn_1(t) - \bar{Q}_1^{z_0}(t;w)\big\} \\
D_2^{z_0}(t;o) &:= -I(t < t_0) H_g^{z_0}(t;o) R_1^{z_0}(t;w)\left\{dn_2(t) - \bar{Q}_2^{z_0}(t;w)\right\} \\
D_W(w) &:= F_1^{z_0}(t_0;w) - F_1^{z_0}(t_0)
\end{aligned}$$

and we have defined

$$\begin{aligned}
H_g^{z_0}(t;o) &:= \frac{I(z = z_0, n(t-1) = c(t-1) = 0)}{g_z(z_0;w) G_C^{z_0}(t;w)} \ , \\
R_1^{z_0}(t;w) &:= \frac{F_1^{z_0}(t_0;w) - F_1^{z_0}(t;w)}{S^{z_0}(t;w)}.
\end{aligned}$$

With this result, it is possible to show that the Aalen-Johansen estimator is an efficient estimator when no prognostic covariates are available and censoring is random. If prognostic covariates are available, an efficient estimator may be constructed by taking a linear combination of stratum-specific Aalen-Johansen estimators with weights provided by the empirical stratum probabilities. This estimator may nevertheless have poor finite sample performance. For example, if W is high-dimensional, even with reasonably large sample sizes, resulting estimates can be poorly behaved. Furthermore, if W is continuously-valued, stratification is not even possible. We might expect better finite sample performance using a smoothed estimate of Q_0 that borrows information across levels of W. However, statistical inference using estimators employing nonparametric smoothing techniques is generally difficult. Estimators naively constructed in this manner are generally overly biased to be asymptotically linear, let alone efficient. This is because the bias-variance trade-off the smoother must make is different when considering optimal estimation of Q_0 and of ψ_0, and the latter trade-off is generally much more delicate to achieve. To correct this, we must clarify how suboptimal estimation of Q_0 for the sake of ψ_0 results in the suboptimal performance of $\Psi(Q_n)$.

Lemma 1. *Under regularity conditions explicitly stated in Appendix C, we have that*

$$\Psi(Q_n) - \Psi(Q_0) = P_n D^*(Q_0, g_0) - P_n D^*(Q_n, g_n) + o_P(n^{-1/2}). \tag{1.3}$$

From the lemma, we see that if we are able to construct estimators Q_n^* and g_n that satisfy regularity conditions and such that $P_n D^*(Q_n^*, g_n) = 0$, then the substitution estimator $\psi_n^* := \Psi(Q_n^*)$ is asymptotic linear and efficient. Targeted minimum loss-based estimation (TMLE) is designed precisely to produce such an estimate Q_n^*. In the following section we provide a brief introduction to the TMLE framework, illustrating how TMLE produces an estimate Q_n^* such that $P_n D^*(Q_n^*, g_n) = 0$. We subsequently present two novel TMLEs for estimation of cumulative incidence.

1.4 TMLE estimators

The general framework of TMLE was formally introduced in van der Laan and Rubin [2006] as a means of constructing efficient substitution estimators in infinite-dimensional models. An early instance of a TMLE estimator appears in Scharfstein et al. [1999]. We refer readers to van der Laan and Rose [2011] for a full treatment of the topic. TMLE works by updating an initial estimator Q_n of Q_0 using an appropriately-chosen parametric submodel and loss function. Suppose that for given g the function $(Q, o) \mapsto L_g(Q; o) \in \mathbb{R}$ is a loss function for Q_0 in the sense that $Q_0 = \text{argmin}_Q P_0 L_{g_0}(Q)$. Suppose also that, given (Q, g), we can construct a parametric fluctuation submodel $\mathcal{Q}(Q; g) := \{Q_g(\epsilon) : \epsilon\} \subseteq Q(\mathcal{M})$ such that $Q_g(0) = Q$ for every g. The TMLE procedure is generally implemented in the following steps, which we subsequently explain in some detail:

1. compute initial estimates Q_n and g_n, and set $Q_{n,0} = Q_n$;
2. given $Q_{n,k}$, find $\epsilon_{k,n} := \text{argmin}_\epsilon P_n L_{g_n}\{Q_{n,k,g_n}(\epsilon)\}$ and set $Q_{n,k+1} := Q_{n,k,g_n}(\epsilon_{k,n})$;
3. repeat step 2 until $\epsilon_{K,n} \approx 0$ and let $Q_n^* = Q_{n,K}$.

Recall that our goal is to find a good estimator Q_n^* such that $P_n D^*(Q_n^*, g_n) = 0$. This is achieved by making an appropriate choice of submodel-loss function combination. Heuristically, at step k the submodel provides a set of possible Q that lie in a parametric neighborhood of $Q_{n,k}$ indexed by ϵ. At each optimization step, provided the value $\epsilon_{k,n}$ of ϵ that minimizes the empirical risk over this submodel lies in the interior of the set of possible ϵ-values, the equation $(d/d\epsilon) P_n L_{g_n}\{Q_{n,k,g_n}(\epsilon)\} = 0$ is satisfied at $\epsilon = \epsilon_{k,n}$. Suppose that Q_n^* denotes the updated estimate of Q_0 at which convergence occurs. By definition, any attempt to update Q_n^* will yield an optimizer value $\epsilon = 0$ and the equation $(d/d\epsilon) P_n L_{g_n}\{Q_{n,g_n}^*(\epsilon)\} = 0$ will be solved at $\epsilon = 0$. The form of this score equation is key. If the submodel and loss function are chosen such that $(d/d\epsilon) L_g\{Q_g(\epsilon)\}|_{\epsilon=0} = D^*(Q, g)$, then it follows that $P_n D^*(Q_n^*, g_n) = 0$. This indicates that by iterating until the optimizer $\epsilon_{K,n}$ is approximately zero, we arrive at an estimate Q_n^* for which the efficient influence function estimating equation is solved. Provided regularity conditions are satisfied, in view of (1.3), the estimator $\psi_n^* = \Psi(Q_n^*)$ is then asymptotic linear and efficient. This estimator is referred to as 'targeted' because at each step the current estimate of Q_0 is updated to one that is preferred for the purpose of estimating ψ_0.

1.4.1 TMLE in the competing risks setting

A TMLE algorithm for a treatment-specific survival probability in the context of right-censored data was first introduced in Moore and van der Laan [2009]. We seek to generalize these results to the competing risks setting. In order to motivate the development of two distinct TMLEs in the present work, we note that many target parameters can be written as functions of multiple choices of $Q = Q(P)$. The cumulative incidence function is an example of such. In addition to the cause-specific hazards representation given in (1.1), we show in Section 1.4.1 that it also has a representation as a function of iteratively-defined conditional means. As a general rule, whenever possible, it is preferable to select a more parsimonious representation since then it may be easier to construct estimators of the involved Q meeting required regularity conditions. In practice, this may often translate in improved finite-sample performance – see, for example, the discussion in van der Laan and Gruber [2012].

For both TMLEs, we take a component-wise approach. In general, if Q is a function of $(Q_1, Q_2, ..., Q_M)$, we can write $\Psi(Q) = \Psi_1(Q_1, Q_2, ..., Q_M)$ for some mapping Ψ_1. The TMLE procedure can be constructed by first defining a loss function, say $Q_m \mapsto L_{m,Q_{\bar{m}}}(Q_m)$, for each component $m \in \{1, ..., M\}$, possibly indexed by other components of Q, which we denote by $Q_{\bar{m}}$. We also require an appropriate submodel $\mathcal{Q}_m = \{Q_{m,g}(\epsilon) : \epsilon\}$ for $m = 1, ..., M$. These loss functions and submodels must be chosen such that the linear span of the generalized scores of the submodel for $(Q_1, Q_2, \ldots, Q_M)$ contains the efficient influence function. van der Laan and Gruber [2012] provided further details on iteratively defined TMLEs in longitudinal data structures.

Hazard-based TMLE

We first outline a TMLE for the representation of the cumulative incidence function given in (1.1). This representation views the parameter as a function of three components: Q_W,

$\bar{Q}_1^{z_0}$ and $\bar{Q}_2^{z_0}$. We require initial estimates of these three components, as well as estimates of $g_z(z_0)$ and $G_C^{z_0}$. We will use the empirical distribution $Q_{W,n}$ as initial estimator of the baseline covariate distribution. If the treatment mechanism is known to be, say, $g_{z,0}$, as it would be in a randomized trial, then we let $g_{z_0,n} = g_{z_0,0}$. To obtain initial estimates for $\bar{Q}_{1,0}^{z_0}$, $\bar{Q}_{2,0}^{z_0}$, $G_{C,0}^{z_0}$, and (possibly) $g_{z,0}$, we make use of ensemble machine-learning learning techniques, such as the Super Learner, and employ a wide library of candidate estimators, including both parametric and nonparametric estimators [van der Laan and Dudoit, 2003, van der Laan and Polley, 2007].

We note that, for a given time t, $\bar{Q}_1^{z_0}(t;w)$ and $\bar{Q}_2^{z_0}(t;w)$ cannot sum to more than one, implying that they are not variationally independent. Thus, if these two parameters are estimated independently, it may well be that no underlying $P \in \mathcal{M}$ is compatible with estimates of both $\bar{Q}_1^{z_0}$ and $\bar{Q}_2^{z_0}$. This would stymie our efforts to develop a true substitution estimator. Nonetheless, a reparametrization involving variationally independent components is possible. Indeed, for each $t = 1, 2, \ldots, \tau$, defining pointwise $\check{Q}_2^{z_0}(t;w) := E\{dN_2(t) \mid Z = z_0, C(t-1) = 0, W = w, N_1(t) = 0, N_2(t-1) = 0\}$, we observe that $\check{Q}_2^{z_0}$ is variationally independent of $\bar{Q}_1^{z_0}$, and furthermore, $\bar{Q}_2^{z_0}(t;w)$ can be recovered using the fact that

$$\bar{Q}_2^{z_0}(t;w) = \check{Q}_2^{z_0}(t;w)\left\{1 - \bar{Q}_1^{z_0}(t;w)\right\} \ . \tag{1.4}$$

This is simply the reparametrization of a joint probability as the product of marginal and conditional probabilities.

Now, we must define appropriate parametric submodels and loss functions in order to update the estimators $\bar{Q}_{1,n}^{z_0}$, $\bar{Q}_{2,n}^{z_0}$ and $Q_{W,n}$ we have in hand. For Q_W, we define the submodel $\mathcal{Q}_W(Q_W) = \{(1 + \epsilon_W D_W)Q_W : \epsilon_W\}$, where D_W is as in Theorem 1. The score for ϵ_W using the negative log-likelihood loss at $\epsilon_W = 0$ is simply D_W. For this choice of submodel and loss function, $\epsilon_{W,n} = 0$ at the first step as a result of having used the empirical distribution function $Q_{W,n}$ and thus the nonparametric maximum likelihood estimator of $Q_{W,0}$. Thus, our estimate of $Q_{W,0}$ requires no updating at all.

For type 1 hazards at each time, we may define pointwise the negative log-likelihood loss

$$L(\bar{Q}_1^{z_0}; o) := -I\{n(t-1) = c(t-1) = 0, Z = z_0\}\times \\ \big[dn_1(t)\,\log\{\bar{Q}_1^{z_0}(t;w)\} + \{1 - dn_1(t)\}\,\log\{1 - \bar{Q}_1^{z_0}(t;w)\}\big]$$

and submodel

$$\bar{Q}_1^{z_0}(t;w)(\epsilon_1;g) := \operatorname{expit}\big[\operatorname{logit}\{\bar{Q}_{1,n}^{z_0}(t;w)\} + \epsilon_1 H_g^{z_0}(t;o)\big\{1 - R_1^{z_0}(t;w)\big\}\big]$$

. It is easy to verify that the true conditional hazard for type 1 events minimizes $P_0 L(\bar{Q}_1^{z_0})$ and that

$$\frac{d}{d\epsilon_1} L\{\bar{Q}_1^{z_0}(t;w)(\epsilon_1)\}\bigg|_{\epsilon_1=0} = D_1^{z_0}(t;o)\ .$$

An appropriate loss function for $\check{Q}_2^{z_0}$ indexed by $\bar{Q}_1^{z_0}$ is given by

$$L_{\bar{Q}_1^{z_0}}(\check{Q}_2^{z_0}; o) := -I\{n(t-1) = c(t-1) = 0, z = z_0\}\times \tag{1.5}$$
$$\left[\frac{dn_2(t)}{1-\bar{Q}_1^{z_0}(t;w)}\log\left\{\frac{\bar{Q}_2^{z_0}(t;w)}{1-\bar{Q}_1^{z_0}(t;w)}\right\} + \left\{1 - \frac{dn_2(t)}{1-\bar{Q}_1^{z_0}(t;w)}\right\}\log\left\{1 - \frac{\bar{Q}_2^{z_0}(t;w)}{1-\bar{Q}_1^{z_0}(t;w)}\right\}\right]$$

with $\bar{Q}_2^{z_0}$ defined in terms of $\bar{Q}_1^{z_0}$ and $\check{Q}_2^{z_0}$ via (1.4), whereas an appropriate submodel is

$$\check{Q}_2^{z_0}(t;w)(\epsilon_2;g) = \operatorname{expit}\big[\operatorname{logit}\{\check{Q}_2^{z_0}(t;w)\} + \epsilon_2 H_g^{z_0}(t;o) R_1^{z_0}(t;w)\big]$$

, where $\bar{Q}_1^{z_0}$ is considered fixed. We can verify that when $\bar{Q}_1^{z_0}$ equals the true value $\bar{Q}_{1,0}^{z_0}$, the minimizer of the expected loss is indeed the true value of $\check{Q}_2^{z_0}$ and that

$$\frac{d}{d\epsilon_2} L_{\bar{Q}_1^{z_0}}\{\check{Q}_2^{z_0}(t;w)(\epsilon_2)\}\bigg|_{\epsilon_2=0} \propto D_2^{z_0}(t;o)\ .$$

It follows that when we aggregate the three loss functions above and sum them over time, the generated score at zero will indeed span the full efficient influence function, as required.

A useful feature of these particular choices of loss functions and submodels is that easily accessible software may be utilized in the minimization step. Indeed, $\epsilon_{1,n}$ can be obtained as the estimated coefficient in a logistic regression of the binary outcome $dn_1(t)$ on covariate $H_{1,n}(t;o)$ with offset $\operatorname{logit}\{\bar{Q}_{1,n}^{z_0}(t;w)\}$. Similarly, $\epsilon_{2,n}$ can be obtained by fitting a logistic regression of the binary outcome $dn_2(t)$ on covariate $H_{2,n}(t;o)$ with offset $\operatorname{logit}[\check{Q}_{2,n}^{z_0}(t;w)\{1 - \bar{Q}_{1,n}^{z_0}(t;w)\}]$. Both regressions are fit using only the subset of the data with $c(t-1) = 0$ and $z = z_0$. A detailed summary of how to implement this TMLE is provided in Appendix D.

Mean-based TMLE

We now present a TMLE based on an alternative representation of cumulative incidence as a function of iteratively-defined conditional means. This representation stems from general results on identifying causal parameters in the context on longitudinal data structures [Bang and Robins, 2005]. First define $\tilde{Q}_1^{z_0}(t_0; w)$ as the conditional mean of $N_1(t_0)$ given $N(t_0-1) = n(t_0-1)$, $W = w$, $C(t_0-1) = 0$ and $Z = z_0$. Then, for $r = 1, 2, \ldots, t_0 - 1$, recursively set $\tilde{Q}_1^{z_0}(t_0-r; w)$ equal to the conditional mean of $\tilde{Q}^{z_0}(t_0; w)$ given $N(t_0) = n(t_0-r-1)$, $W = w$, $C(t_0-r-1) = 0$ and $Z = z_0$. Finally, define $\tilde{Q}_1^{z_0}(0)$ to be the average of $\tilde{Q}_1^{z_0}(1; w)$ over the distribution of W. We find that $\tilde{Q}_1^{z_0}(1; w) = F_1^{z_0}(t_0; w)$ and $\tilde{Q}_1^{z_0}(0) = F_1^{z_0}(t_0)$; an explicit derivation is provided in Appendix B.2. Note that, in order to keep notation consistent, we have suppressed the dependence of on $\tilde{Q}_1^{z_0}(t; w)$ on $n(t_0-1)$. Trivially, one may set $\tilde{Q}_1^{z_0}(t; w) = 1$ whenever $n_1(t-1) = 1$ and $\tilde{Q}_1^{z_0}(t; w) = 0$ whenever $n_2(t-1) = 1$. See (B.5) for a useful representation for these means; this representation will be utilized extensively in Chapter 2.

We have thus represented the cumulative incidence of type 1 failures at the t_0-th time is a function of t_0 iteratively-defined conditional means and the distribution of baseline covariates. We now adapt the general framework of van der Laan and Gruber [2012] to define a TMLE based on this representation. We again require loss functions and parametric submodels yielding scores that span the efficient influence function. For the distribution of baseline covariates, we use the same submodel and loss as for the hazard-based approach, which again yields the score D_W. For the conditional mean at each time $t = t_0, ..., 1$, we define the loss function indexed by $\tilde{Q}^{z_0}(t+1; w)$ and given by

$$\begin{aligned} &L_{\tilde{Q}_1^{z_0}(t+1)}(\tilde{Q}_1^{z_0}(t); o) := \\ &\quad - I\{z = z_0, c(t-1) = 0\} \left[\tilde{Q}_1^{z_0}(t+1; w) \log\bigl\{\tilde{Q}_1^{z_0}(t; w)\bigr\} + \right. \\ &\qquad\qquad\qquad \left. \{1 - \tilde{Q}_1^{z_0}(t+1; w)\} \log\bigl\{1 - \tilde{Q}_1^{z_0}(t; w)\bigr\}\right]. \end{aligned}$$

The true risk from this loss function will be minimized by the true value $\tilde{Q}_{1,0}^{z_0}(t)$ if the index parameter equals the true value $\tilde{Q}_{1,0}^{z_0}(t+1; w)$. Here, we can use the submodel

$\tilde{Q}_1^{z_0}(t;w)(\epsilon_t;g) = \text{expit}[\text{logit}\{\tilde{Q}_1^{z_0}(t;w)\} + \epsilon_t H_g(t;o)]$, which under the loss specified above, yields a generalized score at zero equal to $D_1^{z_0}(t;o) + D_2^{z_0}(t;o)$ at each time t. Thus, the sum over all time points along with D_W gives a score equal to the efficient influence function. We note that again we have chosen computationally convenient loss functions. At any t, we obtain $\epsilon_{t,n}$ as the estimated coefficient in a logistic regression with outcome $\tilde{Q}_1^{z_0}(t+1;w)$, offset $\tilde{Q}_1^{z_0}(t;w)$, and covariate $H_g(t;o)$ fit using observations with $z = z_0$ and $c(t-1) = n_1(t-1) = n_2(t-1) = 0$. A description of this TMLE algorithm is detailed in Appendix D, where we also compare the implementation of the two proposed TMLEs.

1.5 Statistical Properties and Inference

The TMLE estimators discussed share many statistical properties, so in this section we will use Q_0 to generically denote either the conditional cause-specific hazards or iteratively-defined conditional means. First, both TMLEs presented are examples of double-robust estimators and will be consistent if either of the relevant estimators Q_n and g_n of Q_0 and g_0 is consistent [van der Laan and Gruber, 2012]. This is a useful property for an estimator, ensuring some robustness against misspecification.

Both estimators will be asymptotically linear under certain regularity conditions, as stated in Appendix C. If the treatment and censoring mechanisms are known exactly, the influence curve of ψ_n is given by $D^*(Q_0^*, g_0)$, where Q_0^* is the (possibly misspecified) limit of the TMLE Q_n^*. In this case, the asymptotic variance of $n^{1/2}(\psi_n^* - \psi_0)$ can be consistently estimated by $\sigma_n^2 = P_n D^*(Q_n^*, g_0)^2$. If g_0 is unknown, the asymptotic variance is more complicated. However, if g_n is an asymptotically efficient estimator within a parametric model $\mathcal{G}$, we may use σ_n^2 as a conservative estimate of the asymptotic variance (see, e.g., Theorem 2.3 of van der Laan and Robins, 2003). The Wald interval $\left(\psi_n - z_{1-\alpha/2}\sigma_n n^{-1/2}, \psi_n + z_{1-\alpha/2}\sigma_n n^{-1/2}\right)$, where z_β is the β-quantile of the standard normal distribution, will have asymptotic coverage no smaller than $1-\alpha$, with equality when $g_n = g_0$. Similarly, given a fixed $\psi^\circ \in [0,1]$, a two-sided test of the null hypothesis $\psi_0 = \psi^\circ$ of asymptotic size no larger than α can be constructed by rejecting the null hypothesis whenever $|n^{1/2}(\psi_n - \psi^\circ)/\sigma_n| > z_{1-\alpha/2}$.

An additional strength of the influence curve-based approach is that the form of the asymptotic variance of a function of multiple estimators is readily available. This is useful for assessing whether treatment efficacy against the two types of events is different, a common question in competing risks analyses. We can easily define a Wald test of the hypothesis that

$$\frac{F_1^1(t_0)}{F_1^0(t_0)} = \frac{F_2^1(t_0)}{F_2^0(t_0)} \tag{1.6}$$

via simple application of the delta method.

1.6 Simulation Study

1.6.1 Design

We simulated data to mimic a randomized placebo-controlled trial, where n subjects are allocated one-to-one to treatment or placebo at baseline. A covariate predictive of both censoring and events, $W_1 \sim N(0,1)$, and a binary effect modifier of treatment efficacy, $W_2 \sim Bernoulli(\theta)$, are also measured at baseline. Subjects are followed for two years, with events and censoring determined every three months, so that $\tau = 8$. We assume that there are two causes of failure and focus on studying type 1 failures. At each time $t = 1, ..., \tau$ the hazard for an event of type j is given by $\bar{Q}^z_{j,0}(t; w) = \text{expit}(\beta_0 + \beta_1 z + \beta_2 w_1 + \beta_3 z w_2)/2 \;\; j = 1, 2$. The censoring curvival distribution at each time $t = 1, ..., \tau - 1$ is given by $G^{z_0}_{C,0}(t; w) = \prod_{s=1}^{t} \big\{1 - \text{expit}(\gamma_0 + \gamma_1 w_1)\big\}$. We generated a failure time based on the combined hazard for events of either type. A cause of failure was then randomly assigned with the probability that the event was of type j equal to the ratio of the j-specific hazard to the total hazard at that time. This resulted in event types being assigned with equal probability. We then generated a censoring time and let the observed data be the minimum of the censoring and failure time. Ties between event and censoring times were deemed to be an event.

In both of our simulation studies, we fixed the expected number of events at 250 and the two-year cumulative incidence of both event types combined in the placebo arm at 15% by selecting appropriate values for n and β_0. These event rates match the expectation for a typical HIV vaccine efficacy trial. Initial estimates for both TMLEs were generated using the Super Learner algorithm [Polley and van der Laan, 2013, van der Laan and Polley, 2007]. Tables E.1 and E.1 give values for all simulation parameters and Table E.2 shows all candidate estimators included in the Super Learner library.

1.6.2 Simulation Study 1

The first simulation study addressed the performance of the two TMLEs relative to the Aalen-Johansen estimator for estimating the cumulative incidence function in the placebo arm at a fixed time. For each scenario considered, we generated 1,000 data sets and calculated estimates of $F_1^0(8)$ using the Aalen-Johansen estimator and both TMLEs. We compared Monte Carlo bias, coefficient of variation, and mean squared error (MSE) of the estimators.

We considered four censoring scenarios consisting of pairwise combinations of low/high levels of censoring and non-/strongly informative covariates for censoring. The low and high censoring levels respectively corresponded to 10% and 30% dropout rate over the two year study; this was achieved by varying γ_0. The predictiveness of censoring was determined by γ_1, which was set equal to 0 and $\log(3.2)$ to generate non- and strongly-informative censoring. For each of the four censoring scenarios, we varied the predictiveness of W_1 on events by letting $\beta_1 = 0$, $\log(1.6)$, and $\log(3.2)$, corresponding to non-, moderately-, and strongly-informative covariates. For simplicity and because we were only examining estimation in one treatment arm, we set $\beta_2 = 0$ so that there was no effect modification.

Results

Figures 1.1, 1.2, 1.3 show the results of the first simulation study. In Figure 1.1 we see the bias of the estimators as a percentage of the truth. Both TMLE estimators are approximately unbiased in every setting, while the Aalen-Johansen estimator is biased in the scenarios with high levels of dependent censoring. The bias in the estimator is quite substantial when confounding of T and C by W_1 is strongest. Figure 1.2 shows the relative coefficient of variation for the two TMLE estimators relative to the Aalen-Johansen estimator. Thus, a value lower than one indicates less variability of the TMLE estimator relative to its mean than the Aalen-Johansen estimator. We see that as covariate predictiveness of events increases, the relative variability of the TMLE estimates becomes substantially smaller than that of the Aalen-Johansen estimator in each case. Figure 1.3 shows the relative mean squared error for the two TMLEs compared to the Aalen-Johansen estimator, showing that the TMLE is performing at least as well as Aalen-Johansen in every scenario except one. Figure E.3 shows a plot illustrating the performance of the influence curve-based standard error estimates, which are accurate.

Overall, the simulations support the theory developed herein. When covariates informative of events are available, TMLE is more efficient than Aalen-Johansen, even when censoring is non-informative. When censoring is informative, the performance of the Aalen-Johansen estimator suffers more; the estimator is no longer consistent and thus shows substantial finite sample bias. It is also of note that even in cases where covariates are non-informative of both censoring and events we find that for the most part TMLE performs as well as Aalen-Johansen. Recall that these are the scenarios where Aalen-Johansen is the efficient estimator. We are thus reassured that there is little to lose by using TMLE even when covariates turn out to be non-informative.

Additional simulations that explore the performance of TMLE for data generating mechanisms when covariates are measured subject to measurement error are included in Appendix F. Overall, the TMLE still performs well in this situation.

1.6.3 Simulation Study 2

In the second simulation, we examined the power of TMLE-based Wald tests to reject the null hypothesis of no vaccine efficacy against type 1 events at time t_0, defined as

$$VE(t_0, 1) := 1 - \frac{F_1^1(t_0)}{F_1^0(t_0)}. \tag{1.7}$$

We constructed one-sided, level 0.025, Wald tests using influence curve-based standard errors as described in Section 1.5. The standard error of the log ratio of cumulative incidences was calculated using the delta method so that the null hypothesis was rejected whenever

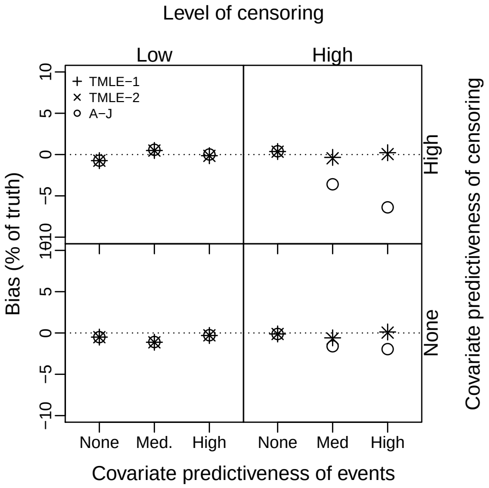

Figure 1.1: Bias of the hazard-based TMLE (TMLE-1), mean-based TMLE (TMLE-2), and the Aalen-Johansen (A-J) estimator. The bias is presented as a percentage of the true cumulative incidence function.

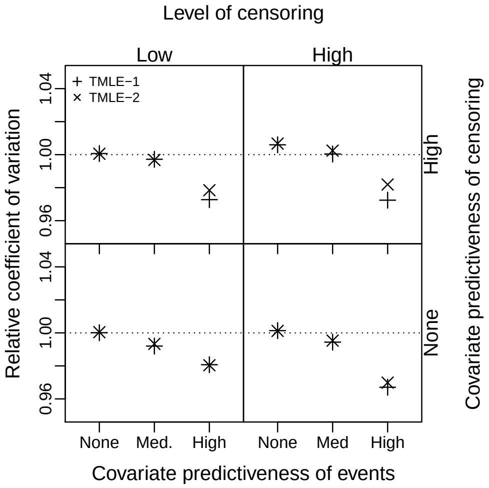

Figure 1.2: Relative coefficient of variation of the hazard-based (TMLE-1) and mean-based (TMLE-2) TMLE estimators compared to the Aalen-Johansen (A-J) estimator. Values below 1.0 indicate a smaller coefficient of variation for the TMLE estimator.

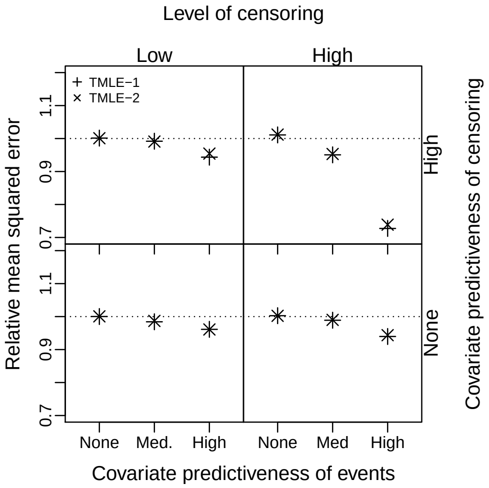

Figure 1.3: Relative mean squared error of the hazard-based (TMLE-1) and mean-based (TMLE-2) estimators compared to the Aalen-Johansen (A-J) estimator. Values below 1.0 indicate a smaller mean squared error for the TMLE estimator.

$log(\widehat{RR})/\widehat{SE}\{log(\widehat{RR})\} < z_{0.025}$, where $log(\widehat{RR}) := log\{\hat{F}_1^1(t_0)/\hat{F}_1^0(t_0)\}$.

We examined the power of the Wald test under true vaccine efficacies of 0, 0.25, 0.50, and 0.75. The data were generated with a two year censoring rate of 10% and covariates strongly-informative of events and moderately informative of censoring, as described above. We conjectured that accounting for effect modifiers would provide greater efficiency gain than accounting for prognostic covariates and we examined this conjecture by allowing for effect modification. This was accomplished by varying β_1 and β_3, while setting $\theta = 0$ so that the treatment was more effective for half of the population on average. The conditional vaccine efficacies in the two subgroups were {0.5,-0.5}, {0.5,0}, {0.8,0.15}, {0.90,0.60} in the $W_2 = 0$ and $W_2 = 1$ subgroups, respectively. These values yielded unconditional vaccine efficacies of 0, 0.25, 0.50, and 0.75, corresponding with the same unconditional VEs in the no effect modification scenario.

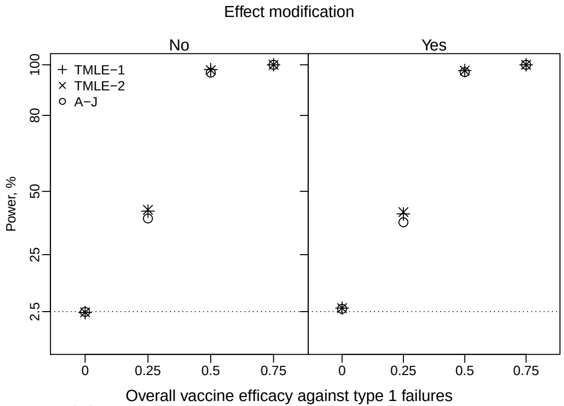

Power relative to A−J

TMLE−1	1.07	1.01	1.00	1.09	1.01	1.00
TMLE−2	1.09	1.01	1.00	1.11	1.01	1.00

Figure 1.4: Power of a one-sided level $\alpha = 0.025$ (dotted line) Wald test of the null hypothesis of no vaccine efficacy against type 1 events at the end of follow up: $VE(8,1) := 1 - F_1^1(8 \mid Z = 1)/F_1^0(8) = 0$. Tests are based on hazard-based TMLE (TMLE-1), mean-based TMLE (TMLE-2), and the Aalen-Johansen (A-J) estimators. The left panel shows power with no effect modification; the right panel shows power with effect modification.

Results

Figure 1.4 shows the results of the second simulation study. We see that the TMLE-based tests give correct sizes and offer a slight gain in power over A-J based tests for moderately effective treatments. We see that indeed a slightly larger gain in power is achieved by accounting for effect modifiers (right panel) rather than prognostic covariates alone.

1.7 Examples

The RV144 ALVAC/AIDSVAX vaccine efficacy trial was a randomized, placebo-controlled HIV-1 vaccine efficacy trial [Rerks-Ngarm et al., 2009]. Vaccine efficacy was estimated to be 31.2% (95% CI: 1.1-51.2; p=0.04) using a Cox-proportional hazards-based definition of vaccine efficacy. Secondary objectives were to assess whether the vaccine efficacy differed based on the genotype of the infecting virus. A previous analysis based on the cause-specific Cox model [Lunn and McNeil, 1995, Prentice et al., 1978] showed differential protection for

HIV genotypes defined by sites 169 and 181 of the V2 Envelope protein [Rolland et al., 2012]. Given that waning vaccine efficacy is a general concern [Durham et al., 1999] and occurred for overall vaccine efficacy in RV144 [Robb et al., 2012], it is of interest to assess the vaccine's genotype-specific efficacy over time. The cumulative incidence vaccine efficacy parameter defined in (1.7) above is therefore very relevant. We reexamined the RV144 data using our proposed methods to estimate vaccine efficacy against HIV genotypes that were matched and mismatched to the vaccine at site 169.

We present results for the mean-based TMLE; the results for the hazard-based TMLE were very similar. Initial estimates for censoring and iteratively-defined conditional means were generated using Super Learner. Table G shows the models included in each library. These models adjusted for various combinations of gender, age, and behavioral risk. Figure 1.5 shows the estimated cumulative incidences in the vaccine and placebo groups at six month time intervals as well as the estimated vaccine efficacy at the end of the study. The estimated ratio of cumulative incidence ratios is 2.43 (95% CI: 0.97-6.09, p=0.06 for test of hypothesis defined in (1.6)) supporting the hypothesis that the vaccine was more effective against 169-matched infections. The TMLE methods provided almost identical answers as Aalen-Johansen in terms of point estimates as well as confidence interval widths. This likely reflects the fact that only weak predictors of HIV-infection and censoring were collected in this trial that did not capture the true heterogeneity of risk across subjects.

We also analyzed a more recent preventive HIV vaccine efficacy trial (HVTN 505) with design very similar to RV144 with randomization 1:1 to a candidate vaccine versus placebo [Hammer et al., 2013]. We again used Super Learner to generate initial estimates and censoring adjusting for self-reported sexual risk behaviors (details in Appendix H). Results are shown in Figure H in the same format as Figure 1.5. The 95% confidence intervals all include zero, supporting that this vaccine did not confer efficacy against either 169 matched nor 169 mismatched HIV infection at any time-point. The estimated standard errors for the TMLE estimator were generally 1-3 percent smaller than the Aalen-Johansen estimator (Figure H), which is consistent with our simulation results showing modest efficiency gains

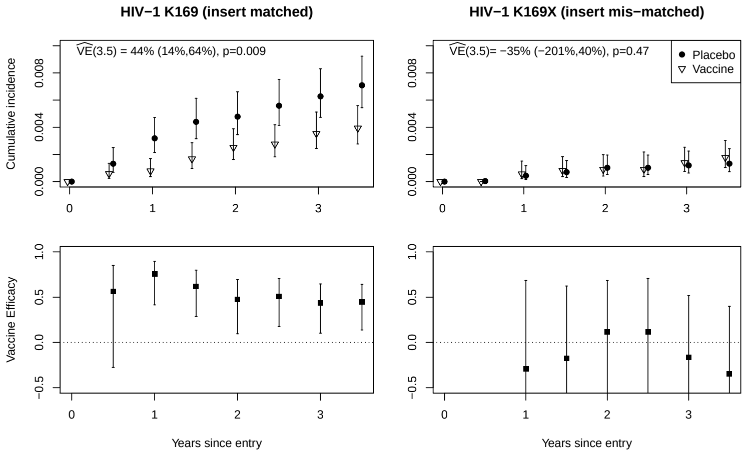

Figure 1.5: Results from the RV 144 HIV vaccine efficacy trial. The top row shows estimated cumulative incidence over time using the hazard-based TMLE for 169-matched and mismatched HIV-infections with point-wise 95% confidence intervals. The bottom row shows estimated vaccine efficacy (VE) and 95% confidence interval (CI) for each type of infection.

when covariates are moderately informative.

1.8 Discussion

In this article, we proposed two novel estimators of cumulative incidence in the setting of competing risks. The estimators developed overcome several shortcomings of previous estimators. Our estimator is able to account for prognostic covariates and allows censoring to depend on covariates while still targeting an estimand with a desirable model-agnostic interpretation. This is advantageous compared to existing semiparametric methods that require correct model specification for interpretability. Our simulations demonstrate superior performance to the nonparametric Aalen-Johansen estimator in settings where prognostic covariates and dependent censoring are possibly present. Furthermore, we illustrated that our estimators perform well even when covariates are not informative of events or censoring, providing confirmation that there is no risk (but only potential gain) to using the more involved methods proposed. Though the efficiency gains seen in simulations were modest, we emphasize that this is not an indication of inadequate asymptotic performance of the estimators – so long as regularity conditions are satisfied, they will achieve the efficiency bound. The magnitude of the efficiency gains should thus be interpreted as the amount by which prognostic covariates modulate the efficiency bound.

The two proposed TMLE estimators performed equally well in simulations. However, in practice we find the mean-based TMLE more appealing than the hazard-based TMLE. The hazard-based TMLE requires consistent estimation of conditional hazards, which are difficult quantities to estimate. The mean-based TMLE only requires estimates of conditional means, which are of lower dimension and are easier to estimate well in practice.

Acknowledgements

Research reported in this publication was supported by the National Cancer Institute and National Institute Of Allergy And Infectious Diseases (NIAID) of the National Institutes of Health (NIH) under Award Numbers T32CA09168, R37AI054165 and UM1AI068635. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH. The authors thank the participants, investigators, and sponsors of the HVTN 505 and RV144 trials, for the latter including Morgane Rolland and the U.S. Military HIV Research Program (MHRP); U.S. Army Medical Research and Materiel Command; NIAID; U.S. and Thai Components, Armed Forces Research Institute of Medical Science; Ministry of Public Health, Thailand; Mahidol University; Sanofi Pasteur; and Global Solutions for Infectious Diseases.

Chapter 2

IMPROVED ESTIMATION OF THE CUMULATIVE INCIDENCE OF RARE OUTCOMES

2.1 Introduction

In many areas of research, outcomes may be very uncommon in the population of interest. This can pose statistical problems for estimation of the causal effect of a treatment on the outcome. For example, in preventive HIV vaccine efficacy trials, researchers often compare the cumulative incidence of HIV between the vaccine and placebo arms at some meaningful time after trial enrollment. However, the incidence of HIV in the target population may be quite low; for example, in a recent landmark HIV vaccine trial in Thailand, less than 1% of subjects in both the placebo and vaccine arms were infected over three and a half years of follow-up [Rerks-Ngarm et al., 2009]. More difficult yet, subsequent studies were interested in comparing the cumulative incidence of HIV with specific genetic patterns between vaccine and placebo [Rolland et al., 2012]. Some of these genetic patterns are extremely rare, and yielded an estimated cumulative incidence of less than 0.1% in both the vaccine and placebo arms. Rare events are also common in observational data with longitudinal follow-up. Cancer registries collect information on many rare cancers and researchers are often interested in, for example, comparing survival between different groups of patients [Gatta et al., 2006]. Observational data poses even greater statistical challenges for estimation of causal effects owing to the necessity of adjusting for a potentially large number of confounders. In some cases, the number of putative confounders may nearly equal the number of events in the observed data. This can lead to many problems with standard estimators, such as instability, incorrect type I error rates, and incorrect confidence interval coverage [Vittinghoff and McCulloch, 2007]. This leaves researchers in a difficult position, forced to consider a trade-off between the validity and stability of an estimator.

It is often true that researchers have a-priori knowledge that the outcome of interest is uncommon; this could be due to previous research in similar populations, the etiology of a disease, surveillance data, or any number of other sources. Heuristically, such knowledge might be useful in constructing estimators of incidence. However, few estimators have been proposed that take this information into account. Recent work by Gruber and van der Laan [2010] proposed improved estimators of treatment effects on a bounded continuous outcome. This work was extended to the binary endpoint setting by Balzer and van der Laan [2013], where it was shown that knowledge of an upper bound on the conditional probability of an outcome can lead to large gains in stability and efficiency in finite samples. The present work discusses an extension of these methods to the setting of survival analysis and estimation of survival and cumulative incidence probabilities. Specifically, we provide two new estimators designed to incorporate bounds in the survival setting. These estimators build directly on the results of Chapter 1.

The remainder of this chapter is organized as follows. In Section 2.2, we discuss the effect of a bounded statistical model on the estimation problem. In Section 2.3, we show how the TMLE methodology developed in Chapter 1 may be extended to account for bounded statistical models, while in Section 2.4 we discuss practical considerations for choosing bounds and estimators. Section 2.5 presents two simulation studies to evaluate the performance of the new estimators and Section 2.6 shows their application to the RV144 data set that was introduced in Chapter 1.

Before proceeding, it bears mentioning several key points. First, we discuss the problem of estimation of cumulative incidence at a given time point of interest, t_0. As in Chapter 1, the selection of this timepoint is arbitrary, so that our results directly apply to point-wise estimation of the cumulative incidence function. Second, although we will work within the general framework of competing risks, our results trivially apply to special case where there is only one active failure type. Throughout this chapter, it is thus convenient to work in the general competing risks setting. When multiple failure types are active, the interpretation of a cause-specific survival probability is often subject to extremely strong assumptions [Prentice et al., 1978]. Thus, we focus our discussion on estimation of cumulative incidence rather than survival. Note that if one type of failure is present (or one is willing to make the untestable assumptions discussed, for example, in Prentice et al. [1978]), the results apply equally well to estimation of a (cause-specific) survival probability at a given time point, defined simply as one minus the cumulative incidence at that time point.

2.2 Implications of a Bounded Statistical Model

In the previous chapter, we showed that cumulative incidence could be written as a functional of one of two types of intermediate quantities: cause-specific hazards or iteratively-defined conditional means. In chapter 1, we first defined the cumulative incidence as a mapping $\bar{\Psi} : \bar{\mathcal{Q}} \to [0,1]$, where $\bar{\mathcal{Q}}$ is a nonparametric model for the cause-specific hazards and distribution of baseline covariates. We also showed that it may be expressed as a different mapping $\tilde{\Psi} : \tilde{\mathcal{Q}} \to [0,1]$, where $\tilde{\mathcal{Q}}$ is the nonparametric model for the iteratively-defined conditional means. For convenience, we use Q to generically refer to either the cause-specific hazards or iteratively-defined means when making points that are relevant to both cases.

The present problem is to describe how incorporating bounds in the model for Q affects the estimation problem. We relegate the discussion of practical considerations for bound selection to Section 2.4. Consider the model $\bar{\mathcal{Q}}^b$ that imposes bounds on the cause-specific hazards:

$$\bar{\mathcal{Q}}^b := \left\{ \bar{Q} \in \bar{\mathcal{Q}} \ : \bar{l}_j(t) \le \bar{Q}_j^{z_0}(t;w) \le \bar{u}_j(t) \ \forall \ w, \ t = 1, ..., t_0, \text{ and } j = 1,2 \right\}.$$

Similarly, we define $\tilde{\mathcal{Q}}^b$ as a model that imposes bounds on the iteratively-defined means, $\tilde{Q}_1^{z_0}$:

$$\tilde{\mathcal{Q}}^b := \left\{ \tilde{Q} \in \tilde{\mathcal{Q}} \ : \tilde{l}_1(t) \le \tilde{Q}_1^{z_0}(t;w) \le \tilde{u}_1(t) \ \forall \ w \text{ and } t = 1, ..., t_0 \right\} .$$

Naïvely, it might be expected that estimation in a bounded statistical model, $\mathcal{Q}^b$, would be an easier proposition than estimation in $\mathcal{Q}$. However, asymptotically this is not the case;

the efficiency bound for estimation of ψ_0 is exactly the same in $\mathcal{Q}^b$ as in $\mathcal{Q}$. A heuristic justification of this is that with an infinite sample size, one observes many rare events, so the knowledge that the incidence of an event is rare does not aid estimation. However, in finite samples, we might expect very different performance between estimators that are constructed to obey the constraints implied by $\mathcal{Q}^b$ and estimators that do not respect these constraints. If estimators perform erratically when events are rare, we may expect the enforcement of bounds to decrease the variance, by preventing nonsensical outlying values.

Another argument that has been previously made in favor of using bounded statistical models when constructing estimators is that the variance of the efficient influence function at a misspecified $Q \in \mathcal{Q}^b$ will be smaller than the variance of the efficient influence function at a misspecified $Q \in \mathcal{Q}$ [Balzer and van der Laan, 2013]. This gives asymptotic motivation for bounded statistical models in practice; however, several caveats apply. Obviously, this result requires that the unconstrained estimator in question converges to a $Q \notin \mathcal{Q}^b$, which may not be the case. Furthermore, when Q is misspecified, asymptotic linearity of a standard efficient estimator (e.g., the TMLE presented in Chapter 1) requires that g_n be a $n^{1/2}$-consistent estimator of g_0. This may occur when, for example, $C \perp W$ or when W is low-dimensional and discrete. However, if censoring is dependent or W is continuous-valued, one could assume a finite-dimensional parametric model and use maximum likelihood estimation to construct a $n^{1/2}$-consistent estimator of g_0. However, as noted in Chapter 1, it is often the case that the state of scientific knowledge precludes construction of such finite-dimensional models. Thus, one is motivated to use nonparametric models and data-adaptive (e.g., ensemble machine learning) estimation techniques for constructing estimators of g_0. Using these methods, one may stand a better chance of constructing a consistent estimator of g_0; however, the rates achieved by these methods will be slower than $n^{1/2}$. This in turn prevents asymptotic linearity of naïvely constructed efficient estimators. Thus, in all but the simplest cases, the putative gains in asymptotic efficiency at a misspecified Q are rendered obsolete by the far-greater concern that estimators will no longer be asymptotically linear in such cases. We explore this situation in great detail in Chapter 3 and show how asymptotically linear estimators may be constructed under model misspecification.

We now proceed to describing how the TMLE algorithms presented in Chapter 1 may be modified to incorporate these bounds.

2.3 Constructing TMLE Estimators with Bounds

Recall from Chapter 1 that targeted minimum loss-based estimation is a two-step procedure. The first step consists of constructing initial estimates of relevant portions of P_0, while the second step consists of a bias-reduction procedure wherein these initial estimates are revised in such a manner as to both improve the empirical risk of the estimate, in addition to solving the efficient influence function estimating equation. The fluctuation submodel must be constructed such that the generalized score at zero fluctuation yielded a portion of the canonical gradient, so that, upon convergence of the bias-reduction algorithm, a portion of the efficient influence function estimating equation is solved. The asymptotics of the resulting plug-in estimator rely on this condition to achieve asymptotic linearity (1.3).

We are now faced with the task of modifying the TMLE algorithms presented in Chapter 1 in such a way that both initial estimates and fluctuation submodels respect the bounds of the statistical model. We will accomplish this by modifying the loss functions presented in Gruber and van der Laan [2010] and Balzer and van der Laan [2013] to the longitudinal setting. We will see that these clever loss functions suffice for both constructing and fluctuating initial estimates in such a way as to yield the correct generalized score. Thus, the asymptotic behavior of these restricted substitution estimators will be identical to the unrestricted TMLE estimators established in Chapter 1.

2.3.1 Hazard-based TMLE

For simplicity, we will again consider the case that $J = 2$ and we are interested in estimating the incidence of type-1 failures, $j_0 = 1$, at some relevant time, t_0. We begin by discussing a loss function that may be used to construct and flucutate estimates of $\bar{Q}_1^{z_0}$. Consider the modified pointwise negative log-likelihood loss

$$L(\bar{Q}_1^{z_0,b}; o) := -I\{n(t-1) = c(t-1) = 0, Z = z_0\}\times \\ \left[dn_1^b(t)\log\{\bar{Q}_1^{z_0,b}(t;w)\} + \{1 - dn_1^b(t)\}\log\{1 - \bar{Q}_1^{z_0,b}(t;w)\}\right] \ , \qquad (2.1)$$

where

$$dn_1^b(t) := \frac{dn_1(t) - \bar{l}_1(t)}{\bar{u}_1(t) - \bar{l}_1(t)}$$

and

$$\bar{Q}_1^{z_0,b}(t;w) := \frac{\bar{Q}_1^{z_0}(t;w) - \bar{l}_1(t)}{\bar{u}_1(t) - \bar{l}_1(t)}$$

is the conditional mean of $dN_1^b(t)$. Note that

$$dn_1^b(t) \in \left\{\frac{-\bar{l}_1(t)}{\bar{u}_1(t) - \bar{l}_1(t)}, \frac{1 - \bar{l}_1(t)}{\bar{u}_1(t) - \bar{l}_1(t)}\right\} \ , \qquad (2.2)$$

while $\bar{Q}_1^{z_0,b}(t;w) \in [0,1]$, and we have the mapping

$$\bar{Q}_1^{z_0}(t;w) = \bar{l}_1(t) + \left\{\bar{u}_1(t) - \bar{l}_1(t)\right\}\bar{Q}_1^{z_0,b}(t;w) \ . \qquad (2.3)$$

The form of (2.1) is similar to the standard negative log-likelihood loss function; however, the outcome now assumes values outside of $\{0,1\}$, as shown in (2.2). Gruber and van der Laan [2010] calls this a pseudo-log-likelihood loss function and establish that this is a valid loss function for a binary variable. Thus, we may use the loss function (2.1) to estimate $\bar{Q}_1^{z_0,b}$ and in turn map those into estimates of $\bar{Q}_1^{z_0}$ via (2.3). We may minimize this loss function by considering a modified logistic regression that regresses the outcome $n_1^b(t)$ on functions of w as predictors, and fit this regression in the subgroup with $z = z_0, n(t-1) = c(t-1) = 0$. Standard regression software will generally not be equipped to handle this minimization; however, all purpose optimizers, such as `R`'s `optim` function, may be used with relative ease. One could also consider more data-adaptive algorithms, such as the Super Learner, for constructing initial estimators of $\bar{Q}_1^{z_0,b}$; however, such algorithms must respect the bounded statistical model, which may limit the number of available learning algorithms. An area of future research will be adapting existing log-likelihood loss-based learning algorithms to minimize loss functions of the form (2.1).

A valid fluctuation submodel is given by

$$\bar{Q}_1^{z_0,b}(t;w)(\epsilon_1;g) := \text{expit}\big[\text{logit}\big\{\bar{Q}_{1,n}^{z_0,b}(t;w)\big\} + \epsilon_1 H_1^{z_0}(t;o)\big] \ ,$$

where $H_1^{z_0}(t;o) := H_g^{z_0}(t;o)\big\{1 - R_1^{z_0}(t;w)\big\}$ and we remind readers that

$$H_g^{z_0}(t;o) := \frac{I(z = z_0, n(t-1) = c(t-1) = 0)}{g^{z_0}(w)G_C^{z_0}(t;w)}$$

$$R_j^{z_0}(t;w) := \frac{F_j^{z_0}(t_0;w) - F_j^{z_0}(t;w)}{S^{z_0}(t;w)} \ .$$

One easily checks that the generalized score of the loss function at $\epsilon_1 = 0$ is proportional to the efficient influence function. To fluctuate initial estimators in practice, one would perform the one-dimensional minimization of the loss along the submodel directly using standard optimization software. This gives the minimizer, $\epsilon_{1,n}$, which we may use to construct updated estimators of $\bar{Q}_1^{z_0,b}(t;w)$.

We now present a loss function for constructing and fluctuating estimators of $\bar{Q}_2^{z_0}$. We take a similar approach as for the type-1 hazards; however, as in Chapter 1, we must again account for the variational dependence between the hazards of the different types at a given time point. We define a loss function that is indexed by the type-1 hazards, which appear in the upper bound:

$$\begin{aligned} L_{\bar{Q}_1^{z_0}}(\bar{Q}_2^{z_0,b};o) := &-I\{n_1(t) = c(t-1) = 0, Z = z_0\}\times \\ &\big[dn_2^b(t)\log\{\bar{Q}_2^{z_0,b}(t;w)\} + \big\{1 - dn_2^b(t)\big\}\log\big\{1 - \bar{Q}_2^{z_0,b}(t;w)\big\}\big] \ , \end{aligned}$$

where

$$dn_2^b(t) := \frac{dn_2(t) - \bar{l}_2(t)}{\bar{u}_2(t) \wedge \bar{Q}_1^{z_0}(t;w) - \bar{l}_2(t)}$$

and

$$\bar{Q}_2^{z_0,b}(t;w) := \frac{\bar{Q}_2^{z_0}(t;w) - \bar{l}_2(t)}{\bar{u}_2(t) \wedge \bar{Q}_1^{z_0}(t;w) - \bar{l}_2(t)}$$

is the conditional mean of $dN_2^b(t)$. We note a similarity between this loss function and that presented in (1.5), which accounted only for $\bar{Q}_1(t;w)$. The present loss function additionally accounts for the upper bound on $\bar{Q}_2^{z_0}(t;w)$. One could again consider the minimization of this loss using a modified logistic regression of outcome $n_1^b(t)$ onto functions of w computed in the subgroup with $z = z_0, n_1(t) = c(t-1) = 0$. General purpose optimization software can once more be employed for this task.

A valid fluctuation submodel for the type-2 hazards is given by

$$\bar{Q}_2^{z_0,b}(t;w)(\epsilon_2;g) := \text{expit}\left[\text{logit}\left\{\bar{Q}_{2,n}^{z_0,b}(t;w)\right\} + \epsilon_2 H_2^{z_0}(t;o)\right] ,$$

where $H_2^{z_0}(t;o) = H_g^{z_0}(t;o) R_1^{z_0}(t;w)$, as in Theorem 1. One again easily checks that the generalized score of the loss function at $\epsilon_2 = 0$ is proportional to the efficient influence function.

Implementing the restricted hazard-based TMLE algorithm follows the exact steps described in Appendix D.1 with loss functions and fluctuation submodels replaced with these versions. In fact, the algorithm described there is just a special case of the current algorithm, with all lower bounds set to 0 and upper bounds set to 1.

2.3.2 Mean-based TMLE

We now describe an algorithm for implementing a closed-form restricted TMLE for the iterative-means representation of cumulative incidence. As noted in Appendix B.2 the iteratively-defined mean at time t may be represented as

$$\tilde{Q}_1^{z_0}(t;w) = I\left\{n_1(t-1) = n_2(t-1) = 0\right\} R_1^{z_0}(t-1;w) + I\left\{n_1(t-1) = 1\right\} . \qquad (2.4)$$

From this representation, it is clear that placing bounds on $\tilde{Q}_1^{z_0}(t;w)$ would not make much sense, as its value surely must be 1 whenever $n_1(t-1) = 1$. We may instead place bounds on the estimation of $R_1^{z_0}(t-1;w)$, which, as discussed in Section 1.4.1, is the portion of $\tilde{Q}_1^{z_0}$ that requires estimation in practice.

Here, we pause to note a nice feature of this representation. Recall that

$$R_1^{z_0}(t-1;w) := \frac{F_1^{z_0}(t_0) - F_1^{z_0}(t)}{1 - F_1^{z_0}(t) - F_2^{z_0}(t)} \ ,$$

or the probability of failure due to type 1 between time t and t_0 conditional on remaining at risk until time t. Defining $l_{j,F}(t)$ and $u_{j,F}(t)$ as the lower and upper bounds on $F_j^{z_0}(t;w)$ at time t for type j failures, we have that

$$\frac{\big\{l_{1,F}(t_0) - u_{1,F}(t)\big\} \vee 0}{1 - u_{1,F}(t) - u_{2,F}(t)} \leq R_1^{z_0}(t;w) \leq \frac{u_{1,F}(t_0) - l_{1,F}(t)}{\big\{1 - u_{1,F}(t) - u_{2,F}(t)\big\} \vee 0} \wedge 1 \ , \tag{2.5}$$

so that one might define

$$\tilde{l}_1(t) := \frac{\big\{l_{1,F}(t_0) - u_{1,F}(t)\big\} \vee 0}{1 - u_{1,F}(t) - u_{2,F}(t)} \ , \tag{2.6}$$

$$\tilde{u}_1(t) := \frac{u_{1,F}(t_0) - l_{1,F}(t)}{\big\{1 - u_{1,F}(t) - u_{2,F}(t)\big\} \vee 0} \wedge 1 \tag{2.7}$$

as the lower and upper bounds for estimation of $R_1^{z_0}(t-1;w)$. Thus, it is straightforward to move between bounds on the iteratively-defined condtional means, which may be difficult to comprehend in practice, and bounds on the cumulative incidence, a quantity that may be more readily understood scientifically. It is interesting to note in (2.7) that the upper bound at time t will only be smaller than one if $u_{2,F}(t) < 1$; that is, if the cumulative incidence for type-2 events can be bounded above.

Now that we have established bounds, we may define pointwise a loss function for $\tilde{Q}_1^{z_0}(t;w)$ indexed by $\tilde{Q}_1^{z_0}(t+1;w)$:

$$\begin{aligned} &L_{\tilde{Q}_1^{z_0}(t+1)}\big(\tilde{Q}_1(t);o\big) := \\ &\quad - I\{z = z_0, c(t-1) = 0\}\big[\tilde{Q}_1^{z_0,b}(t+1;w) \log\big\{\tilde{Q}_1^{z_0,b}(t;w)\big\} + \\ &\qquad\qquad \big\{1 - \tilde{Q}_1^{z_0}(t+1;w)\big\} \log\big\{1 - \tilde{Q}_1^{z_0}(t;w)\big\}\big] \ . \end{aligned}$$

An appropriate fluctuation submodel is

$$\tilde{Q}_1^{z_0,b}(t;w)(\epsilon_t;g) = \operatorname{expit}[\operatorname{logit}\{\tilde{Q}_1^{z_0,b}(t;w)\} + \epsilon_t H_g(t;o)] \ ,$$

which indeed yields a generalized score at zero fluctuation proportional to a component of the efficient influence function at each time. Thus, the sum of the generalized scores of the losses at each time point along their respective submodels at zero fluctuation span the relevant portions of the efficient influence function.

Implementing the restricted mean-based TMLE algorithm follows the exact steps described in Appendix D.2 with loss functions and fluctuation submodels replaced by their bounded counterparts.

2.4 Practical Considerations for Determining Bounds

In some situations, determining bounds on the scale of cause-specific hazards may be quite natural. For example, in preventive HIV vaccine trials, subjects attend routinely-scheduled clinic visits to be tested for HIV. If the time scale of the analysis is clinic visits, then the hazards have a natural interpretation as a probability of infection between two clinic visits conditional on being uninfected at the previous clinic visit. Even if bounds are not known for every subgroup defined by W, we may be willing to assume that the conditional hazard of infection is bounded above by a multiple of the marginal hazard of infection. If, for example, only 0.5% of the at-risk population tests positive for HIV at any given clinic visit, we may be willing to assume that the conditional hazard is bounded above by $c \times 0.005$, where c is some large number, for example $c = 10$ or 20.

Another appealing aspect of the hazard-based mapping is its relationship to the familiar nonparametric maximum likelihood estimators (NPMLE), the so-called Kaplan-Meier and Aalen-Johansen estimators Aalen [1978], Johansen [1978], Kaplan and Meier [1958]. If the grid of follow-up times is fine enough to only include a single failure time in each bin (assuming no ties), these estimators can be constructed by plugging in empirical hazard estimates to the mapping (1.1). However, because empirical estimates are used, these estimators may behave erratically when risk sets become small during the course of follow-up. This may be a problem at the end of follow-up when, for example, a large proportion of the subjects under study have been administratively censored. Thus, imposing bounds – even potentially wide bounds – on these empirical estimates may improve finite sample behavior of the estimator. This could prove especially useful for studies featuring a large amount of left truncation. In such cases, too-large empirical estimates of hazards at early times will negatively impact estimates of incidence at all subsequent timepoints. Previous works have presented regularization techniques for these NPMLEs. One such estimator simply discards risk sets that are deemed too small [Lai and Ying, 1991]. This of course, requires tuning parameters to determine the definition of "too small", which limits the estimators practical utility. The definition of "too small" in practice is likely to be ad-hoc and chosen so that the estimator gives a "reasonable" estimate of the quantity of interest. Further, this method of regularization is not very appealing in the rare events setting; one could imagine the unfortunate situation that a rare event occurs when only few subjects are under observation. The regularized estimator would have the researcher throw out this valuable source of scientific information only for the sake of statistical stability. In contrast, not only is a bounded statistical model transparent and scientifically motivated, it also allows the retention of all subject information regardless of the size of risk sets under observation.

In other cases, determining bounds on the hazard scale will be difficult. If the time scale of the analysis is continuous or consists of a large number of discrete time points, hazards may be very small by definition, making the determination of appropriate bounds a potential challenge. Another drawback to the hazard-based approach to estimation is that there is not a one-to-one correspondence with a bounded model for the conditional cumulative incidence functions if there is more than one type of active failure. An intuitive explanation for this is that a low incidence of events of type 1 could be due to small hazards of type 1 events or large hazards of type 2 events. Thus, bounds on the scale of the target parameter do not map into a unique set of bounds for the cause-specific hazards.

Determining the bounds on the iteratively-defined conditional means initially seems less natural than doing so on the hazard scale. However, as noted in the previous section, these means have an intuitive representation and we are able to map bounds on $F_j^{z_0}(t; w)$ into bounds on $\tilde{Q}_1^{z_0}(t; w)$. Researchers often have some notion of, for example, the expected annual incidence of events in the population. If the annual incidence in the population of interest is bounded above by some small constant f_0, one may be willing to assume that the monthly incidence in any subgroup is bounded above by $cf_0/12$ for some constant c, e.g. $c = 10$ or 20.

We conclude this section by noting that the methodology could be easily adopted to account for discrete W-specific bounds. One would simply generate initial estimates stratified on the relevant components of W tailoring the bounded loss function to appropriately account for the bounds in each group. Similarly, initial estimates would be fluctuated using parametric submodels stratified on those same covariates. As an example of a situation where this may be useful is malaria vaccine trials, where there is substantial variation in the genetic characteristics of the Plasmodium falciparum parasite between West and East Africa. Thus, one may wish to impose bounds on the cumulative incidence of a given genotype only for subjects enrolled at clinics in regions where this genotype is rare. Another example of discrete-covariate specific bounds is in preventive HIV vaccine trials, where women are infected far less often than men. One can imagine many other discrete covariates for which a certain outcome is more uncommon in some groups than others. On the other hand, one can also easily imagine situations where W is high-dimensional or features continuous covariates and reasonable bounds cannot be ascertained from existing data or scientific knowledge. In these situations, one could resort to cross-validation to data-adaptively select bounds, though this will obviously be difficult if sample sizes are very small. Future work will be devoted to implementing and evaluating the performance of such cross-validated estimators.

2.5 Simulation Studies

2.5.1 Simulation 1

Our first simulation study assesses the performance of the proposed estimators for estimating the additive effect of a treatment on the cumulative incidence of an event when there is only one type of event active (e.g., death). An analogous simulation considering multiple failure types is placed in Appendix I. We consider estimation of the difference in cumulative incidence by time $t_0 = 10$ between the treatment and placebo group. We considered the data-generating mechanism where n subjects are assigned treatment with probability one half and three independent baseline covariates are measured: $W_1 \sim N(0, 0.25^2)$, $W_2 \sim U(0,1)$, and $W_3 \sim Bernoulli(0.5)$. The censoring survival function is given by

$$G_C^{z_0}(t;w) = \prod_{s=1}^{t} \{1 - expit(-4 + w_1 - w_2 + w_3)\} \ ,$$

while the hazard for an event at each time is given by

$$\bar{Q}_1^{z_0}(t;w) = expit(\beta_0 - z + 2w_1 + w_2 - 3w_3) \ .$$

Note that β_0 controls the rarity of the outcome and in order to study how the performance of our proposed estimator depends on the rarity of the event, we considered $\beta_0 \in \{-7, -6, -2\}$, which corresponded to a marginal incidence by t_0 of $\{0.003, 0.009, 0.30\}$ and $\{0.009, 0.02, 0.48\}$ in the treatment and placebo arms, respectively. We also varied sample size, considering $n = \{500, 750, 1000, 2000, 5000\}$. We consider estimation of the risk difference and compare our proposed restricted TMLE estimator to the unrestricted TMLE presented in Chapter 1, as well as the Kaplan-Meier estimator. We expect that the most benefit of the restricted TMLE over the unconstrained TMLE should be seen for rare events and in small sample sizes. This benefit should become negligible as the number of events increases. Both TMLE estimators are expected to outperform the Kaplan-Meier estimator due to the presence of informative censoring.

For simplicity, we focus on results for the mean-based TMLE; results for the hazard-based TMLE are in Appendix **??**. Both the restricted and unrestricted TMLE used a correctly specified logistic regression model to estimate the censoring mechanism. The unrestricted TMLE used a correctly specified logistic regression to estimate the iteratively-defined means, while the restricted TMLE minimized the bounded binary log-likelihood loss function according to the correctly specified main terms model. In this simulation, the restricted TMLE imposed bounds equal to the exact upper bound of the iterative means. This allows us to assess “best-case” scenario for finite-sample gains that can be achieved by the estimator. The second simulation study considers how the performance depends on how close the posited bounds are to the true bounds.

Figure 2.1 illustrates the performance of the estimators. First, at the bottom of the figure, note the number of observed events is quite small in the “very rare” and “rare” situations. In fact, in the “very rare” case almost 6% of simulated data sets had no observed events when $n = 500$. Even with $n = 5000$, only thirty events were observed on average. Moving to the top row of the figure, we see the performance of the estimators in terms of bias, expressed as a percentage of the truth. We find that unrestricted TMLE performs best in terms of bias when only a few events observed. However, the restricted TMLE performs at least as well in every other case. The NPMLE does not account for the dependent censoring and thus exhibits substantial bias at all sample sizes. The middle row presents the variance of the restricted TMLE relative to the unrestricted and NPMLE estimators. Note that in this row values below one indicate a lower variance for the restricted TMLE. We see that, as expected, the variance of the restricted TMLE is considerably smaller than its competitors when only few events are observed. As the sample size grows, the benefits of enforcing bounds are diminished. These trends in the first two rows are reiterated by the behavior of the mean-squared error in the bottom row. By this criteria, we find that the NPMLE is doomed to poor performance owing to its bias, while the restricted TMLE suffers due to its variance. Very similar results were seen for the case of multiple failure types (Appendix I).

2.5.2 Simulation Study 2

In this simulation, we examined the question of how the tightness of bounds affects the performance of the restricted TMLE relative to the unrestricted TMLE. We repeated the first simulation, but rather than using the exact bound – which would be unknown in practice – we use a multiple thereof. We considered making the bound too tight and too loose both

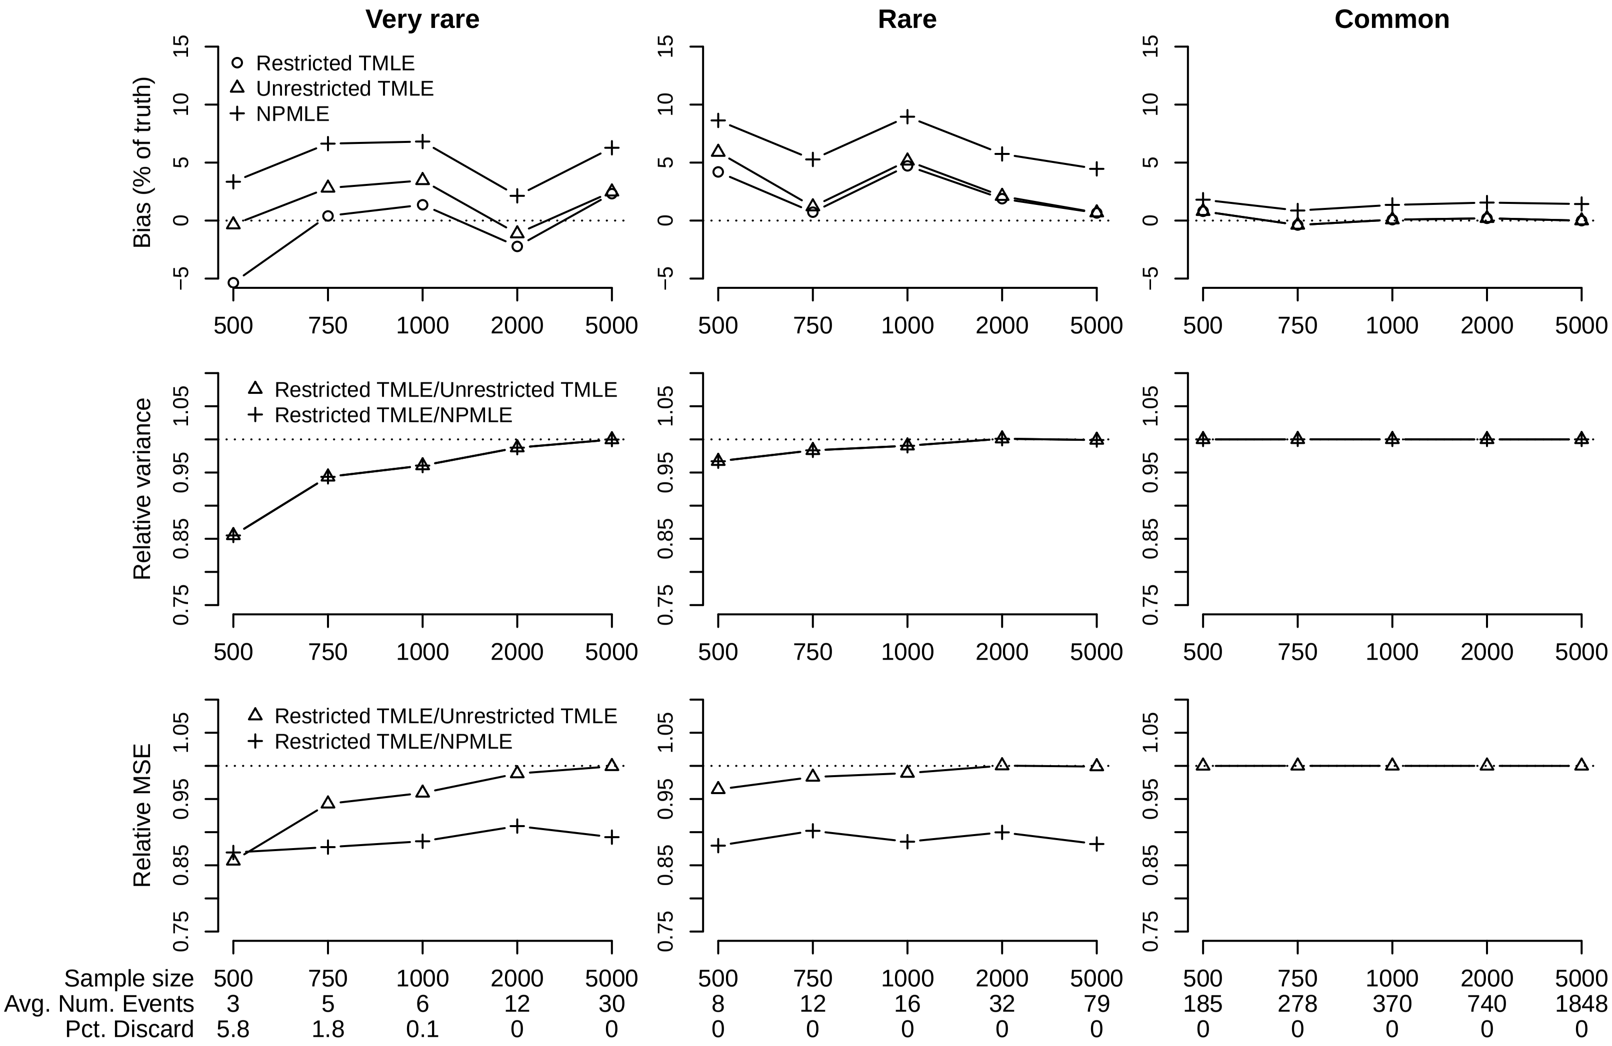

Figure 2.1: Results from Simulation 1. The bias, relative variance, and relative mean-squared error (MSE) of the three estimators is shown in the top, middle, and bottom rows, respectively.

by factors of 10% and 25%. These estimators were compared to the restricted TMLE using the correctly specified bound and the unrestricted TMLE using the same criteria as the first simulation. We expect that bounds that are too tight will lead to improvements in the variance in small samples, but at the cost of bias in large samples. As bounds become larger than the true bound, we expect decreasing improvements in variance relative to the unrestricted estimator. Indeed, as the upper bound approaches one, the restricted TMLE becomes more similar to the unrestricted TMLE.

Figure 2.5.2 shows the performance of the estimators as a function of the tightness of the bounds imposed. In the "very rare" setting, the restricted TMLE is outperformed in terms of bias when $n = 500$; however, at all other sample sizes, the performance of the restricted TMLE is at least as good as the unrestricted TMLE irrespective of bound tightness. In the "rare" event setting, the bias of the unrestricted and restricted TMLE estimators are quite similar for all sample sizes. However, in the "common" event setting, we find that overzealous bounds hurt the asymptotic performance. To some extent, this bias may be mitigated due to double-robustness of the estimator, which assures consistency even if the bounds are misspecified. However, as illustrated by the upper-right-most plot, in finite samples, we may still expect considerable bias if the bounds are much too tight. Moving to the middle row, we find that for small sample sizes, tighter bounds indeed translate to smaller variances. However, the variance of the restricted TMLE with the tightest bounds is seen to grow with sample size relative to the unrestricted TMLE in the "common" event setting. This is also reflected in the mean-squared error shown in the bottom row. Overall, the restricted TMLEs that do not vastly underestimate the bounds perform quite well. This illustrates that one need not have precise knowledge on the bound to expect reasonable gains in practice. An analogous simulation with multiple failure types is featured in Appendix J.

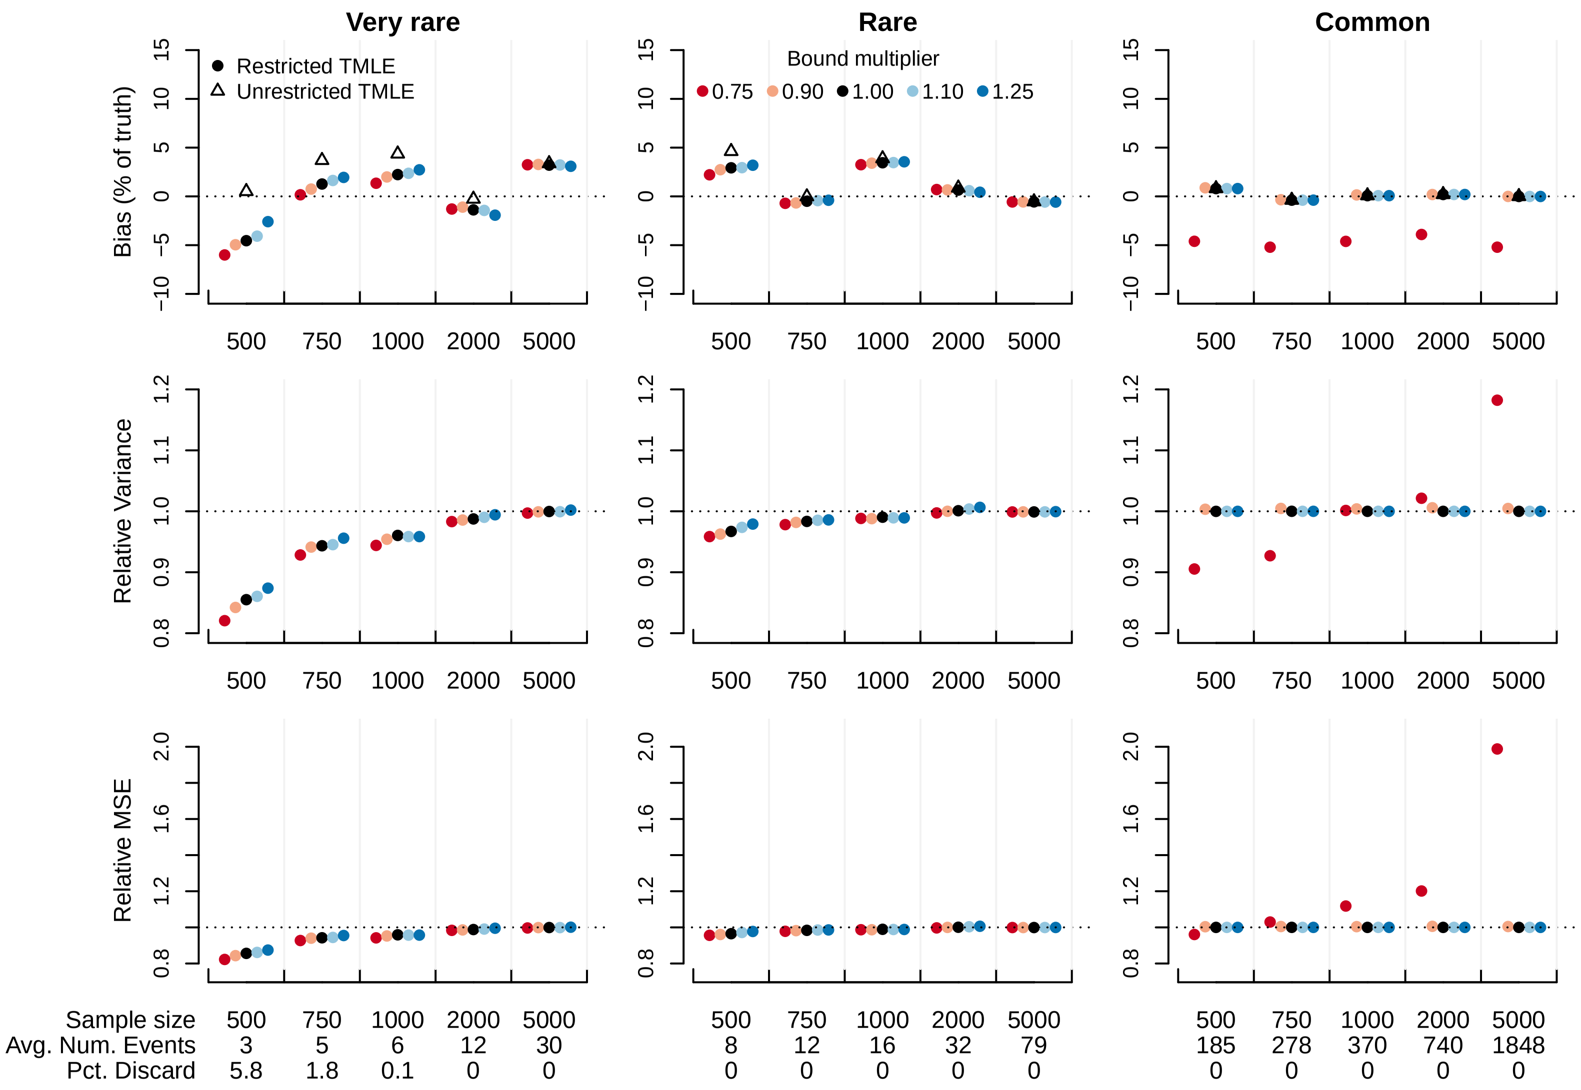

Figure 2.2: Results from Simulation 2. The bias, relative variance, and relative mean-squared error (MSE) of the restricted TMLE compared to the unrestricted TMLE are shown in the top, middle, and bottom rows, respectively. The black circle represents the restricted TMLE that uses the exact bounds. Red circles to the left indicate bounds that are smaller than the truth, while blue circles to the right indicate bounds that are larger than the truth.

2.6 Discussion

In this chapter, we examined the gains afforded by incorporating known bounds into the estimation of cumulative incidence. The targeted minimum loss-based framework is well-suited to handle this modification and we presented appropriate loss functions and fluctuation submodels that could be used in the procedure. Of course, the most important consideration for utilizing these estimators in practice is whether there is enough existing information that plausible bounds can be constructed. Lacking this information, cross-validation could be employed to select bounds that depend on baseline covariates and this will be an important area of future research. It will be interesting to consider relevant objective criteria for the cross validation procedure, as well as the structure of the cross-validation itself (e.g., V-fold vs. leave-one-out).

Our simulations here demonstrated that the restricted estimators perform well in practice, both when only one cause of failure is active and when multiple causes are active. Furthermore, we demonstrated that one need not be incredibly precise in bound selection to see performance benefits in finite samples. However, a substantial risk does exist if one selects bounds that are much too small. Further simulation studies are warranted to better understand the driving factors for this poor behavior. One future question to address is whether poor behavior is driven more by the absolute magnitude of bound misspecification or the proportion of the population that truly has risk above the incorrect bound. For example, it would be interesting to consider whether it is worse to have one percent of the population with incidence much higher than the selected bound, or instead to have ten percent of the population with only slightly higher incidence than the selected bound? Such information would prove useful for researchers in various disease areas when selecting bounds in practice.

Chapter 3

TOWARDS DOUBLY-ROBUST INFERENCE FOR COUNTERFACUTAL MEANS IN LONGITUDINAL SETTINGS

3.1 Introduction

In recent years, doubly-robust (DR) estimators have gained immense popularity, particularly in the field of causal inference. By definition, DR estimators are consistent for the target parameter of interest if one of two nuisance parameters is consistently estimated. This gives a natural appeal of DR estimators: bias induced by misspecification of one nuisance parameter may be mitigated by the consistent estimation of the other. In many contexts, DR estimators arise naturally in pursuit of a locally efficient estimator. Such an estimator must have influence function equal to the canonical gradient of the parameter of interest at the true data generating mechanism in the assumed statistical model when both nuisances are estimated consistently. If an estimator is constructed to have this particular gradient as influence function, double-robustness may follow naturally; this all depends on the form of the gradient. In this sense, often double-robustness may not simply be a property of an estimator itself, but rather a property of the canonical gradient of the target parameter. Double-robustness can typically be shown in two steps: (1) the gradient at the limit of the nuisance parameter estimators must have mean zero; and (2) the linearization of the target parameter at the estimator must yield remainder terms that converge to zero in probability. DR canonical gradients are frequently found for statistical parameters relevant to causal inference, which helps explain the ubiquity of DR estimators in this field. For example, under standard causal assumptions, the statistical parameter identifying the average counterfactual response under a point treatment yields a doubly-robust canonical gradient in a nonparametric model [Robins et al., 1994]. Thus, locally efficient estimators of this statistical target parameter will naturally be DR.

While the conceptual appeal of DR estimators is apparent, questions remain about how these estimators should be constructed in practice. The DR literature has long noted that finite-dimensional models will likely not be sufficient to guarantee consistency of nuisance parameter estimators [Bang and Robins, 2005]. Nevertheless, much of the current work on DR estimation has recommended the use of parametric models and maximum likelihood for estimation of nuisance parameters. This approach has come under scrutiny since Kang and Schafer [2007] showed DR estimators can be arbitrarily poorly behaved if both nuisance parameters are inconsistently estimated. This motivated a recent proposal that developed estimators that minimize the first-order bias resulting from misspecification [Vermeulen and Vansteelandt, 2014]. Others have argued in favor of employing infinite-dimensional models and aggressive data-adaptive estimation in an attempt to ensure consistency of at least one of the nuisance parameter estimates [van der Laan and Rose, 2011].

Regardless of the nuisance parameter estimators employed, studying the behavior of DR estimators under misspecification is of interest. However, due to the widespread use of parametric models and maximum likelihood estimation in the literature, this issue has often been overlooked. It is easy to understand why this is the case: the problems arising from misspecification in finite-dimensional models are somewhat inconsequential in practice. Under misspecification of one nuisance parameter, DR estimators will still be asymptotically linear, though their asymptotic variance may be difficult to calculate explicitly. Nevertheless, resampling techniques, such as the nonparametric bootstrap, may be employed to construct intervals and perform statistical hypothesis tests. In contrast, complications arising from misspecification when using nonparametric estimates of nuisance parameters are much more drastic. Misspecification adversely affects the convergence rate of the estimator of the target parameter and in particular, the bias of this estimator will not vanish sufficiently quickly. This has dire implications for performing inference under misspecification. The coverage of naïvely constructed confidence intervals of all levels will converge to zero and the type-1 error rate for hypothesis tests will converge to one. Furthermore, we may no longer rely on the nonparametric bootstrap as our saving grace. Due to the use of data-adaptive procedures and the irregularity of the resulting estimator, this technique will in general not be valid for constructing intervals and tests.

To our knowledge, the only redress for this problem is presented in van der Laan [2014]. The author proposes a corrected estimator of the average causal effect of a single intervention using targeted minimum loss-based estimation (TMLE). The resulting estimators are doubly-robust not only with respect to consistency, but also with respect to rates of convergence and thereby to inference. This important development opens the door to several new and interesting questions. In the present work, we explore whether the general correction strategy of van der Laan [2014] can be scaled to estimation of a counterfactual mean response under multiple treatments occurring at different points in time. We also examine whether the computational complexity of the corrected estimators may be mitigated by the use of the simpler one-step Newton-Raphson framework for constructing estimators as an alternative to TMLE.

The remainder of the chapter is organized as follows. In Section 3.2, we introduce notation, define our parameter of interest, and examine the theoretical underpinnings DR estimators. We present a linear expansion of the parameter that is the basis for studying the asymptotic properties of DR estimators under misspecification. In Section 3.3, we review the general strategy of van der Laan [2014] for establishing the theory necessary to construct estimators that account for nuisance parameter misspecification. We show how this strategy may be scaled to the longitudinal setting. The section concludes by presenting two one-step estimators and establishing their statistical properties. We contrast the properties of these one-step estimators to estimators constructed using the TMLE framework. In Section 3.4, we examine the proposed estimators' performance in a simple simulation study. We conclude with a discussion.

3.2 DR estimation of counterfactual mean

3.2.1 Notation and parameter of interest

Suppose the observed data consist of n independent copies of the longitudinal data structure $O := (L_0, A_0\ , L_1, A_1, ..., L_{K-1}, A_{K-1}, L_K)$, where L_0 denotes a vector of baseline covariates, A_t a binary treatment at time t and L_t a time-dependent confounder for $t = 1, ..., K-1$, and L_K denotes the outcome of interest. We will use the bar notation to denote the history of a variable, e.g., $\bar{L}_k := (L_0, ..., L_k)$. We define $Q^{\bar{a}}_{k,0}$ as the distribution of L_k conditional on $\bar{L}_{k-1}$ and $\bar{A}_{k-1} = \bar{a}_{k-1}$. Similarly, we define $g^{\bar{a}}_{k,0}$ as the distribution of A_k conditional on $\bar{L}_k$ and $\bar{A}_{k-1} = \bar{a}_{k-1}$. For notational simplicity, we subsequently drop the superscript $\bar{a}$ index so that $Q_{k,0} = Q^{\bar{a}}_{k,0}$ and $g_{k,0} = g^{\bar{a}}_{k,0}$.

Because it is unlikely that the true nuisance parameters fall in a parametric model, we consider a nonparametric statistical model $\mathcal{M}$ for the observed data. We place no restriction on the conditional distributions $Q_{k,0}, k = 0, ..., K$ and define $\mathcal{Q} := \{Q : Q = Q(P) \text{ for some } P \in \mathcal{M}\}$. Similarly, we place no restriction on $g_{k,0}, k = 0, ..., K$ and define $\mathcal{G} := \{g : g = g(P) \text{ for some } g \in \mathcal{M}\}$. We note, however, that our results remain unchanged even if there are restrictions on $\mathcal{G}$. The statistical model can be represented as $\mathcal{M} = \{P = Qg : Q \in \mathcal{Q}, g \in \mathcal{G}\}$. We will use the notation $P_n f := n^{-1} \sum_{i=1}^{n} f(o_i)$ and $P_0 f := \int f(o) dP_0(o)$. We subsequently assume without loss of generality that the treatment assignment of interest consists of setting all treatment nodes to 1. For simplicity, we restrict attention to the case that $K = 2$ and place details of the generalization to an arbitrary K in the supplementary material.

Suppose that O represents a coarsened data structure on the full data

$$X := \big\{L_0, L_1(\bar{a}_0), L_2(\bar{a}_1)\big\} \sim P_0^X\ ,$$

where $L_k(\bar{a}_{k-1})$ is the (potentially unobserved) value of L_k when assigned treatment $\bar{a}_{k-1}$. We are interested in estimation of the counterfactual response under treatment, $E_{P_0^X}\big\{L_2(\bar{1})\big\}$. Bang and Robins [2005] gave assumptions under which this quantity is identified by the statistical parameter $\Psi : \mathcal{Q} \to [0,1]$, where Ψ is defined as

$$\Psi(Q) := \iiint l_2 dQ_2(l_2) dQ_1(l_1) dQ_0(l_0) \ . \tag{3.1}$$

Note that Ψ depends only on P through the conditional means $\bar{Q}_2(\bar{l}_1) := E_{Q_2}\big(L_2 \mid \bar{A}_1 = \bar{1}, \bar{L}_1 = \bar{l}_1\big)$, $\bar{Q}_1(l_0) := E_{Q_1}\big\{\bar{Q}_2(l_0, L_1) \mid A_0 = 1, L_0 = l_0\big\}$ and Q_0, the unconditional distribution of L_0. As we will see in the next section, asymptotically efficient estimation of $\psi_0 := \Psi(Q_0)$ additionally requires estimation of $\bar{g}_{k,0} = E_{g_{k,0}}\big(A_k \mid \bar{A}_{k-1} = \bar{1}_{k-1}, \bar{L}_k\big)$, for $k = 0, 1$, i.e. the conditional probability of treatment at each time point. We henceforth use $Q := (\bar{Q}_2, \bar{Q}_1, Q_0)$ to refer to two outcome regressions and the distribution of L_0, and use $g := (\bar{g}_0, \bar{g}_1)$ to refer to the two propensity regressions. We also use $Q_n := (\bar{Q}_{2,n}, \bar{Q}_{1,n}, Q_{0,n})$ and $g_n := (\bar{g}_{0,n}, \bar{g}_{1,n})$ to denote estimators of these quantities, while $Q_+ := (\bar{Q}_{2,+}, \bar{Q}_{1,+}, Q_{0,+})$ and $g_n := (\bar{g}_{0,+}, \bar{g}_{1,+})$ denote their (possibly misspecified) limit.

3.2.2 Analysis

We begin this section with a brief overview of efficiency theory and a review of several methods for constructing locally efficient estimators. We then present a lemma that is useful for the study of the asymptotic behavior of DR estimators of ψ_0. This lemma is used to study the behavior of estimators under misspecification, which reveals the particularly problematic behavior when employing data-adaptive estimators of Q_0 and g_0. We finish this section by providing useful representations of the asymptotic first-order bias of DR estimators that will be utilized in the next section.

A regular estimator ψ_n is asymptotically linear if and only if it has representation $\psi_n = \psi_0 + P_n D(P_0) + o_P(n^{-1/2})$, where $D(P_0)$ is a gradient of Ψ at P_0 in the Hilbert space $L_2^0(P_0)$ of mean zero finite variance functions endowed with the covariance inner product. Furthermore, such estimators are efficient if and only if their influence function is given by the canonical gradient $D^*(P_0)$ [Bickel et al., 1997]. The canonical gradient is the unique gradient that lies in the tangent space $T_{\mathcal{M}}(P_0) \subseteq L_2^0(P_0)$ of $\mathcal{M}$ at P_0. Thus, gradients and in particular the canonical gradient serve as crucial ingredients in constructing asymptotically linear and efficient estimators.

Several methodologies exist for constructing regular asymptotically linear estimators with a user-specified gradient as influence function. The earliest seems to have appeared in the work of Ibragimov et al. [1981] and Pfanzagl [1982]. This work represented an extension of the well-known Newton-Raphson estimator from finite- to infinite-dimensional models. These estimators are constructed by adding the empirical mean of the gradient evaluated at an estimate of P_0 to an initial plug-in estimator based on this same estimate, yielding the estimator $\psi_n := \Psi(\hat{P}) + P_n D(\hat{P})$. This simple construction makes these estimators generally computationally convenient in practice. Targeted minimum loss-based estimation (TMLE) is a more recent method for constructing regular asymptotically linear estimators. In this methodology, an initial estimate $\hat{P}$ is fluctuated along a carefully selected parametric submodel within which the MLE $\hat{P}^*$ is obtained, and the final estimator is defined as the resulting plug-in, $\Psi(\hat{P}^*)$. The first example of such an estimator appears in Scharfstein et al. [1999], though the full generality of this procedure was first described and rigorously studied in van der Laan and Rubin [2006], and the methodology was fully reviewed and discussed in van der Laan and Rose [2011].

We now turn our attention to the asymptotic behavior of DR estimators. We first present a lemma that provides a useful linear expansion of the parameter Ψ that can be used to study DR estimators constructed using the above frameworks. For that we require the canonical gradient of the pathwise derivative of Ψ at (Q_0, g_0) in $\mathcal{M}$; namely

$$\begin{aligned} &D^*(Q_0, g_0)(o) \\ &\quad := \frac{a_0 a_1}{\bar{g}_{0,0}(l_0)\bar{g}_{1,0}(\bar{l}_1)}\{l_2 - \bar{Q}_{2,0}(\bar{l}_1)\} + \frac{a_0}{\bar{g}_{0,0}(l_0)}\{\bar{Q}_{2,0}(\bar{l}_1) - \bar{Q}_{1,0}(l_0)\} + \bar{Q}_{1,0}(l_0) - \Psi(Q_0)\ . \end{aligned}$$

This particular representation of the canonical gradient was first provided in Bang and Robins [2005]. We now have the following lemma.

Lemma 1. *The first-order expansion of* Ψ *as defined in (3.1) at* Q_n *is given by*

$$\Psi(Q_n) - \Psi(Q_+) = (P_n - P_0)D^*(Q_+, g_+) - B_n^* + M_n^* + R_{1,n} + R_{2,n} \tag{3.2}$$

where

$$
\begin{aligned}
B_n^* &:= B_n^*(Q_n, g_n) := P_n D^*(Q_n, g_n) \\
M_n^* &:= M_n^*(Q_n, Q_+, g_n, g_+) := \big(P_n - P_0\big)\big\{D^*(Q_n, g_n) - D^*(Q_+, g_+)\big\} \\
R_{1,n} &:= R_{1,n}(Q_n, Q_0, g_n, g_0) := P_0\left\{\left(\frac{\bar{g}_{0,0} - \bar{g}_{0,n}}{\bar{g}_{0,n}}\right)\left(\bar{Q}_{1,0} - \bar{Q}_{1,n}\right)\right\} \\
R_{2,n} &:= R_{2,n}(Q_n, Q_0, g_n, g_0) := P_0\left\{\frac{A_0}{\bar{g}_{0,n}}\left(\frac{\bar{g}_{1,0} - \bar{g}_{1,n}}{\bar{g}_{1,n}}\right)\left(\bar{Q}_{2,0} - \bar{Q}_{2,n}\right)\right\} .
\end{aligned}
$$

The lemma organizes the study of the asymptotics of the plug-in estimator $\Psi(Q_n)$ into the study of four kinds of terms: (1) the difference $P_n - P_0$ applied to the canonical gradient at the limit of the nuisance parameter estimates; (2) a first-order bias term, B_n^*; (3) an empirical process term, M_n^*; and (4) two remainder terms, $R_{1,n}$ and $R_{2,n}$.

We now use this lemma to illustrate how double-robustness and asymptotic linearity arise. This is a useful exercise for introducing several key types of conditions that will be used in the remainder of the text. By definition, double-robustness requires the right-hand side of (3.2) to converge to zero in probability if either $Q_+ = Q_0$ or $g_+ = g_0$. This first requires that $P_0 D^*(Q_+, g_+) = 0$, which can be shown to be true in this case. This combined with the weak law of large numbers gives that the first term converges to zero in probability. The bias term, B_n^*, must be accounted for in some way by the estimator itself. This can be accomplished by either making B_n an additive correction to the plug-in estimate $\Psi(Q_n)$ (one-step estimator) or constructing Q_n and g_n in such a way that $B_n^* = 0$ (TMLE). Assuming B_n^* has been accounted for in the estimator, we proceed to the empirical process term M_n^*. In the remainder, we will assume that such empirical process terms are asymptotically negligible, i.e., $M_n^* = o_P(n^{-1/2})$. Sufficient, but not necessary, conditions establishing this are:

1. $D^*(Q_n, g_n)$ falls into a P_0-Donsker class with probability tending to one; and

2. $P_0\big\{D^*(Q_n, g_n) - D^*(Q_+, g_+)\big\}^2$ converges to zero in probability.

For a full development of the theory of empirical processes, we refer readers to van der Vaart and Wellner [1996]. The final step in showing double-robustness is demonstrating that the remainder terms $R_{1,n}$ and $R_{2,n}$ converge to zero in probability. It is clear this will be the case, as each term involves two differences, one of which will converge to zero if $Q_+ = Q_0$ or $g_+ = g_0$. We have thus shown double-robustness.

Asymptotic linearity is a more stringent requirement than double-robustness and thus requires further development. We again assume that $P_0 D^*(Q_+, g_+) = 0$, that B_n^* has been accounted for in the estimator, and that $M_n^* = o_P(n^{-1/2})$. The focus then turns back to the remainder terms. Note that for the purpose of asymptotic linearity, it is no longer sufficient for the remainder terms to converge to zero in probability: they must do so faster than $n^{-1/2}$. If both $Q_+ = Q_0$ and $g_+ = g_0$, then the remainder terms are second-order in that they consist of the product of two terms, both going to zero. Thus, it may be reasonable to assume that $R_{1,n} = o_P(n^{-1/2})$ and $R_{2,n} = o_P(n^{-1/2})$. Sufficient, but not necessary conditions to establish this for $R_{1,n}$, for example, are that $\bar{Q}_{1,n} - \bar{Q}_{1,0} = o_P(n^{-1/4})$ and $\bar{g}_{0,n} - \bar{g}_{0,0} = o_P(n^{-1/4})$ in some appropriate norm. We will subsequently assume negligibility of second-order terms and note that similar quarter-rate arguments will be sufficient for these assumptions to hold. In view of the above, we find that $\psi_n^* - \psi_0 = P_n D^*(Q_0, g_0) + o_P(n^{-1/2})$, where we have used ψ_n^* to denote an estimator that accounts for B_n^* in some way. This establishes not only asymptotic linearity, but also asymptotic efficiency. By the central limit theorem, $n^{1/2}(\psi_n^* - \psi_0)$ converges to a mean-zero normally distributed variate with variance $P_0 D^*(Q_0, g_0)^2$. However, this argument clearly relies on the fact that $Q_+ = Q_0$ and $g_+ = g_0$, which allowed us to claim $R_{1,n}$ and $R_{2,n}$ were second-order and thus asymptotically negligible. However, if either is misspecified, this term is not second order, as one of these differences is no longer converging to zero. Thus, establishing asymptotic linearity requires more effort.

To help understand the behavior of $R_{1,n}$ and $R_{2,n}$ under misspecification, we now present two useful representations of these terms. Define

$$\begin{aligned}
\Phi_{1,0}(\bar{Q}_1) &:= \Phi_{1,g_0,g_+}(\bar{Q}_1) := P_0 \left\{ \left(\frac{\bar{g}_{0,+} - \bar{g}_{0,0}}{\bar{g}_{0,+}} \right) \bar{Q}_1 \right\} \\
\Gamma_{0,0}(\bar{g}_0) &:= \Gamma_{0,g_+,Q_+,Q_0}(\bar{g}_0) := P_0 \left\{ \left(\frac{\bar{Q}_{1,+} - \bar{Q}_{1,0}}{\bar{g}_{0,+}} \right) \bar{g}_0 \right\} ,
\end{aligned}$$

and note that these mappings can be used to describe the first order behavior of $R_{1,n}$:

$$R_{1,n} = \big\{\Gamma_0(\bar{g}_{0,n}) - \Gamma_0(\bar{g}_{0,0})\big\} + \big\{\Phi_1(\bar{Q}_{1,n}) - \Phi_1(\bar{Q}_{1,0})\big\} + \\ S^*_{1,n}(Q_n, Q_0, g_n, g_0, g_+) + S^*_{2,n}(Q_n, Q_+, g_n, g_+)\ , \tag{3.3}$$

where the second-order terms

$$S^*_{1,n}(Q_n, Q_0, g_n, g_0, g_+) := P_0\Bigg\{\left(\frac{\bar{g}_{0,0} - \bar{g}_{0,n}}{\bar{g}_{0,+}\bar{g}_{0,n}}\right)\big(\bar{g}_{0,n} - \bar{g}_{0,+}\big)\ \big(\bar{Q}_{1,0} - \bar{Q}_{1,n}\big)\Bigg\}\ \text{, and}$$
$$S^*_{2,n}(Q_n, Q_+, g_n, g_+) := P_0\Bigg\{\left(\frac{\bar{g}_{0,+} - \bar{g}_{0,n}}{\bar{g}_{0,+}}\right)\big(\bar{Q}_{1,+} - \bar{Q}_{1,n}\big)\Bigg\}$$

can be considered to be $o_P(n^{-1/2})$. Note that $S^*_{1,n}$ is truly second-order in that it consists of either a squared difference between $\bar{g}_{0,n}$ and $\bar{g}_{0,0}$ (if $g_+ = g_0$) or the cross-product between $\bar{g}_{0,n} - \bar{g}_{0,+}$ and $\bar{Q}_{1,n} - \bar{Q}_{1,0}$ (if $Q_+ = Q_0$). Thus, (3.3) uses the combined estimation error of parameters $\Phi_{1,0}$ and $\Gamma_{0,0}$ to approximate the first-order behavior of $R_{1,n}$. We produce a similar representation for $R_{2,n}$, by defining

$$\Phi_{2,0n}(\bar{Q}_2) := \Phi_{2,g_n,g_+,g_0}(\bar{Q}_2) := P_0\Bigg\{\frac{A_0}{g_{0,n}}\left(\frac{\bar{g}_{1,+} - \bar{g}_{1,0}}{\bar{g}_{1,+}}\right)\bar{Q}_2\Bigg\}$$
$$\Gamma_{1,0n}(\bar{g}_1) := \Gamma_{1,g_n,Q_+,Q_0}(\bar{g}_1) := P_0\Bigg\{\frac{A_0}{g_{0,n}}\left(\frac{\bar{Q}_{2,+} - \bar{Q}_{2,0}}{\bar{g}_{1,0}}\right)\bar{g}_1\Bigg\}\ ,$$

and noting that

$$R_{2,n} = \big\{\Phi_2(\bar{Q}_{2,n}) - \Phi_2(\bar{Q}_{2,+})\big\}\ +\ \big\{\Gamma_1(\bar{g}_{1,n}) - \Gamma_1(\bar{g}_{1,0})\big\} \\ + S^*_{3,n}(Q_n, Q_0, g_n, g_0, g_+) + S^*_{4,n}(Q_n, Q_+, g_n, g_+)\ , \tag{3.4}$$

where we define the second-order terms

$$S^*_{3,n}(Q_n, Q_0, g_n, g_0, g_+) := P_0\Bigg\{\frac{A_0}{\bar{g}_{0,0}}\left(\frac{\bar{g}_{1,+} - \bar{g}_{1,n}}{\bar{g}_{1,+}\bar{g}_{1,n}}\right)\big(\bar{g}_{1,0} - \bar{g}_{1,n}\big)\ \big(\bar{Q}_{2,0} - \bar{Q}_{2,n}\big)\Bigg\}\ ,\ \text{and}$$
$$S^*_{4,n}(Q_n, Q_+, g_n, g_+) := P_0\Bigg[\frac{A_0}{\bar{g}_{0,n}}\left(\frac{\bar{g}_{1,+} - \bar{g}_{1,n}}{\bar{g}_{1,+}}\right)\big(\bar{Q}_{2,+} - \bar{Q}_{2,n}\big)\Bigg]\ .$$

A similar argument to that for $S^*_{1,n}$ can be made to show that $S^*_{3,n}$ is truly second-order, and we will again consider both $S^*_{3,n}$ and $S^*_{4,n}$ to be $o_P(n^{-1/2})$.

We note that the mappings defined in (3.3) and (3.4) are indexed by unknown portions of P_0 (e.g., Q_0) and sometimes additionally indexed by the data (e.g., g_n). Such representations may seem unusual, but will prove useful for our analysis. Considering the first remainder, we see that if the first outcome regression is consistently estimated, $R_{1,n}$ behaves in first-order as $\Phi_{1,0}(\bar{Q}_{1,n}) - \Phi_{1,0}(\bar{Q}_{1,0})$. Conversely, if the first propensity regression is consistently estimated, $R_{1,n}$ behaves in first-order as $\Gamma_{1,0}(\bar{g}_{0,n}) - \Gamma_{1,0}(\bar{Q}_{1,0})$. Similarly for $R_{2,n}$, if the outcome regression is consistently estimated, $R_{2,n}$ behaves as $\Phi_{2,0n}(\bar{Q}_{2,n}) - \Phi_{2,0n}(\bar{Q}_{2,0})$. On the other hand, if the second propensity is correctly specified, $R_{2,n}$ behaves as $\Gamma_{1,0n}(\bar{g}_{1,n}) - \Gamma_{1,0n}(\bar{g}_{1,n})$.

Consider the implications of (3.3) and (3.4) with respect to choice of nuisance parameter estimators. If $\mathcal{M}$ is a parametric model and maximum likelihood estimation is used to estimate Q_0 and g_0, then by the invariance property of maximum likelihood estimators, we have that, for example, $\phi_{1,n} := \Phi_1(\bar{Q}_{1,n})$ is the maximum likelihood estimator of $\phi_{1,0} := \Phi_1(\bar{Q}_{1,0})$. Thus, under suitable regularity conditions (e.g., van der Vaart [2000], pg 67-71), $n^{1/2}(\phi_{1,n} - \phi_{1,0})$ converges in distribution to Z_{ϕ_1}, a normally distributed mean zero random variate. The variance of Z_{ϕ_1} is given by the CR bound and could be computed using the delta method. This illustrates the convenience of employing parametric models and maximum likelihood estimation: under misspecification, no further correction to ψ_n is required to establish its asymptotic linearity. Furthermore, while estimating the variance of ψ_n may be difficult in practice, resampling techniques, such as the nonparametric boostrap, may be used to construct confidence intervals and perform hypothesis tests.

However, as noted previously, it is unlikely that Q_0 or g_0 fall into established parametric models. Thus, in practice, it may be more appealing to use infinite-dimensional models and employ data-adaptive estimators. Under misspecification of one of the nuisance parameters, the above representation clearly illustrates why problems arise. Consider again the case that $g_+ \neq g_0$; however, instead of using maximum likelihood, we use instead a kernel regression with cross-validated bandwidth selection to estimate $\bar{Q}_{1,0}$. If L_0 is bivariate, under regularity conditions, a kernel regression optimizing mean-squared error with respect to $\bar{Q}_{1,0}$ will have an optimal bandwidth of the order $n^{-1/6}$ and bias on the order of $n^{-1/3}$ based on the optimal bandwidth. This large bias will be inherited by the functional evaluation $\Phi_1(\bar{Q}_{1,n})$, and, as a consequence, ψ_n^* will not achieve asymptotic linearity without further correction.

This section has presented a clear picture of the statistical problem at hand. Maximum likelihood estimation in parametric models is convenient statistically in that establishing asymptotic linearity of DR estimators requires no further effort. However, the use of finite dimensional models is often questionable scientifically, as it is unlikely that either true nuisance parameter falls in a known parametric model. We are therefore motivated to consider nonparametric models and data-adaptive estimation of nuisance parameters. This presents a problem: inconsistent estimation of a nuisance parameter prevents our estimator from achieving a proper convergence rate and a regular asymptotic limit and thereby prevents us from performing inference with this estimator. In the next section, we outline the general strategy of van der Laan [2014] for constructing corrections that account for this misspecification and thereby allow the construction of doubly-robust confidence intervals and hypothesis tests.

3.3 Correcting for misspecification

Using the case that $Q_+ = Q_0$ and $g_+ \neq g_0$ as an example, the general approach for correcting for misspecification is as follows:

1. Approximate the first-order behavior of each remainder term, $R_{k,n}(Q_n, Q_0, g_n, g_0) = \Phi_{k,0n}(\bar{Q}_{k,n}) - \Phi_{k,0n}(\bar{Q}_{k,0}) + o_P(n^{-1/2})$ for some mapping $\Phi_{k,0n}$ that depends on P_0 and potentially also on the data, for $k = 1, 2$.

2. Approximate the first-order behavior of $\Phi_{k,0n}$ using mappings that can be computed based only on the data, i.e. identify $\Phi_{k,n}$ such that $\Phi_{k,0n}(\bar{Q}_{k,n}) - \Phi_{k,0n}(\bar{Q}_{k,0}) = \Phi_{k,n}(\bar{Q}_{k,n}) - \Phi_{k,n}(\bar{Q}_{k,0}) + o_P(n^{-1/2})$.

3. Linearize $\Phi_{k,n}$ around $\bar{Q}_{k,0}$,

$$\Phi_{k,n}(\bar{Q}_{k,n}) - \Phi_{k,n}(\bar{Q}_{k,0}) = $$
$$P_n D_{Q_k,0}(Q_0, g_0^r) - B_{Q_k,n}(Q_n, g_n^r) + M_{Q_k,n}(Q_n, Q_0, g_n^r, g_0^r) + S_{Q_k,n}(Q_n, Q_0, g_n^r, g_0^r) \ ,$$

   where $D_{Q_k,n}$ is the canonical gradient of the pathwise derivative of $\Phi_{k,n}$, while $D_{Q_k,0}$ is its limit and the canonical gradient of $\Phi_{k,0}$. This gradient depends on P_0 through Q_0 and nuisance parameter g_0^r. The bias term is given by $B_{Q_k,n}(Q_n, g_n^r) := P_n D_{Q_k,n}(Q_n, g_n^r)$, while the empirical process (M_n) and second-order term ($S_{Q_k,n}$) can be considered to be $o_P(n^{-1/2})$.

4. Construct an estimator ψ_n^* that accounts in some way for the bias term $B_{Q_k,n}(Q_n, g_n^r)$, thus establishing the asymptotic linearity.

In the previous section, we showed how the remainder terms may be approximated in first order (Step 1). The present section deals with how these approximations may themselves be approximated in first-order by simpler mappings (Step 2) and how to apply the first-order expansion to these mappings (Step 3). The following section describes how to construct estimators that account for bias terms (Step 4).

Describing Step 2 was the key insight of van der Laan [2014]. The necessity of this step becomes apparent if, for example, one attempts to apply a first-order expansion to $\Phi_{k,0n}$. In this case, one finds that the canonical gradient of the pathwise derivative of $\Phi_{k,0n}$ at P_0 depends on P_0 through Q_0 and g_0. Thus, the bias term resulting from this expansion would involve g_0, which was inconsistently estimated to begin with. Therefore, one would be unable to construct an estimator that adequately corrects for the bias term. In contrast, the problem is simplified greatly when considering the mappings constructed in Step 2. The canonical gradient of this mapping at P_0 is indexed by g_0^r rather than g_0. We will show that g_0^r may be constructed to be of lower dimension than g_0; in fact, g_0^r takes the form of (possibly several) univariate regression(s). This allows one to utilize nonparametric smoothing techniques with reasonable rates of convergence for estimation of g_0^r. Remarkably, we will show that the dimension of this reduced-dimension nuisance parameter does not grow with the number of time points, though the number of such univariate regressions does increase. This allows the technique of van der Laan [2014] to be feasibly scaled to the longitudinal setting.

We first consider the mappings that result when $Q_+ \neq Q_0$: $\Gamma_{0,0n}$ and $\Gamma_{1,0n}$. We begin by re-writing $\Gamma_{0,0n}$ as

$$\begin{aligned}\Gamma_{0,0n}(\bar{g}_{0,n}) - \Gamma_{0,0n}(\bar{g}_{0,0}) &= P_0\left\{\left(\frac{\bar{Q}_{1,+} - \bar{Q}_{1,0}}{\bar{g}_{0,0}}\right)\left(\bar{g}_{0,n} - \bar{g}_{0,0}\right)\right\} \\ &= -P_0\left\{\frac{A_0 A_1}{\bar{g}_{0,0}\bar{g}_{1,0}}\left(\frac{L_2 - \bar{Q}_{1,+}}{\bar{g}_{0,0}}\right)\left(\bar{g}_{0,n} - \bar{g}_{0,0}\right)\right\} \\ &= -P_0\left\{\frac{\bar{Q}^r_{1,0n}}{g^2_{0,0}}\left(\bar{g}_{0,n} - \bar{g}_{0,0}\right)\right\},\end{aligned}$$

where

$$\bar{Q}^r_{1,0n}(l_0) := E_0\left[\frac{A_0 A_1}{\bar{g}_{1,0}(\bar{L}_1)}\left\{L_2 - \bar{Q}_{1,+}(\bar{L}_0)\right\}\,\middle|\, \bar{g}_{0,n}(l_0), \bar{g}_{0,0}(l_0)\right].$$

In this set of equalities, we first replaced the unknown quantity $\bar{Q}_{1,0}$ by L_2; the explicit justification of this step is shown in Appendix K. We next took an inner expectation conditioning on $\bar{g}_{0,n}(L_0)$ and $\bar{g}_{0,0}(L_0)$, which yielded the final equality. We introduced the superscript r to denote the dimension reduction in $\bar{Q}^r_{1,0n}$ achieved by conditioning on summaries of L_0 rather than all of L_0. We note, however, that we cannot directly estimate $\bar{Q}^r_{1,0n}$ in practice because it involves unknown quantities. Thus, we proceed by approximating the first-order behavior of this quantity using a mapping that can be computed based only on the data. Before doing so, note that the estimation error is only non-zero for the case that $\bar{Q}_{1,+} \neq \bar{Q}_{1,0}$, in which case we are assuming that $\bar{g}_{0,+} = \bar{g}_{0,0}$, i.e.

$$\Gamma_{0,0n}(\bar{g}_{0,n}) - \Gamma_{0,0n}(\bar{g}_{0,0}) = I\big(\bar{Q}_{1,+} \neq \bar{Q}_{1,0}, \bar{g}_{0,+} = \bar{g}_{0,0}\big)\big\{\Gamma_{0,0n}(\bar{g}_{0,n}) - \Gamma_{0,0n}(\bar{g}_{0,0})\big\}.$$

We suppress this indicator for notational simplicity; however, we will find that it is indeed needed to generate a mapping $\Gamma_{0,n}$ that (1) can be computed based only on data and (2) approximates $\Gamma_{0,0n}$ in first order.

We proceed as follows:

$$
\begin{aligned}
&- P_0\left\{ \frac{\bar{Q}^r_{1,0n}}{\bar{g}^2_{0,0}} \left(\bar{g}_{0,n} - \bar{g}_{0,0}\right) \right\} \\
&\quad = -P_0\left\{ \frac{\bar{Q}^r_{1,n}}{\bar{g}^2_{0,n}} \left(\bar{g}_{0,n} - \bar{g}_{0,0}\right) \right\} + S_{g_0,1}(\bar{Q}^r_{1,0n}, \bar{Q}^r_{1,0}, \bar{g}_{0,0}, \bar{g}_{0,n}) + S_{g_0,2}(\bar{g}_{0,n}, \bar{g}_{1,n}, \bar{Q}^r_{1,n}, \bar{Q}^r_{1,0}) \ ,
\end{aligned}
\tag{3.5}
$$

where we defined the terms

$$
\begin{aligned}
S_{g_0,1}(\bar{Q}^r_{1,0n}, \bar{Q}^r_{1,0}, \bar{g}_{0,0}, \bar{g}_{0,n}) &:= -P_0\left\{ \left(\frac{\bar{Q}^r_{1,0n} - \bar{Q}^r_{1,0}}{\bar{g}^2_{0,0}} \right) \left(\bar{g}_{0,n} - \bar{g}_{0,0}\right) \right\} \\
S_{g_0,2}(\bar{g}_{0,n}, \bar{g}_{1,n}, \bar{Q}^r_{1,n}, \bar{Q}^r_{1,0}) &:= P_0\left\{ \left(\frac{\bar{Q}^r_{1,n}}{\bar{g}^2_{0,n}} - \frac{\bar{Q}^r_{1,0}}{\bar{g}^2_{0,0}} \right) \left(\bar{g}_{0,n} - \bar{g}_{0,0}\right) \right\} \ ,
\end{aligned}
$$

with

$$
\begin{aligned}
\bar{Q}^r_{1,0}(L_0) &:= E_0\left[\frac{A_0 A_1}{\bar{g}_{1,0}(\bar{L}_1)} \left\{L_2 - \bar{Q}_{1,+}(\bar{L}_0)\right\} \,\middle|\, \bar{g}_{0,0}(L_0) \right] \\
\bar{Q}^r_{1,n}(L_0) &:= E_0\left[\frac{A_0 A_1}{\bar{g}_{1,n}(\bar{L}_1)} \left\{L_2 - \bar{Q}_{1,n}(\bar{L}_0)\right\} \,\middle|\, \bar{g}_{0,n}(L_0) \right] \ .
\end{aligned}
$$

Before proceeding, we make a quick remark on notation. For the sake of brevity, we henceforth only provide definitions for the reduced-dimension regressions that condition on both a true and estimated quantity, in which case we will use $0n$ in the subscript (e.g., $\bar{Q}_{2,0n}$). The same reduced-dimension regressions that involve only estimated quantities will be denoted with subscript n (e.g., $\bar{Q}_{2,n}$), while the regressions involving only true quantities will be denoted with subscript 0 (e.g., $\bar{Q}_{2,0}$).

Note that in the situation that $\bar{Q}_{1,+} \neq \bar{Q}_{1,0}$ and $\bar{g}_{0,+} = \bar{g}_{0,0}$, the terms $S_{g_0,1}$ and $S_{g_0,2}$ are second-order and thus we may consider both terms to be $o_P(n^{-1/2})$. Thus, we may define

$$
\Gamma_{0,n}(\bar{g}_0) := P_0\left(\frac{\bar{Q}^r_{1,n}}{\bar{g}^2_{0,n}} \ \bar{g}_0 \right) \tag{3.6}
$$

and we have shown that

$$
\begin{aligned}
&\Gamma_{0,0n}(\bar{g}_{0,n}) - \Gamma_{0,0n}(\bar{g}_{0,0}) = \\
&\qquad - I\big(\bar{Q}_{1,+} \neq \bar{Q}_{1,0}, \bar{g}_{0,+} = \bar{g}_{0,0}\big)\big\{\Gamma_{0,n}(\bar{g}_{0,n}) - \Gamma_{0,n}(\bar{g}_{0,0}) + o_P(n^{-1/2})\big\} \ .
\end{aligned}
\tag{3.7}
$$

This concludes Step 2 of the general strategy for $\Gamma_{0,0n}$. We proceed to Step 3, performing a first-order expansion of $\Gamma_{0,n}$ at $\bar{g}_{0,n}$:

$$\begin{aligned}&\Gamma_{0,n}(\bar{g}_{0,n}) - \Gamma_{0,n}(\bar{g}_{0,0}) = \\ &\qquad \big(P_n - P_0\big) D_{g_0,0}(\bar{Q}^r_{1,0}, g_{0,0}) + B_{g_0,n}(\bar{Q}^r_{1,n}, \bar{g}_{0,n}) \ + M_{g_0,n}(\bar{Q}^r_{1,n}, \bar{Q}^r_{1,0}, \bar{g}_{0,n}, \bar{g}_{0,0}) \ ,\end{aligned}$$

where we have defined the canonical gradient of $\Gamma_{0,n}$ at $\bar{Q}^r_1, \bar{g}_0$ as

$$D_{g_0,0}(\bar{Q}^r_1, \bar{g}_0)(o) := \frac{\bar{Q}^r_1(l_0)}{\bar{g}^2_0(l_0)} \ \big(a_0 - \bar{g}_0(l_0)\big) \ ,$$

the first-order bias as $B_{g_0,n} := B_{g_0,n}(\bar{Q}^r_{1,n}, \bar{g}_{0,n}) := P_n D_{g_0,n}(\bar{Q}^r_{1,n}, \bar{g}_{0,n})$ and the empirical process term

$$M_{g_0,n}(\bar{Q}^r_{1,n}, \bar{Q}^r_{1,0}, \bar{g}_{0,n}, \bar{g}_{0,0}) := \big(P_n - P_0\big) \ \big\{D_{g_0,n}(\bar{Q}^r_{1,n}, \bar{g}_{0,n}) - D_{g_0,0}(\bar{Q}^r_{1,0}, \bar{g}_{0,0})\big\} \ ,$$

which by assumption is $o_P(n^{-1/2})$. Through this effort, we were able to isolate the first-order bias term $B_{g_0,n}$. We will use this term to construct asymptotically linear estimators in the next section. However, we first present the reduced-dimension parameter mappings and expansions for the remaining parameters defined in the previous section.

Considering $\Gamma_{1,0n}$, we have

$$\begin{aligned}\Gamma_{1,0n}(\bar{g}_{1,n}) - \Gamma_{1,0n}(\bar{g}_{1,0}) &= P_0\bigg\{\frac{A_0}{\bar{g}_{0,n}} \left(\frac{\bar{Q}_{2,+} - \bar{Q}_{2,0}}{\bar{g}_{1,0}}\right) \big(\bar{g}_{1,n} - \bar{g}_{1,0}\big)\bigg\} \\ &= -P_0\bigg\{\frac{A_0 A_1}{\bar{g}_{0,n}\bar{g}_{1,0}} \left(\frac{Y - \bar{Q}_{2,+}}{\bar{g}_{1,0}}\right) \big(\bar{g}_{1,n} - \bar{g}_{1,0}\big)\bigg\} \\ &= -P_0\bigg\{\frac{\bar{Q}^r_{2,0n}}{\bar{g}^2_{1,0}} \ \big(\bar{g}_{1,n} - \bar{g}_{1,0}\big)\bigg\} \ ,\end{aligned}$$

where

$$\bar{Q}^r_{2,0n}(A_0, L_0) := E_0\left[\frac{A_0 A_1}{\bar{g}_{0,n}(L_0)\bar{g}_{1,0}(\bar{L}_1)} \ \big\{Y - \bar{Q}_{2,+}(\bar{L}_1)\big\}\middle| A_0, \bar{g}_{1,n}(\bar{L}_1), \bar{g}_{1,0}(\bar{L}_1)\right] \ .$$

Similar to the first time point, we begin with a series of equalities that allow us to define a reduced-dimension regression that is, in this case, of lower dimension that $\bar{Q}_{2,0}$. Recalling our convention for the subscripts of these reduced-dimension regressions, we proceed as follows:

$$\begin{aligned}
&- P_0\left\{ \frac{\bar{Q}^r_{2,0n}}{\bar{g}^2_{1,0}} \ \big(\bar{g}_{1,n} - \bar{g}_{1,0}\big) \right\} = \\
&\qquad - P_0\left\{ \frac{\bar{Q}^r_{2,n}}{g_{1,n}} \ \big(\bar{g}_{1,n} - \bar{g}_{1,0}\big) \right\} + S_{g_1,1}(\bar{Q}^r_{2,0n}, \bar{Q}^r_{2,0}, \bar{g}_{1,n}, \bar{g}_{1,0}) + S_{g_1,2}(\bar{Q}^r_{2,n}, \bar{Q}^r_{2,0}, \bar{g}_{1,n}, \bar{g}_{1,0}) \ ,
\end{aligned}$$

where

$$\begin{aligned}
S_{g_1,1}(\bar{Q}^r_{2,0n}, \bar{Q}^r_{2,0}, \bar{g}_{1,n}, \bar{g}_{1,0}) &:= -P_0\left\{ \left(\frac{\bar{Q}^r_{2,0n} - \bar{Q}^r_{2,0}}{\bar{g}^2_{1,0}} \right) \ \big(\bar{g}_{1,n} - \bar{g}_{1,0}\big) \right\} \ , \\
S_{g_1,2}(\bar{Q}^r_{2,n}, \bar{Q}^r_{2,0}, \bar{g}_{1,n}, \bar{g}_{1,0}) &:= -P_0\left\{ \left(\frac{\bar{Q}^r_{2,n}}{\bar{g}_{1,n}} - \frac{\bar{Q}^r_{2,0}}{\bar{g}^2_{1,0}} \right) \ \big(\bar{g}_{1,n} - \bar{g}_{1,0}\big) \right\} \ ,
\end{aligned}$$

will be second-order terms whenever $\bar{Q}_{2,+} \neq \bar{Q}_{2,0}$ and $\bar{g}_{1,+} = \bar{g}_{1,0}$ and thus considered to be $o_P(n^{-1/2})$ in this case. Defining

$$\Gamma_{1,n}(\bar{g}_1) := P_0\left(\frac{\bar{Q}^r_{2,n}}{g_{1,n}} \ \bar{g}_1 \right),$$

we have shown that

$$\begin{aligned}
&\Gamma_{1,0n}(\bar{g}_{1,n}) - \Gamma_{1,0n}(\bar{g}_{1,0}) = \\
&\qquad - I\big(\bar{Q}_{2,+} \neq \bar{Q}_{2,0}, \bar{g}_{1,+} = \bar{g}_{1,0}\big)\big\{\Gamma_{1,n}(\bar{g}_{1,n}) - \Gamma_{1,n}(\bar{g}_{1,0}) + o_P(n^{-1/2})\big\} \ . \qquad (3.8)
\end{aligned}$$

We now write a first-order expansion to $\Gamma_{1,n}$ at $\bar{g}_{1,n}$:

$$\begin{aligned}
&\Gamma_{1,n}(\bar{g}_{1,n}) - \Gamma_{1,n}(\bar{g}_{1,0}) = \\
&\qquad \big(P_n - P_0\big) D_{g_1,0}(\bar{Q}^r_{2,0}, g_{1,0}) + B_{g_1,n}(\bar{Q}^r_{2,n}, \bar{g}_{1,n}) \ + M_{g_1,n}(\bar{Q}^r_{2,n}, \bar{Q}^r_{2,0}, \bar{g}_{1,n}, \bar{g}_{1,0}) \ ,
\end{aligned}$$

where we have defined the canonical gradient of $\Gamma_{1,n}$ at $(\bar{Q}^r_{2,0}, \bar{g}_{1,0})$ as

$$D_{g_1,0}(\bar{Q}^r_{2,0}, \bar{g}_{1,0})(o) := \frac{\bar{Q}^r_{2,0}(a_0, \bar{l}_1)}{\bar{g}^2_{1,0}(\bar{l}_1)} \ \big\{a_1 - \bar{g}_{1,0}(\bar{l}_1)\big\} \ ,$$

the first-order bias as $B_{g_1,n} := B_{g_1,n}(\bar{Q}^r_{2,n}, \bar{g}_{1,n}) := P_n D_{g_1,n}(\bar{Q}^r_{2,n}, \bar{g}_{1,n})$, and the empirical process term

$$M_{g_1,n}(\bar{Q}^r_{2,n}, \bar{Q}^r_{2,0}, \bar{g}_{1,n}, \bar{g}_{1,0}) := \big(P_n - P_0\big) \ \big\{D_{g_1,n}(\bar{Q}^r_{2,n}, \bar{g}_{1,n}) - D_{g_1,0}(\bar{Q}^r_{2,0}, \bar{g}_{1,0})\big\} \ ,$$

which is by assumption $o_P(n^{-1/2})$.

We now proceed to the same operations for the mappings that appear when $g_+ \neq g_0$. We begin with $\Phi_{1,0n}$:

$$\begin{aligned}
\Phi_{1,0n}(\bar{Q}_{1,n}) - \Phi_{1,0n}(\bar{Q}_{1,0}) &= -P_0\left\{ \left(\frac{\bar{g}_{0,0} - \bar{g}_{0,+}}{\bar{g}_{0,+}} \right) \left(\bar{Q}_{1,n} - \bar{Q}_{1,0} \right) \right\} \\
&= -P_0\left\{ \left(\frac{A_0 - \bar{g}_{0,+}}{\bar{g}_{0,+}} \right) \left(\bar{Q}_{1,n} - \bar{Q}_{1,0} \right) \right\} \\
&= -P_0\left\{ \bar{g}^r_{0,0n} \left(\bar{Q}_{1,n} - \bar{Q}_{1,0} \right) \right\} \\
&= -P_0\left\{ \frac{A_0}{\tilde{g}^r_{0,0n}} \ \bar{g}^r_{0,0n} \left(\bar{Q}_{1,n} - \bar{Q}_{1,0} \right) \right\} .
\end{aligned}$$

Here, we have defined two reduced-dimension regressions:

$$\begin{aligned}
\bar{g}^r_{0,0n}(L_0) &:= E_0\left\{ \frac{A_0 - \bar{g}_{0,+}(L_0)}{\bar{g}_{0,+}(L_0)} \ \middle| \ \bar{Q}_{1,0}(L_0), \bar{Q}_{0,n}(L_0) \right\} \text{ , and} \\
\tilde{g}^r_{1,0n}(L_0) &:= E_0\left\{ A_0 \ \middle| \ \bar{Q}_{1,0}(L_0), \bar{Q}_{1,n}(L_0) \right\} \ ,
\end{aligned}$$

which will be referred to collectively as g^r_{0n}. We proceed as follows:

$$\begin{aligned}
&- P_0\left\{ \frac{A_0}{\tilde{g}^r_{0,0n}} \ \bar{g}^r_{0,0n} \left(\bar{Q}_{1,n} - \bar{Q}_{1,0} \right) \right\} \\
&\quad = -P_0\left\{ \frac{A_0}{\tilde{g}^r_{0,n}} \ \bar{g}^r_{0,n} \left(\bar{Q}_{1,n} - \bar{Q}_{1,0} \right) \right\} + S_{Q_1,1}(g^r_{0,0n}, g^r_{0,0}, \bar{Q}_{1,n}, \bar{Q}_{1,0}) + S_{Q_1,1}(g^r_{0,n}, g^r_{0,0}, \bar{Q}_{1,n}, \bar{Q}_{1,0}) \ ,
\end{aligned}$$

where we used the shorthand $g^r_{0,0n} := (\bar{g}_{0,0n}, \tilde{g}_{0,0n})$ to refer to both reduced-dimension regression collectively and we have defined

$$\begin{aligned}
S_{Q_1,1}(g^r_{0,0n}, g^r_{0,0}, \bar{Q}_{1,n}, \bar{Q}_{1,0}) &:= -P_0\left\{ A_0 \left(\frac{\bar{g}^r_{0,0n}}{\tilde{g}^r_{0,0n}} - \frac{\bar{g}^r_{0,0}}{\tilde{g}^r_{0,0}} \right) \left(\bar{Q}_{1,n} - \bar{Q}_{1,0} \right) \right\} \\
S_{Q_1,2}(g^r_{0,n}, g^r_{0,0}, \bar{Q}_{1,n}, \bar{Q}_{1,0}) &:= P_0\left\{ A_0 \left(\frac{\bar{g}^r_{0,n}}{\tilde{g}^r_{0,n}} - \frac{\bar{g}^r_{0,0}}{\tilde{g}^r_{0,0}} \right) \left(\bar{Q}_{1,n} - \bar{Q}_{1,0} \right) \right\} .
\end{aligned}$$

These terms will be second-order whenever $\bar{Q}_{1,+} = \bar{Q}_{1,0}$ and $\bar{g}_{0,+} \neq \bar{g}_{0,0}$ and thus will be considered to be $o_P(n^{-1/2})$ in this case. Defining

$$\Phi_{1,n}(\bar{Q}_1) := P_0\left(\frac{A_0}{\tilde{g}^r_{0,n}} \ \bar{g}^r_{0,n} \ \bar{Q}_1 \right),$$

we thus have that

$$\begin{aligned}&\Phi_{1,0n}(\bar{Q}_{1,n}) - \Phi_{1,0n}(\bar{Q}_{1,0}) = \\ &\qquad I(\bar{Q}_{1,+} = \bar{Q}_{1,0}, \bar{g}_{0,+} \neq \bar{g}_{0,0})\big\{\Phi_{1,n}(\bar{Q}_{1,n}) - \Phi_{1,n}(\bar{Q}_{1,0}) + o_P(n^{-1/2})\big\} \ .\end{aligned}$$

We now perform an expansion of $\Phi_{1,n}$ at $\bar{Q}_{1,n}$ and find that

$$\begin{aligned}&\Phi_{1,n}(\bar{Q}_{1,n}) - \Phi_{1,n}(\bar{Q}_{1,0}) \\ &\quad = -(P_n - P_0)D_{Q_{1},0}(Q_0, g_0^r) + B_{Q_{1},n}(Q_n, g_n^r) + M_{Q_{1},n}(Q_n, Q_0, g_n^r, g_0^r) + R_{Q_{1},n}(Q_n, g_n^r) \ ,\end{aligned}$$

where we define the canonical gradient of the pathwise derivative of $\Phi_{1,n}$ at (Q_0, g_0^r) as

$$D_{Q_{1},0}(Q_0, g_0^r)(o) := \frac{a_0}{\tilde{g}_{0,0}^r(l_0)} \ \bar{g}_{0,0}^r(l_0) \ \big(\bar{Q}_{2,0}(\bar{l}_1) - \bar{Q}_{1,0}(l_0)\big) \ ,$$

the first-order bias term as $B_{Q_{1},n} := B_{Q_{1},n}(Q_n, g_n^r) := P_n D_{Q_{1},n}(Q_n, g_n^r)$, the empirical process term as

$$\begin{aligned}&M_{Q_{1},n}(Q_n, Q_0, g_n^r, g_0^r) \\ &\qquad = \big\{P_n - P_0\big\} \ \big\{D_{Q_{1},n}(Q_n, g_n^r) - D_{Q_{1},0}(Q_0, g_0^r)\big\},\end{aligned}$$

and the remainder term as

$$R_{Q_{1},n}(Q_n, Q_0, g_n^r) := -P_0\bigg\{\frac{A_0}{\tilde{g}_{0,n}^r} \ \bar{g}_{0,n}^r \ \big(\bar{Q}_{2,n} - \bar{Q}_{2,0}\big)\bigg\} \ .$$

Here, we note an interesting finding. In contrast to the previous cases studied, the first-order expansion of $\Phi_{1,n}$ yields a first-order remainder $R_{Q_{1},n}$. This appears to be due to the iterative construction of $\bar{Q}_{1,0}$. Nevertheless, we may apply the same strategy, approximating the first-order behavior of $R_{Q_{1},n}$ using reduced-dimension regressions. Specifically, we have that

$$\begin{aligned}-P_0\bigg\{\frac{A_0}{\tilde{g}_{0,n}^r} \ \bar{g}_{0,n}^r \ \big(\bar{Q}_{2,n} - \bar{Q}_{2,0}\big)\bigg\} &= -P_0\big\{\bar{g}_{0,0n}^R \ \big(\bar{Q}_{2,n} - \bar{Q}_{2,0}\big)\big\} \\ &= -P_0\bigg\{\frac{A_0 A_1}{\tilde{g}_{1,0n}^r}\bar{g}_{0,0n}^R \ \big(\bar{Q}_{2,n} - \bar{Q}_{2,0}\big)\bigg\} \ ,\end{aligned}$$

where we define

$$\bar{g}^R_{0,0n}(\bar{L}_1) := E_0\left\{ \frac{A_0}{\tilde{g}^r_{0,n}(L_0)}\ \bar{g}^r_{0,n}(L_0) \ \middle|\ \bar{Q}_{2,n}(\bar{L}_1), \bar{Q}_{2,0}(\bar{L}_1) \right\}$$

$$\tilde{g}^r_{\bar{1},0n}(\bar{L}_1) := E_0\big\{A_0A_1 \ \big|\ \bar{Q}_{2,n}(\bar{L}_1), \bar{Q}_{2,0}(\bar{L}_1)\big\}\ . \tag{3.9}$$

While the notation is becoming cumbersome, we have simply applied the same strategy as before. However, $\bar{g}^R_{0,0n}$ now features as outcome the reduced-dimension regressions from the first step. This mapping may also look strange in that the outcome of the regression is a function of L_0, while we are conditioning on summaries of a measurement, L_1, obtained later in time. However, we are merely using these data-adaptive mappings as tools to isolate bias terms that must be corrected for in estimation. It is for this reason that these parameters are of interest and we make no attempt to provide an interpretation for the limiting value of these data-adaptive parameters. Now, we may proceed by noting

$$\begin{aligned} &-P_0\left\{\frac{A_0A_1}{\bar{g}^r_{\bar{1},0n}}\bar{g}^R_{0,0n}\ \big(\bar{Q}_{2,n}-\bar{Q}_{2,0}\big)\right\} \\ &\qquad = -P_0\left\{\frac{A_0A_1}{\bar{g}^r_{\bar{1},n}}\bar{g}^R_{0,n}\ \big(\bar{Q}_{2,n}-\bar{Q}_{2,0}\big)\right\} + S^R_{Q_1,1}(Q_n,Q_0,g^R_{0n},g^R_0) + S^R_{Q_1,2}(Q_n,Q_0,g^R_{0n},g^R_0)\ , \end{aligned}$$

where we defined the second-order terms

$$\begin{aligned} S^R_{Q_1,1}(Q_n,Q_0,g^R_{0n},g^R_0) &:= -P_0\left\{A_0A_1\left(\frac{\bar{g}^R_{0,0n}}{\tilde{g}^r_{\bar{1},0n}} - \frac{\bar{g}^R_{0,0}}{\tilde{g}^r_{\bar{1},0}}\right)\ \big(\bar{Q}_{2,n}-\bar{Q}_{2,0}\big)\right\} \\ S^R_{Q_1,2}(Q_n,Q_0,g^R_{n},g^R_0) &:= P_0\left\{A_0A_1\left(\frac{\bar{g}^R_{0,n}}{\tilde{g}^r_{\bar{1},n}} - \frac{\bar{g}^R_{0,0}}{\tilde{g}^r_{\bar{1},0}}\right)\ \big(\bar{Q}_{2,n}-\bar{Q}_{2,0}\big)\right\}\ , \end{aligned}$$

which are considered to be $o_{n^{-1/2}}$. Thus, defining

$$\Phi^R_{1,n}(\bar{Q}_2) := P_0\left\{\frac{A_0A_1}{\tilde{g}^r_{\bar{1},0n}}\bar{g}^R_{0,0n}\ \bar{Q}_2\right\}\ ,$$

we have the representation

$$\begin{aligned} &R_{Q_1,n}(\bar{Q}_{2,0},\bar{Q}_{2,n},g^r_{1,n}) \\ &\qquad = I(\bar{Q}_{2,+}=\bar{Q}_{2,0})\big\{\Phi^R_{1,n}(\bar{Q}_{2,n}) - \Phi^R_{1,n}(\bar{Q}_{2,0}) + o_P(n^{-1/2})\big\}\ . \end{aligned}$$

We may now apply a first order expansion to $\Phi^R_{1,n}$ at $\bar{Q}_{2,n}$:

$$\begin{aligned}
&\Phi^R_{1,n}(\bar{Q}_{2,n}) - \Phi^R_{1,n}(\bar{Q}_{2,0}) \\
&\qquad = -(P_n - P_0) D^R_{Q_1,0}(Q_0, g^R_0) + B^R_{Q_1,n}(Q_n, g^R_n) + M^R_{Q_1,n}(Q_n, Q_0, g^R_n, g^R_0) \ ,
\end{aligned}$$

where we defined the canonical gradient of $\Phi^R_{1,n}$ at (Q_0, g^R_0) as

$$D^R_{Q_1,0}(Q_0, g^R_0)(o) := \frac{a_0 a_1}{\tilde{g}^r_{1,0}(\bar{l}_1)} \ \bar{g}^R_{0,0}(\bar{l}_1) \ \left\{ l_2 - \bar{Q}_{2,0}(\bar{l}_1) \right\} \ ,$$

the first-order linear bias as $B^R_{Q_1,n} := B^R_{Q_1,n}(Q_n, g^R_n) := P_n D^R_{Q_2,n}(Q_n, g^R_n)$, and the empirical process term as

$$M^R_{Q_1,n}(Q_n, Q_0, g^R_n, g^R_0) = \left\{ P_n - P_0 \right\} \left\{ D^R_{Q_2,n}(Q_n, g^R_n) - D^R_{Q_2,0}(Q_0, g^R_0) \right\} \ ,$$

and we consider that $M_{Q_1,n} = o_P(n^{-1/2})$.

Finally, we come to the analysis of $\Phi_{2,0n}$:

$$\begin{aligned}
\Phi_{2,0n}(\bar{Q}_{2,n}) - \Phi_{2,0n}(\bar{Q}_{2,0}) &= P_0 \left\{ \frac{A_0}{\bar{g}_{0,n}} \left(\frac{\bar{g}_{1,+} - \bar{g}_{1,0}}{\bar{g}_{1,+}} \right) \left(\bar{Q}_{2,n} - \bar{Q}_{2,0} \right) \right\} \\
&= -P_0 \left\{ \frac{A_0}{\bar{g}_{0,n}} \left(\frac{A_1 - \bar{g}_{1,+}}{\bar{g}_{1,+}} \right) \left(\bar{Q}_{2,n} - \bar{Q}_{2,0} \right) \right\} \\
&= -P_0 \left\{ \bar{g}^r_{1,0n} \left(\bar{Q}_{2,n} - \bar{Q}_{2,0} \right) \right\} \\
&= -P_0 \left\{ \frac{A_0 A_1}{\tilde{g}^r_{1,0n}} \ \bar{g}^r_{1,0n} \left(\bar{Q}_{2,n} - \bar{Q}_{2,0} \right) \right\} \ ,
\end{aligned}$$

where we defined

$$\bar{g}^r_{1,0n}(\bar{l}_1) := E_0 \left[\frac{A_0}{\bar{g}_{0,n}(\bar{L}_1)} \left\{ \frac{A_1 - \bar{g}_{1,+}(\bar{L}_1)}{\bar{g}_{1,+}(\bar{L}_1)} \right\} \ \middle| \ \bar{Q}_{2,n}(\bar{l}_1), \bar{Q}_{2,0}(\bar{l}_1) \right] \ ,$$

and $\tilde{g}^r_{1,0n}$ is as defined in (3.9). The process should now be familiar and we continue by generating second-order terms:

$$\begin{aligned}
&- P_0 \left\{ \frac{A_0 A_1}{\tilde{g}^r_{1,0n}} \ \bar{g}^r_{1,0n} \left(\bar{Q}_{2,n} - \bar{Q}_{2,0} \right) \right\} \\
&\qquad = -P_0 \left\{ \frac{A_0 A_1}{\tilde{g}^r_{1,n}} \ \bar{g}^r_{1,n} \left(\bar{Q}_{2,n} - \bar{Q}_{2,0} \right) \right\} + S_{Q_2,1}(Q_n, Q_0, g^r_{0n}, g^r_0) + S_{Q_2,2}(Q_n, Q_0, g^r_n, g^r_0) \ ,
\end{aligned}$$

with

$$S_{Q_{2,1}}(Q_n, Q_0, g^r_{0n}, g^r_0) := -P_0\Big\{ A_0 A_1 \left(\frac{\bar{g}^r_{1,0n}}{\tilde{g}^r_{1,0n}} - \frac{\bar{g}^r_{1,0}}{\tilde{g}^r_{1,0}}\right) \left(\bar{Q}_{2,n} - \bar{Q}_{2,0}\right)\Big\} \text{ , and}$$

$$S_{Q_{2,2}}(Q_n, Q_0, g^r_n, g^r_0) := P_0\Big\{ A_0 A_1 \left(\frac{\bar{g}^r_{1,n}}{\tilde{g}^r_{1,n}} - \frac{\bar{g}^r_{1,0}}{\tilde{g}^r_{1,0}}\right) \left(\bar{Q}_{2,n} - \bar{Q}_{2,0}\right)\Big\}$$

both considered to be $o_P(n^{-1/2})$ whenever $\bar{Q}_{2,+} = \bar{Q}_{2,0}$ and $\bar{g}_{1,+} \neq \bar{g}_{1,0}$. Defining

$$\Phi_{2,n}(\bar{Q}_2) := P_0\bigg(\frac{A_0 A_1}{\tilde{g}^r_{1,n}}\ \bar{g}^r_{1,n}\ \bar{Q}_2 \bigg) \ ,$$

we have thus shown that

$$\begin{aligned} &\Phi_{2,0n}(\bar{Q}_{2,n}) - \Phi_{2,0n}(\bar{Q}_{2,0}) \\ &\qquad = I\big(\bar{Q}_{2,+} = \bar{Q}_{2,0}, \bar{g}_{1,+} \neq \bar{g}_{1,0}\big)\big\{\Phi_{2,n}(\bar{Q}_{2,n}) - \Phi_{2,n}(\bar{Q}_{2,0}) + o_P(n^{-1/2})\big\} \ . \end{aligned}$$

We now write a first-order expansion of $\Phi_{2,n}$ at $\bar{Q}_{2,n}$:

$$\begin{aligned} &\Phi_{2,n}(\bar{Q}_{2,n}) - \Phi_{2,n}(\bar{Q}_{2,0}) \\ &\qquad = -(P_n - P_0) D_{Q_2,0}(Q_0, g^r_0) + B_{Q_2,n}(Q_n, g^r_n)\ + M_{Q_2}(Q_n, Q_0, g^r_n, g^r_0) \ , \end{aligned}$$

where the canonical gradient of $\Phi_{2,n}$ at (Q_0, g^r_0) is defined as

$$D_{Q_2,0}(Q_0, g^r_0)(o) := \frac{a_0 a_1}{\tilde{g}^r_{1,0}(\bar{l}_1)}\ \bar{g}^r_{1,0}(\bar{l}_1)\ \big\{l_2 - \bar{Q}_{2,0}(\bar{l}_1)\big\} \ ,$$

the first-order bias is given by $B_{Q_2,n} := B_{Q_2,n}(Q_n, g^r_n) := P_n D_{Q_2,n}(Q_n, g^r_n)$, and the empirical process term

$$M_{Q_2}(Q_n, Q_0, g^r_{1,n}, g^r_{1,0}) := \big\{P_n - P_0\big\}\ \big\{D_{Q_2,n}(Q_n, g^r_n) - D_{Q_2,n}(Q_0, g^r_0)\big\} \ ,$$

is considered to be $o_P(n^{-1/2})$.

To conclude this section, we summarize the results we have developed in this section in the following key representation:

$$\begin{aligned}
&\Psi(Q_n) - \Psi(Q_+) \\
&\quad = P_n D^*(Q_+, g_+) - P_n D^*(Q_n, g_n) + o_P(n^{-1/2}) \\
&\qquad - I(Q_+ \neq Q_0, g_+ = g_0)\big[P_n\big\{D_{g_0,0}(Q_0^r, g_0) + D_{g_1,0}(Q_0^r, g_0)\big\} - \\
&\qquad\qquad\qquad\qquad \big\{B_{g_0,n}(Q_n^r, g_n) + B_{g_1,n}(Q_n^r, g_n)\big\} + o_P(n^{-1/2})\big] \\
&\qquad - I(g_+ \neq g_0, Q_+ = Q_0)\big[P_n\big\{D_{Q_1,0}(Q_0, g_0^r) + D^R_{Q_1,0}(Q_0^r, g_0^R) + D_{Q_2,0}(Q_0, g_0^r)\big\} - \\
&\qquad\qquad\qquad\qquad \big\{B_{Q_1,n}(Q_n, g_n^r) + B^R_{Q_1,n}(Q_n, g_n^R) + B_{Q_2,n}(Q_n, g_n^r)\big\} + o_P(n^{-1/2})\big]\ .
\end{aligned} \tag{3.10}$$

This equation will serve as the basis for constructing estimators that correct for misspecification. Appendix L contains the same operations developed in this section for general remainder terms arising when the longitudinal process involves an arbitrary number of time points.

3.4 Corrected estimators, intervals, and hypothesis tests

The previous section developed theory that is extremely useful for constructing estimators. In this section, we consider how estimators may be constructed using a one-step Newton-Raphson framework. We will show the limitations that arise from utilizing this framework, which motivates a discussion of how the TMLE framework may instead be used. We conclude the section with a discussion of practical considerations for estimator construction.

3.4.1 One-step Newton-Raphson estimators

Recall that one-step estimators are constructed by adding an approximation of the bias B_n^* resulting from the first-order expansion of Ψ at Q_n. This was necessary because B_n^* is not $o_P(n^{-1/2})$ and therefore the naïve plug-in estimator $\Psi(Q_n)$ is overly biased and not asymptotically linear. From (3.10), we see that it may be possible to take a similar approach to correct for the bias that arises in the remainder terms due to misspecification. In fact, we can read directly from (3.10) that the corrected estimator

$$\psi_n^{*,+} := \Psi(Q_n) + B_n^* - I(Q_+ \neq Q_0, g_+ = g_0)B_{g,n} - I(g_+ \neq g_0, Q_+ = Q_0)B_{Q,n} \ , \qquad (3.11)$$

is asymptotically linear with influence function

$$\begin{aligned} &D^{*,+}(Q_+, Q_0^r, g_+, g_0^r) \\ &\quad := D^*(Q_+, g_+) + I(Q_+ \neq Q_0, g_+ = g_0)\{D_{g_0,0}(Q_0^r, g_0) + D_{g_1,0}(Q_0^r, g_0)\} + \\ &\qquad I(g_+ \neq g_0, Q_+ = Q_0)\{D_{Q_1,0}(Q_0, g_0^r) + D^R_{Q_1,0}(Q_0^r, g_0^R) + D_{Q_2,0}(Q_0, g_0^r)\} \ . \end{aligned}$$

Here, we used $B_{g,n} := B_{g_0,n} + B_{g_1,n}$ to denote the sum of the bias terms resulting from misspecification of Q and $B_{Q,n} := B_{Q_1,n} + B^R_{Q_1,n} + B_{Q_2,n}$ to similarly denote the sum of the bias terms resulting from misspecification of g. Consistent estimators of the asymptotic variance of $\psi_n^{*,+}$ can thus be constructed as

$$\sigma_n^2 := P_n D^{*,+}(Q_n, Q_n^r, g_n, g_n^r)^2 \ ,$$

and may be used to construct asymptotic $(1-\alpha)\%$ Wald confidence intervals of the form

$$\left(\psi_n^{*,+} - z_{1-\alpha/2}\frac{\sigma_n}{n^{1/2}} \ , \ \psi_n^{*,+} + z_{1-\alpha/2}\frac{\sigma_n}{n^{1/2}}\right) \ ,$$

as well as two-sided asymptotic level α Wald tests that reject the hypothesis that $\psi_0 = \psi^\circ$ whenever

$$\left|\frac{n^{1/2}(\psi_n^{*,+} - \psi^\circ)}{\sigma_n}\right| > z_{1-\alpha/2} \ ,$$

where $z_{1-\alpha/2}$ is the $(1-\alpha/2)$-quantile of a standard normal distribution.

Note that $\psi_n^{*,+}$ only offers protection from the consequences of misspecification if we have guarded against misspecification of a nuisance parameter that has in fact been misspecified. This limits the practical utility of such an estimator; however, it is not difficult to imagine situations where this estimator may still prove useful. For example, in many settings outcome regressions may be more difficult to consistently estimate in practice, owing to a lack of scientific knowledge on the outcome process. In contrast, in many settings useful a priori information may be available for the propensity regression. For example, it may be known that the probability of receiving treatment at a given time point is only a function of a small number of predictors, and so it may be possible to construct a consistent nonparametric estimator of the propensity regressions converging at sufficient rates. In other cases, the propensity is known by design. In such cases, one could guard against the misspecification of the outcome regressions by subtracting $B_{g,n}$ from the standard one-step estimator, thereby assuring asymptotic linearity of the estimator in spite of the misspecification of the outcome regressions.

We now examine the question of whether one-step estimators may be used to simultaneously guard against misspecification of either Q or g without needing to know which. Consider the estimator

$$\psi_n^{*,\diamond} := \Psi(Q_n) + B_n^* - B_{g,n} - B_{Q,n} \ , \tag{3.12}$$

which may be re-written as

$$\psi_n^{*,\diamond} = \psi_n^{*,+} - \big\{1 - I(Q_+ \neq Q_0, g_+ = g_0)\big\} B_{g,n} - \big\{1 - I(g_+ \neq g_0, Q_+ = Q_0)\big\} B_{Q,n} \ . \tag{3.13}$$

The expression (3.13) combined with (3.10) imply that to understand the behavior of $\psi_n^{*,\diamond}$, we must study of the bias correction terms $B_{g,n}$ and $B_{Q,n}$ in the situations when the corrections were unnecessary. For the sake of brevity, we relegate the study of these terms in this situation to Appendix M and provide the conclusions of that analysis here. We find that $\psi_n^{*,\diamond}$ has the following properties:

1. If $g_+ = g_0$ and $Q_+ \neq Q_0$ but we nevertheless perform the correction for misspecification of g (i.e. involving $B_{Q,n}$), $\psi_n^{*,\diamond}$ is still DR in the classical sense. However, $\psi_n^{*,\diamond}$ will only be asymptotically linear if g_n^r converges to g_0^r faster than $n^{-1/2}$.

2. If $Q_+ = Q_0$ and $g_+ \neq g_0$, but we nevertheless perform the correction to guard against misspecification of Q (i.e. involving $B_{g,n}$), then $\psi_n^{*,\diamond}$ will not be DR in the classical

sense, nor will it be asymptotically linear. This is true even if Q_n^r converges to Q_0^r faster than $n^{-1/2}$.

Both points above are somewhat problematic for the use of the estimator $\psi_n^{*,\diamond}$ in practice. Together these points establish that the confidence intervals and tests developed for $\psi_n^{*,+}$ above will only be correct for $\psi_n^{*,\diamond}$ if we are in the very special case that neither Q nor g was misspecified! The first point additionally tells us that if we incorrectly adjust for misspecification of g, the reduced-dimension regressions g_n^r must converge at a parametric rate in order to retain asymptotic linearity of $\psi_n^{*,\diamond}$. However, correctly specifying a parametric regression model for these reduced-dimension regressions will be no easy task, if it is indeed at all possible. Furthermore, even if the parametric regression model is correctly specified, the inference developed above will no longer be valid as there will be a contribution to the influence function of $\psi_n^{*,\diamond}$ resulting from the incorrect correction for $B_{Q,n}$. The second point is far more disconcerting than the first. Here, we see that incorrectly adjusting for misspecification of Q, not only prevents $\psi_n^{*,\diamond}$ from achieving asymptotic linearity, but also prevents its consistency as well.

Based on the discussion above, it appears as though one-step estimators are capable of correcting for misspecification only so long as the user knows which nuisance parameter has been misspecified. The attempt to guard against either form of misspecification without knowing which actually occurs comes with high risks, including the loss of consistency and hence of double-robustness.

3.4.2 TMLE estimators

We recall from the introduction and previous chapters that TMLE is a general methodology for updating initial estimators of nuisance parameters. These updates are constructed in such a way that the revised estimator solves a set of user-specified estimating equations in P; typically, this is the efficient influence function estimating equation, i.e., the revised estimator Q_n^* is constructed such that $P_n D^*(Q_n^*, g_n) = 0$. Now, suppose we were able to construct estimators Q_n^*, g_n^*, $Q_n^{r,*}$, and $g_n^{r,*}$ such that not only $P_n D^*(Q_n^*, g_n^*) = 0$, but also that $B_{g,n}(Q_n^{r,*}, g_n^*) = 0$ and $B_{Q,n}(Q_n^*, g_n^{r,*}) = 0$. Then by (3.10), one would immediately have that $\psi_{n,TMLE}^* := \Psi(Q_n^*)$ would be asymptotically linear for arbitrary patterns of misspecification, so long as the updated estimators (Q_n^*, g_n^*, $Q_n^{r,*}$, $g_n^{r,*}$) still satisfied the second-order and empirical process conditions outlined in the previous section. An iterative TMLE algorithm certainly could be used to construct estimators that solve the necessary equations; however, it remains to be seen whether solving additional unnecessary equations would have any negative impact of the performance of the estimator in practice. In van der Laan [2014], this was not the case and therefore this appears to be a very promising avenue. Such questions are beyond the scope of this dissertation but will be the focus of future research.

3.4.3 Practical considerations

We now make a few remarks on considerations for implementing estimators in practice. Asymptotic linearity relies on consistent estimation of the reduced-dimension regressions outlined in the previous sections. These terms are not easily interpretable and one may be concerned that consistent estimation would prove difficult in practice. However, by construction these terms involve only univariate regressions, so that nonparametric smoothing methods may be successfully employed to generate consistent estimates. For example, it is well known that a univariate kernel regression estimator optimizing mean-squared error will have bias on the order of $n^{-2/5}$ based on an optimal bandwidth of the order $n^{-1/5}$. Thus, if cross-validation is used to select the optimal bandwidth, the bias correction terms will converge at a relatively fast rate to their true value, helping to ensure relevant second-order conditions are satisfied. On the other hand, in finite samples parametric and semiparametric regression techniques may lead to superior performance. An appealing practical solution is thus to employ the Super Learner, an ensemble machine learning technique that uses cross-validation to select the optimal combination of a given collection of candidate learners [van der Laan and Polley, 2007]. The candidate learners could include cross-validated kernel regressions, as well as semiparametric (e.g., general additive models) or parametric (e.g., splines or simple univariate regression). The oracle results proven in van der Laan and Dudoit [2003] imply that the ensemble Super Learner model will perform asymptotically as well as the best combination. Thus, by including nonparametric smoothers as candidates, we ensure consistency while allowing for potential finite sample improvements that may be seen with more structured learners.

3.5 Simulation Study

3.5.1 Design

We consider the finite-sample performance of $\psi_n^{*,+}$ and $\psi_n^{*,\diamond}$ in a simple survival analysis setting. Let L_0 be a two-dimensional vector consisting of independent components $L_{01} \sim U(-5,5)$ and $L_{02} \sim Bernoulli(0.5)$. We use L_1 to describe a subject's failure status at the first time point, with $L_1 = 0$ denoting the subject is still at risk, while $L_1 = 1$ denotes a subject has already failed. Similarly, L_2 denotes failure status at the second time point. As such, it is known that $\bar{Q}_{2,0}(l_0, 1) = 1$ for every l_0. We consider a binary baseline treatment A_0 and a binary censoring indicator A_1. We use $A_1 = 1$ to denote remaining uncensored at the first time point. We parametrize the distribution of the L-process using the hazard function:

$$E_0(L_1 \mid A_0 = 1, L_0 = l_0) = E_0(L_2 \mid \bar{A}_1 = \bar{1}, L_1 = 0, L_0 = l_0) = expit(-0.2l_{01} + 0.4l_{01}l_{02}) \ ,$$

Similarly we set

$$E_0(A_0 \, | L_0 = l_0) = E_0(A_1 \mid A_0 = 1, L_1 = 0, L_0 = l_0) = expit(-0.2l_{01} + 0.4l_{01}l_{02}) \ .$$

We examine the behavior of the proposed estimators under misspecification of both Q and g. In each case, the misspecified regression estimator employed is a logistic regression that includes only main terms for L_{01} and L_{02}, while the correctly specified regression estimator is a bivariate kernel regression with cross-validated bandwidth selection based on the mean-squared error criterion. Univariate kernel regression with cross-validated bandwidth selection was used to estimate the reduced-dimension regression needed for the bias correction terms.

We compare our proposed estimators to the standard one-step Newton-Raphson estimator that does not correct for misspecification. In particular, we study two outcomes: the convergence rate of the bias of the estimators and the nominal 95% coverage probabilities of confidence intervals. We expect that the bias of the uncorrected one-step estimator will not converge faster than $n^{-1/2}$, as would be required for asymptotic linearity, due to misspecification of the nuisance parameters. This in turn should be reflected in the coverage of confidence intervals, which should converge to zero asymptotically. We expect that the bias of the one-step estimator that is able to properly account for misspecification $\psi_n^{*,+}$ should tend to zero faster than $n^{-1/2}$ rate and achieve nominal coverage levels asymptotically. As discussed in the previous section, we expect the bias of $\psi_n^{*,\diamond}$ not to converge at the proper rate, though we may expect to find bias smaller than ψ_n^{*}. However, in the case that $Q_+ = Q_0$ and $g_+ \neq g_0$, we expect the inconsistency of $\psi_n^{*,\diamond}$ to adversely affect its asymptotic performance.

3.5.2 Results

Figure 3.5.2 shows the results from the simulation. Starting with the left column, we examine the situation where the propensity scores have been misspecified. In the top plot, we find that the bias of the standard one-step estimator is not converging to zero faster than $n^{-1/2}$, as evidenced by the fact that $n^{1/2}$ times the bias of the estimator is not converging to zero. We see that the estimators that make additional corrections have bias converging to zero, but that the inconsistency of the estimator that incorrectly adjusts for the additional possible misspecification of Q begins to appear at large sample sizes. Figure N.3 (Supplement) indeed shows that the additional correction term is not converging to zero, thus confirming this estimators inconsistency. Nevertheless, the estimator still far outperforms the standard one-step for all sample sizes considered. The bottom row shows the nominal 95% coverage of the estimators, where we see that the bias has a strong negative impact on the coverage probability for the standard one-step estimator. In contrast, the corrected estimators achieve

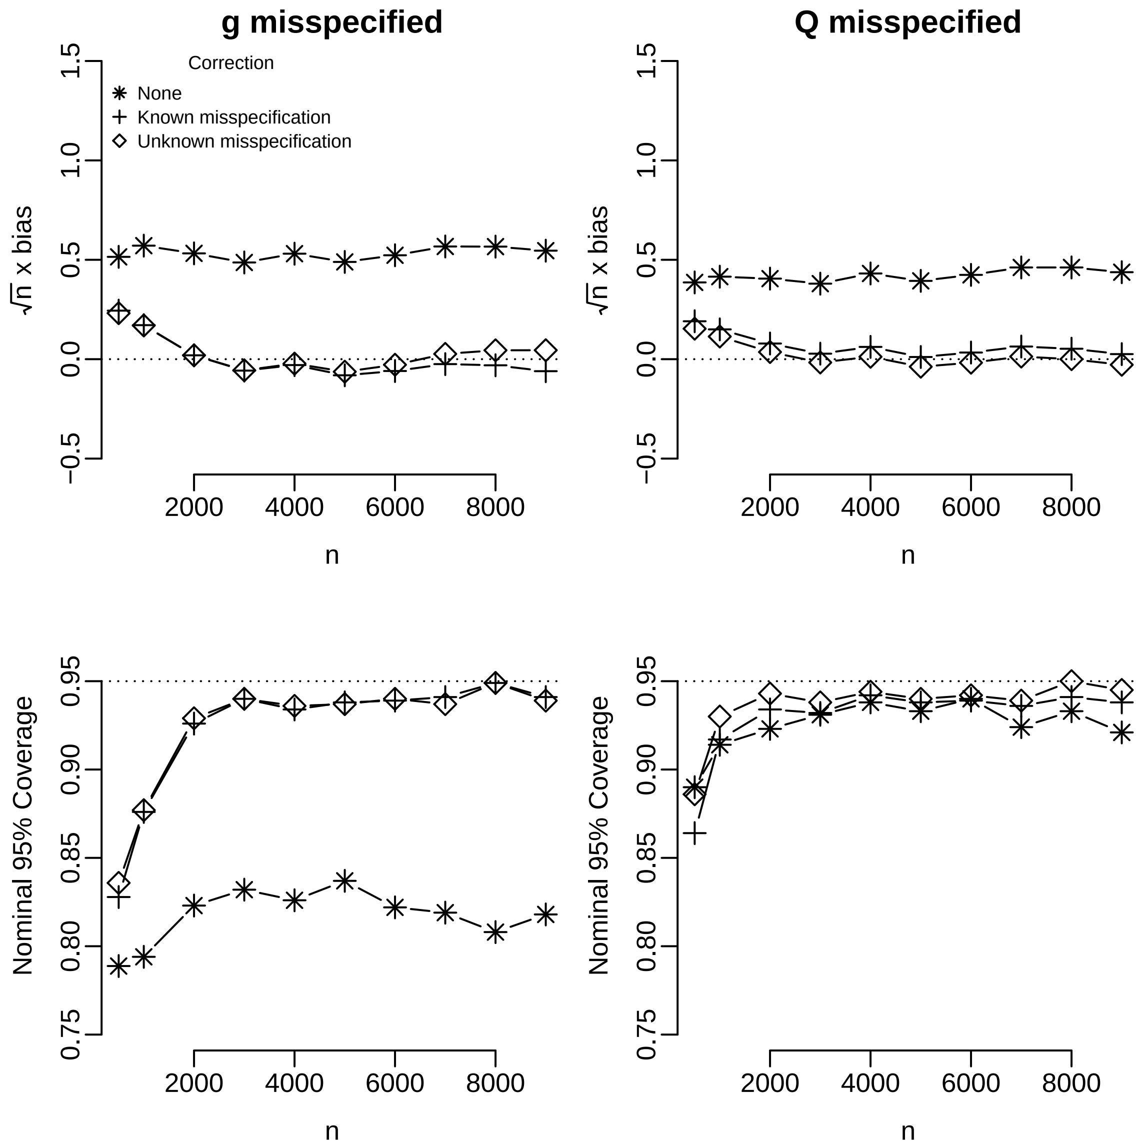

Figure 3.1: Simulation results for the standard vs. corrected one-step estimators. The top row shows $n^{1/2}$ times the bias of the estimator, while the bottom row shows the nominal 95% coverage probability. The left column is the case where the propensity regressions (g) have been misspecified, while the right column is the case where the outcome regressions (Q) have been misspecified.

a near nominal coverage at a reasonable sample size. Moving to the right column, we examine the case that Q is misspecified, while g is correctly specified. Here, we see again that the bias of the corrected estimators appears to be converging to zero close to the correct rate, while the bias of standard one-step is not converging to zero fast enough. However, the bottom panel shows that this does not appear to have a strong adverse effect on its coverage. This is explained by the fact that the efficient influence function-based standard errors are conservative in this case, as illustrated in Figure N.5 (Supplement).

3.6 Discussion

As machine learning and data-adaptive estimation techniques become more popular in practice, the statistical theory necessary to perform valid inference must be developed to keep pace. In this chapter, we discussed the issues that arise in doubly-robust estimation due to nuisance parameter misspecification. We found these issues to be particularly problematic when nonparametric estimators had been employed due to the slow convergence rate of the bias of standard doubly-robust estimators. The finding that one-step estimation methodology may not be fully equipped to develop estimators that are doubly-robust with respect to inference (rather than consistency alone) is very interesting. This points to the necessity of utilizing the TMLE framework for constructing estimators with such properties. However, substantial work remains to be done in that arena to examine whether constructing estimators to solve the additional, unnecessary estimating equations – and in particular, the non-doubly robust estimating equation discussed in Appendix M.5 – will adversely affect the finite-sample and asymptotic behavior of the estimator. This will be a very interesting topic of future research. Additionally, it will be interesting to examine whether a TMLE algorithm for achieving doubly-robust inference in the longitudinal setting is computationally stable. Such an algorithm will be highly iterative and may have convergence difficulties in practice. This motivates the usage of universally least favorable submodels, a recent idea proposed in van der Laan [2015] that would obviate the need for iteration in the TMLE algorithm. This too is an area of interest for future research.

BIBLIOGRAPHY

Aalen, O. (1978). Nonparametric Estimation of Partial Transition Probabilities in Multiple Decrement Models. *The Annals of Statistics*, pages 534–545.

Balzer, L. B. and van der Laan, M. J. (2013). Estimating Effects on Rare Outcomes: Knowledge is Power. *UC Berkeley Division of Biostatistics Working Paper Series.*

Bang, H. and Robins, J. M. (2005). Doubly Robust Estimation in Missing Data and Causal Inference Models. *Biometrics*, 61(4):962–973.

Bickel, P., Klaassen, C., Ritov, Y., and Wellner, J. (1997). *Efficient and Adaptive Estimation for Semiparametric Models.* Johns Hopkins University Press Baltimore.

Chen, K. N., Jin, Z. Z., and Ying, Z. L. (2002). Semiparametric Analysis of Transformation Models with Censored Data. *Biometrika*, 89(3):659–668.

Durham, L. K., Halloran, M. E., Longini, I. M., and Manatunga, A. K. (1999). Comparison of Two Smoothing Methods for Exploring Waning Vaccine Effects. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 48(3):395–407.

Fine, J. P. and Gray, R. J. (1999). A Proportional Hazards Model for the Subdistribution of a Competing Risk. *Journal of the American Statistical Association*, 94(446):496–509.

Gatta, G., Ciccolallo, L., Kunkler, I., Capocaccia, R., Berrino, F., Coleman, M. P., De Angelis, R., Faivre, J., Lutz, J. M., Martinez, C., et al. (2006). Survival from Rare Cancer in Adults: A Population-based Study. *The Lancet Oncology*, 7(2):132–140.

Gelman, A. and Su, Y.-S. (2014). *arm: Data Analysis Using Regression and Multilevel/Hierarchical Models.* R package version 1.7-03.

Gruber, S. and van der Laan, M. J. (2010). A Targeted Maximum Likelihood Estimator of a Causal Effect on a Bounded Continuous Outcome. *The International Journal of Biostatistics*, 6(1).

Hammer, S. M., Sobieszczyk, M. E., Janes, H., Karuna, S. T., Mulligan, M. J., Grove, D., Koblin, B. A., Buchbinder, S. P., Keefer, M. C., and Tomaras, G. D. (2013). Efficacy Trial of a DNA/rAd5 HIV-1 Preventive Vaccine. *New England Journal of Medicine*, 369(22):2083–2092.

Hastie, T. (2013). *gam: Generalized Additive Models*. R package version 1.09.1.

Ibragimov, I. A., Has' minskii, R. Z., and Kotz, S. (1981). *Statistical Estimation: Asymptotic Theory*, volume 2. Springer-Verlag New York.

Johansen, S. (1978). The Product Limit Estimator as Maximum Likelihood Estimator. *Scandinavian Journal of Statistics*, pages 195–199.

Kang, J. D. and Schafer, J. L. (2007). Demystifying Double Robustness: A Comparison of Alternative Strategies for Estimating a Population Mean from Incomplete Data. *Statistical science*, pages 523–539.

Kaplan, E. L. and Meier, P. (1958). Nonparametric Estimation from Incomplete Observations. *Journal of the American statistical association*, 53(282):457–481.

Lai, T. L. and Ying, Z. (1991). Estimating a distribution function with truncated and censored data. *The Annals of Statistics*, pages 417–442.

Lunn, M. and McNeil, D. (1995). Applying Cox regression to competing risks. *Biometrics*, 51:524–532.

Moore, K. and van der Laan, M. (2009). Increasing power in randomized trials with right censored outcomes through covariate adjustment. *Journal of Biopharmaceutical Statistics*, 19(6):1099–1131.

Petersen, M. L., Porter, K. E., Gruber, S., Wang, Y., and van der Laan, M. J. (2010). Diagnosing and responding to violations in the positivity assumption. *Statistical Methods in Medical Research*, 21(1):31–54.

Pfanzagl, J. (1982). *Contributions to a General Asymptotic Statistical Theory*. Springer.

Polley, E. and van der Laan, M. (2013). *SuperLearner: Super Learner Prediction*. R package version 2.0-10.

Prentice, R. L., Kalbfleisch, J. D., Peterson Jr, A. V., Flournoy, N., Farewell, V., and Breslow, N. (1978). The analysis of failure times in the presence of competing risks. *Biometrics*, 34(4):541–554.

Rerks-Ngarm, S., Pitisuttithum, P., Nitayaphan, S., Kaewkungwal, J., Chiu, J., Paris, R., Premsri, N., Namwat, C., de Souza, M., and Adams, E. (2009). Vaccination with ALVAC and AIDSVAX to prevent HIV-1 infection in Thailand. *New England Journal of Medicine*, 361(23):2209–2220.

Robb, M. L., Rerks-Ngarm, S., Nitayaphan, S., Pitisuttithum, P., Kaewkungwal, J., Kunasol, P., Khamboonruang, C., Thongcharoen, P., Morgan, P., Benenson, M., et al. (2012). Risk behaviour and time as covariates for efficacy of the HIV vaccine regimen ALVAC-HIV (vCP1521) and AIDSVAX B/E: a post-hoc analysis of the Thai phase 3 efficacy trial RV 144. *The Lancet Infectious Diseases*, 12(7):531–537.

Robins, J. M., Rotnitzky, A., and Zhao, L. P. (1994). Estimation of regression coefficients when some regressors are not always observed. *Journal of the American Statistical Association*, 89(427):846–866.

Rolland, M., Edlefsen, P. T., Larsen, B. B., Tovanabutra, S., Sanders-Buell, E., Hertz, T., Carrico, C., Menis, S., Magaret, C. A., and Ahmed, H. (2012). Increased HIV-1 vaccine efficacy against viruses with genetic signatures in Env V2. *Nature*, 490(7420):417–420.

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

Scheike, T., Zhang, M.-J., and Gerds, T. (2008). Predicting cumulative incidence probability by direct binomial regression. *Biometrika*, 1(95):205–220.

Stitelman, O. M., Wester, C. W., De Gruttola, V., and van der Laan, M. J. (2011). Targeted maximum likelihood estimation of effect modification parameters in survival analysis. *The International Journal of Biostatistics*, 7(1):1–34.

Struthers, C. A. and Kalbfleisch, J. D. (1986). Misspecified proportional hazard models. *Biometrika*, 73(2):363–369.

Tsiatis, A. (2007). *Semiparametric theory and missing data*. Springer.

van der Laan, M. and Rubin, D. (2006). Targeted maximum likelihood learning. *The International Journal of Biostatistics*, 2(1):1–40.

van der Laan, M. J. (2014). Targeted estimation of nuisance parameters to obtain valid statistical inference. *The International Journal of Biostatistics*, 10:29–57.

van der Laan, M. J. (2015). One-step targeted minimum loss-based estimation based on universal least favorable one-dimensional submodels.

van der Laan, M. J. and Dudoit, S. (2003). Unified cross-validation methodology for selection among estimators and a general cross-validated adaptive epsilon-net estimator: Finite sample oracle inequalities and examples. *UC Berkeley Division of Biostatistics Working Paper Series*.

van der Laan, M. J. and Gruber, S. (2012). Targeted minimum loss based estimation of causal effects of multiple time point interventions. *The International Journal of Biostatistics*, 8(1):1–34.

van der Laan, M. J. and Polley, E. C. (2007). Super learner. *Statistical Applications in Genetics and Molecular Biology*, 6(1):1–23.

van der Laan, M. J. and Robins, J. M. (2003). *Unified Methods for Censored Longitudinal Data and Causality*. Springer.

van der Laan, M. J. and Rose, S. (2011). *Targeted Learning: Causal Inference for Observational and Experimental Data*. Springer.

van der Vaart, A. W. (2000). *Asymptotic Statistics*, volume 3. Cambridge university press.

van der Vaart, A. W. and Wellner, J. A. (1996). *Weak Convergence*. Springer.

Vermeulen, K. and Vansteelandt, S. (2014). Biased-reduced doubly robust estimation. *Journal of the American Statistical Association*, (forthcoming):00–00.

Vittinghoff, E. and McCulloch, C. E. (2007). Relaxing the Rule of Ten Events per Variable in Logistic and Cox Regression. *American Journal of Epidemiology*, 165(6):710–718.

Women's Health Initiative Investigators, W. G. (2002). Risks and benefits of estrogen plus progestin in healthy postmenopausal women: principal results from the Women's Health Initiative randomized controlled trial. *JAMA*, 288(3):321–333.

Appendix A

CAUSAL ASSUMPTIONS

We use a counterfactual framework for defining causal effects. Let $T(z)$ be the potential time to the first event when subject i is assigned treatment z, so that $T = (1-Z)T(0)+ZT(1)$; similarly, define $J(z)$ as the potential type of event under this same treatment assignment. We assume that the treatment assignment of the ith individual does not affect the potential outcomes of other individuals and that there are no multiple forms of treatment, i.e. that the single unit treatment value assumption (SUTVA) holds. We also adopt the axiom of consistency: if a unit is assigned to treatment z, then it must be that $(T, J) = \{T(z), J(z)\}$. The full data we would like to see for each subject are denoted $X(z) = \{T(z), J(z), W\}$.

Our goal is estimation of the cumulative incidence of events of type $J = 1$ at a fixed time t_0 under fixed treatment assignment z_0:

$$F_{z_0}(t_0, 1) = P\{T(z_0) \leq t_0, J(z_0) = 1\} \quad \text{for} \quad z_0 = 0, 1. \tag{A.1}$$

To identify this quantity based on the observed data, we must make several non-testable assumptions:

1. Conditionally independent treatment assignment: $\{Z \perp T(z)\} \mid W$ for $z = 0, 1$
2. Positivity: $P\{U(z) > t_0 \mid Z = z, W = w\} > 0$ for $P_0 -$ almost every w and for $z = 0, 1$
3. Coarsening at random (CAR): $P_{O|Z=z,X(z)} = P_{O|Z=z,W}$ for $z = 0, 1$

Assumption 1 would be satisfied in any trial where treatment probabilities are known based on baseline covariates W, as would be in a randomized trial. Assumption 2 calls for a positive probability of remaining uncensored up to time t_0 within each strata of W; this ensures counterfactual event times are well defined [Petersen et al., 2010]. Assumption 3 requires there be no unmeasured confounders of $T(z)$ and $U(z)$ [Tsiatis, 2007, van der Laan and Robins, 2003]. If we are unwilling to make untestable assumptions, as perhaps we would be in an observational study, the statistical parameter defined in Section 2 Equation (1) may still be an interesting parameter for assessing treatment efficacy.

Appendix B

PROOF OF THEOREM 1

We provide two derivations of the efficient influence curve: the first via projections onto tangent spaces generated by fluctuation submodels for Q, the second as a direct application of Theorem 1 in van der Laan and Gruber [2012]. These two derivations arrive at different, but equivalent representations of the efficient influence curve.

B.1 Representation 1

The first method of obtaining the canonical gradient of cumulative incidence involves projecting an initial gradient, $D_{IPCW}(P)$, onto the tangent space of the model at P, denoted $T(P)$. The tangent space for a random variable $O(t)$ is defined as the mean-square closure of the space spanned by the scores of regular parametric submodels through $P[O(t) \mid Pa\{O(t)\}]$. Because cumulative incidence is only a function of P through the Q portion of the likelihood, we can consider the submodel $\mathcal{M}(g)$ of $\mathcal{M}$ where g is treated as fixed and known. Due to the factorization of $P = Qg$, it follows that the canonical gradient in this smaller model will be the same as in the full model [van der Laan and Robins, 2003]. The tangent space in the reduced model is a sum space consisting of the tangent spaces generated by submodels $P(\epsilon)$ that only vary Q_W, $\bar{Q}_1^{z_0}$, and $\bar{Q}_2^{z_0}$; we denote these tangent spaces as $T_W(P)$, $T_{dN_1(t)}(P), T_{dN_2(t)}(P), t = 1, ..., \tau$, respectively.

Let $L_2^0(P)$ be the Hilbert space of mean zero functions with finite variance with respect to P equipped with inner product $< f, g >:= Pfg$. For the competing risks problem the relevant tangent spaces are

$$\begin{aligned}
T_W(P) &= \{v(W) \in L_2^0(P) : E\{v(W)\} = 0\}, \\
T_{dN_1(t)}(P) &= \{v[dN_1(t), Pa\{dN_1(t)\}] \in L_2^0(P) : \\
&\qquad E(v[dN_1(t), Pa\{dN_1(t)\}] \mid Pa\{dN_1(t)\} = 0)\}, \quad \text{and} \\
T_{dN_2(t)}(P) &= \{v[dN_2(t), Pa\{dN_2(t)\}] \in L_2^0(P) : \\
&\qquad E(v[dN_2(t), Pa\{dN_2(t)\}] \mid Pa\{dN_2(t)\}) = 0\}.
\end{aligned}$$

To project an initial gradient onto the tangent space for variable $O(t)$, we use the projection operator given by

$$\begin{aligned}
&\Pi\{D(P) \mid T_{O(t)}(P)\}[o(t), pa\{o(t)\}] \\
&\quad = E[D(P;O) \mid O(t) = o(t), Pa\{O(t)\} = pa\{o(t)\}] - E[D(P;O) \mid Pa\{O(t)\} = pa\{o(t)\}].
\end{aligned}$$

For the binary variables $dN_1(t), dN_2(t)$ we have the special form

$$\Pi\{D(P) \mid T_{O(t)}(P)\} = H_{O(t)}[Pa\{O(t)\}](O(t) - P[O(t) = 1 \mid Pa\{O(t)\}]), \tag{B.1}$$

where

$$\begin{aligned}
H_{O(t)}[pa\{o(t)\}] &:= E[D(P;O) \mid O(t) = 1, Pa\{O(t)\} = pa\{o(t)\}] \\
&\quad -E[D(P;O) \mid O(t) = 0, Pa\{O(t)\} = pa\{o(t)\}]. \qquad \text{(B.2)}
\end{aligned}$$

See van der Laan and Rose [2011] (p 538-542) for proofs of these results.

Now consider the gradient in $\mathcal{M}(g)$ at P evaluated at a standard data unit o,

$$D_{IPCW}(P;o) = H_g(t_0;o)n_1(t_0) - F_1^{z_0}(t_0), \tag{B.3}$$

where $H_g(t_0;o)$ is as defined in Theorem 1.

The projection of this gradient onto $T_W(P)$ is easily seen to be

$$\begin{aligned}
\Pi\{D_{IPCW}(P) \mid T_W(P)\}(o) &= E\{D_{IPCW}(P;O) \mid W = w\} \\
&= F_1^{z_0}(t_0;w) - F_1^{z_0}(t_0).
\end{aligned}$$

To find the projection onto $T_{dN_1(t)}(P)$, we calculate (B.2) for $O(t) = dN_1(t), t = 1, ..., \tau$. Note that (B.1) will be 0 if $N_1(s-1) = 1$, $N_2(s-1) = 1$, or $C(s-1) = 1$, for any $s < t$. It follows that we can condition on $N_1(t-1) = N_2(t-1) = C(t-1) = 0$ in our calculations; we denote this condition as $\tilde{T} \geq t$.

The first term in (B.2) for $t \leq t_0$ is

$$\begin{aligned} &E[D_{IPCW}(P;O) \mid dN_1(t) = 1, \tilde{T} \geq t, Pa\{dN_1(t)\} = pa\{dn_1(t)\}] \\ &\quad = H_g(t;o) - F_1^{z_0}(t_0), \end{aligned}$$

For $t > t_0, E[D_{IPCW}(P;O) \mid dN_1(t) = 1, Pa\{dN_1(t)\} = pa\{dn_1(t)\}] = -F(t_0, 1 \mid Z = z_0)$. The second term in (B.2) for $t < t_0$ is,

$$\begin{aligned} &E[D_{IPCW}(P;O) \mid dn_1(t) = 0, \tilde{T} \geq t, pa\{dn_1(t)\}] \\ &\quad = H_g(t;o)\{1 - \check{Q}_2^{z_0}(t;w)\} \sum_{s=t+1}^{t_0} \bar{Q}_1^{z_0}(s;w) \prod_{m=t+1}^{s-1} \{1 - \bar{Q}_1^{z_0}(m;w) - \bar{Q}_2^{z_0}(m;w)\} \\ &\quad = H_g(t;o)\{1 - \check{Q}_2^{z_0}(t;w)\} R_1^{z_0}(t;w). \end{aligned}$$

For $t \geq t_0, E[D_{IPCW}(P;O) \mid dN_1(t) = 0, \tilde{T} \geq t, Pa\{dN_1(t)\} = pa\{dn_1(t)\}] = -F_1^{z_0}(t_0)$. Combining these results, we conclude that

$$H_{dN_1}(t;o) = H_g(t;o)\{I(t < t_0)[1 - \{1 - \check{Q}_2^{z_0}(t;w)\} R_1^{z_0}(t;w)] + I(t = t_0)\}.$$

We can similarly determine the projection of $D_{IPCW}(P)$ onto $T_{dN_2(t)}(P)$. The projection will be 0 when $N_1(t) = 1, N_2(t-1) = 1$, or $C(t-1) = 1$; we use $\{\tilde{T} \geq t, dN_1(t) = 0\}$ to denote the complement of this condition. The first term in (B.2) is $-F(t_0, 1 \mid Z = z_0) \ \forall \ t \in \{1, ..., t_0\}$. The second term in (B.2) for $t < t_0$ is,

$$\begin{aligned} &E[D_{IPCW}(P;O) \mid dN_2(t) = 0, Pa\{dN_2(t)\} = pa\{dN_2(t)\}, \tilde{T} \geq t, dN_1(t) = 0] \\ &\quad = H_g(t;o) \sum_{s=t+1}^{t_0} \bar{Q}_1^{z_0}(s;w) \prod_{m=t+1}^{s-1} \{1 - \bar{Q}_1^{z_0}(m;w) - \bar{Q}_2^{z_0}(m;w)\} - F_1^{z_0}(t_0) \\ &\quad = H_g(t;o) R_1^{z_0}(t;w) - F_1^{z_0}(t_0). \end{aligned}$$

For $t \geq t_0$, this term is $-F_1^{z_0}(t_0)$. Combining these we conclude that

$$H_{dN_2}(t;o) = -I(t < t_0) H_g\{t, (o-1)\} R(t;w)$$

Now, by summing over all the projections, we have the first form of the efficient influence curve:

$$\begin{aligned} D^*(P;o) &= \sum_{t=1}^{\tau} H_g(t;o) H_{dN_1}(t;o) \big\{ dn_1(t) - \bar{Q}_1^{z_0}(t;w) \big\} \\ &\quad + \sum_{t=1}^{\tau} H_g(t;o) H_{dN_2}(t;o) \{ dn_2(t) - \check{Q}_2^{z_0}(t;w) \} \\ &\quad + F_1^{z_0}(t_0;w) - F_1^{z_0}(t_0). \end{aligned} \tag{B.4}$$

It can easily be shown that this is equivalent to the form given in the theorem.

B.2 Representation 2

Theorem 1 in van der Laan and Gruber [2012] gives the form of the efficient influence curve for longitudinal parameters under multiple time point interventions as established in Bang and Robins [2005]. By viewing censoring as a sequentially randomized treatment (equivalent to assuming CAR, see e.g. van der Laan and Rose [2011] Appendix 5) and the event process N_2 as a time-varying confounder, we may directly apply the results of this theorem.

We begin by showing that cumulative incidence can indeed be written as a function of iteratively defined conditional means. Note that for $t = 1, ..., t_0$, $\tilde{Q}_1^{z_0}(m;w)$ can be written

$$\tilde{Q}_1^{z_0}(m;w) = I\{n_1(m-1) = n_2(m-1) = 0\} \tag{B.5}$$

$$\times \left[\sum_{t=m}^{t_0} \bar{Q}_1^{z_0}(t;w) \prod_{s=m}^{t-1} \{1 - \bar{Q}_1^{z_0}(s;w) - \bar{Q}_2^{z_0}(s;w)\} \right] + I\{n_1(m-1) = 1\} \tag{B.6}$$

$$= I\{n_1(m-1) = n_2(m-1) = 0\} R_1^{z_0}(m-1) + I\{n_1(m-1) = 1\}, \tag{B.7}$$

where we take empty products to be 1 and empty sums to be 0 as necessary so that this form applies for all $m = 1, ..., t_0$. Letting $m = 1$ we see that $\tilde{Q}_1^{z_0}(1;w) = F_1^{z_0}(t_0;w)$. Taking the expected value of this with respect to Q_W gives the unconditional cumulative incidence.

Now we show the form of the efficient influence curve for this representation of the parameter. A direct application of Theorem 1 in van der Laan and Gruber [2012] gives the efficient influence curve at P to be $\sum_{t=0}^{t_0} \tilde{D}_1^{z_0}(t)$, where

$$\begin{aligned}
\tilde{D}_1^{z_0}(t_0; o) &:= H_g(t_0; o)\{dn_1(t_0) - \bar{Q}_1^{z_0}(t_0; w)\} \\
\tilde{D}_1^{z_0}(t; o) &:= H_g(t; o)\{\tilde{Q}_1^{z_0}(t+1; w) - \tilde{Q}_1^{z_0}(t; w)\}, \quad \text{for } t = t_0 - 1, ..., 1 \\
\tilde{D}_W(w) &:= F_1^{z_0}(t_0; w) - F_1^{z_0}(t_0)
\end{aligned}$$

First note that for $t = t_0$, we have $D^*_{t_0}(P; o) = D_1^{z_0}(t_0; o) + D_2^{z_0}(t_0; o) = D_1^{z_0}(t_0; o)$ from Theorem 1. Now for all $t = t_0 - 1, ..., 1$, the following equalities will be used:

$$\begin{aligned}
R_1^{z_0}(t; w) - R_1^{z_0}(t-1; w) &= -\{1 - R_1^{z_0}(t; w)\}\bar{Q}_1^{z_0}(t; w) + \bar{Q}_2^{z_0}(t; w)R_1^{z_0}(t; w) \\
I\{N_1(t) = 1\} - I\{N_1(t-1) = 1\} &= dN_1(t); \\
I\{N_1(t) = N_2(t) = 0\} &= \{1 - dN_1(t) - dN_2(t)\}I\{N(t-1) = 0\}; \\
I\{N(t-1) = 0\}dN_1(t) &= dN_1(t).
\end{aligned}$$

For $t = t_0 - 1, ..., 1$, we can then write

$$\begin{aligned}
&(\tilde{Q}_1^{z_0}(t+1; w) - \tilde{Q}_1^{z_0}(t; w)) \\
&\quad = I\{n(t) = 0\}R_1^{z_0}(t; w) + I\{n_1(t) = 1\} - I\{n_1(t-1) = 0\}R_1^{z_0}(t-1; w) - I\{n_1(t-1) = 1\} \\
&\quad = dn_1(t) + \{1 - dn_1(t) - dn_2(t)\}I\{n(t-1) = 0\}R_1^{z_0}(t; w) - I\{n(t-1) = 0\}R_1^{z_0}(t-1; w) \\
&\quad = I\{n(t-1) = 0\}\{dn_1(t) - R_1^{z_0}(t; w)dn_1(t) - R_1^{z_0}(t; w)dN_2(t) + R_1^{z_0}(t; w) - R_1^{z_0}(t-1; w)\} \\
&\quad = I\{n(t-1) = 0\}[\{1 - R_1^{z_0}(t; w)\}\{dn_1(t) - \bar{Q}_1^{z_0}(t; w)\} - R_1^{z_0}(t; w)\{dn_2(t) - \bar{Q}_2^{z_0}(t; w)\}],
\end{aligned}$$

which is equivalent to $D_1^{z_0}(t; o) + D_2^{z_0}(t; o)$ from Theorem 1.

Appendix C

REGULARITY CONDITIONS FOR LEMMA 1

Lemma 1 will hold for estimators Q_n and g_n that satisfy the following conditions:

1. $(P_n - P_0)\{D^*(Q_n, g_n) - D^*(Q_0, g_0)\} = o_p(n^{-1/2})$

2. $U(Q_0, Q_n, g_0, g_n) = o_p(n^{-1/2})$,

where

$$U(Q_0, Q_n, g_0, g_n) = \Psi(Q_n) - \Psi(Q_0) + P_0 D^*(Q_n, g_n),$$

is the remainder term resulting from the linearization of $\Psi(Q)$ and involves second- and higher-order differences between Q_n and Q_0 and between g_n and g_0. Condition (1) would be satisfied if there exists a P_0-Donsker class $\mathcal{F}$ such that $D^*(Q_n, g_n) \in \mathcal{F}$ with probability tending to 1 and $P_0\{D^*(Q_n, g_n) - D^*(Q_0, g_0)\}^2 \rightarrow_p 0$. For a detailed proof see van der Laan and Rose [2011] Appendix 18, Theorem A5.

Appendix D

CONSTUCTION OF TMLES

D.1 Hazard-based TMLE

The hazard-based TMLE can be constructed as follows:

1. Generate initial estimates, $g_{z_0,n}$ and $G_{C,n}$, of the conditional treatment and censoring mechanisms. These may be obtained through standard methods (e.g. parametric regression for $g_{z_0,n}$; Kaplan-Meier for $G_{C,n}$), or more ideally using machine learning techniques, possibly combined data-adaptively using the Super Learner.

2. Let $k = 0$. Generate initial estimates, $\bar{Q}^{z_0}_{1,n,k}$ and $\check{Q}^{z_0}_{2,n,k}$, of the cause-specific hazards. Again, these could be obtained through standard logistic regression, but ideally would make use of more data adaptive methods.

3. Use the current estimates $\bar{Q}^{z_0}_{1,n,k}$ and $\check{Q}^{z_0}_{2,n,k}$ to compute $H_{1,n,k}(o)$.

4. Obtain $\epsilon_{1,n,k}$ as the estimated coefficient in a logistic regression model of outcome $dn_1(t)$ on covariate $H_{1,n,k}(o)$ with offset $\bar{Q}^{z_0}_{1,n,k}(t;w)$ in the subset of data with $z = z_0, c(t-1) = 0$. Set $\bar{Q}^{z_0}_{1,n,k+1}(t;w) := \bar{Q}^{z_0}_{1,n,k}(t;w)(\epsilon_{1,n,k})$.

5. Use $\bar{Q}^{z_0}_{1,n,k+1}$ and $\check{Q}^{z_0}_{2,n,k}$ to compute $H_{2,n,k+1}$.

6. Obtain $\epsilon_{2,n,k}$ as the estimated coefficient in a logisitic regression model of outcome $dn_2(t)$ on covariate $H_{2,n,k+1}$ with offset $\text{logit}[\check{Q}^{z_0}_{2,n,k}(t;w)\{1-\bar{Q}^{z_0}_{1,n,k+1}(t;w)\}]$ in the subset of data with $Z = z_0, C(t-1) = 0$. Set $\check{Q}^{z_0}_{2,n,k+1}(t;w) := \check{Q}^{z_0}_{2,n,k}(t;w)(\epsilon_{2,n,k})$.

7. Iterate steps 3-6 to a large number K such that $\epsilon_{1,n,K} \approx 0$ and $\epsilon_{2,n,K} \approx 0$. Let $\bar{Q}^{z_0,*}_{1,n}(t;w)$, $\check{Q}^{z_0,*}_{2,n}(t;w)$ denote the estimates at the final iteration.

8. Apply the mapping in (1) using $\bar{Q}^{z_0,*}_{1,n}(t;w)$, $\check{Q}^{z_0,*}_{2,n}(t;w)$, and $Q_{W,n}$ to obtain the estimator $\psi^*_n := F^*_n(t_0, 1|Z = z_0)$.

D.2 Iterative mean-based TMLE

We begin by briefly discussing how to obtain initial estimators of $\tilde{Q}^{z_0}_1(t;w)$. Consider the conditional mean at the last time point, $\tilde{Q}^{z_0}_1(t_0;w)$. The initial estimate of this quantity should trivially assign 1 to subjects who have failed due to cause 1 prior to t_0 and 0 to subjects who have failed due to cause 2. It thus only remains to estimate the conditional hazard of a failure at time t_0. This can be achieved through standard parametric regression or more data adaptive methods, but in each case we only use data for subjects at risk and uncensored at $t_0 - 1$. Moving to the next time point, $t_0 - 1$, the initial estimate of $\tilde{Q}^{z_0}_1(t_0 - 1;w)$ now assigns 1 to subjects who have failed due to cause 1 prior to $t_0 - 1$ and 0 to subjects who have failed due to cause 2. We then estimate the probability of a type 1 event at either $t_0 - 1$ or t_0 in the subset of subjects who have not failed due to any cause by time $t_0 - 2$. Again, this can be done using standard parametric regression or more data adaptive methods. An example of a standard regression approach would be a logistic regression model with $\tilde{Q}^{z_0}_{1,n}(t_0 - 1;w)$ as the outcome and functions of w as predictors, fit using the subset of the data for which $Z = z_0, C(t_0 - 2) = N(t_0 - 2) = 0$. We iterate this estimation process at each time. Note that the subset of uncensored subjects used in the non-trivial estimation at each step is getting larger until eventually at $t = 1$, all subjects are used to obtain an estimate of $F^{z_0}_1(t_0;w)$. We then average over covariate levels to obtain an estimate of the unconditional cumulative incidence function using e.g. the empirical distribution of baseline covariates as our estimate of Q_W.

The TMLE procedure follows the estimation procedure outlined above, but adds in a targeting step at each time point. Thus, the procedure is as follows:

1. Generate initial estimates, $g_{z_0,n}(w)$ and $G_{C,n}(\cdot, w)$, of the treatment and censoring mechanisms. Use these to compute $H_{g,n}$.

2. Generate an initial estimate $\tilde{Q}^{z_0}_{1,n}(t_0; w)$ of the first conditional mean as outlined above.

3. Obtain $\epsilon_{t_0,n}$ by fitting logistic regression with $dn_1(t_0)$ as outcome, $logit\{\tilde{Q}^{z_0}_{1,n}(t_0; w)\}$ as offset, and $H_{g,n}(t_0, o)$ as covariate in the subset of data with $z = z_0, c(t_0 - 1) = n_1(t_0 - 1) = n_2(t_0 - 1) = 0$. Set $\tilde{Q}^{z_0,*}_{1,n}(t_0; w) = \tilde{Q}^{z_0}_{1,n}(t_0; w)(\epsilon_{t_0,n})$. Let $t = t_0 - 1$.

4. Generate an initial estimate $\tilde{Q}^{z_0}_{1,n}(t; w)$ of the t-th conditional mean as outlined above using $\tilde{Q}^{z_0}_{1,n}(t + 1; w)$ as the outcome in e.g. a logistic regression model or a data adaptive learning algorithm.

5. Obtain $\epsilon_{t,n}$ by fitting logistic regression with $\tilde{Q}^{z_0,*}_{1,n}(t + 1; w)$ as outcome, $\tilde{Q}^{z_0}_{1,n}(t; w)$ as offset, and $H_{g,n}(t, o)$ as covariate in the subset of data with $z = z_0, c(t-1) = n_1(t-1) = n_2(t - 1) = 0$.

6. Iterate steps 4-5 for $t = t_0 - 2, ..., 1$.

7. Use the empirical distribution of covariates to average $\tilde{Q}^{z_0,*}_{1,n}(1; w) = F^{z_0}_1(t_0; w)$ over the distribution of W, obtaining estimator $\psi^*_n = F^{z_0,*}_{1,n}(t_0)$.

D.3 Practical implementation

We first note that the targeting in the mean-based TMLE is completed in a single step, as opposed to the hazard-based approach, which requires iteration. However, the mean-based TMLE requires estimation of $\tilde{Q}^{z_0}_{1,0}(t; w)$ at each time, whereas there were only two steps of initial estimation in the hazard-based approach. If the initial estimation is being done with a computationally intensive method, e.g. Super Learner, then the mean-based approach might be much slower to implement in practice. However, for a modest number of discrete time points, as in the simulation studies presented, we found that the methods take approximately the same amount of time to execute.

Appendix E

SIMULATION DETAILS

E.1 Simulation parameters

For the first simulation study, we set $\beta_0 = -3.876$, $exp(\beta_1) = 0.474$, and $exp(\beta_3) = 1$ for all scenarios. Table E.1 contains values for all other parameters. The exponentiated parameters can be interpreted as a cause-specific hazard ratio for a unit difference in the associated covariate. For the second simulation study, we set $\beta_0 = -3.876$, $exp(\beta_1) = 3.123$, $exp(\gamma_1) = 1.557$ for all scenarios and include additional parameter values in Table 2.

E.2 Super Learner libraries

Table E.2 shows the algorithms included in the Super Learner library for cause-specific hazards (TMLE-1), iterated conditional means (TMLE-2) and censoring (both TMLE-1 and TMLE-2) in the simulation studies. The generalized additive models were computed using the gam R package [Hastie, 2013] and the Super Learner algorithm was executed using the SuperLearner R package [Polley and van der Laan, 2013].

E.3 Standard error estimates

Figure E.3 shows the performance of the influence curve-based standard error estimates of $\hat{F}_1^{z_0}(8)$ for $z_0 = 0, 1$ from simulation study 1. The estimates perform well in all settings.

Pred. cens.	$exp(\gamma_1)$	Cens. rate	γ_0	Pred. event	$exp(\beta_2)$	n
None	1	Low	-4.125	None	1	2180
None	1	Low	-4.115	Medium	1.557	2150
None	1	Low	-4.100	High	3.123	1570
None	1	High	-2.887	None	1	2630
None	1	High	-2.875	Medium	1.557	2450
None	1	High	-2.840	High	3.123	1780
High	1.557	Low	-4.210	None	1	2340
High	1.557	Low	-4.200	Medium	1.557	2180
High	1.557	Low	-4.130	High	3.123	1610
High	1.557	High	-2.935	None	1	2650
High	1.557	High	-2.920	Medium	1.557	2530
High	1.557	High	-2.850	High	3.123	1900

Table E.1: Parameters for simulation study 1 used to generate data for all scenarios considered: combinations of none/high covariate predictiveness of censoring (Pred. cens.), low/high censoring rates (Cens. rate), none/medium/high covariate predictiveness of events (Pred. event).

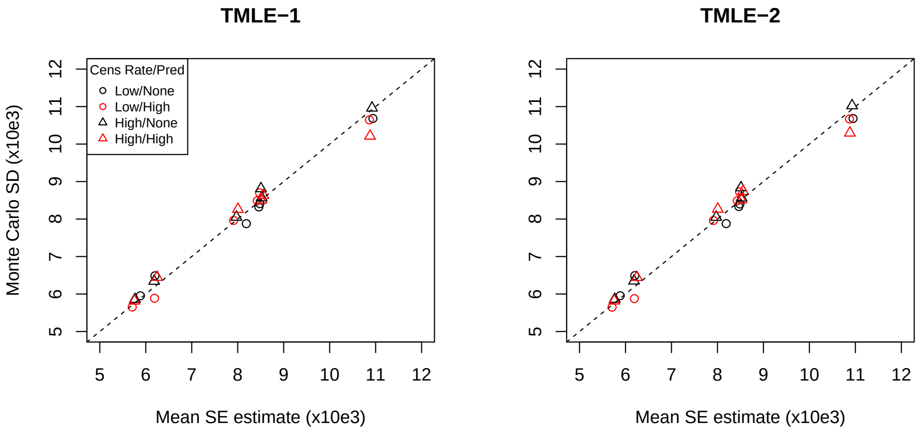

Figure E.1: Monte Carlo standard deviation (SD) plotted against the mean standard error (SE) estimate for simulation study 1 for the hazard-based TMLE (TMLE-1) and the mean-based TMLE (TMLE-2). The dashed line is the identity line, along which points should fall if estimates of standard error accurately reflect the true standard deviation of the estimates. Circles indicate the low censoring rate setting and triangles represent high censoring rate. Red points indicate high covariate predictiveness of censoring, while black points indicate low covariate predictiveness of censoring.

Eff. Mod.	$exp(\beta_3)$	VE	$exp(\beta_1)$	γ_0	n
No	1	0	1	-4.170	1640
No	1	0.25	0.726	-4.175	1880
No	1	0.5	0.474	-4.190	2190
No	1	0.75	0.225	-4.206	2630
Yes	0.280	0	1.620	-4.120	1200
Yes	0.474	0.25	1	-4.170	1880
Yes	0.220	0.5	0.830	-4.200	2150
Yes	0.250	0.75	0.375	-4.215	2630

Table E.2: Parameters for simulation study 2 used to generate data for all scenarios considered: combinations of effect modification yes/no (Eff. mod.), and vaccine efficacy 0,0.25,0.5,0.75 (VE).

Model type	Time	Covariates
Hazard estimates (TMLE-1)		
`glm`	$\emptyset$	Z
`glm`	$\emptyset$	$Z + W_1 + W_2$
`glm`	$\emptyset$	$Z + W_1 + W_2 + Z * W_2$
`glm`	`factor(t)`	$Z + W_1 + W_2$
`glm`	`factor(t)`	$Z + W_1 + W_2 + Z * W_2$
`gam`	$\emptyset$	$Z + \texttt{s}(W_1, df = 3) + W_2$
`gam`	$\texttt{s}(t, df = 3)$	$Z + W_1 + W_2$
`gam`	$\emptyset$	$Z + \texttt{s}(W_1, df = 3) + W_2 + Z * W_2$
`gam`	$\texttt{s}(t, df = 3)$	$Z + W_1 + W_2 + Z * W_2$
Conditional mean estimates (TMLE-2)		
`glm`	$\emptyset$	Z
`glm`	$\emptyset$	$Z + W_1 + W_2$
`glm`	$\emptyset$	$Z + W_1 + W_2 + Z * W_2$
`gam`	$\emptyset$	$Z + \texttt{s}(W_1, df = 3) + W_2$
`gam`	$\emptyset$	$Z + \texttt{s}(W_1, df = 3) + W_2 + Z * W_2$
Censoring estimates (TMLE-1 & TMLE-2)		
`glm`	$\emptyset$	$\emptyset$
`glm`	$\emptyset$	$Z + W_1 + W_2$
`glm`	$\emptyset$	$Z + W_1 + W_2 + Z * W_2$
`glm`	`factor(t)`	$Z + W_1 + W_2$
`glm`	`factor(t)`	$Z + W_1 + W_2 + Z * W_2$
`gam`	$\emptyset$	$Z + \texttt{s}(W_1, df = 3) + W_2$
`gam`	$\texttt{s}(t, df = 3)$	$Z + W_1 + W_2$
`gam`	$\emptyset$	$Z + \texttt{s}(W_1, df = 3) + W_2 + Z * W_2$
`gam`	$\texttt{s}(t, df = 3)$	$Z + W_1 + W_2 + Z * W_2$

Table E.3: Models included in Super Learner libraries for simulation studies. The columns indicate model (`glm`=generalized linear model, `gam` = generalized additive model), time ($\emptyset$ denotes time was omitted from the model, `factor(t)` dummy variables), and covariates ($x * y$ indicates a cross product between covariates x and y). $\texttt{s}(x, df = d)$ denotes that variable x

Appendix F

ADDITIONAL SIMULATION STUDIES

We examined the performance of our estimators when covariates are measured with error, e.g. self-reported risk behavior in an HIV vaccine efficacy trial. To generate data with measurement error, we suppose there exists a latent variable $L_1 \sim \mathrm{N}(0, \sigma^2_{L_1})$ and W_1 is a version of L_1 subject to measurement error,

$$W_1 = L_1 + e_1, \quad e_1 \sim \mathrm{N}(0, \sigma^2_{e_1}), \quad e_1 \perp L_1,$$

so that $W_1 \sim \mathrm{N}(0, \sigma^2_1)$, with $\sigma^2_1 = \sigma^2_{L_1} + \sigma^2_{e_1}$.

We suppose there exists a second latent variable, $L_2^\star$, with

$$L_2^\star = L_2^\dagger + e_2, \quad L_2^\dagger \sim N(0, \sigma^2_{L_2}), \quad e_2 \sim \mathrm{N}(0, \sigma^2_{e_2}), \quad e_2 \perp L_2^\dagger,$$

so that $L_2^\star \sim \mathrm{N}(0, \sigma^2_2)$, with $\sigma^2_2 = \sigma^2_{L_2} + \sigma^2_{e_2}$. W_2 is a dichotomized version of $L_2^\star$, i.e. $W_2 \equiv I(L_2^\star > \theta)$, which implies

$$W_2 \sim \text{Bernoulli}\{\Phi(-\theta/\sigma_{obs})\},$$

where Φ is the standard normal CDF. We let $L_2 \equiv I(L_2^\dagger > \theta)$ be the true variable we would like to have measured. We denote $\rho_1 \equiv 1 - \sigma^2_{e_1}/\sigma^2_1$ as the fraction of variability of W_1 explained by L_1; similarly, $\rho_2 \equiv 1 - \sigma^2_{e_2}/\sigma^2_2$ is the fraction of variability in $L_2^\star$ explained by $L_2^\dagger$. We let $\rho_1 = 0.9$ and repeated simulation study 1. We then let $\rho_1 = \rho_2 = 0.9$ and $\theta = 0$ and repeated simulation study 2.

F.1 Simulation 1 Results

Figures F.1, F.1, and F.1 show the results for the first simulation study when covariates are measured with measurement error. We see that when covariates are measured with error

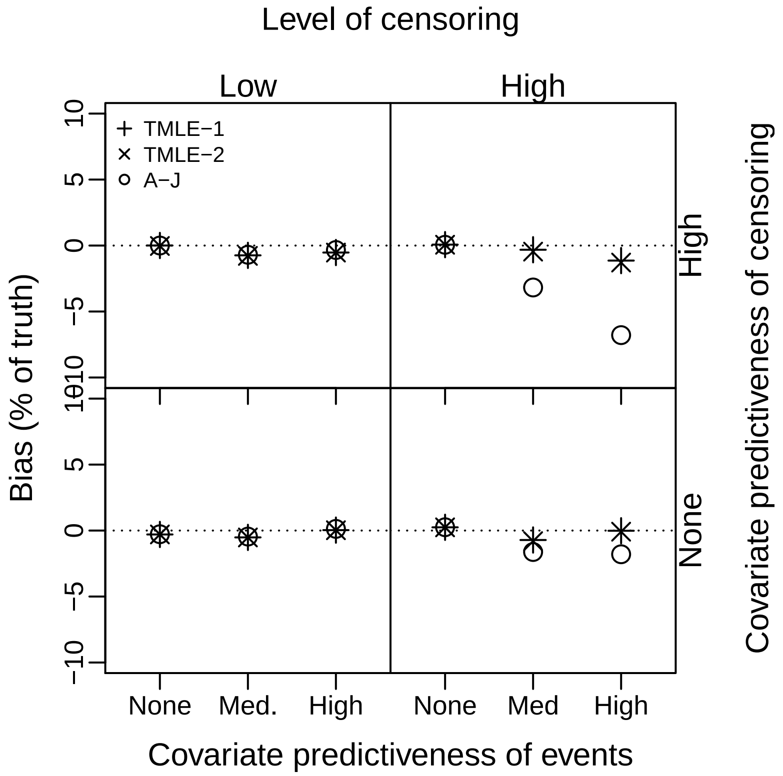

Figure F.1: Bias of the hazard-based TMLE (TMLE-1), mean-based TMLE (TMLE-2), and the Aalen-Johansen (A-J) estimator when covariates are measured with error. The bias is presented as a percentage of the true cumulative incidence function.

TMLE estimates have some bias and efficiency gains are less than when there is no error in covariate measurement. In general, when covariates are not prognostic TMLE still performs as well as Aalen-Johansen.

F.2 Simulation 2 Results

Figure F.2 shows the results for simulation study 2 when covariates are measured with error. We see that TMLE still provides some gains in power over Aalen-Johansen for moderately effective vaccines. When there was no measurement error, TMLE offered a larger gain in efficiency in the effect modification scenario. However, we see that the gain in power in the effect modification scenario is smaller than the gain in the no effect modification scenario.

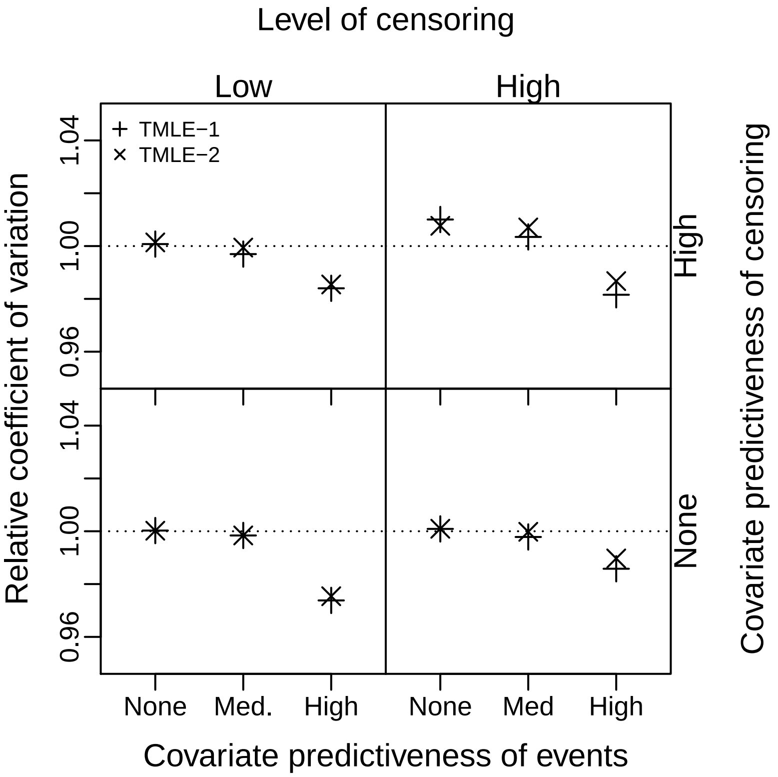

Figure F.2: Relative coefficient of variation of the hazard-based (TMLE-1) and mean-based (TMLE-2) TMLE estimators compared to the Aalen-Johansen (A-J) estimator when covariates are measured with error. Values below 1.0 indicate a smaller coefficient of variation for the TMLE estimator.

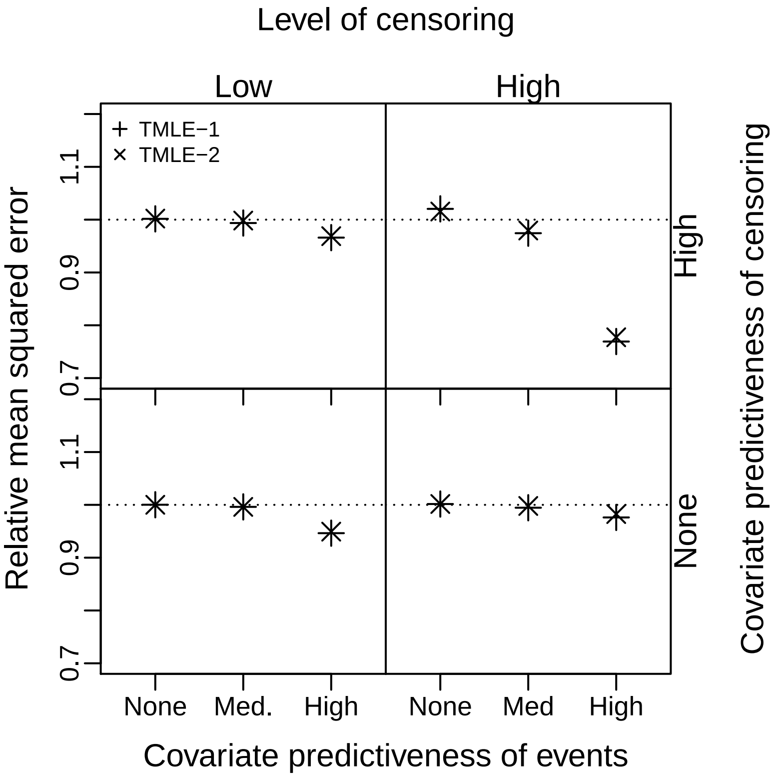

Figure F.3: Relative mean squared error of the hazard-based (TMLE-1) and mean-based (TMLE-2) estimators compared to the Aalen-Johansen (A-J) estimator when covariates are measured with error. Values below 1.0 indicate a smaller mean squared error for the TMLE estimator.

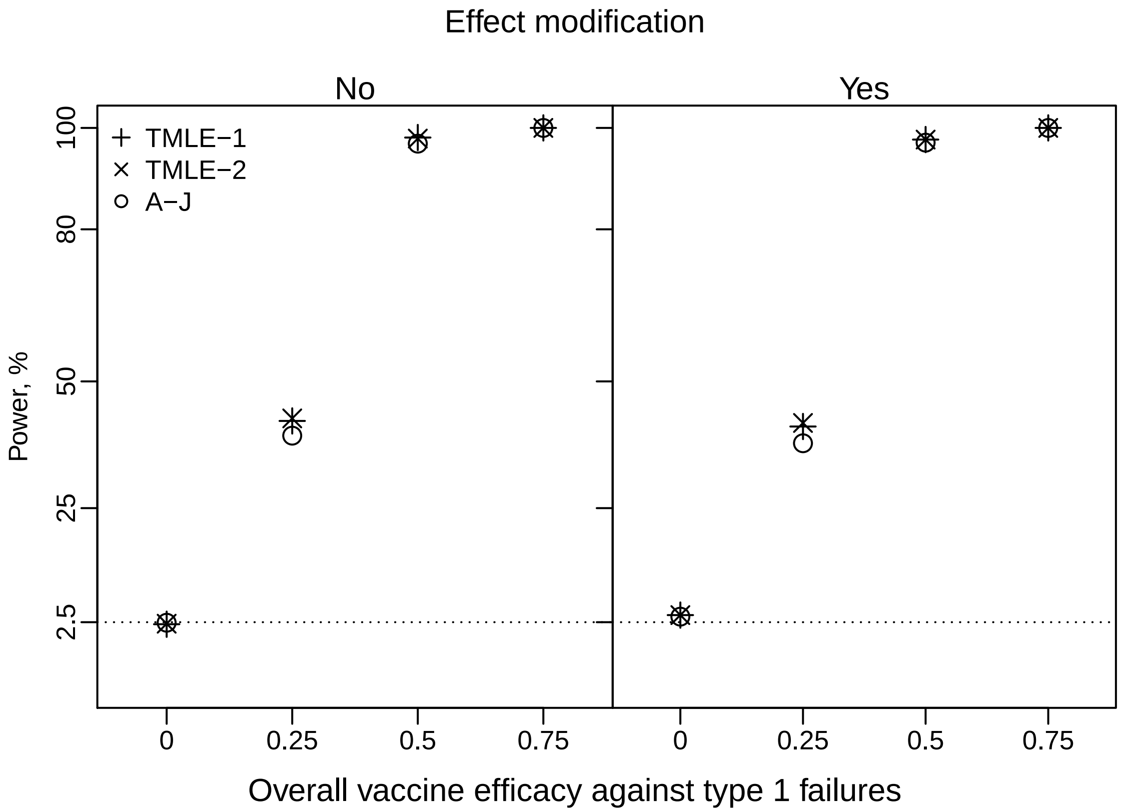

Power relative to A−J

TMLE−1	1.07	1.01	1.00	1.09	1.01	1.00
TMLE−2	1.09	1.01	1.00	1.11	1.01	1.00

Figure F.4: Power of a one-sided level $\alpha = 0.025$ (dotted line) Wald test of the null hypothesis of no vaccine efficacy against type 1 events at the end of follow up when covariates are measured with error. Tests are based on hazard-based TMLE (TMLE-1), mean-based TMLE (TMLE-2), and the Aalen-Johansen (A-J) estimators. The left panel shows power with no effect modification; the right panel shows power with effect modification.

Appendix G

RV144 ANALYSIS DETAILS

Table G shows the models that were included in the Super Learner library for the RV144 analysis. Bayesian generalized linear models were calculated using the arm package [Gelman and Su, 2014]. We use B to denote baseline behavioral risk score (low/medium/high), G to denote gender, and A to denote age category ($\leq$ 20, 21-25, $\geq$ 26). These covariates were included in models using dummy variables.

Model type	Time	Covariates
Conditional mean estimates (TMLE-2)		
`glm`	$\emptyset$	Z
`glm`	$\emptyset$	$Z + B + G + A$
`glm`	$\emptyset$	$Z * B$
`bayesglm`	$\emptyset$	Z
`bayesglm`	$\emptyset$	$Z + B + G + A$
`bayesglm`	$\emptyset$	$Z * B$
`step`	$\emptyset$	$Z * B + G + A$
Censoring estimates (TMLE-1 & TMLE-2)		
`glm`	`factor(t)`	$\emptyset$
`glm`	`factor(t)`	$Z * t$
`glm`	$\emptyset$	Z
`glm`	$\emptyset$	$Z + B + G + A$
`glm`	t	Z
`glm`	t	$Z + B + G + A$
`glm`	$log(t)$	Z
`glm`	$log(t)$	$Z + B + G + A$
`step`	`factor(t)`	$Z + B + G + A$
`step`	t	$Z + B + G + A$
`step`	$log(t)$	$Z + B + G + A$
`gam`	$s(t, 3)$	Z
`gam`	$s(t, 3)$	$Z + B + G + A$
`gam`	$s(t, 2)$	Z
`gam`	$s(t, 2)$	$Z + B + G + A$

Table G.1: Models included in Super Learner libraries for RV144 analysis. The columns indicate what type of model was used (`glm`=generalized linear model, `bayesglm` = bayesian glm with default priors, `step` = stepwise glm using both AIC and BIC as selection criteria), how time was modeled ($\emptyset$ denotes time was omitted from the model, `factor(t)` indicates dummy variables were used), and what covariates were included ($x * y$ indicates a cross product between covariates x and y). We use $\mathtt{s}(x, df = d)$ to denote that variable x was modeled using a polynomial spline of degree d.

Appendix H

HVTN 505 RESULTS

Figure H shows results from an analysis of the HVTN 505 HIV vaccine efficacy trial using TMLE. We implemented the mean-based TMLE using a the Super Learner to generate our initial estimates at each time point as well as the censoring. We considered self-reported sexual risk behavior (number of sexual partners, unprotected sex, drug/alcohol use, and a derived sexual risk score). The Super Learner library included generalized linear models, Bayesian generalized linear models, generalized additive models, and stepwise selection procedures very similar to those listed in Table G. The point estimates given by TMLE were very similar to those obtain via Aalen-Johansen, while the estimated variances were generally 1-3% smaller (Figure H).

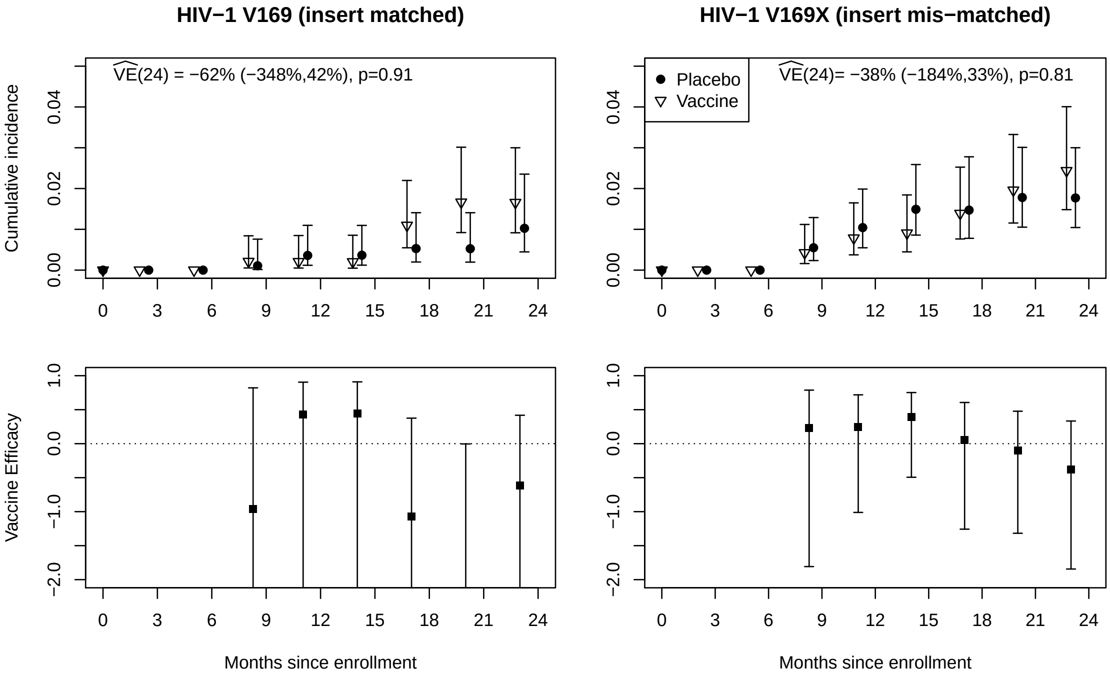

Figure H.1: Results from the HVTN 505 HIV vaccine efficacy trial. The top row shows estimated cumulative incidence over time using the hazard-based TMLE for 169-matched and mismatched HIV-infections with point-wise 95% confidence intervals. The bottom row shows estimated vaccine efficacy (VE) and 95% confidence interval (CI) for each type of infection.

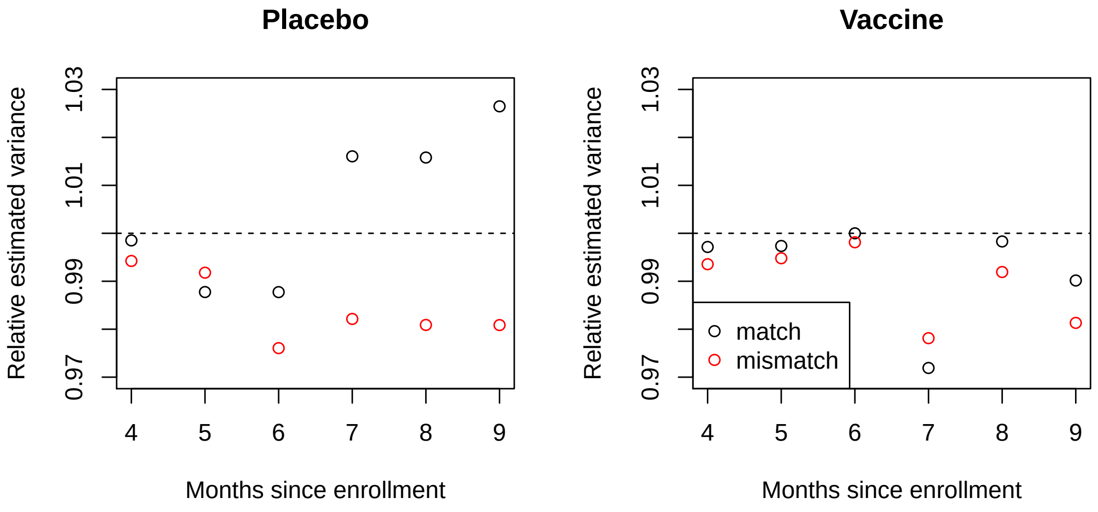

Figure H.2: Comparison of the estimated variance of Aalen-Johansen estimator and TMLE estimator across time points in HVTN-505. Points below 1 indicate a smaller estimated variance for TMLE estimators.

Appendix I

RARE EVENTS SIMULATION 1 - MULTIPLE FAILURE TYPES

I.1 Design

We now consider estimation of cumulative incidence when multiple failure types are present. Again, we assume that at baseline either a treatment and placebo is assigned with equal probability and baseline covariates are measured with the same distribution as the first simulation. The censoring distribution is also the same as in the first simulation. The cause-specific hazard for events at each time is given by $\bar{Q}_j^z(t;w) = expit(\beta_{0,j} - z + 2w_1 + w_2 - 3w_3)$, $j = 1, 2$. We evaluate three situations: (1) both event types are common $\beta_{0,1:2} = (-2, -2)$; (2) one common event and one rare event $\beta_{0,1:2} = (-2, -6)$; and (3) both event types are rare $\beta_{0,1:2} = (-6, -6)$. In each case we focus on estimating the additive effect of Z on the cumulative incidence of type-2 events. We again compare our proposed restricted estimator to the unrestricted TMLE and the nonparametric maximum likelihood estimator, the Aalen-Johannsen estimator.

Similar to the first simulation, in both TMLEs the censoring mechanism is estimated using a correctly specified logistic regression model. In the unconstrained TMLE, the iteratively-defined means are estimated using a correctly specified logistic regression, while the restricted TMLE minimizes the bounded binary log-likelihood loss function according to the correctly specified main terms model. In this simulation, the restricted TMLE imposes bounds equal to the exact upper limit of the iterative means. This allows us to assess the "best-case" scenario for finite sample gains that can be achieved by the estimator. The second simulation study in the following section considers how the performance depends on how close the posited bounds are to the true bounds.

I.2 Results

The results are shown in Figure I.2. Here, we see a similar story to the case of a single failure type. The bias of the restricted TMLE is typically at least as good as the other estimators, while its variance is far better for small sample sizes. As the sample size grows, the unrestricted TMLE's performance begins to match that of the restricted TMLE.

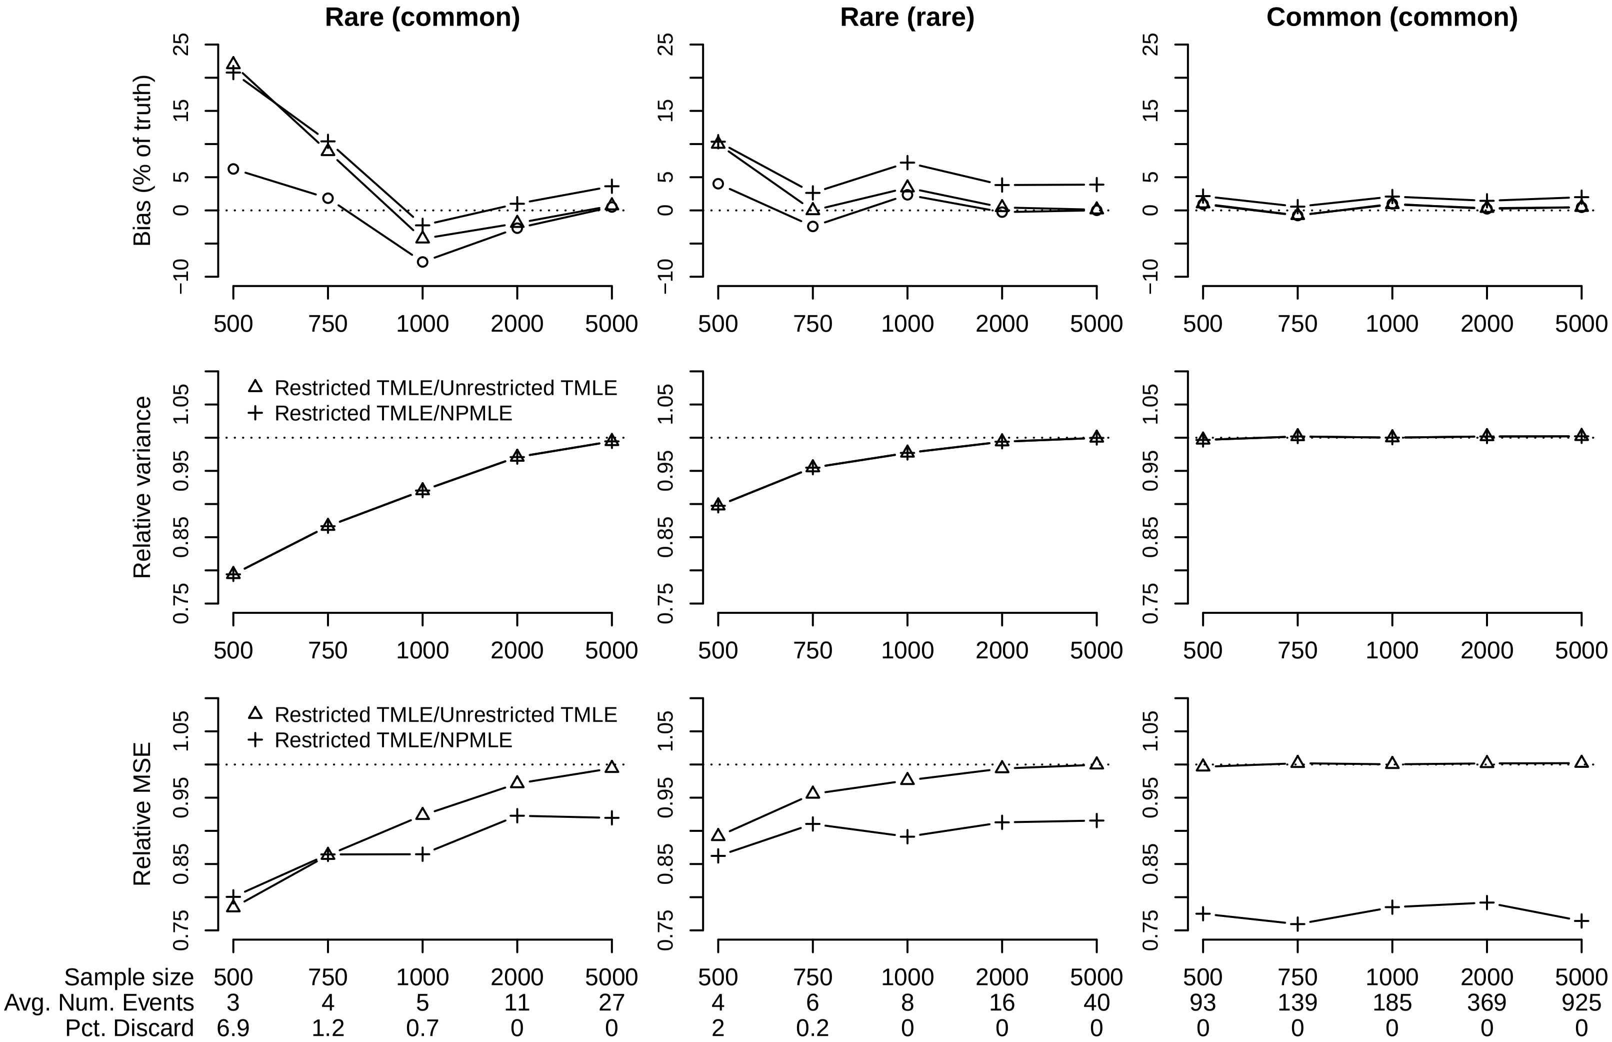

Figure I.1: Results from Simulation 1 with multiple failure types. The bias, relative variance, and relative mean-squared error (MSE) of the three estimators are shown in the top, middle, and bottom rows, respectively.

Appendix J

RARE EVENTS SIMULATION 2 - MULTIPLE FAILURE TYPES

J.1 Design

We now consider how the tightness of the bounds on the iteratively-defined conditional means affects the performance of the restricted TMLE estimator of cumulative incidence. To that end, we consider the same data-generating mechanism described in Appendix I. However, rather than using the exact bound, we use a multiple thereof. We considered increasing and decreasing the bound by 10 and 25% and compare the bias, variance, and mean-squared error of the restricted TMLE to the unrestricted TMLE. We again expect that bounds that are too tight will lead to improvements in finite sample variance but at the cost of asymptotic bias. As bounds become larger than the true bound, decreasing improvements in variance relative to the unrestricted estimator are expected while asymptotic bias should be negligible.

J.2 Results

Figure J.2 shows the performance of the estimators relative to the NPMLE (the Aalen-Johansen estimator). For very small sample sizes, we find that the bias of the restricted TMLE compares favorably to that of the unrestricted TMLE though no clear trend is evident. If more than about ten events of the rare failure type are observed, the estimators tend to perform approximately the same in terms of their bias. We see that the choice of bounds that are too tight comes at the cost of increased bias in large samples (top right figure). However, we see an indication that the double-robustness property of the estimator kicking in at large sample sizes. This is reflected by the decreasing bias as sample size gets large. Similar to the

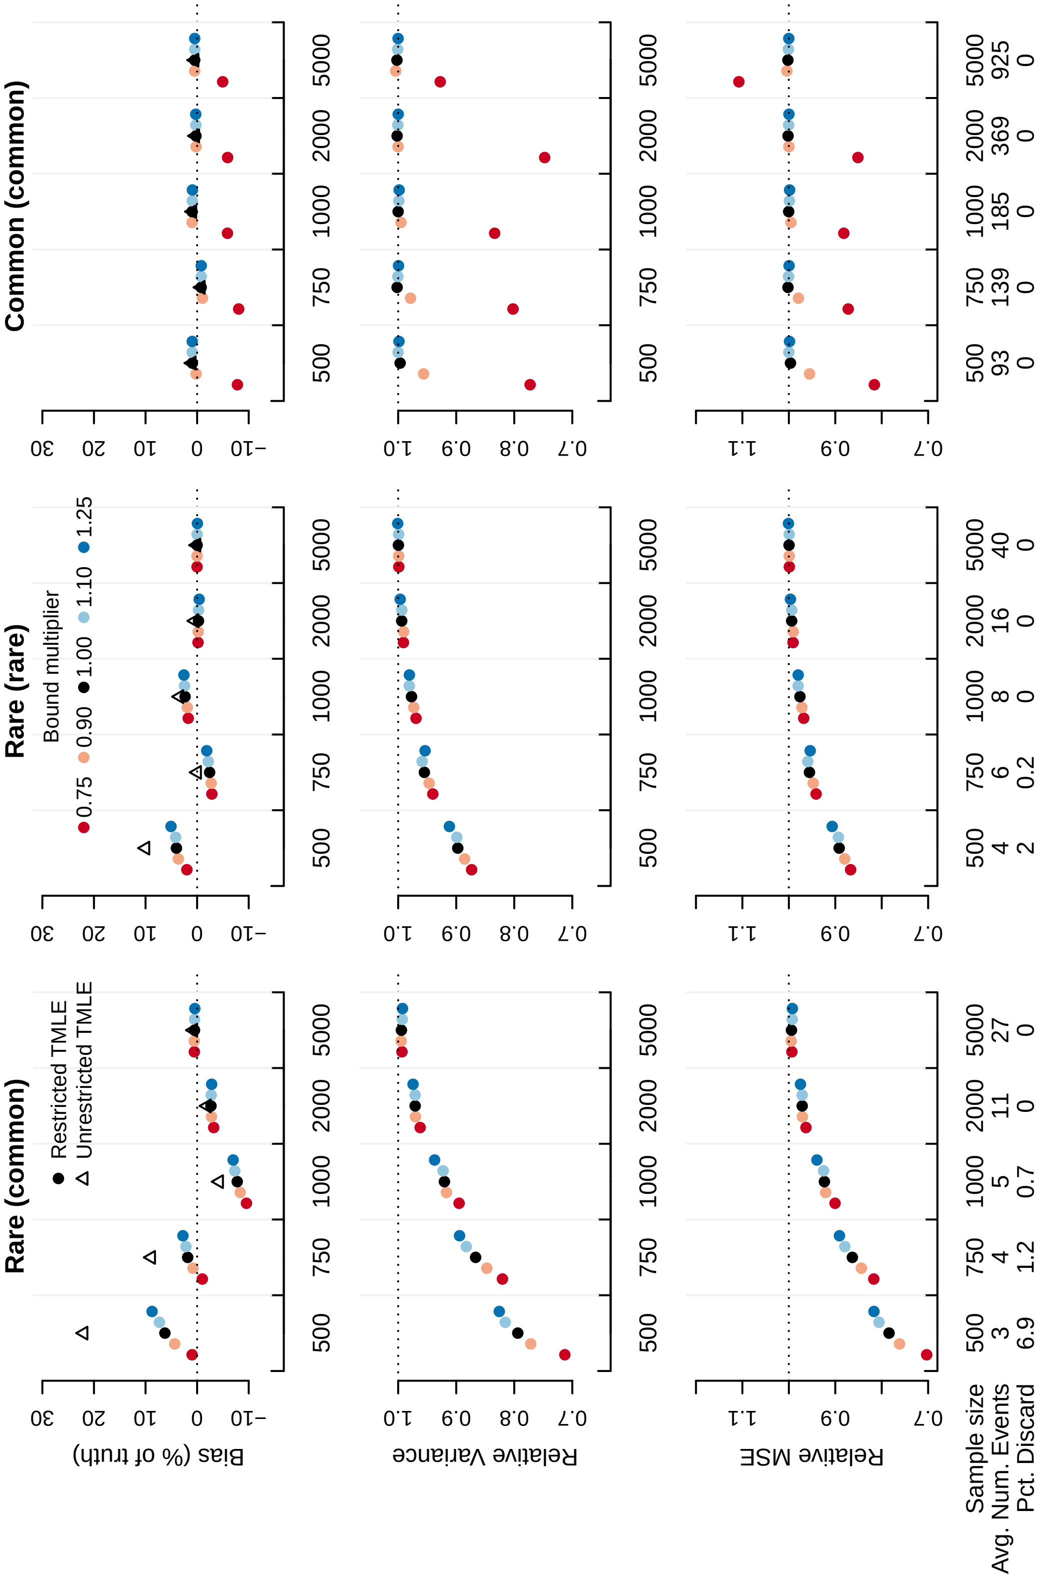
Rare (common)
Rare (rare)
Common (common)
Restricted TMLE
Unrestricted TMLE
Bound multiplier
0.75
0.90
1.00
1.10
1.25
Bias (% of truth)
Relative Variance
Relative MSE
30
20
10
0
−10
1.0
0.9
0.8
0.7
1.1
Sample size
Avg. Num. Events
Pct. Discard
500 3 6.9
750 4 1.2
1000 5 0.7
2000 11 0
5000 27 0
500 4 2
750 6 0.2
1000 8 0
2000 16 0
5000 40 0
500 93 0
750 139 0
1000 185 0
2000 369 0
5000 925 0

single failure type setting, the middle row shows that, unsurprisingly the tighter the bounds the smaller the variance in finite samples. Even for the loosest bounds, we find improved variance relative to the unrestricted TMLE. This is also reflected in the mean-squared error, shown in the bottom row.

Appendix K

PROOFS FOR CHAPTER 3

We provide a justification of the claim that

$$P_0\left\{\left(\frac{\bar{Q}_{1,+}-\bar{Q}_{1,0}}{\bar{g}_{0,0}}\right)\left(\bar{g}_{0,n}-\bar{g}_{0,0}\right)\right\}=P_0\left\{\frac{A_0A_1}{\bar{g}_{0,0}\bar{g}_{1,0}}\left(\frac{L_2-\bar{Q}_{1,+}}{\bar{g}_{0,0}}\right)\left(\bar{g}_{0,n}-\bar{g}_{0,0}\right)\right\}.$$

Working backwards, we have that

$$\begin{aligned}
&E_0\left[\frac{A_0A_1}{\bar{g}_{0,0}(L_0)\bar{g}_{1,0}(\bar{L}_1)}\left\{\frac{L_2-\bar{Q}_{1,+}(\bar{L}_0)}{\bar{g}_{0,0}(L_0)}\right\}\left\{\bar{g}_{0,n}(L_0)-\bar{g}_{0,0}(L_0)\right\}\right]\\
&\qquad=E_0\left(E_0\left[\frac{A_0A_1}{\bar{g}_{0,0}(L_0)\bar{g}_{1,0}(\bar{L}_1)}\left\{\frac{L_2-\bar{Q}_{1,+}(\bar{L}_0)}{\bar{g}_{0,0}(L_0)}\right\}\left\{\bar{g}_{0,n}(L_0)-\bar{g}_{0,0}(L_0)\right\}\;\middle|\;\bar{A}_1,\bar{L}_1\right]\right)\\
&\qquad=E_0\left(\frac{A_0A_1}{\bar{g}_{0,0}(L_0)\bar{g}_{1,0}(\bar{L}_1)}\left\{\frac{E_0(L_2\mid\bar{A}_1,\bar{L}_1)-\bar{Q}_{1,+}(\bar{L}_0)}{\bar{g}_{0,0}(L_0)}\right\}\left\{\bar{g}_{0,n}(L_0)-\bar{g}_{0,0}(L_0)\right\}\right)\\
&\qquad=E_0\left[\frac{A_0A_1}{\bar{g}_{0,0}(L_0)\bar{g}_{1,0}(\bar{L}_1)}\left\{\frac{\bar{Q}_{2,0}(\bar{L}_1)-\bar{Q}_{1,+}(\bar{L}_0)}{\bar{g}_{0,0}(L_0)}\right\}\left\{\bar{g}_{0,n}(L_0)-\bar{g}_{0,0}(L_0)\right\}\right]\\
&\qquad=E_0\left(E_0\left[\frac{A_0A_1}{\bar{g}_{0,0}(L_0)\bar{g}_{1,0}(\bar{L}_1)}\left\{\frac{\bar{Q}_{2,0}(\bar{L}_1)-\bar{Q}_{1,+}(\bar{L}_0)}{\bar{g}_{0,0}(L_0)}\right\}\left\{\bar{g}_{0,n}(L_0)-\bar{g}_{0,0}(L_0)\right\}\middle|\;\bar{L}_1,A_0\right]\right)\\
&\qquad=E_0\left[\frac{A_0}{\bar{g}_{0,0}(L_0)}\left\{\frac{\bar{Q}_{2,0}(\bar{L}_1)-\bar{Q}_{1,+}(\bar{L}_0)}{\bar{g}_{0,0}(L_0)}\right\}\left\{\bar{g}_{0,n}(L_0)-\bar{g}_{0,0}(L_0)\right\}\right]\\
&\qquad=E_0\left(\left[\frac{A_0}{\bar{g}_{0,0}(L_0)}\left\{\frac{\bar{Q}_{2,0}(\bar{L}_1)-\bar{Q}_{1,+}(\bar{L}_0)}{\bar{g}_{0,0}(L_0)}\right\}\left\{\bar{g}_{0,n}(L_0)-\bar{g}_{0,0}(L_0)\right\}\;\middle|\;A_0,L_0\right]\right)\\
&\qquad=E_0\left(\frac{A_0}{\bar{g}_{0,0}(L_0)}\left\{\frac{E_0\left[\bar{Q}_{2,0}(\bar{L}_1)\mid A_0,L_0\right]-\bar{Q}_{1,+}(\bar{L}_0)}{\bar{g}_{0,0}(L_0)}\right\}\left\{\bar{g}_{0,n}(L_0)-\bar{g}_{0,0}(L_0)\right\}\right)\\
&\qquad=E_0\left[\frac{A_0}{\bar{g}_{0,0}(L_0)}\left\{\frac{\bar{Q}_{1,0}(\bar{L}_0)-\bar{Q}_{1,+}(\bar{L}_0)}{\bar{g}_{0,0}(L_0)}\right\}\left\{\bar{g}_{0,n}(L_0)-\bar{g}_{0,0}(L_0)\right\}\right]\\
&\qquad=E_0\left(E_0\left[\frac{A_0}{\bar{g}_{0,0}(L_0)}\left\{\frac{\bar{Q}_{1,0}(\bar{L}_0)-\bar{Q}_{1,+}(\bar{L}_0)}{\bar{g}_{0,0}(L_0)}\right\}\left\{\bar{g}_{0,n}(L_0)-\bar{g}_{0,0}(L_0)\right\}\;\middle|L_0\right]\right)\\
&\qquad=E_0\left[\left\{\frac{\bar{Q}_{1,0}(\bar{L}_0)-\bar{Q}_{1,+}(\bar{L}_0)}{\bar{g}_{0,0}(L_0)}\right\}\left\{\bar{g}_{0,n}(L_0)-\bar{g}_{0,0}(L_0)\right\}\right]
\end{aligned}$$

Appendix L

ANALYSIS OF A GENERAL REMAINDER TERM

Consider the case of K timepoints and the remainder term involving the cross-product of the outcome and propensity regressions at time k. We introduce the shorthand notation $A_{0:k} := \prod_{m=0}^{k} A_m$ and $\bar{g}_{0:k} := \prod_{m=0}^{k} \bar{g}_m$. The form of the remainder term is

$$R_{k,n} := R_{k,n}(Q_n, Q_0, g_n, g_0) := P_0 \left\{ \frac{A_{0:k-2}}{\bar{g}_{0:k-2,n}} \left(\frac{\bar{g}_{k-1,0} - \bar{g}_{k-1,n}}{\bar{g}_{k-1,n}} \right) \left(\bar{Q}_{k,0} - \bar{Q}_{k,n} \right) \right\} .$$

This may be written as

$$\begin{aligned} R_{k,n} = & \left\{ \Gamma_{k-1,0n}(\bar{g}_{k-1,n}) - \Gamma_{k-1,0n}(\bar{g}_{k-1,0}) \right\} + \left\{ \Phi_{k,0n}(\bar{Q}_{k,n}) - \Phi_{k,0n}(\bar{Q}_{k,0}) \right\} + \\ & S_{k,1}(Q_n, Q_0, g_n, g_0) + S_{k,2}(Q_n, Q_0, g_n, g_0) \ , \end{aligned}$$

where we have defined

$$\begin{aligned} \Gamma_{k-1,0n}(\bar{g}_{k-1}) &:= P_0 \left\{ \frac{A_{0:k-2}}{\bar{g}_{0:k-2,n}} \left(\frac{\bar{Q}_{k,+} - \bar{Q}_{k,0}}{\bar{g}_{k-1,0}} \right) \bar{g}_{k-1} \right\} , \\ \Phi_{k,0n}(\bar{Q}_k) &:= P_0 \left\{ \frac{A_{0:k-2}}{\bar{g}_{0:k-2,n}} \left(\frac{\bar{g}_{k-1,+} - \bar{g}_{k-1,0}}{\bar{g}_{k-1,+}} \right) \bar{Q}_k \right\} , \\ S_{k,1}(Q_n, Q_0, g_n, g_0) &:= P_0 \left\{ \frac{A_{0:k-2}}{\bar{g}_{0:k-2,n}} \left(\frac{\bar{g}_{k-1,+} - \bar{g}_{k-1,n}}{\bar{g}_{k-1,+} \ \bar{g}_{k-1,n}} \right) \left(\bar{g}_{1,0} - \bar{g}_{1,n} \right) \left(\bar{Q}_{k,0} - \bar{Q}_{k,n} \right) \right\} \\ S_{k,2}(Q_n, Q_0, g_n, g_0) &:= P_0 \left\{ \frac{A_{0:k-2}}{\bar{g}_{0:k-2,n}} \left(\frac{\bar{g}_{k-1,+} - \bar{g}_{k-1,n}}{\bar{g}_{k-1,+}} \right) \left(\bar{Q}_{k,+} - \bar{Q}_{k,n} \right) \right\} . \end{aligned}$$

We start with the analysis of $\Gamma_{k-1,0n}$. We can write

$$\begin{aligned}
\Gamma_{k-1,0n}(\bar{g}_{k-1,n}) - \Gamma_{k-1,0n}(\bar{g}_{k-1,0}) &= P_0\left\{ \frac{A_{0:k-2}}{\bar{g}_{0:k-2,n}} \left(\frac{\bar{Q}_{k,+} - \bar{Q}_{k,0}}{\bar{g}_{k-1,0}} \right) \big(\bar{g}_{k-1,n} - \bar{g}_{k-1,0} \big) \right\} \\
&= -P_0\left\{ \frac{A_{0:K-1}}{\bar{g}_{0:k-2,n}\ \bar{g}_{k-1:K-1,0}} \left(\frac{Y - \bar{Q}_{k,+}}{\bar{g}_{k-1,0}} \right) \big(\bar{g}_{k-1,n} - \bar{g}_{k-1,0} \big) \right\} \\
&= -P_0\left\{ \frac{\bar{Q}^r_{k,0n}}{\bar{g}^2_{k-1,0}} \big(\bar{g}_{k-1,n} - \bar{g}_{k-1,0} \big) \right\} \\
&= -P_0\left\{ \frac{\bar{Q}^r_{k,n}}{\bar{g}^2_{k-1,n}} \big(\bar{g}_{k-1,n} - \bar{g}_{k-1,0} \big) \right\} \\
&\qquad + S_{Q_k,1}(Q^r_{0n}, Q^r_0, g_n, g_0) + S_{Q_k,2}(Q^r_n, Q^r_0, g_n, g_0)\ ,
\end{aligned}$$

with the reduced-dimension regression

$$\bar{Q}^r_{k,0n}(\bar{l}_k, \bar{a}_{k-1}) := E_0\left[\frac{A_{0:K-1}}{\bar{g}_{0:k-2,n}\ \bar{g}_{k-1:K-1,0}} \left(\frac{Y - \bar{Q}_{k,+}}{\bar{g}_{k-1,0}} \right) \middle| \bar{A}_{k-1} = \bar{a}_{k-1}, \bar{g}_{k-1,n}(\bar{l}_k), \bar{g}_{k-1,0}(\bar{l}_k) \right]$$

where we defined the second-order terms

$$\begin{aligned}
S_{Q_k,1}(Q^r_{0n}, Q^r_0, g_n, g_0) &:= -P_0\left\{ \frac{\bar{Q}^r_{k,0n} - \bar{Q}^r_{k,0}}{\bar{g}^2_{k-1,0}} \big(\bar{g}_{k-1,n} - \bar{g}_{k-1,0} \big) \right\} \\
S_{Q_k,2}(Q^r_n, Q^r_0, g_n, g_0) &:= P_0\left\{ \left(\frac{\bar{Q}^r_{k,n}}{\bar{g}^2_{k-1,n}} - \frac{\bar{Q}^r_{k,0}}{\bar{g}^2_{k-1,0}} \right) \big(\bar{g}_{k-1,n} - \bar{g}_{k-1,0} \big) \right\}\ ,
\end{aligned}$$

which are considered to be $o_P(n^{-1/2})$. We thus have shown that

$$\begin{aligned}
&\Gamma_{k-1,0n}(\bar{g}_{k-1,n}) - \Gamma_{k-1,0n}(\bar{g}_{k-1,0}) \\
&\qquad = -I(\bar{g}_{k-1,+} = \bar{g}_{k-1,0}, \bar{Q}_{k,+} \neq \bar{Q}_{k,0}) \big\{ \Gamma_{k-1,n}(\bar{g}_{k-1,n}) - \Gamma_{k-1,n}(\bar{g}_{k-1,0}) + o_P(n^{-1/2}) \big\}\ ,
\end{aligned}$$

where

$$\Gamma_{k-1,n}(\bar{g}_{k-1}) := P_0\left(\frac{\bar{Q}^r_{k,n}}{\bar{g}^2_{k-1,n}}\ \bar{g}_{k-1} \right)\ .$$

We then perform a first-order expansion of $\Gamma_{k-1,n}$ at $\bar{g}_{k-1,n}$:

$$\begin{aligned}
&\Gamma_{k-1,n}(\bar{g}_{k-1,n}) - \Gamma_{k-1,n}(\bar{g}_{k-1,0}) \\
&\qquad = -(P_n - P_0) D_{g_{k-1},0}(Q^r_0, g_0) + B_n(Q^r_n, g_n) + M_{g_{k-1}}(Q^r_n, Q_0, g_n, g_0)\ ,
\end{aligned}$$

where the canonical gradient of $\Gamma_{k-1,n}$ at (Q_0^r, g_0) is

$$D_{g_{k-1},0}(Q_0^r, g_0)(o) := \frac{\bar{Q}^r_{k,0}(\bar{a}_{k-1}, \bar{l}_{k-1})}{\bar{g}^2_{k-1,0}(\bar{l}_{k-1})} \left\{a_{k-1} - \bar{g}_{k-1,0}(\bar{l}_{k-2})\right\} ,$$

the first-order bias is given by $B_n(Q_n^r, g_n) := P_n D_{g_{k-1},n}(Q_n^r, g_n)$, and the empirical process term is given by

$$M_{g_{k-1}}(Q_n^r, Q_0, g_n, g_0) := \big(P_n - P_0\big) \left\{D_{g_{k-1},n}(Q_n^r, g_n) - D_{g_{k-1},0}(Q_0^r, g_0)\right\}$$

and is considered to be $o_P(n^{-1/2})$.

Continuing to the analysis of $\Phi_{k,0n}$:

$$\begin{aligned}
\Phi_{k,0n}(\bar{Q}_{k,n}) - \Phi_{k,0n}(\bar{Q}_{k,0}) &= P_0\left\{\frac{A_{0:k-2}}{\bar{g}_{0:k-1,n}} \left(\frac{\bar{g}_{k-1,+} - \bar{g}_{k-1,0}}{\bar{g}_{k-1,+}}\right) \big(\bar{Q}_{k,n} - \bar{Q}_{k,0}\big)\right\} \\
&= -P_0\left\{\frac{A_{0:k-2}}{\bar{g}_{0:k-1,n}} \left(\frac{A_{k-1} - \bar{g}_{k-1,+}}{\bar{g}_{k-1,+}}\right) \big(\bar{Q}_{k,n} - \bar{Q}_{k,0}\big)\right\} \\
&= -P_0\left\{\bar{g}^r_{k-1,0n} \, \big(\bar{Q}_{k,n} - \bar{Q}_{k,0}\big)\right\} \\
&= -P_0\left\{\frac{A_{0:k-1}}{\tilde{g}^r_{k-1,0n}} \, \bar{g}^r_{k-1,0n} \, \big(\bar{Q}_{k,n} - \bar{Q}_{k,0}\big)\right\} \\
&= -P_0\left\{\frac{A_{0:k-1}}{\tilde{g}^r_{k-1,n}} \, \bar{g}^r_{k-1,n} \, \big(\bar{Q}_{k,n} - \bar{Q}_{k,0}\big)\right\} \\
&\qquad + S_{Q_k,1}(Q_n, Q_0, g^r_{k,0n}, g^r_{k,0}) + S_{Q_k,2}(Q_n, Q_0, g^r_{k,n}, g^r_{k,0}) ,
\end{aligned}$$

where we defined the reduced dimension regression

$$\begin{aligned}
\bar{g}^r_{k-1,0n}(\bar{L}_{k-1}) &:= E_0\left[\frac{A_{0:k-2}}{\bar{g}_{0:k-1,n}(\bar{L}_{k-1})} \left\{\frac{A_{k-1} - \bar{g}_{k-1,+}(\bar{L}_{k-1})}{\bar{g}_{k-1,+}(\bar{L}_{k-1})}\right\} \,\middle|\, \bar{Q}_{k,n}(\bar{L}_{k-1}), \bar{Q}_{k,0}(\bar{L}_{k-1})\right] \text{, and} \\
\tilde{g}^r_{k-1,0n}(\bar{L}_{k-1}) &:= E_0\left\{A_{0:k-1} \,\middle|\, \bar{Q}_{k,n}(\bar{L}_{k-1}), \bar{Q}_{k,0}(\bar{L}_{k-1})\right\} ,
\end{aligned}$$

as well as the second-order terms

$$\begin{aligned}
S_{Q_k,1}(Q_n, Q_0, g^r_{k,0n}, g^r_{k,0}) &:= -P_0\left\{A_{0:k-1}\left(\frac{\bar{g}^r_{k-1,0n}}{\tilde{g}^r_{k-1,0n}} - \frac{\bar{g}^r_{k-1,0}}{\tilde{g}^r_{k-1,0}}\right) \big(\bar{Q}_{k,n} - \bar{Q}_{k,0}\big)\right\} \text{and} \\
S_{Q_k,2}(Q_n, Q_0, g^r_{k,0n}, g^r_{k,0}) &:= P_0\left\{A_{0:k-1}\left(\frac{\bar{g}^r_{k-1,n}}{\tilde{g}^r_{k-1,n}} - \frac{\bar{g}^r_{k-1,0}}{\tilde{g}^r_{k-1,0}}\right) \big(\bar{Q}_{k,n} - \bar{Q}_{k,0}\big)\right\} ,
\end{aligned}$$

which are considered to be $o_P(n^{-1/2})$. Thus, defining

$$\Phi_{k,n}(\bar{Q}_k) := P_0\left(\frac{A_{0:k-1}}{\tilde{g}^r_{\overline{k-1},n}} \ \bar{g}^r_{k-1,n} \ \bar{Q}_k \right) ,$$

we have shown that

$$\Phi_{k,0n}(\bar{Q}_{k,n}) - \Phi_{k,0n}(\bar{Q}_{k,0}) = I(\bar{Q}_{k,+} = \bar{Q}_{k,0}, \bar{g}_{k-1,+} \neq \bar{g}_{k-1,0}) \ \big\{ \Phi_{k,n}(\bar{Q}_{k,n}) - \Phi_{k,n}(\bar{Q}_{k,0}) \big\} \ .$$

Now, we consider a first-order expansion of $\Phi_{k,n}$ at $\bar{Q}_{k,n}$:

$$\begin{aligned} &\Phi_{k,n}(\bar{Q}_{k,n}) - \Phi_{k,n}(\bar{Q}_{k,0}) \\ &\quad = -(P_n - P_0) D_{Q_k,0}(Q_0, g^r_0) + B_{Q_k,n}(Q_n, g^r_n) + M_{Q_k}(Q_0, Q_n, g^r_0, g^r_n) + R_{Q_k}(Q_n, Q_0, g^r_0, g^r_n) \ , \end{aligned}$$

where the canonical gradient of $\Phi_{k,n}$ at (Q_0, g^r_0) is given by

$$D_{Q_k,0}(Q_0, g^r_0)(o) := \frac{a_{0:k-1}}{\tilde{g}^r_{k-1,n}(\bar{l}_{k-1})} \ \bar{g}^r_{k-1,n}(\bar{l}_{k-1}) \ \big\{ \bar{Q}_{k+1,0}(\bar{l}_k) - \bar{Q}_{k,0}(\bar{l}_{k-1}) \big\} \ ,$$

the first-order bias term is given by $B_{Q_k,n}(Q_n, g^r_n) := P_n D_{Q_k,n}(Q_n, g^r_n)$, and the empirical process term is given by

$$M_{Q_k}(Q_0, Q_n, g^r_0, g^r_n) = \big(P_n - P_0\big) \ \big\{ D_{Q_k,n}(Q_n, g^r_n) - D_{Q_k,0}(Q_0, g^r_0) \big\} \ ,$$

and the first-order remainder term is given by

$$R^k_{Q_k}(Q_n, Q_0, g^r_0, g^r_n) := P_0 \left\{ \frac{A_{0:k-1}}{\tilde{g}^r_{\overline{k-1},n}} \ \bar{g}^r_{k-1,n} \ \big(\bar{Q}_{k+1,0} - \bar{Q}_{k+1,n}\big) \right\} .$$

Now, we apply an iterative process, where for $m = k, k+1, ..., K-1$ we have

$$\begin{aligned} R^k_{Q_m}(Q_n, Q_0, g^r_0, g^r_n) &= P_0 \left\{ \frac{A_{0:k-1}}{\tilde{g}^r_{\overline{m-1},n}} \ \bar{g}^r_{m-1,n} \ \big(\bar{Q}_{m+1,0} - \bar{Q}_{m+1,n}\big) \right\} \\ &= P_0 \big\{ \bar{g}^R_{m,0n} \ \big(\bar{Q}_{k+1,0} - \bar{Q}_{k+1,n}\big) \big\} \\ &= P_0 \left\{ \frac{A_{0:m}}{\tilde{g}^r_{\bar{m},0n}} \bar{g}^R_{m,0n} \ \big(\bar{Q}_{m+1,0} - \bar{Q}_{m+1,n}\big) \right\} \\ &= P_0 \left\{ \frac{A_{0:m}}{\tilde{g}^r_{\bar{m},n}} \bar{g}^R_{m,n} \ \big(\bar{Q}_{m+1,0} - \bar{Q}_{m+1,n}\big) \right\} \\ &\qquad + S^m_{Q_k}(Q_n, Q_0, g^r_{0n}, g^r_0) + S^m_{Q_k}(Q_n, Q_0, g^r_n, g^r_0) \end{aligned}$$

where we defined the reduced dimension regressions

$$\bar{g}^R_{m,0n}(\bar{L}_m) := E_0\left\{ \frac{A_{0:k-1}}{\tilde{g}^r_{\overline{m-1},n}}\ \bar{g}^r_{m-1,n} \ \middle|\ \bar{Q}_{m+1,0}(\bar{L}_m), \bar{Q}_{m+1,n}(\bar{L}_m) \right\}\ , \text{ and}$$

$$\tilde{g}^r_{\bar{m},0n}(\bar{L}_m) := E_0\big\{ A_{0:m} \ \big|\ \bar{Q}_{m+1,0}(\bar{L}_m), \bar{Q}_{m+1,n}(\bar{L}_m) \big\}\ ,$$

as well as the second-order terms

$$S^m_{Q_k}(Q_n, Q_0, g^r_{0n}, g^r_0) := -P_0\left\{ A_{0:m}\left(\frac{\bar{g}^R_{m,0n}}{\tilde{g}^r_{\bar{m},0n}} - \frac{\bar{g}^R_{m,0}}{\tilde{g}^r_{\bar{m},0}} \right)\ \left(\bar{Q}_{m+1,0} - \bar{Q}_{m+1,n} \right) \right\}$$

$$S^m_{Q_k}(Q_n, Q_0, g^r_{0n}, g^r_0) := P_0\left\{ A_{0:m}\left(\frac{\bar{g}^R_{m,n}}{\tilde{g}^r_{\bar{m},n}} - \frac{\bar{g}^R_{m,0}}{\tilde{g}^r_{\bar{m},0}} \right)\ \left(\bar{Q}_{m+1,0} - \bar{Q}_{m+1,n} \right) \right\}\ ,$$

which are both considered to be $o_P(n^{-1/2})$. Now, defining

$$\Phi^{R,m}_{k,n}(\bar{Q}_{m+1}) := P_0\left(\frac{A_{0:m}}{\tilde{g}^r_{\bar{m},n}} \bar{g}^R_{m,n}\ \bar{Q}_{m+1} \right)\ ,$$

we have shown that

$$R^k_{Q_m}(Q_n, Q_0, g^r_0, g^r_n) = I(\bar{Q}_{m+1,+} = \bar{Q}_{m+1,0})\big\{ \Phi^{R,m}_{k,n}(\bar{Q}_{m+1,n}) - \Phi^{R,m}_{k,n}(\bar{Q}_{m+1,0}) + o_P(n^{-1/2}) \big\}\ .$$

We write a first-order expansion of $\Phi^{R,m}_{k,n}$ at $\bar{Q}_{m+1,n}$:

$$\begin{aligned} &\Phi^{R,m}_{k,n}(\bar{Q}_{m+1,n}) - \Phi^{R,m}_{k,n}(\bar{Q}_{m+1,0}) \\ &\qquad = -(P_n - P_0) D^{R,m}_{Q_k,0}(Q_n, Q_0, g^R_0, g^R_n) + B^{R,m}_{Q_k,n}(Q_n, g^R_n) + \\ &\qquad\quad + M^{R,m}_{Q_k,n}(Q_n, Q_0, g^R_0, g^R_n) + R^{m+1}_{Q_k}(Q_n, Q_0, g^R_0, g^R_n) \end{aligned}$$

where canonical gradient of $\Phi^{R,m}_{k,n}$ at (Q_0, g^R_0) is defined as

$$D^{R,m}_{Q_k,0}(Q_0, g^R_0)(o) = \frac{a_{0:m}}{\tilde{g}^r_{\bar{m},0}(\bar{l}_m)} \bar{g}^R_{m,0}(\bar{l}_m)\ \left\{ \bar{Q}_{m+2,0}(\bar{l}_{m+1}) - \bar{Q}_{m+1,0}(\bar{l}_m) \right\}\ ,$$

the first-order bias term as $B^{R,m}_{Q_k,n}(Q_n, g^R_n) := P_n D^{R,m}_{Q_k,n}(Q_n, g^R_n)$, and the empirical process term (considered to be $o_P(n^{-1/2})$) as

$$M^{R,m}_{Q_k,n}(Q_n, Q_0, g^R_0, g^R_n) := \big\{ P_n - P_0 \big\}\ \big\{ D^{R,m}_{Q_k,n}(Q_n, g^R_n) - D^{R,m}_{Q_k,0}(Q_0, g^R_0) \big\}\ ,$$

and the first-order remainder as

$$R^{m+1}_{Q_k}(Q_n, Q_0, g^R_0, g^R_n) := P_0\left\{ \frac{A_{0:m}}{\tilde{g}^r_{\bar{m},0}} \bar{g}^R_{m,0}\ \left(\bar{Q}_{m+2,0} - \bar{Q}_{m+2,n} \right) \right\}\ .$$

Appendix M

ANALYSIS OF BIAS CORRECTION TERMS UNDER CORRECT SPECIFICATION

M.1 $B_{Q_1,n}$

We begin with the analysis of $B_{Q_1,n} = P_n D_{Q_1,n}(g_n^r, Q_n)$ when $g_+ = g_0$ and $Q_+ \neq Q_0$. We have

$$\begin{aligned} P_n D_{Q_1,n}(g_n^r, Q_n) &= (P_n - P_0) D_{Q_1,n}(g_n^r, Q_n) + P_0 D_{Q_1,n}(g_n^r, Q_n) \\ &= (P_n - P_0) D_{Q_1,n}(g_+^r, Q_+) + P_0 D_{Q_1,n}(g_n^r, Q_n) + M_{Q_1}^c(Q_n, Q_+, g_n^r, g_+^r) \ , \end{aligned} \tag{M.1}$$

where we define the empirical process term

$$M_{Q_1}^c(Q_n, Q_+, g_n^r, g_+^r) := \big(P_n - P_0\big) \ \big\{ D_{Q_1,n}(g_n^r, Q_n) - D_{Q_1,n}(g_+^r, Q_+) \big\} \ ,$$

which is considered to be $o_P(n^{-1/2})$. Now, note that

$$\begin{aligned} \bar{g}_{0,+}^r(l_0) &= E_0 \bigg\{ \frac{A_0 - g_{0,0}(L_0)}{g_{0,0}(L_0)} \ \bigg| \ \bar{Q}_{1,+}(l_0) \bigg\} \\ &= E_0 \bigg[E_0 \bigg\{ \frac{A_0 - g_{0,0}(L_0)}{g_{0,0}(L_0)} \ \bigg| \ L_0 = l_0, \bar{Q}_{1,+}(l_0) \bigg\} \ \bigg| \ \bar{Q}_{1,+}(l_0) \bigg] \\ &= 0 \ . \end{aligned}$$

Thus, the first quantity in (M.1) is equal to 0 and it remains to study

$$P_0 D_{Q_1,n}(g_n^r, Q_n) = P_0 \bigg\{ \frac{A_0}{\tilde{g}_{0,+}^r} \ \bar{g}_{0,n}^r \ \big(\bar{Q}_{2,n} - \bar{Q}_{1,n} \big) \bigg\} \ .$$

We proceed by noting that

$$\begin{aligned} P_0 D_{Q_1,n}(g_n^r, Q_n) = P_0 &\bigg\{ \frac{A_0}{\tilde{g}_{0,n}^r} \ \bar{g}_{0,n}^r \ \big(\bar{Q}_{2,+} - \bar{Q}_{1,+} \big) \bigg\} \\ &+ S_{Q_1,1}^c(Q_n, Q_+, g_n^r) + S_{Q_1,2}^c(Q_n, Q_+, g_n^r) + S_{Q_1,3}^c(Q_n, Q_+, g_n^r) \ , \end{aligned}$$

where

$$
\begin{aligned}
S^c_{Q_1,1}(Q_n, Q_+, g^r_n) &:= P_0\left\{ \frac{A_0}{\tilde{g}^r_{0,n}}\ \bar{g}^r_{1,n}\ \left(\bar{Q}_{2,n} - \bar{Q}_{2,+}\right) \right\} \\
S^c_{Q_1,2}(Q_n, Q_+, g^r_n) &:= -P_0\left\{ \frac{A_0}{\tilde{g}^r_{0,n}}\ \bar{g}^r_{1,n}\ \left(\bar{Q}_{1,n} - \bar{Q}_{1,+}\right) \right\} \\
S^c_{Q_1,3}(Q_n, Q_+, g^r_n) &:= P_0\left\{ A_0\ \left(\bar{Q}_{2,+} - \bar{Q}_{1,+}\right)\ \bar{g}^r_{0,n}\ \left(\frac{\tilde{g}^r_{0,n} - \tilde{g}^r_{0,+}}{\tilde{g}^r_{0,n}\tilde{g}^r_{0,+}} \right) \right\}
\end{aligned}
$$

and these terms are second-order because $\bar{g}^r_{1,n} \to \bar{g}^r_{1,+} = 0$. Thus, we may consider both terms to be $o_P(n^{-1/2})$. We have thus arrived at the representation

$$
P_0 D_{Q_1,n}(g^r_n, Q_n) = \Theta_{1,n}(\bar{g}^r_{0,n}) - \Theta_{1,n}(\bar{g}^r_{0,+}) + o_P(n^{-1/2})\ ,
$$

where we define

$$
\Theta_{1,n}(\bar{g}^r_0) := P_0\left\{ \frac{A_0}{\tilde{g}^r_{0,n}}\ \left(\bar{Q}_{2,+} - \bar{Q}_{1,+}\right)\ \bar{g}^r_0 \right\}\ ,
$$

and note that $\Theta_{1,n}(\bar{g}^r_{0,+}) = 0$. This representation makes clear that the bias correction term originally meant to correct for misspecification of g_0 will behave as $\Theta_{1,n}(\bar{g}^r_{0,n}) - \Theta_{1,n}(\bar{g}^r_{0,+})$, when in fact $g_+ = g_0$. We conclude that $\psi^{*,\diamond}_n$ will still be DR; however, unless $\bar{g}^r_{0,n}$ converges to zero at a parametric rate, $\psi^{*,\diamond}_n$ will not generally be asymptotically linear. Nevertheless the order of the bias is likely to still be reduced relative to ψ^*_n in these cases. The bias of ψ^*_n will be on the order of the bias of Q_n, which is likely to be converging much more slowly than the bias of $\bar{g}^r_{0,n}$, a univariate regression. We conclude by noting that the representation of this term as a data-adaptive parameter $\Theta_{1,n}$ begs the question as to whether $\Theta_{1,n}$ itself may be approximated in first-order by a further reduced-dimension parameter. This is beyond the scope of this dissertation but may be a direction of future research.

M.2 $B^R_{Q_1,n}$

We proceed to the analysis of $B^R_{Q_1,n} = P_n D^R_{Q_1,n}(g^R_n, Q_n)$, the second bias term arising from the putative misspecification of $g_{0,0}$. Consider again the case $g_+ = g_0$ and $Q_+ \neq Q_0$. We have

$$\begin{aligned} P_n D^R_{Q_1,n}(g^R_n, Q_n) &= (P_n - P_0) D^R_{Q_1,n}(g^R_n, Q_n) + P_0 D^R_{Q_1,n}(g^R_n, Q_n) \\ &= (P_n - P_0) D^R_{Q_1,n}(g^R_+, Q_+) + P_0 D^R_{Q_1,n}(g^R_n, Q_n) + M^{R,c}_{Q_1}(Q_n, Q_+, g^R_n, g^R_+) \ , \end{aligned} \tag{M.2}$$

where we define the empirical process term

$$M^{R,c}_{Q_1}(Q_n, Q_+, g^R_n, g^R_+) := \big(P_n - P_0\big) \ \big\{D^R_{Q_1,n}(g^R_n, Q_n) - D_{Q_1,n}(g^R_+, Q_+)\big\} \ ,$$

which is considered to be $o_P(n^{-1/2})$. Now, we recall that

$$\bar{g}^r_{1,+}(\bar{l}_1) := E_0\left\{ \frac{A_0}{\tilde{g}^r_{0,+}} \bar{g}^r_{0,+}(L_0) \ \middle| \ \bar{Q}_{2,+}(\bar{l}_1) \right\} \ ,$$

and that we argued in the previous section that $\bar{g}^r_{0,+}(L_0) = 0$, so we also have that $\bar{g}^r_{1,+}(\bar{L}_1) = 0$. Thus, the first term in (M.2) will be zero. It remains then to examine $P_0 D^R_{Q_1,n}(g^R_n, Q_n)$, which we accomplish in the following:

$$P_0 D^R_{Q_1,n}(g^R_n, Q_n) = P_0\left\{ \frac{A_0 A_1}{\tilde{g}^r_{\bar{1},+}} \ \big(L_2 - \bar{Q}_{2,+}\big) \ \bar{g}^R_{0,n} \right\} + S^R_{Q_1,1}(Q_n, Q_+, \bar{g}^R_{0,n}) + S^R_{Q_1,2}(Q_n, Q_+, \bar{g}^R_{0,n}),$$

where we define the second-order terms

$$\begin{aligned} S^R_{Q_1,1}(Q_n, Q_+, \bar{g}^R_{0,n}) &:= -P_0\left\{ \frac{A_0 A_1}{\tilde{g}^r_{\bar{1},n}} \ \big(\bar{Q}_{2,n} - \bar{Q}_{2,+}\big) \ \bar{g}^R_{0,n} \right\} \text{ and} \\ S^R_{Q_1,2}(Q_n, Q_+, \bar{g}^R_{0,n}) &:= P_0\left\{ A_0 A_1 \ \big(L_2 - \bar{Q}_{2,+}\big) \ \bar{g}^R_{0,n} \ \left(\frac{\tilde{g}^r_{\bar{1},n} - \tilde{g}^r_{\bar{1},+}}{\tilde{g}^r_{\bar{1},n} \tilde{g}^r_{\bar{1},+}} \right) \right\} \end{aligned}$$

which are considered to be $o_P(n^{-1/2})$. Thus, we have the representation that

$$B^R_{Q_1,n} = \Theta^R_{1,n}(\bar{g}^R_{0,n}) - \Theta^R_{1,n}(\bar{g}^R_{0,+}) + o_P(n^{-1/2}) \ ,$$

and all the notes made in the previous section with regard to the behavior or $\Theta_{1,n}(\bar{g}^r_{0,n})$ apply here as well.

M.3 $B_{Q_2,n}$

We now move to the analysis of $B_{Q_2,n} := P_n D_{Q_2,n}(Q_n, g_n^r)$, the term correcting for the possible misspecification of $\bar{g}_{1,0}$ when in fact $g_+ = g_0$ and $Q_+ \neq Q_0$. We have that

$$\begin{aligned} P_n D_{Q_2,n}(Q_n, g_n^r) &= (P_n - P_0) D_{Q_2,n}(Q_n, g_n^r) + P_0 D_{Q_2,n}(Q_n, g_n^r) \\ &= (P_n - P_0) D_{Q_1,+}(Q_+, g_+^r) + P_0 D_{Q_2,n}(Q_n, g_n^r) + M_{Q_2}^c(Q_n, Q_+, g_n^r, g_+^r) \ , \end{aligned} \tag{M.3}$$

where we define the empirical process term

$$M_{Q_2}^c(Q_n, Q_+, g_n^r, g_+^r) := \big(P_n - P_0\big) \ \big\{ D_{Q_2,n}(Q_n, g_n^r) - D_{Q_1,+}(Q_+, g_+^r) \big\} \ ,$$

which is considered to be $o_P(n^{-1/2})$. Now, note that

$$\begin{aligned} \bar{g}_{1,+}^r(\bar{L}_1) &:= E_0\left[\frac{A_0}{\bar{g}_{0,0}(L_0)} \left\{ \frac{A_1 - \bar{g}_{1,0}(\bar{L}_1)}{\bar{g}_{1,0}(\bar{L}_1)} \right\} \ \middle| \ \bar{Q}_{2,+}(\bar{L}_1) \right] \\ &= E_0\left(E_0\left[\frac{A_0}{\bar{g}_{0,0}(L_0)} \left\{ \frac{A_1 - \bar{g}_{1,0}(\bar{L}_1)}{\bar{g}_{1,0}(\bar{L}_1)} \right\} \ \middle| \ \bar{Q}_{2,+}(\bar{L}_1), A_0, L_0, \bar{g}_{1,0}(\bar{L}_1) \right] \ \middle| \ \bar{Q}_{2,+}(\bar{L}_1) \right) \\ &= E_0\left(E_0\left[\frac{A_0}{\bar{g}_{0,0}(L_0)} \left\{ \frac{A_1 - \bar{g}_{1,0}(\bar{L}_1)}{\bar{g}_{1,0}(\bar{L}_1)} \right\} \ \middle| \ \bar{Q}_{2,+}(\bar{L}_1), A_0 = 1, L_0, \bar{g}_{1,0}(\bar{L}_1) \right] \ \middle| \ \bar{Q}_{2,+}(\bar{L}_1) \right) \\ &= 0 \ . \end{aligned}$$

Thus, the first quantity in (M.3) will be zero and it remains to study $P_0 D_{Q_2,n}(Q_n, g_n^r)$. Note that

$$\begin{aligned} P_0 D_{Q_2,n}(Q_n, g_n^r) &= P_0\left\{ \frac{A_0 A_1}{\tilde{g}_{\bar{1},0}} \ \bar{g}_{1,n}^r \ \big(L_2 - \bar{Q}_{2,n}\big) \right\} \\ &= P_0\left\{ \frac{A_0 A_1}{\tilde{g}_{\bar{1},0}} \ \bar{g}_{1,n}^r \ \big(L_2 - \bar{Q}_{2,+}\big) \right\} + S_{Q_2,1}(Q_n, Q_0, g_n^r) + S_{Q_2,2}(Q_n, Q_0, g_n^r) \ , \end{aligned}$$

where

$$\begin{aligned} S_{Q_2,1}(Q_n, Q_0, g_n^r) &:= -P_0\left\{ \frac{A_0 A_1}{\tilde{g}_{\bar{1},0}} \ \bar{g}_{1,n}^r \ \big(\bar{Q}_{2,n} - \bar{Q}_{2,+}\big) \right\} \text{ and} \\ S_{Q_2,2}(Q_n, Q_0, g_n^r) &:= -P_0\left\{ A_0 A_1 \ \big(L_2 - \bar{Q}_{2,+}\big) \ \bar{g}_{1,n}^r \ \left(\frac{\tilde{g}_{\bar{1},n}^r - \tilde{g}_{\bar{1},+}^r}{\tilde{g}_{\bar{1},n}^r \tilde{g}_{\bar{1},+}^r} \right) \right\} \end{aligned}$$

are considered to be $o_P(n^{-1/2})$. Thus, we have the representation that

$$B_{Q_2,n} = \Theta_{2,n}(\bar{g}^r_{1,n}) - \Theta_{2,n}(\bar{g}^R_{0,+}) + o_P(n^{-1/2}) \ ,$$

and we may expect, as previously, that $B_{Q_2,n}$ will converge to zero in probability, but will not be asymptotically linear unless $\bar{g}^r_{1,n}$ falls in a parametric model and is estimated via maximum likelihood.

M.4 $B_{g_1,n}$

We continue with the analysis of $B_{g_1,n} := P_n D_{g_1,n}(Q^r_n, g_n)$, the correction term correcting for the putative misspecification of $\bar{Q}_{2,0}$ when in fact $Q_+ = Q_0$ and $g_+ \neq g_0$. Proceeding in a similar manner to the previous two sections, we have

$$\begin{aligned} P_n D_{g_1,n}(Q^r_n, g_n) &= (P_n - P_0) D_{g_1,n}(Q^r_n, g_n) + P_0 D_{g_1,n}(Q^r_n, g_n) \\ &= (P_n - P_0) D_{g_1,+}(Q^r_+, g_+) + P_0 D_{g_1,n}(Q^r_n, g_n) + M^c_{g_1}(Q^r_n, Q^r_+, g_n, g_+) \ , \end{aligned} \tag{M.4}$$

where we define the empirical process term

$$M^c_{g_1}(Q^r_n, Q^r_+, g_n, g_+) := \big(P_n - P_0\big) \left\{ D_{g_1,n}(Q^r_n, g_n) - D_{g_1,+}(Q^r_+, g_+) \right\} \ ,$$

which is considered to be $o_P(n^{-1/2})$. Now we note that

$$\begin{aligned} \bar{Q}^r_{2,+}(\bar{l}_1, a_0) &= E_0\left[\frac{A_0 A_1}{\bar{g}_{0,+}} \left\{ L_2 - \bar{Q}_{2,+}(\bar{L}_1) \right\} \middle| a_0, \bar{g}_{1,+}(\bar{l}_1) \right] \\ &= E_0\Bigg(E_0\left[\frac{A_0 A_1}{\bar{g}_{0,+}(L_0)} \left\{ L_2 - \bar{Q}_{2,0}(\bar{L}_1) \right\} \middle| \bar{A}_1, L_0 = l_0, \bar{Q}_{2,0}(\bar{l}_1), \bar{g}_{1,+}(\bar{l}_1) \right] \Bigg| \ A_0 = a_0, \bar{g}_{1,+}(\bar{l}_1) \Bigg) \\ &= E_0\Bigg(E_0\left[\frac{A_0 A_1}{\bar{g}_{0,+}(L_0)} \left\{ L_2 - \bar{Q}_{2,0}(\bar{L}_1) \right\} \middle| \bar{A}_1 = \bar{1}, L_0, \bar{Q}_{2,0}(\bar{l}_1), \bar{g}_{1,+}(\bar{l}_1) \right] \Bigg| \ A_0 = a_0, \bar{g}_{1,+}(\bar{l}_1) \Bigg) \\ &= 0 \ , \end{aligned}$$

where the final line follows from the fact that

$$E_0\big\{ L_2 \ \big| \ \bar{A}_1 = \bar{1}, \bar{Q}_{2,0}(\bar{l}_1), L_0 = l_0, \bar{g}_{1,+}(\bar{l}_1) \big\} = \bar{Q}_{2,0}(\bar{l}_1) \ .$$

This implies that the first quantity of (M.4) will be 0. Now, note that

$$\begin{aligned} P_0 D_{g_1,n}(Q_n^r, g_n) &= P_0 \left\{ \frac{\bar{Q}_{2,n}^r}{\bar{g}_{1,n}^2} \, \left(A_1 - \bar{g}_{1,n} \right) \right\} \\ &= P_0 \left\{ \frac{\bar{Q}_{2,n}^r}{\bar{g}_{1,+}^2} \, \left(A_1 - \bar{g}_{1,+} \right) \right\} + S_{g_1,1}^c(Q_n^r, Q_+^r, g_n, g_+) + S_{g_1,2}^c(Q_n^r, Q_+^r, g_n, g_+) \ , \end{aligned}$$

where the second-order terms

$$\begin{aligned} S_{g_1,1}^c(Q_n^r, Q_+^r, g_n, g_+) &:= -P_0 \left\{ \frac{\bar{Q}_{2,n}^r}{\bar{g}_{1,n}^2} \, \left(\bar{g}_{1,n} - \bar{g}_{1,+} \right) \right\} \text{ and} \\ S_{g_1,1}^c(Q_n^r, Q_+^r, g_n, g_+) &:= -P_0 \left\{ \bar{Q}_{2,n}^r \, \left(A_1 - \bar{g}_{1,+} \right) \, \left(\frac{\bar{g}_{1,n}^2 - \bar{g}_{1,+}^2}{\bar{g}_{1,n} \bar{g}_{1,+}} \right) \right\} \end{aligned}$$

are considered to be $o_P(n^{-1/2})$. Thus, we have the representation

$$B_{g_1,n} := \eta_{1,n}(\bar{Q}_{1,n}^r) - \eta_{1,n}(\bar{Q}_{1,+}^r) + o_P(n^{-1/2}) \ ,$$

where we defined

$$\eta_{1,n}(\bar{Q}_1^r) := P_0 \left\{ \left(\frac{A_1 - \bar{g}_{1,+}}{\bar{g}_{1,+}} \right) \, \bar{Q}_1^r \right\} \ .$$

Again, the comments of the previous two sections apply to this term as well.

M.5 $B_{g_0,n}$

We saved the most interesting term for last, $B_{g_0,n} := P_n D_{g_0,n}(Q_n^r, g_n)$. This is the correction term for the putative misspecification of $\bar{Q}_{1,0}$ when in fact $Q_+ = Q_0$ and $g_+ \neq g_0$. We begin as previously:

$$\begin{aligned} P_n D_{g_0,n}(Q_n^r, g_n) &= (P_n - P_0) D_{g_0,n}(Q_n^r, g_n) + P_0 D_{g_0,n}(Q_n^r, g_n) \\ &= (P_n - P_0) D_{g_0,+}(Q_+^r, g_+) + P_0 D_{g_0,n}(Q_n^r, g_n) + M_{g_0}^c(Q_n^r, Q_+^r, g_n, g_+) \ , \end{aligned}$$

where the empirical process term

$$M_{g_0}^c(Q_n^r, Q_+^r, g_n, g_+) := \left(P_n - P_0 \right) \, \left\{ D_{g_0,n}(Q_n^r, g_n) - D_{g_0,+}(Q_+^r, g_+) \right\} \ ,$$

is considered to be $o_P(n^{-1/2})$. Now, we note that, contrary to the previous reduced-dimension terms,

$$\begin{aligned}\bar{Q}^r_{1,+}(l_0) &:= E_0\left[\frac{A_0A_1}{\bar{g}_{1,+}(\bar{L}_1)}\ \left\{L_2-\bar{Q}_{1,0}(L_0)\right\}\ \middle|\ \bar{g}_{0,+}(l_0)\right]\\ &= E_0\left(E_0\left[\frac{A_0A_1}{\bar{g}_{1,+}(\bar{L}_1)}\ \left\{L_2-\bar{Q}_{1,0}(L_0)\right\}\ \middle|\ \bar{g}_{0,+}(l_0),\bar{L}_1,\bar{A}_1\right]\ \middle|\ \bar{g}_{0,+}(l_0)\right)\\ &= E_0\left[\frac{A_0\bar{g}_{1,0}(\bar{L}_1)}{\bar{g}_{1,+}(\bar{L}_1)}\ \left\{\bar{Q}_{2,0}(\bar{L}_1)-\bar{Q}_{1,0}(L_0)\right\}\ \middle|\ \bar{g}_{0,+}(L_0)\right]\ ,\end{aligned}$$

which does not necessarily equal zero. It thus appears that $D_{g_0,n}$, the canonical gradient of $\Gamma_{0,n}$, is not doubly-robust; it will have mean zero if either (1) $Q_+ = Q_0$ and $g_+ = g_0$ or (2) $Q_+ \neq Q_0$ and $g_+ = g_0$, but not when $Q_+ = Q_0$ and $g_+ \neq g_0$. Thus, the estimator $\psi_n^{*,\diamond}$ will be locally efficient but not doubly-robust in the classical sense.

Appendix N

ADDITIONAL SIMULATION RESULTS

In this section we include figures examining the behavior of the bias correction terms discussed in the previous section. In particular, we examine $B_{Q_2,n}$ and $B_{g_0,n}$. Recall that in the previous section, we argued that the former would converge to zero, but not at a parametric rate, owing to the fact that $\bar{g}^r_{1,n}$ was estimated nonparametrically. This is clearly illustrated in Figures N.1 and N.2; the former illustrating the convergence in probability to zero, the latter that the term is not $o_P(n^{-1/2})$. Now consider Figures N.3 and N.4, which show the same figures for $B_{g_0,n}$. Recall that we argued in the previous section that this term would not be converging to zero at all, let alone at a parametric rate due to its lack of double-robustness. This too is borne out by the figures.

Finally, Figure N.5 shows the performance of the variance estimators of each of the estimators proposed. On the left side, we find that the efficient influence curve-based variance estimators for the standard one-step estimator are hugely anti-conservative, which, along with the bias, helps explain the poor coverage of the uncorrected one-step estimator seen in Figure 3.5.2. Conversely, on the right side, we find that the uncorrected variance estimators yield conservative variance estimators, which helps explain why the coverage was found to be approximately nominal in this case in spite of the excess bias. The standard error estimators for the corrected one-steps perform reasonably well in both settings. That the variance estimator for the estimator correcting for both potential misspecifications is found to behave reasonably well may indicate that the contribution to the influence function due to the unnecessary correction term may be negligible at these sample sizes.

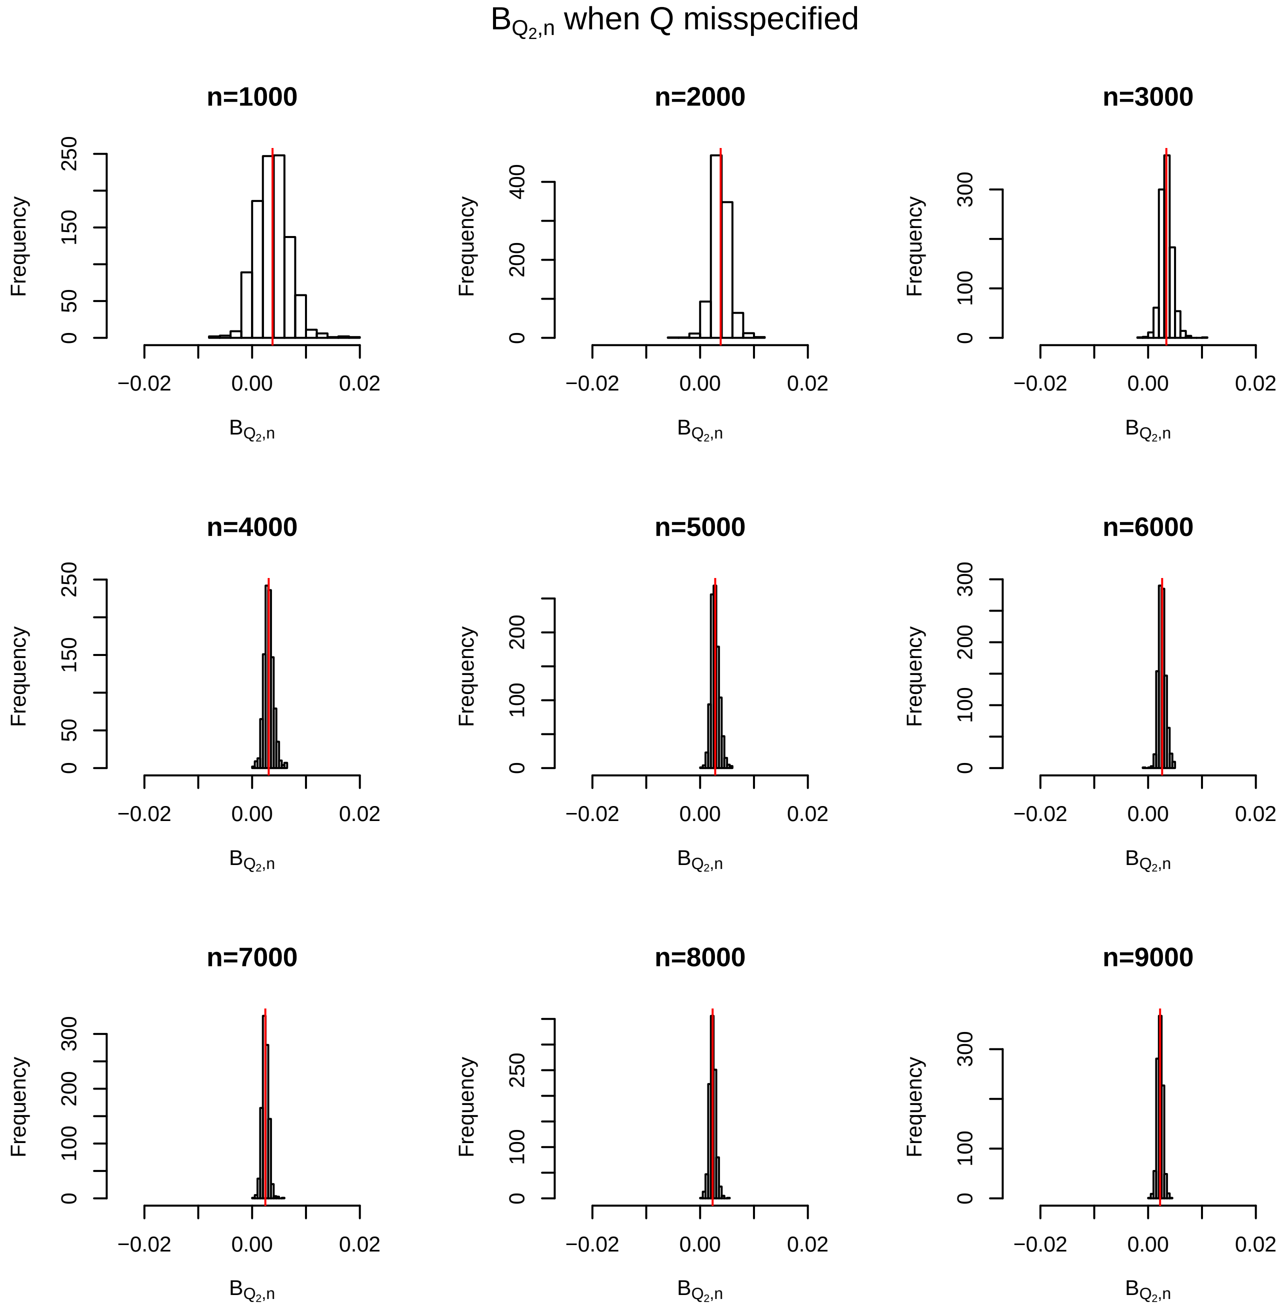

Figure N.1: A histogram of the sampling distribution of $B_{Q_2,n}$ in Simulation 1 when Q is misspecified, and thus $B_{Q_2,n}$ is unneeded. The red vertical line represents the mean over all simulations.

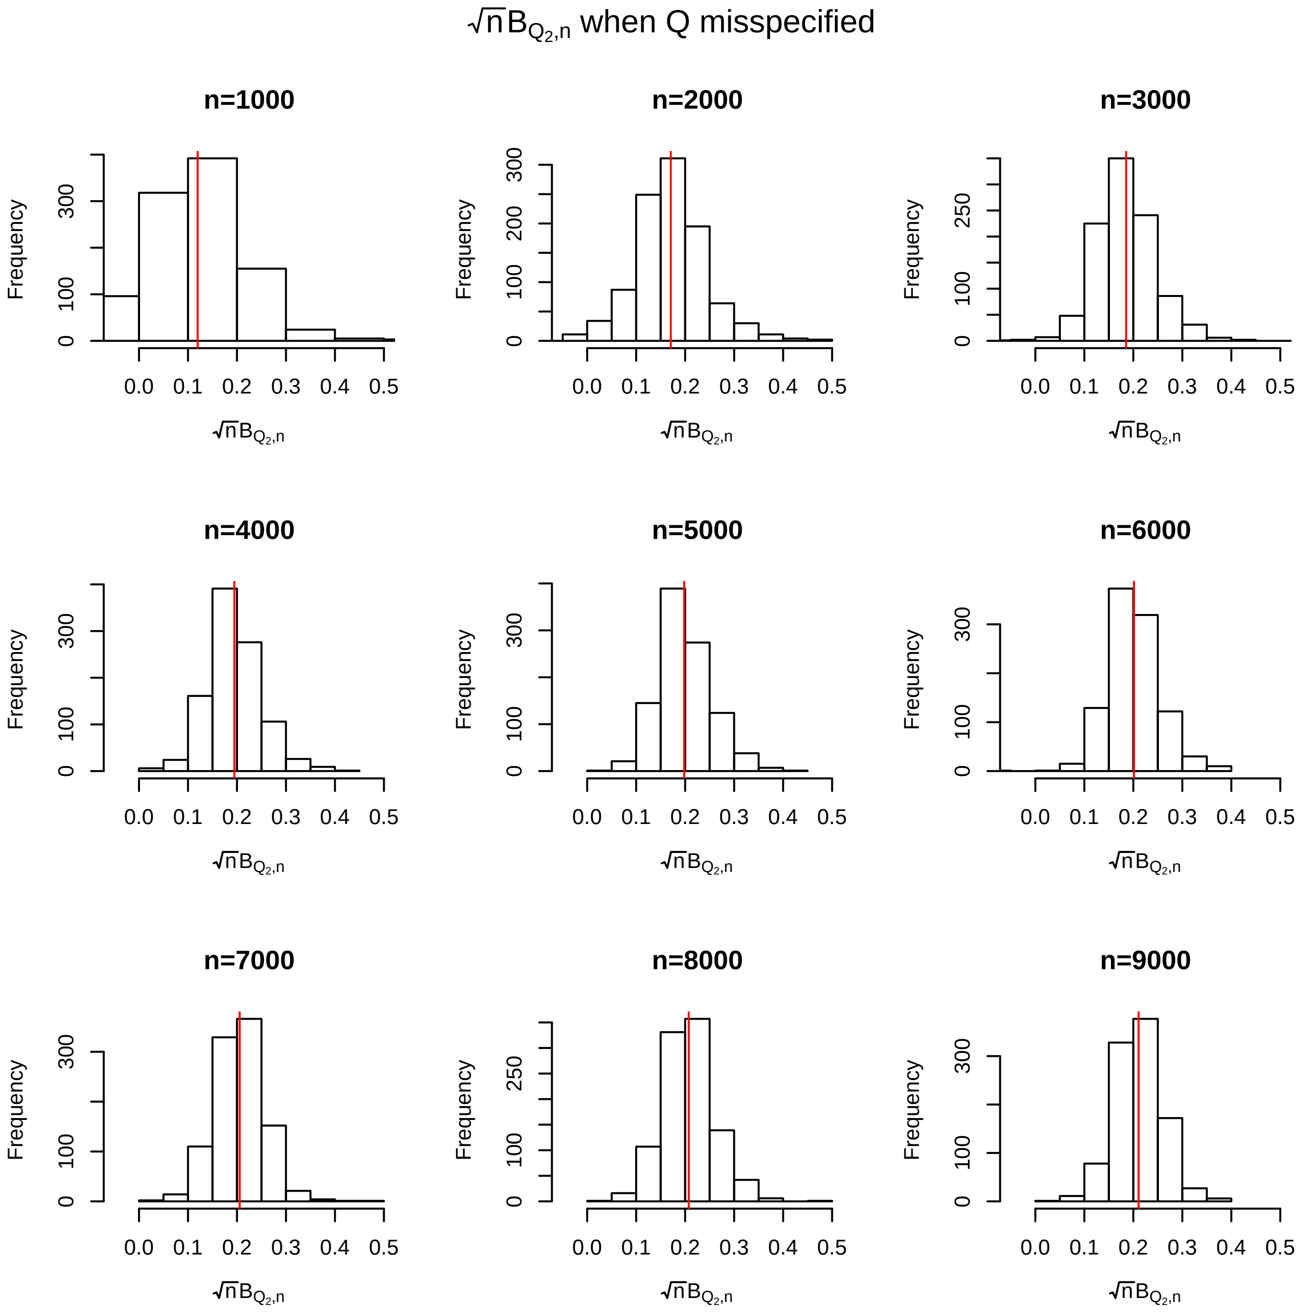

Figure N.2: A histogram of the sampling distribution of $\sqrt{n}B_{Q_2,n}$ in Simulation 1 when Q is misspecified, and thus $B_{Q_2,n}$ is unneeded. The red vertical line represents the mean over all simulations.

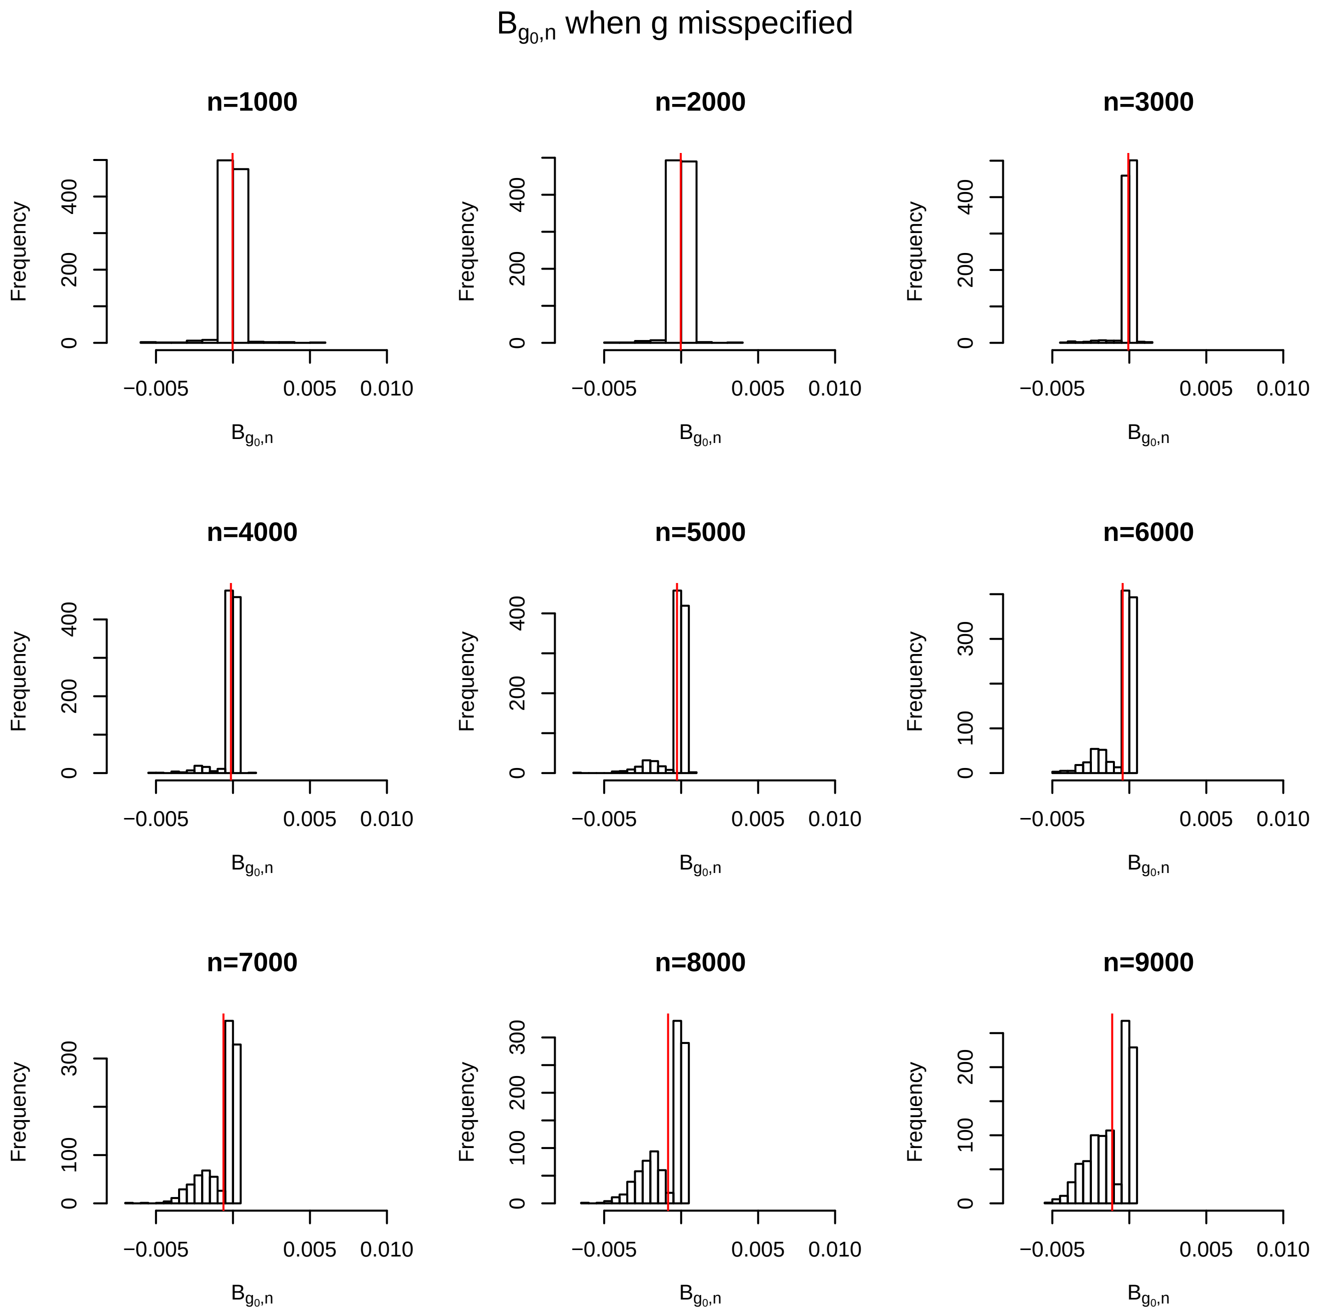

Figure N.3: A histogram of the sampling distribution of $B_{g_0,n}$ in Simulation 1 when g is misspecified, and thus $B_{g_0,n}$ is unneeded. The red vertical line represents the mean over all simulations.

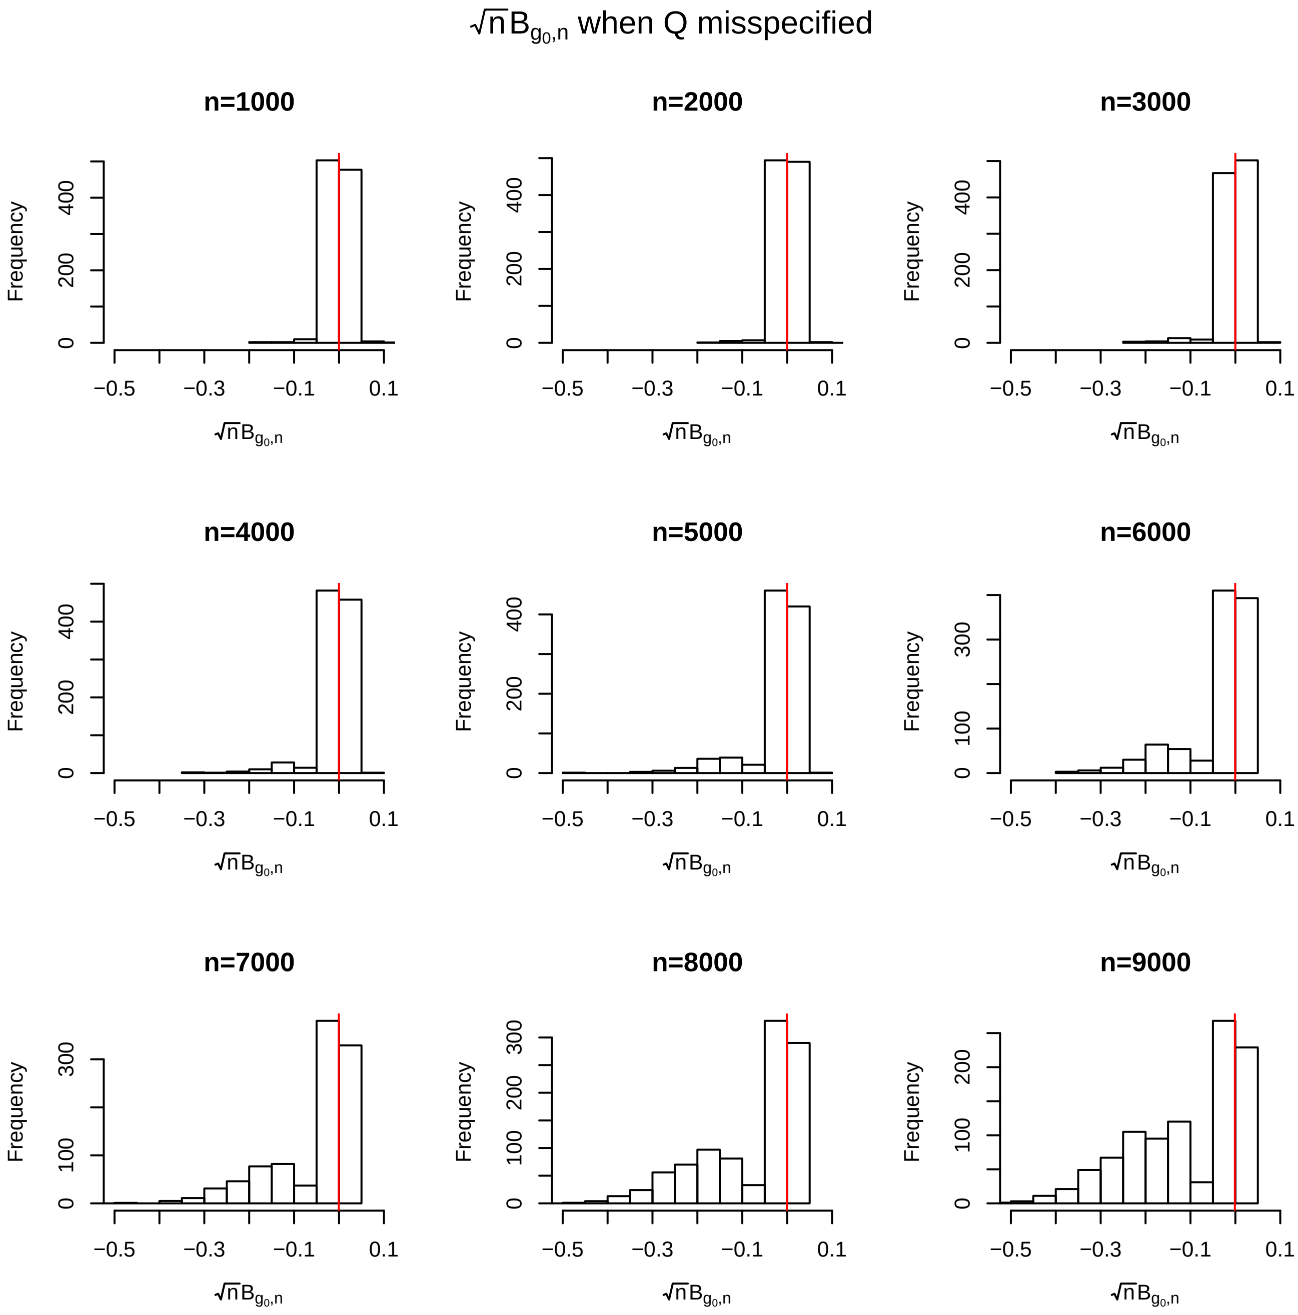

Figure N.4: A histogram of the sampling distribution of $\sqrt{n}B_{g_0,n}$ in Simulation 1 when g is misspecified, and thus $B_{g_0,n}$ is unneeded. The red vertical line represents the mean over all simulations.

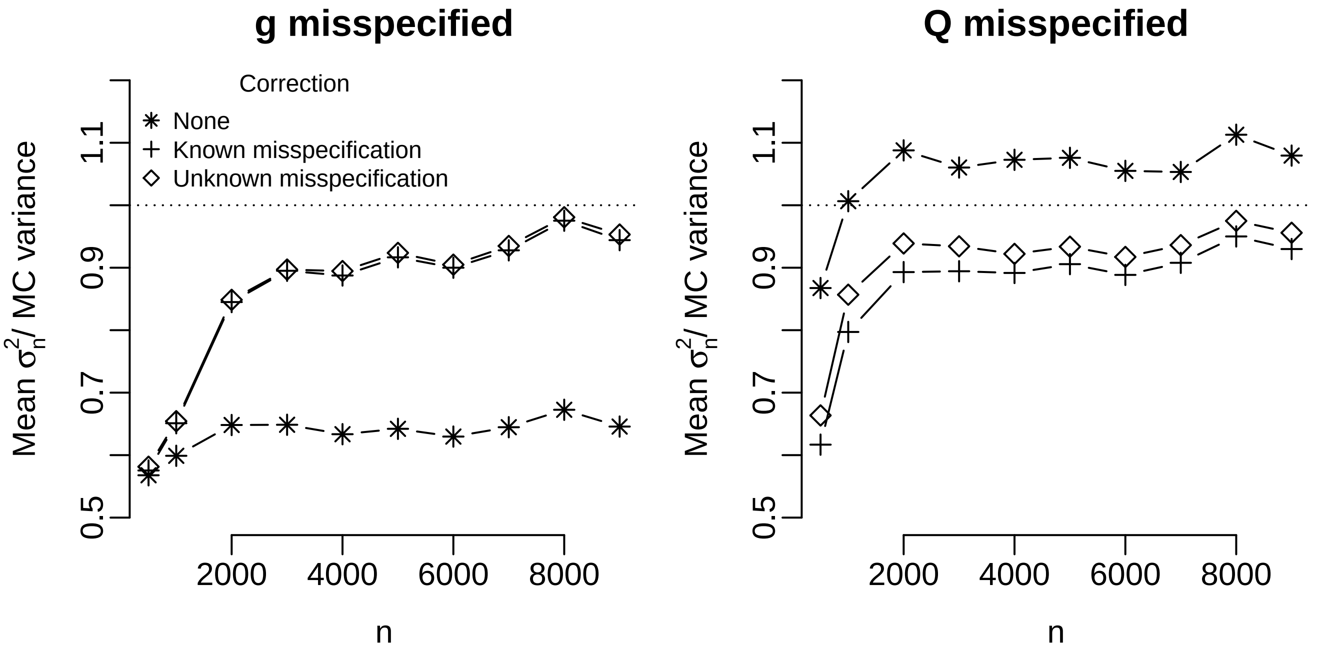

Figure N.5: The performance of influence curve based variance estimators. The ratio of the average variance estimate compared to the Monte Carlo variance of the estimators is shown. Values below one indicate an anti-conservative estimate, while values above one indicate a conservative estimate.