Comparative Causal Effect Estimation and Robust Variance for Longitudinal Data Structures with Applications to Observational HIV Treatment

by

Linh Mai Tran

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Biostatistics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Maya Petersen, Chair
Professor Mark van der Laan
Professor John Colford

Spring 2016

Comparative Causal Effect Estimation and Robust Variance for Longitudinal Data Structures with Applications to Observational HIV Treatment

Abstract

Comparative Causal Effect Estimation and Robust Variance for Longitudinal Data Structures with Applications to Observational HIV Treatment

by

Linh Mai Tran

Doctor of Philosophy in Biostatistics

University of California, Berkeley

Professor Maya Petersen, Chair

This dissertation discusses the application and comparative performance of double robust estimators for estimating the intervention specific mean outcome in longitudinal settings with time-dependent confounding as well as the corresponding estimator variances. Specifically, we focus on carefully defining target causal parameters to avoid known positivity issues, estimating these parameters using the asymptotically efficient and double robust targeted minimum loss-based estimation, comparing this to other double robust estimators of the same causal parameter, and estimating the corresponding variances in a way which demonstrates valid Type I errors while retaining statistical power. Chapter 1 begins by introducing the open problem in statistics. We present the International epidemiologic Databases to Evaluate AIDS, East Africa region (IeDEA-EA) cohort and the implementation of a low risk express care program implemented between $2007-2009$. We continue in Chapter 2 by presenting the targeted learning road map for causal inference. This road map is applied, as a case study, to the IeDEA-EA cohort in evaluating the impact of the low risk express care program. Targeted minimum loss-based estimation is used to estimate the intervention specific mean outcome using data adaptive machine learning candidate estimators for the nuisance parameters. Practical issues are addressed, including carefully defining the causal parameters (and the corresponding causal contrasts) and remaining within the boundaries implied by the statistical model while using the machine learning algorithms. In Chapter 3, we compare additional estimators for the intervention specific mean outcome. The iterated conditional expectation estimator, inverse probability weighted estimator, augmented inverse probability weighted estimator, double robust iterated conditional expectation estimator, and targeted minimum loss-based estimator are presented. Additionally, variations on the double robust iterated conditional expectation estimator and targeted minimum loss-based estimator are reviewed and implemented. Simulations are conducted to analyze the finite sample performance of each estimator, in both correct and mis-specified models. The estimators are also applied to estimating the impact of enrollment into the low risk express care program in the IeDEA-EA cohort. Chapter 4 studies the estimation of estimator variance for estimators solving the efficient influence function. A robust approach of estimating the efficient influence function variance is presented, followed by approaches for estimating the derived expec-

tation of the variance. This robust approach of estimating the EIF variance can be used to raise a red flag for unreliable statistical inference due to sparsity, thereby declaring that the target parameter is practically not identifiable from the data, and that the reported variance estimates (though large) will themselves be imprecise. We additionally present a bootstrap approach based on fitting the initial density of the data once, followed by a non-parametric bootstrap of the targeting step. This bootstrap approach can be used to estimate the variance of substitution based estimators solving the efficient influence function. Simulations are conducted, demonstrating the bias, variance, coverage, and statistical power resulting from each of the variance estimators. Standard errors and confidence intervals are calculated using each of the variance estimators in estimating the impact of enrollment into the low risk express care program in the IeDEA-EA cohort. The primary appendices present relevant proofs for the analyses conducted in this dissertation, while R code for implementing the various estimators are provided in secondary appendices.

To my wife. Forever and ever.

Contents

List of Figures

List of Tables

Acknowledgments

In hindsight, so many things have made this very challenging journey so memorable and worth all of the vested effort. From the amazing professors and colleagues at U.C. Berkeley, to the fascinating and relevant work being conducted, to family, friends, and mentors, I owe so much of what I have done to everything and everyone that surrounds me. I want to sincerely thank the International epidemiologic Databases to Evaluate AIDS, East Africa team for providing the support, data, and research questions that motivated much of the work presented here. To the U.C. Berkeley Biostatistics Group, thank you for putting together such an amazing graduate program. To my advisor Maya Petersen, thank you for your guidance, patience, and endearment. Your lessons on causal inference will be forever ingrained in me. To my professor Mark van der Laan, thank you for your leadership and rigor. You have shown me the true brilliance to continually strive towards. To Sharon Norris, thank you for coping with me through all of the paperwork and deadlines. Your organization has always amazed me and aid has always been appreciated. And lastly, to my wife Alyson, for all of the love and support she has continually provided while I chased my dreams. I really am the luckiest man in the world.

Chapter 1

Introduction

1.1 The open problem

Obtaining a causal effect from observed data is currently one of the greatest open problems of the modern era. In the vast majority of research studies where the researcher is tasked with analyzing an exposure and outcome, the primary interest is usually in estimating a causal effect. That is, how would some parameter of the outcome distribution differ under alternative exposure assignments? For example, how would the probability of death by some time point differ if all subjects versus none had received the exposure? While the exponential growth of computing power has lead to faster and easier data collection and analysis, consequently providing droves of potentially powerful insights, the analysis and resulting inference still varies widely from one analyst to the next. This discrepancy can lead to very different decisions involving potentially large amounts of resources from consumers of the results. It follows that much rigor and care needs to be placed into such analyses so that significant differences (if any) in the results are not due to errors such as incorrect assumptions or inaccurate formulations of the corresponding statistical problem.

A statistical model $\mathscr{M}$ is formally defined as a set of potential probability distributions. Conventionally, the statistical model provides a framework in which parameters can be estimated and inferences drawn from the observational data used in the model. For example, with classical statistics, an analyst may choose a parametric linear regression model (Figure 1.1) with p parameters for the data observed. Once the parameters are estimated from the data, the parameter that corresponds to the treatment variable of interest is interpreted as the treatment effect. In scenarios where these models contain the truth, then efficient estimates can be obtained with standard maximum likelihood estimation software. In practice, however, these simple models rarely (if ever) capture the true relationship of the covariates from the data. An obvious solution, then, is to use larger semi-parametric or non-parametric models. To aid in doing so a systematic road map, backed by theory, could be employed providing a robust and reproducible approach of establishing inference.

van der Laan and Rose [49], Petersen and van der Laan [66], Pearl [61], and others have advocated for the use of such a systematic road map for translating causal questions into statistical analyses and interpreting the results. This road map requires the analyst to learn as much as possi-

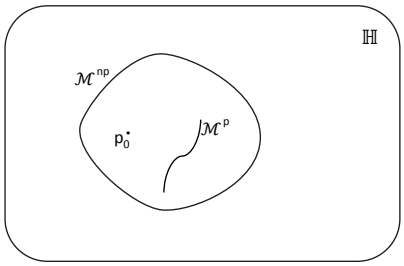

Figure 1.1: The true probability distribution p_0 along with parametric $\mathscr{M}^p$ and non-parametric $\mathscr{M}^{np}$ statistical models, residing within the Hilbert space of probability functions $p \in \mathscr{P}$.

ble about how the data were generated, use the information to posit a realistic statistical model that encompasses these findings, and assign a corresponding estimand that can answer the scientific question of interest. While in some cases this approach can be straightforward, in practice it generally requires much consideration when implementing. This is especially true in observational data, where a common objective is to estimate the joint effect of one or more longitudinal exposures, or series of sequential treatment decisions (for example, Bodnar [4], Bryan [7], and Petersen [63]). For example, one may be interested in contrasting immediate enrollment into a particular program (a single treatment decision) with delayed enrollment (a series of treatment decisions, in which enrollment is sequentially deferred at multiple time points at which it could have been initiated).

In this introduction, we provide an overview of the International epidemiologic Databases to Evaluate AIDS, East Africa region (IeDEA-EA) cohort used in this dissertation. The low risk express care program is presented and detailed. The cohort is described and formalized as n copies of $O \overset{iid}{\sim} P_0$. We review the likelihood and specify the notation for this observed data.

Chapter 2 reviews the framework for causal inference as applied to our setting [49, 62]. This includes, though is not limited to (1) stating the scientific question, (2) specifying the causal model $\mathscr{M}^{\mathscr{F}}$, (3) specifying the causal parameter $\Psi^{\mathscr{F}}(P_d^0)$, and (4) assessing identifiability of our causal parameter as some function of the observed data distribution P_0. We provide a detailed description of this analysis, with an emphasis on several practical challenges likely to arise in similar applications.

Chapter 3 considers additional estimators for the same target parameter as Chapter 2 (i.e. the intervention specific mean outcome). Each estimator considered is described in detail and practical implementations are discussed. The decision between conditioning on treatment and pooling across different treatment regimes is considered. We conduct a simulation study to analyze finite sample performance in the estimators, in addition to applying each estimator to our IeDEA-EA cohort.

Chapter 4 presents on two novel methods of estimating the variance of estimators solving the equations corresponding to their influence functions, particularly in the setting of practical pos-

itivity violations. A robust estimation approach that directly targets the variance of the efficient influence is first presented, along with estimators of this parameter. A bootstrap approach is also presented, which is shown to perform well for substitution based estimators such as targeted minimum loss-based estimation. Simulation studies are conducted using these variance estimators on both a single time-point static treatment setting and longitudinal setting with time-dependent confounding. We additionally apply the variance estimation approaches to the IeDEA-EA data, demonstrating the variations in confidence interval lengths from each approach.

Mathematical proofs and software are reserved for the appendices. Appendix A presents a proof demonstrating the equivalence of influence functions between the full data O and the reduced data $O_r = O/L1(K+1)$ used in the Chapter 2 applied analyses. Appendix B contains proofs regarding the robust approach of estimating the efficient influence function variance presented in Chapter 4, generalized to working marginal structural models. Appendix C contains code used to compare the different estimators in Chapter 3 and Chapter 4.

1.2 The IeDEA-EA cohort

The majority of individuals with HIV live in settings, such as East Africa, where noticeable resource and infrastructure constraints place a limitation on the care these patients can receive [100, 115]. Antiretroviral therapy (ART) medication has been shown to reduce both viral loads and mortality in these patients [23, 59, 119, 15, 12, 54], as well as reduce rates of transmission to persons uninfected [68, 2, 13, 9]. However, the shortage of resources and health care professionals limit the number of patients who can be placed on ART [110, 40, 116, 108].

Due to these limitations, various approaches have been undertaken in an effort to ensure that the maximal number of patients who need care can receive it. One such approach shifts care provision tasks for patients considered to be at low risk from physicians and clinical officers to other care professionals, such as nurses. Consequently, the workload for the physicians and clinical officers is reduced, thereby theoretically increasing the attention that higher risk patients can receive for HIV or other conditions [115].

One such program was implemented between 2007 and 2009 among clinics around eastern Africa. These clinics were followed as part of the Academic Model Providing Access to Healthcare (AMPATH) program and contributed data to the International epidemiologic Databases to Evaluate AIDS, East Africa region (IeDEA-EA). The purpose of this Low Risk Express Care (LREC) program is to shift care tasks for patients considered to be at low risk from physician-centered care models to those utilizing non-physician health workers trained in simplified and standardized approaches to care. Patients were considered to be at low risk if they met the following set of criteria.

1. They were older than 18 years of age.
2. They were stable on ART for at least 6-months.
3. They had no AIDS-defining or AIDS-associated events within the past 180-days.
4. During past 6 months, they reported no missed pills when asked about the 7 days prior to each visit.

5. They were not pregnant within the past 30 days.
6. Their most recent CD4 count$>$ 200 cells/μL within 274 days prior to the current visit.

Once eligible, at each clinical visit, the clinician decided whether or not to enroll the patient into the LREC program. Patients enrolled had part of their care, such as identifying and managing ART side effects and opportunistic infections, shifted to nurses. Table 1.1 shows the differences between the standard of care and the LREC model.

Table 1.1: Proportion of visits responsibility assigned between the (a) standard and (b) LREC models of care provision. (P = Physician, CO = Clinical Officer)

Clinical monitoring	**Standard model**		**LREC model**	
	P / CO	**Nurse**	**P / CO**	**Nurse**
Request CD4 / VL counts	All	None	All	None
Monitor / support ART adherence	All	All	1/3	All
Determine functional status	All	None	1/3	2/3
Identify / manage ART side effects	All	None	1/3	2/3
Identify / manage opportunistic infections	All	None	1/3	2/3

Our study population is comprised of subjects found eligible for the LREC program within each of 15 clinics in Kenya between 2006 and 2009, with each clinic starting the LREC program between 2007 and 2008. Table 1.2 shows the characteristics of the 15 clinics included in our study.

Table 1.2: LREC program clinic characteristics (n=15).

Area No.(%)		
Urban	6	(40.0)
Rural	9	(60.0)
Clinic Type No.(%)		
Referral hospital	3	(20.0)
(Sub) district hospital	7	(46.7)
Rural health center	5	(33.3)
Patients No.(%)		
≤ 500	4	(26.7)
$501-1000$	3	(20.0)
$1001-1500$	4	(26.7)
>1500	4	(26.7)

In an effort to ensure that other unmeasured (or unmeasurable) aspects of care remained roughly constant, the first point of patient eligibility was truncated at 1-year prior to LREC program initiation at each clinic. Our target population is therefore comprised of patients found eligible for LREC within 1-year before the clinic's start date up to the administrative censoring date (5 March 2009). Our baseline timepoint ($t = 0$) was defined to be the first date at which a patient was eligible

for the program within this window. Figure 1.2 shows the distribution of the time from baseline to start of the LREC program, with negative values representing patient eligibility post program initiation. A small number of subjects had more than one year of follow up from baseline to LREC program initiation as a result of transferring to a new clinic with a later LREC start date. As patients generally visited clinics every three months, we discretized follow-up into time intervals of 90 days.

Patients were followed from the baseline time point defined above until one of four possible end points:

1. Death
2. Loss to follow-up (LTFU), defined here to be 6.5 months with no clinical visits
3. Database closure, occurring on 5 March 2009
4. Transfer to a clinic with no LREC program

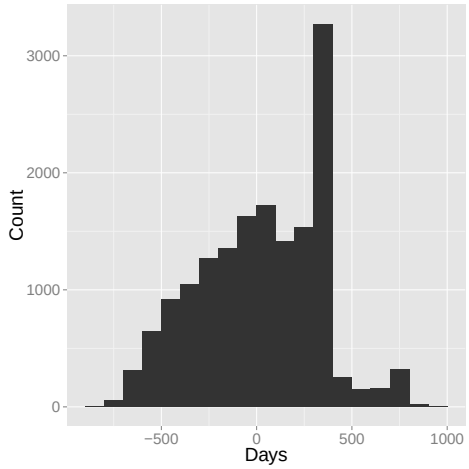

Figure 1.2: Histogram of time from eligibility to LREC program initiation.

Observed data

For notational convenience, we defined variables after the occurrence of any of these end points as equal to their last observed value. Following discretization of the data, we have a longitudinal data set with time-varying covariates, where the time points t correspond to 90-day intervals (e.g. 0, 90, 180, ... days). Let W be the observed baseline time-independent covariates observed at the date the patient was first eligible for LREC (age at eligibility, CD4 count at start of ART, gender, indicator that ART regimen is PI-based at eligibility, treated for tuberculosis at start of ART, indicator at urban clinic at eligibility, and WHO immunologic stage at both the start of ART and maximum stage prior to start). Let the time-varying variables from the observed data for each time point t be:

$$O(t) = (L1(t), Y(t), A1(t), A2(t), C1(t), C2(t)) : t = 0, 1, \ldots, K, \tag{1.1}$$

where

- $L1(t)$ consists of the most recent measures of time-varying covariate values at the start of interval t, inclusive of covariates measured at the clinic level (i.e. calendar date, most recent, nadir, and zenith CD4 count, days since enrolling into the AMPATH program, an indicator of remaining on ART, pregnancy status, indicator in WHO stage III or IV, indicator of being treated for tuberculosis, clinic type (rural or urban), and an indicator of having at least one clinic visit in the previous interval).

- $Y(t)$ is an indicator that the patient is either (a) no longer in care (not seen in the clinic for 6.5-months) or (b) has died by the end of interval t. It jumps to and remains at 1 if either event occurs.

- $A1(t)$ is an indicator of the LREC program availability by the end of interval t. It jumps to and remains at 1 once the program has started.

- $A2(t)$ is an indicator of enrollment into the LREC program by the end of the interval. It jumps to and remains at 1 at time of enrollment and remains at 0 if $A1(t) = 0$.

- $C1(t)$ is an indicator that the patient transfers by the end of interval t to a clinic other than one of the 15 clinics that initiate the LREC program. It also jumps to and remains at 1 once the patient transfers.

- $C2(t)$ is an indicator of data base closure by the end of interval t. It jumps to and remains at 1 at the end of the study. Note that although database closure occurs at a fixed calender date, censoring time due to database closure remains a random variable due to variability in time of eligibility for LREC.

To simplify notation, we refer to the covariate and outcome nodes collectively as $L(t) = (L1(t), Y(t))$. Furthermore, we refer to the treatment and censoring processes together as $A(t) = (A1(t), A2(t), C1(t), C2(t))$. By additionally defining $L(0)$ to include our baseline variables W such that $L(0) = (W, L1(0), Y(0))$, our observable data for each subject i can be expressed as

$$O_i = (L_i(0), A_i(0), L_i(1), A_i(1), \ldots, L_i(K+1)) \tag{1.2}$$

where $K+1$ is our final time point of interest, here equal to 4 (or equivalently 450 days after LREC eligibility). We assume the observed data over all subjects consists of n copies of $O_i \overset{iid}{\sim} P_0 \in \mathscr{M}$, where P_0 is the true underlying distribution (residing in a statistical model $\mathscr{M}$) from which the data are drawn.

Chapter 2

Evaluating the impact of the low-risk express care program

We follow the targeted learning road map as presented by van der Laan and Rose [49] and Petersen and van der Laan [66], as applied to our IeDEA-EA cohort. In doing so, we perform a case study of the targeted learning road map in the longitudinal setting. For organizational purposes, we separate the estimation of our parameter and interpretation of the results into subsequent sections.

In particular, we emphasize two primary issues in this chapter. The first is the translation of the scientific questions into counterfactual causal parameters, including a total effect and controlled direct effect. This is carefully defined as to avoid known and potential positivity violations. The second is the use of recently developed estimation methods, including a double robust and semi-parametric efficient targeted minimum loss-based estimator, the integration of Super Learning for nuisance parameter estimation, and the imposition of global restraints for conditional distributions implied by our statistical model. While previous analyses have applied the longitudinal targeted minimum loss-based estimation (TMLE) method described here [6, 14], to the best of our knowledge, this represents the first such application that integrates advanced machine learning algorithms in an approach known as Super Learning [47]. We discuss several decisions and practical challenges that arise as a result.

2.1 The targeted learning road map

Any causal scientific question must first be framed as a statistical question when being answered with the observed data. For purely statistical questions, this includes properly defining the data, statistical model, and estimand or statistical parameter of interest. Here we consider data that consists of n independent and identically distributed observations of a random (vector) variable O with some underlying probability distribution P_0. The statistical model $\mathscr{M}$ should be chosen to ensure that it contains the true distribution P_0, and thus any model assumptions should be based only on real knowledge about the process which generated the data. As stated in Chapter 1, in practice this generally implies a semi-parametric or nonparametric statistical model.

In order to translate causal questions into such a statistical estimation problem, the road map further requires translating the scientific question into a formal causal query. Such a query is typically defined as a parameter of the counterfactual distribution of the data, or equivalently, the distribution the data would have under some ideal hypothetical experiment or intervention. Here we consider counterfactual experiments indexed by static interventions to set a vector $\bar{A}$ of treatment or exposure variables equal to some fixed value $d(\bar{l})$ as a function of covariates $\bar{l}$, and denote the resulting counterfactual distribution of the observed data P_d^0. The statistical model $\mathscr{M}$ can then be augmented with possible additional non-testable causal assumptions about the data generating process. These causal assumptions are representable using structural causal diagrams or structural equation models [61]. The resulting causal model $\mathscr{M}^{\mathscr{F}}$ therefore represents a model on the counterfactual distribution P_d^0 under each $d(\bar{l})$ of interest.

Finally, one must also determine what additional causal assumptions (if any) are needed in order to obtain identifiability of the target causal parameter from the distribution of the observed data. In other words, what additional assumptions on the data generating process will allow the researcher to express the target causal parameter $\Psi^{\mathscr{F}}(P_d^0)$ as a statistical estimand $\Psi(P_0)$ that is a function of the observed data alone? In this chapter, we focus on the estimand provided by Bang and Robins [3] through the longitudinal g-formula, which identified the causal effects of longitudinal interventions under the sequential randomization assumption.

Once identifiability has been established, the statistical model $\mathscr{M}$ and target parameter of the observed data distribution $\Psi(P_0^d)$ are used to select a corresponding estimator $\hat{\Psi}$ which is a function of the empirical distribution P_n. This process may additionally require nuisance parameter estimation. Here, we describe implementation of a double robust efficient targeted maximum likelihood estimator [46], based on an iterative conditional expectation representation of the longitudinal g-formula proposed by Bang and Robins [3]. The resulting $\hat{\Psi}$ is used in calculating the parameter estimate $\hat{\Psi}(P_n) = \hat{\psi}_n$. Corresponding influence function based standard errors for the estimates are calculated for inference and results are interpreted. In the following sections we illustrate each of these steps in greater detail using the LREC analysis as a case study.

The scientific question

In implementing this task-shifting program, a primary question of interest among health care providers is whether clinic level exposure to and individual level enrollment in the program results in either better or worse clinical outcomes. For example, it could be that enrollment in the program increases loss to follow up and mortality because care is received from individuals that have a lower level of qualification or certification. Alternatively, enrollment in the program might decrease in mortality and loss to follow-up due to more attentive and personal care given to enrolled patients. It could also be that an equivalent level of care is provided and thus, no impact is observed. Furthermore, receiving care at a clinic that has already implemented the LREC program might itself have a direct beneficial or detrimental effect on patient outcomes, in addition to an indirect effect mediated by patient enrollment in the program. Such an effect could result, for example, from changes in clinic practices made possible by shifts in workload.

Likelihood

We first consider the likelihood $L(O)$ for the data structure specified in Chapter 1.2. In particular, we use the following factorization, implied by the assumed time ordering $L(t) \rightarrow A(t)$ such that the likelihood for subject i is

$$\begin{aligned} L(O_i) &= P_0(L_i(0), A_i(0), L_i(1), A_i(1), \ldots, L_i(K+1)) \\ &= P_0(L_i(K+1)|\bar{L}_i(K), \bar{A}_i(K)) \cdot P_0(A_i(K)|\bar{L}_i(K), \bar{A}_i(K-1)) \\ &\quad \cdot P_0(L_i(K)|\bar{L}_i(K-1), \bar{A}_i(K-1)) \cdot P_0(A_i(K-1)|\bar{L}_i(K-1), \bar{A}_i(K-2)) \\ &\quad \cdots P_0(L_i(0)) \\ &= \left[\prod_{t=0}^{K+1} \underbrace{P_0(L_i(t)|\bar{L}_i(t^-), \bar{A}_i(t^-))}_{Q_{0,L(t)}(L(t)|Pa(L(t))} \right] \cdot \left[\prod_{t=0}^{K} \underbrace{P_0(A_i(t)|\bar{L}_i(t), \bar{A}_i(t^-))}_{g_{0,A(t)}(A(t)|Pa(A(t))} \right] \end{aligned} \tag{2.1}$$

where we define $\bar{X}(t) \equiv (X(1), X(2), \ldots, X(t))$ to denote the history of variable X up to time t, $t^- \equiv t-1$ to denote the previous time point, $P_0(\cdot)$ to denote the true probability distribution of O, and $A(-1) = L(-1) = \varnothing$. We assume O is discrete for sake of presentation. We follow convention in using $Pa(X)$ to denote the parents of the node defined as the variables which precede it and denote the conditional probabilities $P_0(L(t)|\cdot)$ and $P_0(A(t)|\cdot)$ as the Q and g factors for the likelihood, respectively, such that $Q_{0,L(t)}(L(t)|Pa(L(t)) = P_0(L(t)|\bar{L}(t^-), \bar{A}(t^-))$ and $g_{0,A(t)}(A(t)|Pa(A(t)) = P_0(A(t)|\bar{L}(t), \bar{A}(t^-))$. For ease of notation, we collectively refer to the entire set of $P_0(L(t)|\cdot)$ and $P_0(A(t)|\cdot)$ for all t respectively as

$$\begin{aligned} Q_0 &\equiv (Q_{0,L(0)}(L(0)|Pa(L(0))), Q_{0,L(1)}(L(1)|Pa(L(1))), \ldots, Q_{0,L(K+1)}(L(K+1)|Pa(L(K+1)))) \\ g_0 &\equiv (g_{0,A(0)}(A(0)|Pa(A(0))), g_{0,A(1)}(A(1)|Pa(A(1))), \ldots, g_{0,A(K)}(A(K)|Pa(A(K)))). \end{aligned} \tag{2.2}$$

Furthermore, we denote the conditional expectation of the $L(t)$ process with an overhead bar notation such that $\bar{Q}_{0,L(t)} \equiv \mathbb{E}_0[L(t)|\bar{L}(t^-), \bar{A}(t^-)]$ and $\bar{Q}_0$ to denote the collection of $\bar{Q}_{0,L(t)}$ across all time points.

The statistical model

With our data and likelihood clear, we first consider a statistical model $\mathscr{M}$ for the true distribution P_0, such that if $\mathscr{Q}$ and $\mathscr{G}$ are the variationally independent sets of all possible values for Q_0 and g_0 respectively, then the statistical model can be represented as $\mathscr{M} = \{P = Q \cdot g : Q \in \mathscr{Q}, g \in \mathscr{G}\}$. We assume a semi-parametric model, which restricts the set of possible distributions for the g_0 and Q_0 components of the likelihood. Specifically, to respect the factual details that we know about the data generating process, we force two model restrictions on the conditional distributions of $g_{0,A(t)}(A(t)|Pa(A(t)) : k = 0, 1, \ldots, K$.

1. Once $A1(t) = 1$, we have that $\underline{A1}(t^+) = 1$.
2. Once $A2(t) = 1$, we have that $\underline{A2}(t^+) = 1$.

where for notational convenience, we define $t^+ \equiv t+1$ to denote the next time point and $\underline{X}(t^+) \equiv (X(t+1), X(t+2), \ldots, X(K+1))$ to denote the remaining history of X from time $t+1$ to $K+1$.

Knowing that our outcome is also a counting process, we additionally force a model restriction on the conditional distributions of $Q_{0,L(t)}(L(t)|Pa(L(t)) : k = 1, \ldots, K+1$.

1. Once $Y(t) = 1$, we have that all last observed values for $O(t)$ are carried forward with probability 1, i.e. $\underline{O}(t) = O(t)$

where we define $\underline{X}(t) \equiv (X(t), X(t+1), \ldots, X(K+1))$ to denote the remaining history of variable X to time $K+1$.

We emphasize again the stark contrast of this approach to the more classical one in which the analyst sees that the dependent variable is binary and, say, chooses a parametric logistic regression model with some "reasonable" set of covariates. The approach taken here makes minimal assumptions for our statistical model that we know to be true, as opposed to the parametric approach that makes many constraints such as linearity, smoothness, and no multi-collinearity.

The causal model

A causal model allows us to represent additional knowledge and assumptions associated with our scientific question that cannot be represented statistically. We present our causal model $\mathscr{M}^{\mathscr{F}}$ by making use of structural equation models to formally present how we assume each variable to be generated. We treat the 15 IeDEA-EA clinics as fixed, rather than sampled, which would describe an experiment in which individual subjects become eligible for the LREC program at random times. Specifically, we define the following non-parametric structural equation model (NPSEM) [61] to represent our knowledge about the causal process that generated the observed data.

$$\begin{aligned}
L1(t) &= f_{L1(t)}(\bar{L}(t^-), \bar{A}(t^-), U_{L1(t)}) \\
Y(t) &= f_{Y(t)}(\bar{L}(t^-), \bar{A}(t^-), U_{Y(t)}) \\
&\qquad \text{for } t = 0, 1, \ldots, K, K+1 \\
A1(t) &= f_{A1(t)}(\bar{L}(t), \bar{A}(t^-), U_{A1(t)}) \\
A2(t) &= f_{A2(t)}(\bar{L}(t), \bar{A}(t^-), \bar{A1}(t), U_{A2(t)}) \\
C1(t) &= f_{C1(t)}(\bar{L}(t), Y(t^-), \bar{A}(t^-), A1(t), A2(t), U_{C1(t)}) \\
C2(t) &= f_{C2(t)}(L(0), Y(t^-), \bar{C}(t^-), C1(t), U_{C2(t)}) \\
&\qquad \text{for } t = 0, 1, \ldots, K
\end{aligned} \tag{2.3}$$

where $U \equiv (U_{L1(t)}, U_{Y(t)}, U_{A1(t)}, U_{A2(t)}, U_{C1(t)}, U_{C2(t)})$ are unmeasured exogenous variables from some underlying probability distribution P_U and $L(-1) = A(-1) = \varnothing$.

This causal model specifies how we believe each of the variables in our data are deterministically generated, with randomness coming only from the unmeasured exogenous variables U. It tells us, for example, that individual enrollment immediately following LREC eligibility $A2(0)$ is generated as a deterministic function $f_{A2(0)}$ of $L(0)$, program availability $A1(0)$, and an error term $U_{A2(0)}$ drawn from some underlying population. Additionally, while not explicitly stated in Equation (2.3), we note that the deterministic function for enrollment $f_{A2(t)}$ sets $A2(t) = 0$ with probability 1 if $\bar{A}1(t) = 1$, i.e. if the program is not yet available.

Note that censoring due to end of study $C2(t)$ is not a function of time updated covariates $\bar{L}(t)$ beyond baseline covariates $L(0)$. This is because, while the values of $L(0)$ may vary due to the calender date at which a subject's baseline eligibility occurs, once this date is set for a given subject, the censoring process due to database closure is deterministic.

Our outcome $Y(t)$ is assumed to be a function of treatment and censoring $A(t)$ only up to the previous time point, $t-1$. This restriction is imposed in order to avoid the possibility of reverse causality, i.e. that death/LTFU ($Y(t)$) that occurs in an interval t affects availability/enrollment ($A(t)$) in the same interval. Consequently, the effects of availability and enrollment within an interval on the composite outcome are only captured beginning in the following interval.

The causal parameter

Recall that we are interested in evaluating the impact of the LREC program for the IeDEA cohort. It follows that we should define a causal parameter that corresponds with this scientific question. More specifically, we are interested in the effect of LREC exposure and enrollment on the probability of both death and remaining in clinical care. Patients who do not return for continuing HIV care are subject to higher risk of complications and health decline [38, 89, 22, 21, 34], placing them at unnecessarily higher mortality rates. To account for this, we define our outcome of interest (Chapter 1.2) as a composite of either the occurrence of death or LTFU. Patients were followed until this "*failure*" or until censoring due to either the end of study or clinic transfer. We aimed to evaluate the impact of (a) implementation of the LREC program at the clinic, and (b) enrollment into the LREC program after implementation on both retention "*in-care*" and survival.

Conceptualizing an ideal hypothetical experiment can help in defining the target counterfactual parameter of interest. In order to evaluate the effect of exposure to and enrollment in the LREC program, we can conceive of an experiment in which we compare survival over time under alternative interventions to set time to program availability, time to enrollment following availability, and under an additional intervention to prevent censoring. As represented above in the causal model, these counterfactual outcome distributions are defined in terms of an intervention on the data-generating mechanism for $A(t) : t = 0, 1, \ldots, K$. In other words, we intervene to set values of program availability, enrollment, and censoring to some fixed values of interest at all time points.

The outcome at $t = 0$ is independent of any potential treatment assignments $A(t)$. As the purpose of this study is to analyze the impact of different levels of treatment on the outcome, we conditioned on survival past $t = 0$. Thus, we had that $Y(0) = 0$ for all subjects included in the study. That is, all subjects in the study survived past the first 90 days.

As the causal parameter, we focus on the mean of the counterfactual outcome at a fixed time point under a specific treatment intervention $d(\bar{l}(K))$. Let $\bar{Y}_d(t^*)$ denote the counterfactual outcome process over time t^* under an intervention to set both time of program availability ($\overline{A1}(t^*-1)$) and time of individual enrollment ($\overline{A2}(t^*-1)$). Our target parameter for a given intervention of interest is $\mathbb{E}Y_d(t^*) : t^* = 1,\ldots,4$, where $\mathbb{E}Y_d(t^*)$ is the cumulative counterfactual probability of failure by time t^* under intervention $d(\bar{l}(t^*-1))$. As we have conditioned on surviving past $t = 0$ (the first 90-days of follow-up), our range represents the cumulative probability of failure from 180 to 450 days post-eligibility. The intervention on the exposure(s) of interest at each time point corresponds to deterministically setting the values of $A(t)$ to some value $d(\bar{l}(t))$ as a function of covariate $\bar{l}(t)$ for all t in the causal model $\mathscr{M}^{\mathscr{F}}$ specified above, resulting in a modified distribution P_d^0. The counterfactual outcome $Y_d(t^*)$ under this intervention can be interpreted as the value of $Y(t^*)$ at some fixed time point $t^* \leq K+1$ that would have been generated had $A(t)$ been deterministically set to $d(\bar{l}(t))$ for all t. We use t^* here to differentiate from the time t in which the data generating process spans.

We focused on static interventions here so that there is only one and exactly one treatment such that $d(\bar{l}(t^*)) = \bar{a}(t^*)$. When contrasting the counterfactual failure probabilities under the distinct interventions, we focused on estimating the absolute risk difference (or average treatment effect). Specifically, we contrasted these intervention specific counterfactual survival probabilities under the three following static interventions: Our first intervention assigns all patients to have no program availability at all time points (set $A1(t) = A2(t) = 0 : t = 1,2,\ldots,t^*$), and forces patients to remain uncensored (set $C1(t) = C2(t) = 0 : t = 1,2,\ldots,t^*$). The corresponding 4 counterfactual failure probabilities $\mathbb{E}Y_{d=00}(t^*) : t^* = 1,\ldots,4$ give us an understanding of survival patterns without the LREC program.

The second intervention of interest is to assign all patients to have immediate program availability (set $A1(t) = 1 : t = 1,2,\ldots,t^*$), but not allow any subjects to enroll into the program (set $A2(t) = 0 : t = 1,2,\ldots,t^*$). Patients would again be forced to remain uncensored and the counterfactual failure probability at each time point $\mathbb{E}Y_{d=1,0}(t^*) : t^* = 1,\ldots,4$ would be calculated. By evaluating $\Psi_{1,0}(P_0) = \mathbb{E}Y_{d=1,0}(t^*) - \mathbb{E}Y_{d=0,0}(t^*) : t^* = 1,...,4$, we target the controlled direct effect of exposure to the program if enrollment were prevented.

The third intervention is to assign all patients to have both immediate availability and enrollment (set $A1(t) = A2(t) = 1 : t = 1,2,\ldots,t^*$). Again, censoring would be prevented and the counterfactual failure probability at each time point $\mathbb{E}Y_{d=1,1}(t^*) : t^* = 1,\ldots,4$ would be calculated. Evaluating $\Psi_{1,1}(P_0) = \mathbb{E}Y_{d=1,1}(t^*) - \mathbb{E}Y_{d=1,0}(t^*) : t^* = 1,\ldots,4$ allows us to target the total effect of enrollment in a scenario where all subjects experienced immediate availability.

Identifiability

It is well recognized that the cumulative effect of longitudinal exposures is often subject to time dependent confounding (for example, Robins et. al. [77], Bodnar [4]). For example, the decision to continue to defer enrollment at post-baseline time points may be affected by covariates that affect the outcome (and are thus confounders), and that are themselves affected by the prior decision not

to enroll. If not properly accounted for, the resulting estimates may be biased for the parameter of interest.

Similar challenges arise in analyses aiming to estimate the effect of one longitudinal exposure, while holding a second longitudinal exposure constant. For example, how would patient outcomes have differed under immediate versus deferred exposure to a clinic level intervention, if individual level enrollment were prevented? Such a controlled direct effect of two longitudinal exposures can (under certain assumptions) be used to investigate the mechanisms by which an exposure is mediated by individual level uptake. In such scenarios, both longitudinal exposures may be subject to time dependent confounding.

Under the untestable assumption of sequential randomization [71, 76] and the partially testable assumption of positivity [10, 64], this parameter $\mathbb{E}Y_d(t^*)$ is identifiable from the observed data using the longitudinal g-computation formula [71]:

$$\mathbb{E}Y_d(t^*) = \sum_{\bar{l}(t^{*-})} \mathbb{E}(Y(t^*)|\bar{L}(t^{*-}) = \bar{l}(t^{*-}), \bar{A}(t^{*-}) = d(\bar{l}(t^{*-}))) \cdot \prod_{j=0}^{t^{*-}} P(l(j)|\bar{L}(j-1) = \bar{l}(j-1), \bar{A}(j-1) = d(\bar{l}(j-1))) \tag{2.4}$$

Sequential randomization

Under the sequential randomization assumption [71], we have that

$$Y_d(t^*) \perp\!\!\!\perp A(t)|Pa(A(t)) : t = 1, 2, \ldots, t^* - 1 \tag{2.5}$$

That is, our treatment is independent of the counterfactual outcome given its parents or informally, that measured covariates are sufficient to control for confounding of treatment and informative censoring. A sufficient condition for this assumption to be met is if all unmeasured exogenous variables affecting the treatment and censoring nodes $U_{A(t)}$ are independent of the exogenous variables affecting future $Y(t)$ nodes given the observed past up to time t. This occurs if all potential confounders for the treatment and outcome of interest are measured and adjusted for appropriately. The randomization assumption or backdoor criteria is often used to formally define whether a given adjustment set of measured covariates is sufficient to control for confounding [86, 71, 60]. In this context, the major concern for violation of this assumption is that among patients classified as clinically stable, some patients are healthier or at lower risk of loss than others in ways not captured by the measured covariates, and these patients are differentially enrolled into the program.

Positivity

Our assumption of positivity [32, 10, 114, 64] states that:

$$P_0(\bar{A}(t^* - 1) = d(\bar{l}(t^* - 1))|\bar{L}(t^* - 1), \bar{A}(t^* - 2) = d(\bar{l}(t^* - 2))) > 0 \; a.e. \tag{2.6}$$

Informally, this states that there is adequate support for each intervention of interest regardless of covariate history. As we demonstrate below in simulations, even in situations where this assumption on P_0 holds, in finite samples certain covariate histories and treatment combinations may be poorly supported, resulting in data sparsity and, consequently, potential increases in estimator bias and variance which threaten valid inference.

Patients losing their eligibility for the LREC program posed a particular threat to this assumption. Our study population is comprised of patients initially deemed eligible for the LREC program due to their low risk. However, a noticeable proportion of the study population (32%) lost their eligibility at some point during follow-up. Once these patients were ineligible, they had a 0 probability of subsequent program enrollment. To circumvent this potential positivity violation, we considered only treatment interventions which avoided enrollment at these time points. For example, consideration of patients who enroll immediately into the program would not encounter this issue, as all patients are eligible at the start of follow-up. Consideration of patients never enrolling into the program is also valid, as patients who are not eligible do not enroll. We further note that patients who lost their eligibility after enrollment into the LREC program were still considered to be enrolled. Similarly, patients who transferred to a new clinic without availability after receiving care at a clinic where LREC was available were considered exposed to the LREC program (in other words, we conducted an intent to treat type analysis of the effect of both availability and enrollment).

We emphasize that while the causal parameter $\mathbb{E}Y_d(t^*)$ is of primary interest to us, in the scenario where the assumptions are unmet, we still have a desirable and well defined statistical parameter. This parameter can be thought of as the association between the intervention of interest and the outcome adjusted for the *measured* subset of potential confounders, both baseline and time varying.

2.2 Estimating $\mathbb{E}Y_d(t^*)$

A range of estimators of our target parameter (or more precisely, of the statistical estimands to which the counterfactual effects correspond under the sequential randomization assumption), have been developed, implemented, and applied (e.g. Robins [70, 75], Robins et. al. [84], Rotnitzky and Robins [87], Hernan and Robins [32], and Pearl [61]). Prominent examples include, inverse probability weighted (IPW) [35, 41, 79, 92, 91, 87, 32], parametric g-computation [84, 32, 76, 102], and double robust estimating equation-based estimators [80, 48, 3, 107, 82].

The efficient influence function

In choosing an estimator, a desirable property to seek is the lowest asymptotic variance among reasonable estimators as this ensures that we have the most statistically powerful estimator asymptotically. All regular asymptotically linear estimators have influence functions (IF) [25, 107, 112] such that the estimator variance is equal to the variance of the influence function divided by n. Estimators with the lowest asymptotic variance will have by definition IF equal to the efficient

influence function (EIF) for $\Psi(P_0)$. The EIF for our target parameter, denoted $D^*(P)(O)$, has been derived previously [93, 80]. It is given by

$$D^*(P)(O) = \sum_{t=0}^{t^*} D_t^*(P)(O), \text{ where} \tag{2.7}$$

$$\begin{aligned} D_0^*(P)(L(0)) &= \bar{Q}_1^d - \bar{Q}_0^d \\ D_t^*(P)(\bar{A}(t^-), \bar{L}(t^-)) &= H_t(g)(\bar{Q}_{t^+}^d - \bar{Q}_t^d) : t = 1, 2, \ldots, t^* \end{aligned}$$

with

$$H_t(g)(\bar{A}(t^-), \bar{L}(t^-)) = \frac{\mathbb{I}(\bar{A}(t^-) = d(\bar{l}(t^-)))}{g_{0:t^-}(\bar{A}(t^-), \bar{L}(t^-))} \tag{2.8}$$

$$\bar{Q}_{t^*+1}^d = Y(t^*) \tag{2.9}$$

$$\bar{Q}_t^d(\bar{L}(t^-)) = \mathbb{E}_P[Y^d(t^*) \mid \bar{L}^d(t^-) = \bar{L}(t^-)] : t = 1, 2, \ldots, t^* \tag{2.10}$$

$$\bar{Q}_0^d = \mathbb{E}_P[Y^d(t^*)]. \tag{2.11}$$

It should be noted that $g_{0:t^-}(\bar{A}(t^-), \bar{L}(t^-))$ represents the cumulative probability of treatment up to time $t-1$ and that $\bar{Q}_t^d(\bar{L}(t^-)) = \mathbb{E}_P[\bar{Q}_{t^+}^d \mid \bar{L}(t^-), \bar{A}(t^-) = d(\bar{l}(t^-))]$ is defined by recursive regression, starting at $t = t^*$ and moving backwards in time. For notational convenience, we let $H_0 = 1$ so that

$$D^*(P)(O) = \sum_{t=0}^{t^*} H_t(g)(\bar{Q}_{t^+}^d - \bar{Q}_t^d).$$

There are a number of points worth re-emphasizing here. Firstly, this EIF is simply a sum of time-specific IFs over all time points. Thus, estimators that solve the estimating equation corresponding to the EIF can be constructed such that they solve the estimating equations individually at each time point. Secondly, when t^* is equal to 1, i.e. there is only one time point, this IF reduces to the well known EIF for the point treatment setting [79, 81, 80]. Lastly, the IF has, in the denominator of the first term, the cumulative probability of treatment up to each time point t. This implies that the variance is highly dependent upon the cumulative probability of treatment given the past, with the magnitude of variance increasing with lower probabilities. Consequently, positivity violations or near violations (Equation 2.6), where the probability of treatment given the past is extremely low, can have large effects on the performance estimators solving this IF.

Targeted minimum loss-based estimation

More recently, Bang and Robins [3] introduced double robust iterated conditional expectation g-computation estimators solving the EIF, which van der Laan and Gruber [46] extended to develop a longitudinal iterated conditional expectation TMLE. The g-computation recursive algorithm for

estimating the cumulative probability of failure, introduced by Robins [75] and expanded in Bang and Robins [3], calls upon the tower rule of conditional expectations for this identity, suggesting an iterative conditional expectation (ICE) approach to estimating our parameters. Using their results, our parameter can therefore instead be expressed as

$$\mathbb{E}[\mathbb{E}[\cdots\mathbb{E}[\mathbb{E}[Y^d(t^*)|\bar{L}^d(t^*-1)]|\bar{L}^d(t^*-2)]\cdots|\bar{L}^d(0)]] \tag{2.12}$$

where $L^d(t)$ is the variable $L(t)$ from the post intervention distribution resulting from setting $\bar{A}(t^-) = d(\bar{l}(t^-))$ and the expectations are taken with respect to this distribution. Thus, given $\bar{L}^d(t)$ is equivalent to conditioning on $\bar{L}(t), \bar{A}(t) = d(\bar{l}(t))$.

This ICE approach has a number of advantages to the more simple standard g-computation procedure. The most noticeable among them is that we are only required to estimate the iteratively defined conditional expectations as opposed to the entire conditional densities of the covariate process $\bar{L}(t^*)$. We therefore use the ICE approach here. A small disadvantage is that these set of regressions must be run separately for each desired treatment rule, $d(\bar{l}(t^*))$, whereas in using the original formulation one only has to estimate the conditional densities once. We note however that this is a very small price to pay when compared against the substantial gain achieved by not having to estimate the entire joint density, especially when dealing with high dimensional data.

While the ICE approach already provides a considerable advantage towards our estimation goals, using targeted minimum loss-based estimation [46, 49] provides a further gain. This approach, which builds upon the double robust ICE estimator of Bang and Robins [3], solves the derived efficient influence function $D(P)$ for our estimand within a substitution based setting and is therefore guaranteed to respect the parameter constraints implied by our statistical model. This removes the bias associated with the untargeted minimization of a global loss function for the density of the data. Furthermore, TMLE does not suffer to the same extent from positivity violations [77, 114, 64] as IPW [77, 32, 64], and allows for full integration of machine learning while retaining valid inference [49, 6, 95, 14]. It is also known to be double robust, in that consistent estimation of the target parameter is obtained if either of the treatment or outcome distributions are estimated consistently. All analyses were conducted on R version 3.1.1 [104]. An R package titled `ltmle` has been developed which implements this estimator [96, 65]. This package takes longitudinal data supplied in wide format and estimates our target parameter for each specified treatment rule $d(l(\bar{t^*}))$. We therefore used this estimator in the analysis on the impact of the LREC program.

A number of options are available to users of the `ltmle` R package. For example, the probabilities of treatment and censoring $A(t)$ at each time point (for the specified treatment rule $d(\bar{l}(t^*))$) can be separately estimated and subsequently fed into the estimation procedure, rather then being fit within. This allows the analyst the option of pooling the observations over all time points and estimating probabilities within this pooled sample, as opposed to fitting separate models for each time point. Doing so provides a lower variance in estimates of the treatment mechanisms at the cost of potentially higher bias.

An additional advantage of pooling over all observations in estimating our treatment and censoring mechanisms is that we can use additional data that is not included in modeling the ICEs. That is, data observed beyond the final time point t^* can be used to aid in estimating the probabili-

ties of treatment for $A(t) : t = 0, 1, \ldots, t^* - 1$. This can be advantageous and promotes stability in the estimates by borrowing information across all time points, also at a cost of potentially higher asymptotic bias. Specifically, it helped in the present study in estimating our censoring from transfer mechanism, due to the extremely small number of transfers observed. For this study, we choose to pool across observations to fit the treatment mechanisms though note that the decision between the two did not significantly affect the parameter estimates.

An additional package option is the ability to pool over all subjects when estimating the ICEs irregardless of observed treatment, as opposed to stratifying the data and using only subjects following the treatment rule specified. This choice implies an analogous bias-variance trade off. Pooling over subjects regardless of treatment history potentially allows for more stable estimates of the ICEs due to the smaller variance. This option is helpful when the number of time points is large and the number of persons following a particular treatment regime over all time points is small.

Use of the `ltmle` R package requires that the data be provided in a wide format and with a time-ordering of the covariates. In doing so, two options are available for the $L(t)$ node at each time point t, i.e. $L1(t), Y(t)$ or $Y(t), L1(t)$. In our causal model in Chapter 2.1, the $L(t)$ node is not affected by the specified order. Therefore, use of either ordering will suffice in our study. We additionally note, however, that even if the time-ordering did matter and our outcome of interest $Y(t^*)$ were, say, dependent upon the covariates $L1(t^*)$, there is still no need to condition on $L1(t^*)$ for the sake of estimating our target parameter as the EIF for our parameter based on the full data O is the same as the efficient influence function based on reduced data $O_r = O/L1(t^*)$. We provide a short proof for this in Appendix A.

Super Learning the nuisance parameters

Consistent, asymptotically linear, and efficient estimation of our target parameter requires that the treatment and censoring $A(t)$ mechanisms as well as ICEs be estimated in a consistent manner and at a fast enough rate. Recall that using parametric models to do this requires correct specification of the functional form for the conditional densities. Given that we do not know a priori the form of the true probability distribution P_0 and the extreme unlikeliness that a simple parametric specification will result in a correctly specified model, use of these models will most likely result in overly biased estimates which will approach statistical significance with probability 1 as sample size increases regardless of whether a treatment effect exists. In other words, the use of mis-specified parametric models elevates the risk of obtaining significant findings even if no true treatment effect is present.

We instead rely on the use of data-adaptive non-parametric methods which reside in a much larger statistical model or set of distributions. Examples include gradient boosting machines [17, 19], neural networks [56], and k-nearest neighbors [1]. In deciding which method to use, we recommend using the ensemble machine learning approach Super Learner, which is based on V-fold cross validation and implemented in the R package titled `SuperLearner` [67]. This algorithm takes a user-supplied loss function (chosen to measure performance) and a library of algorithms, which can include parametric models as well as non-parametric or machine learning algorithms such as those listed above. It uses V-fold cross validation to chose the convex combination of

these algorithms that performs best on independent data (derived from internal data splits). If, as is likely, none of the algorithms in the library achieves the rate of convergence one would have with a correctly specified parametric model, the Super Learner will still perform asymptotically at least as well as the best algorithm in the library. Otherwise, it achieves an almost parametric rate of convergence (i.e. $\log(n)/n$). Furthermore, the derived oracle inequality showing the asymptotic performance also shows that the number of candidate algorithms considered in the cross-validation procedure can be polynomial in size proportional to the number of observations [44, 109, 47]. Therefore, a large number of algorithms can be considered, which can grow with the number of observations, without fear of hampering the Super Learner's performance.

To use Super Learning, a loss function must be chosen and a user-specified library provided. We chose the non-negative log loss function for its desired steeper risk surface, as this function penalizes incorrect probabilities more severely than the more commonly used squared error loss. A number of default candidates are included in the `SuperLearner` package that were used here. These include the overall mean, main terms logistic model, step-wise regression with AIC [33], generalized additive model [27], Bayesian generalized linear model [20], k-nearest neighbors [1], LASSO [105], ridge regression [33], neural net [56], multivariate adaptive polynomial spline [18], generalized boosted regression model [17, 19], and support vector machine [5, 11]. Additionally, most of the algorithms have tuning parameters which can result in better candidate performance. To ensure that we were achieving satisfactory performance, we used different tuning parameters as additional candidates in the ensemble for the generalized additive models, k-nearest neighbors, neural nets, and generalized boosted regression models. We additionally used 4 user-specific parametric models as candidates in the library. The Super Learner fits were constructed using all potential confounders, listed above in Chapter 1.2 as W and $L1(t)$ for each time point t.

Of particular concern to the analyst when deciding on which candidates to include in the Super Learner library is the explicit condition that the candidates not be too data adaptive. This is because the empirical process conditions for the asymptotic linearity of our estimator require that we work within a Donsker class of estimators [50]. Indeed, we have seen that algorithms that tend to overfit the empirical data, such as the machine learning algorithm random forest, will negatively impact our estimators. We therefore excluded these algorithms from our library, though note that such algorithms could still be used in a cross-validated parameter estimation approach such as cross-validated targeted minimum loss-based estimation [117].

Regarding the pooling of observations across time for the treatment mechanism, one further possible option is to use a Super Learner library that is doubled in the number of candidates by including estimates from both the time stratified and pooled approach. Ensemble weights could then be calculated based on the best performing candidate in this larger library and subsequently fed into the `ltmle` package. Consequently, we continue benefiting from the borrowed information at different time points and simultaneously protect ourselves from the asymptotic bias of the previous approach. We opted not to additionally use the stratified approach, due to the computational intensity required of the approach.

Initial ICE fits

In using non-parametric estimators for the estimation of the ICEs, we may potentially unintentionally disregard the global constraints implied by our statistical model. While all the estimators considered here are expected to work well at time $t = t^*$ where the outcome being modelled is binary, their use at $t < t^*$ where the outcome being modelled is known to fall within the interval $[0,1]$ can present issues. For example, continuing to treat the outcome being modelled as binary may result in programmatic errors since many of the algorithms, such as Support Vector Machines or k-nearest neighbors, require that the outcome be supplied in classes or as factors. Alternatively, modelling the outcome continuously may result in extrapolations with estimates greater than 1 or less than 0. As our expectation is known to fall within $[0,1]$, this would result in violations of the constraints of the outcome being modelled.

For each of the algorithms facing this potential issue, we implemented three approaches aimed at ensuring that estimates remained within the constraints. All three were used in the Super Learner library, allowing us to objectively compare the performance of each approach. We define each of the conditional expectations from Equation (2.12) to be $\bar{Q}_{0,L(t)}$ such that, for example, $\bar{Q}_{0,L(t^*)} \equiv \mathbb{E}[Y^d(t^*)|\bar{L}^d(t^*-1)]$ and $\bar{Q}_{0,L(t^*-1)} \equiv \mathbb{E}[\bar{Q}_{0,L(t^*)}|\bar{L}^d(t^*-2)]$. Let each estimator of the conditional expectation within the Super Learner library at time t be denoted as $\bar{Q}^j_{n,L(t)}$ for $j \in 1,2,\ldots,J$ where J is the total number of candidates in the Super Learner library. We considered

1. Simply truncating $\bar{Q}^j_{n,L(t)}$ at both 0 and 1.
2. Taking the logit transformation of the outcome being modelled and truncating at a fixed threshold τ (set here to be 0.0001). We then modelled the transformed outcome on a continuous scale and took the inverse logit transformation on the fitted values.
3. Stratifying the observations by whether they were within the $(0,1)$ open interval or within $\{0,1\}$, i.e. whether they were continuous within the $(0,1)$ interval or dichotomous with only values of 0 or 1. The former were fit on a continuous scale after taking the logit transformation, while the latter were modelled as a binary outcome.

We emphasize that use of Super Learner for estimation of the treatment mechanisms and ICEs provides two important primary benefits. Firstly, its use helps ensure the conditions for the asymptotic linearity of our estimator and the corresponding statistical inference are met by ensuring the consistent estimation of both the intervention mechanism and the iteratively defined conditional expectations. This allows us to establish robust confidence intervals for our estimator. Secondly, we gain efficiency in that we get an asymptotically efficient estimator if both the treatment mechanisms and ICEs are estimated consistently at fast enough rate. Thus, as long as at least one of the library candidates for each of the nuisance parameters achieve this optimal rate, our approach will have the lowest variance among all regular asymptotically linear estimators. Further, even if we fall short of this goal, the improved estimation of both nuisance parameters offered by Super Learner will generally improve finite sample efficiency.

2.3 Results

Among the 15 clinics implementing the LREC program, a total of 16,513 subjects (31% male) were found to be eligible for the program, of which 16,479 survived past $t = 0$. As we only modelled survival up to 450 days, we report figures with follow-up truncated at that point. After discretizing the data, these patients contributed a total of 17,668 person-years of follow-up to the analyses. From these subjects 1,206 failure events were observed, of which 1,102 were losses to follow-up and 104 were deaths. All failure events observed at $t = 1$ were deaths, since our definition of loss to follow-up required at least 6.5 months to pass before a subject could be lost to follow-up. A small number of subjects ($n = 128$) were censored due to transfers to non-LREC clinics, while 3,889 subjects reached the end of study prior to $t = 4$ and prior to experiencing a failure event and were administratively censored. Table 2.1 shows the characteristics of the patients conditioning on survival past $t = 0$.

A small proportion of subjects died (42), were lost to follow-up (286), or were censored (60) before the LREC program became available. Of the 16,050 subjects who were at some point exposed to the program, most (15,294) experienced it by 1-year from baseline. Almost half of the study population began follow-up after the LREC program had already initiated, as indicated by the large spike in the cumulative incidence at time 0 (Figure 2.1). A noticeable spike was also seen at 1-year after baseline, representing the patients who had their first eligibility truncated at 1-year as stated in Chapter 3.1.

Patients who were not exposed to the LREC program could not enroll. Furthermore, once the LREC program was available, decisions on whether to enroll subjects rested upon the treating clinicians or clinical officers. Consequently, only 3,832 subjects were enrolled. As expected, subjects who were healthier were more likely to enroll into the LREC program. For example, univariate analyses showed that subjects who had higher CD4 counts, were receiving ARV, had a WHO stage I or II, were seen in clinics less often, and were not being treated for tuberculosis had higher probabilities of enrolling. Additionally, subjects from (sub) district hospitals and rural health centers (compared to referral hospitals), with fewer missed clinical visits, not on protease-inhibitor based regimen, and who were not pregnant also had higher probabilities of enrolling. We note that despite listing non-pregnancy as a criteria for being at low risk, a small number of subjects (18) enrolled while pregnant. Figure 2.1 shows the cumulative incidence of LREC availability and enrollment. Cumulative incidence of enrollment by 90 and 360 days after baseline was 7% (95% CI: 6.3%, 7.1%) and 19% (95% CI: 18.8%, 20.1%), respectively.

As stated in Chapter 3.1, all patients in our study started follow-up eligible for the LREC program, leading to low or no variance in many of the confounders at early time points with a skewness towards the healthier values. During follow-up, however, many subjects who did not enroll subsequently became less healthy resulting in decreased probabilities of subsequent enrollment. These covariates measuring their health and enrollment probabilities therefore represent classical time-dependent confounding and should be adjusted for appropriately [77]. Indeed, 3,920 subjects were found to have lost their eligibility prior to enrollment and prior to 1-year, precluding interventions to evaluate a range of different enrollment times, as discussed in Chapter 2.1.

Unadjusted analyses using the Kaplan-meier estimator showed overall high probabilities of in-

Table 2.1: Characteristics of 16,479 patients at LREC eligibility (conditioning on survival past $t = 0$).

	Count	(%)
Age (years)		
<30	2,793	(17%)
30-39	7,017	(43%)
≥40	6,669	(40%)
Sex		
Female	11,441	(69%)
Male	5,038	(31%)
CD4 cell count (cells/μL) at ART start		
<200	9,754	(59%)
200-349	2,313	(14%)
350-499	543	(3%)
≥500	387	(2%)
Unknown	3,482	(21%)
CD4 cell count (cells/μL) at baseline		
<200	0	(0%)
200-349	10,125	(61%)
350-499	4,016	(24%)
≥500	2,334	(14%)
Unknown	4	(0%)
PI-based ARV regimen		
No	15,600	(95%)
Yes	879	(5%)
Max WHO stage prior to ARV start		
I/II	7,696	(47%)
III/IV	8,373	(51%)
Unknown	410	(2%)

care survival among all subjects (Figure 2.2). Those with immediate LREC availability who never enrolled had noticeably lower in-care survival than subjects never experiencing LREC availability. For example, at $t = 4$ the proportion of subjects still alive and in care was 0.77 for those not enrolling into LREC, compared to 0.93 for the group of subjects never experiencing LREC availability and 0.94 for subjects enrolling into LREC immediately. Conversely, those with immediate enrollment into the program had the highest survival probabilities. Differences in survival probabilities between treatment groups increased with time, with the largest differences seen between subjects with immediate enrollment and those with LREC availability never enrolling.

The cross-validated risks (using the non-negative log likelihood loss) for the treatment and censoring mechanisms are shown in Figure 2.3 under various models and algorithms. While ob-

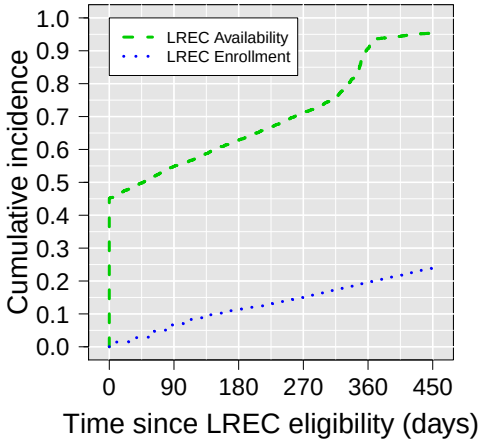

Figure 2.1: Cumulative incidence of LREC availability and enrollment and 95% confidence interval.

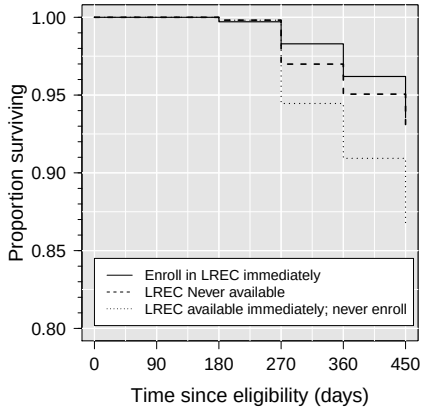

Figure 2.2: Unadjusted Kaplan-Meier survival curves.

servations at all time points were used in the mechanism fits, our interest is only in treatment interventions at $t \leq 3$. We therefore only calculated the risks using those time points. As expected, the Super Learner fit outperformed all of the candidate estimators in the supplied library, as well as the cross-validation selector (i.e. discrete Super Learner, which is equivalent to choosing the single algorithm in the library with the lowest cross validated risk). Compared to the mean model, which assumes no confounding and does not control for any confounders, the Super Learner fits for the LREC availability and end of study mechanisms showed an immense decrease in cross-validated risk. This gain was also noticeable when compared to the candidate model that only

controls for time. The Super Learner fit for the enrollment mechanism also outperformed the mean model, though to a smaller degree. No noticeable gain was seen in the transfer mechanism, presumably due to the extremely low number of transfers observed (218). We did not present the cross-validated risks for the ICE fits as they are too numerous to describe in detail, though note that they were similar to fits for the treatment and censoring mechanisms.

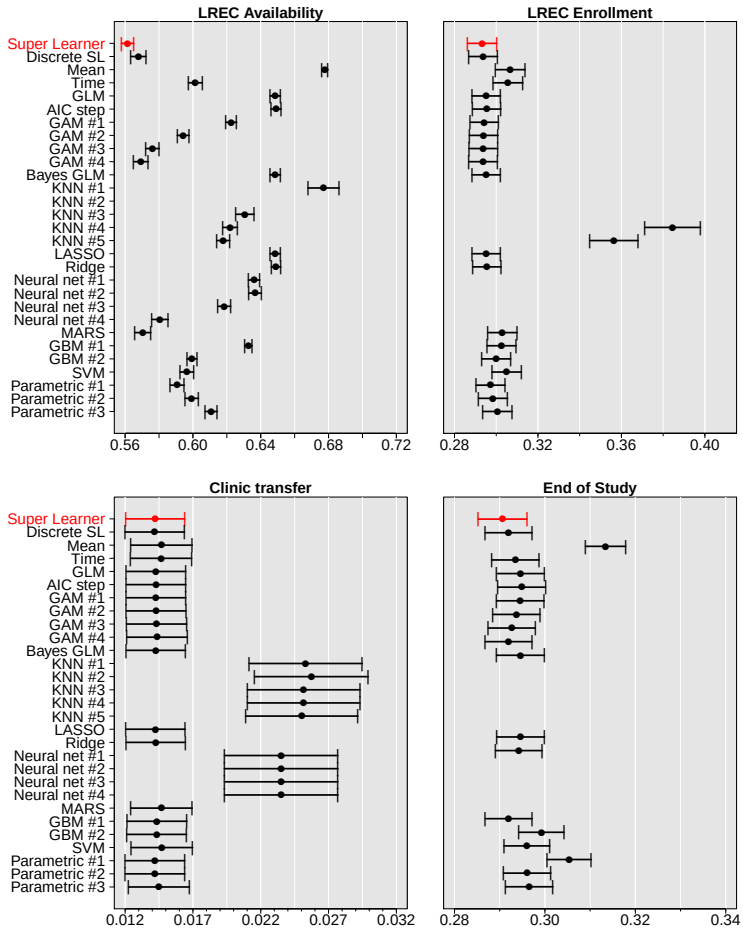

Figure 2.3: Cross-valided risk estimates (using non-negative log-likelihood loss) and 95% confidence interval for the treatment and censoring mechanisms. A number of candidates had cross-validated risks too high to plot in the specified window and consequently are not shown. Mean: marginal probability, Time: logistic regression with time variable only, GLM: logistic regression with all confounders, AIC step: stepwise regression using the Akaike information criterion, GAM: generalized additive model, KNN: k-nearest neighbors, LASSO: least absolute shrinkage and selection operator, MARS: multivariate adaptive regression splines, GBM: generalized boosted regression models, SVM: support vector machines, Parametric: user specified logistic models using a subset of the confounders.

Adjustment for potential confounders using the Super Learner fits resulted in relatively small updates of the survival curves (Figure 2.4), i.e. $1-\hat{\Psi}(P_n)$. Subjects enrolling immediately into the LREC program at eligibility continued to have the highest survival probabilities, while those with immediate availability not enrolling had the lowest. Tables 2.2 and 2.3 show the calculated average treatment effects between the different interventions. Confidence intervals and p-values were calculated based on influence functions. As implied by the survival curves, immediate enrollment into the LREC program at eligibility had a beneficial effect relative to never having LREC available, while having LREC immediately available and never enrolling was adverse. For example, at $t = 4$ the probability of survival for subjects with immediate enrollment was 0.93 (95% CI: 0.91, 0.95) and 0.87 (95% CI: 0.86, 0.87) for subjects with immediate availability never enrolling. For subjects without LREC availability, it was 0.91 (95% CI: 0.90, 0.92). Similar to the unadjusted estimates, the treatment effects increased with time. All estimates after $t = 1$ showed statistical significance.

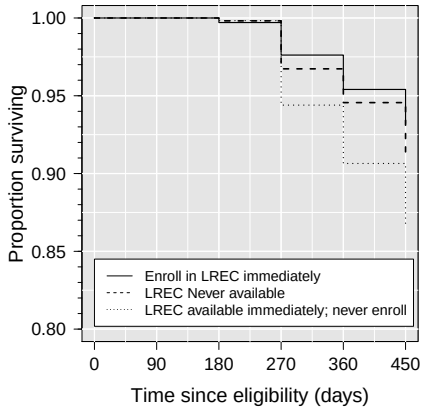

Figure 2.4: Survival curves adjusting for potential confounders.

It is possible that near positivity violations can have large effects on estimates of our parameter. To test for this potential issue, we considered different truncation bounds for our treatment probabilities. Specifically, we considered using a bound of 0.001 and using untruncated probabilities. No differences were seen in the resulting mean outcome estimates.

2.4 Discussion

We have presented a comprehensive approach to applying longitudinal targeted minimum loss-based estimation to evaluate the impact of the LREC program. The results support a somewhat negligible impact of implementation and enrollment, with the lowest survival among patients with immediate LREC availability never enrolling and similar survival among the other two interventions (Figure 2.4). Subjects enrolling immediately into the LREC program have almost identical

Table 2.2: Unadjusted time-specific average treatment effects. (a) compares the intervention immediate LREC availability without enrollment to never having LREC available; (b) compares the intervention immediate LREC availability and enrollment to immediate LREC availability without enrollment.

	(a) $\mathbb{E}Y_{\bar{a}=10}(t^*) - \mathbb{E}Y_{\bar{a}=00}(t^*)$			(b) $\mathbb{E}Y_{\bar{a}=11}(t^*) - \mathbb{E}Y_{\bar{a}=10}(t^*)$		
Time (t^*)	Estimate	(95% CI)	p-value	Estimate	(95% CI)	p-value
1	0.00	(0.00,0.00)	0.93	0.00	(0.00,0.00)	0.56
2	0.03	(0.02,0.03)	0.00	-0.04	(-0.05,-0.03)	0.00
3	0.04	(0.03,0.05)	0.00	-0.05	(-0.07,-0.04)	0.00
4	0.06	(0.05,0.08)	0.00	-0.07	(-0.09,-0.05)	0.00

Table 2.3: Time-specific average treatment effects adjusted for measured potential confounders. (a) compares the intervention immediate LREC availability without enrollment to never having LREC available; (b) compares the intervention immediate LREC availability and enrollment to immediate LREC availability without enrollment.

	(a) $\mathbb{E}Y_{\bar{a}=10}(t^*) - \mathbb{E}Y_{\bar{a}=00}(t^*)$			(b) $\mathbb{E}Y_{\bar{a}=11}(t^*) - \mathbb{E}Y_{\bar{a}=10}(t^*)$		
Time (t^*)	Estimate	(95% CI)	p-value	Estimate	(95% CI)	p-value
1	0.00	(0.00,0.00)	0.81	0.00	(0.00,0.01)	0.58
2	0.02	(0.02,0.03)	0.00	-0.03	(-0.05,-0.02)	0.00
3	0.04	(0.03,0.05)	0.00	-0.05	(-0.07,-0.03)	0.00
4	0.04	(0.03,0.06)	0.00	-0.07	(-0.08,-0.05)	0.00

survival to subjects never being exposed to the program. While the magnitude of difference in survival increased with time, this difference is modest.

It is important to note that our target population is comprised only of subjects at low risk of mortality. Consequently, the majority of our results are driven primarily by subjects not remaining in care, as the number of deaths expected to be observed will be low. In our study, only 104 of the total 1,206 failure events were from deaths. A sensitivity analysis using only loss to follow-up as the outcome resulted in similar estimates.

We chose 90-day intervals for our time points in the current study, due to the understanding that patients would have visits approximately every 3-months. While smaller intervals could have been chosen, doing so can reduce the probability of following a given regime of interest given observed covariates (i.e. increase the extent of practical positivity violations), both by decreasing the probability of availability and enrollment occurring in the first interval, and because the probability of never enrolling given observed covariates involves taking a cumulative probability of not enrolling given the observed past over many more time points. Furthermore, the use of smaller intervals results in more time points, leading to higher computational costs. On the other hand, the use of larger intervals leads to discarding information in order to preserve time ordering, which can result

in a less complete control for confounding as well as failure to capture the full causal effect of the intervention. In order to preserve time ordering, only covariate and outcome values measured at the end of the prior interval are considered possible causes of enrollment and availability in an interval. Longer intervals result in more problems with the assumption. We tested whether there was an effect in our study by re-running the analyses using 30-day intervals. The resulting survival estimates were similar to the ones reported here.

As with all studies, there are limitations that need to be considered. Firstly, it is possible that we did not sufficiently adjust for all the potential confounders. For example, the majority of subjects who had immediate availability and never enrolled had initial eligibility occur after the LREC program had already started. These subjects experiencing incidental eligibility (as opposed to prevalent eligibility from those eligible prior to the LREC program initiation) may have had factors placing them at higher risk. In addition, in defining our composite outcome "dead or lost to follow up" we implicitly assumed that not being seen in clinic for 6.5 months is an undesirable outcome reflecting out of care status. In practice, some of these patients might represent unreported transfers to care in an alternative clinic. If true, however, this would have to occur disproportionately among treatment groups in order to affect the average treatment effect estimates presented here. Lastly, our analysis considered subjects from the same clinics to be causally independent of each other. In specifying our causal model we made a key decision to use an individual level NPSEM despite our interest in both an individual and a clinic level exposure variable. Such a formulation assumes that individuals within a given clinic are causally independent, and in particular, that the exposure received by one patient does not impact the outcome of another (the assumption of no causal interference) [39, 103]. A different formulation is possible that uses a hierarchical or clinic level NPSEM and corresponding hierarchical identification and analysis. We can think of the corresponding experiment as randomizing entire clinics to start the LREC program and within clinics with LREC available, randomizing patients to enroll. However, the sample size then becomes driven by the number of clinics and identification would require adequate variability in the introduction of LREC across clinics [103]. We therefore pursued an individual level formulation, while noting the limitations of this approach. Future research into improved approaches to interference effects in this setting should be undertaken.

We end by stating that, while not conducted here, this framework can be easily generalized to include dynamic interventions that are dependent upon other covariates. For example, there could be interest in intervening to enforce enrollment only on patients who retain eligibility during follow-up while exempting patients who do not. Another option to consider is the use of marginal structural models to smooth treatment effects across time points, as well as availability and enrollment times [84, 77, 75, 65] though care should be taken when implementing as the number of regimes with available data would be limited. These models allow us to project the true underlying dose response curve onto a working parametric model, allowing us to conduct inference on a smaller set of parameters. The `ltmle` package includes a TMLE for causal parameters defined as the projection of the survival or failure curve onto user-specified parametric models.

In summary we applied the targeted learning road map to longitudinal data with a multilevel longitudinal treatment of interest to analyze a nurse-based triage system among HIV patients in East Africa. This included both definition and identification of our causal parameter. Issues with

positivity were handled with careful selection of our target causal parameter. Nuisance parameters were estimated using Super Learner, a cross-validation ensemble algorithm using both parametric and machine learning algorithms. Observations for the estimation of the treatment mechanisms were pooled across time points, which aided us in estimating the censoring mechanism due to clinical transfers. Various approaches were implemented aimed at ensuring the machine learning estimates of the ICEs would respect the underlying statistical model. Estimates of survival at each time point were then contrasted by their differences and inference derived using the empirical influence functions. The results show a somewhat negligible impact of both availability and enrollment in the LREC program on in-care survival.

Chapter 3

Double robust efficient estimators of longitudinal treatment effects

In settings where the exposure of interest is longitudinal, or in other words, is comprised of multiple treatment decisions over time, identification of causal parameters requires non-traditional statistical estimands, such as that provided by the longitudinal g-computation formula under an assumption of sequential randomization [71, 76]. Many estimators have been developed that target these estimands. Available estimators differ in their efficiency, in the nuisance parameters they require estimators for, and in their robustness and consistency of these estimands. For example, the "classical" longitudinal g computation approach [71] relies on consistent estimation of a series of conditional densities while, inverse probability weighted estimators [35] rely on consistent estimation of the treatment mechanism.

A number of these estimators are double robust (DR). In other words, they have the appealing property that if either of the two nuisance parameters are estimated consistently, then the resulting estimator will be consistent for the true parameter value. They are also efficient in a semi-parametric statistical model that makes assumptions, if any, only on the exposure assignment mechanism if both nuisance parameters are estimated consistently at reasonable rates [48, 107, 49]. Such DR semi-parametric efficient estimators include, among others, those based on estimating equations [81, 87], sequential regression approaches based on iterated conditional expectations [93, 75, 3], and TMLE [46, 65, 95].

Important differences also exist within the class of double robust efficient estimators. First, a subset of estimators in this class are based on an alternative iterative conditional expectation representation of the longitudinal g-computation formula [93, 75, 3], an approach that can improve performance by dramatically reducing the dimensionality of one of the two nuisance parameters. We focus here on this class of DR estimators. Second, DR efficient estimators may be defined either as solutions to an estimating equation, or instead as a substitution estimator. The latter approach can improve stability in the face of data sparsity (near positivity violations [32, 10, 64]) and ensure that resulting estimates respect the parameter space. Finally, as we discuss, existing DR estimators differ in their ability to fully incorporate machine learning approaches to nuisance parameter estimation while maintaining the basis of statistical inference.

While a number of sophisticated estimators of longitudinal effects have been proposed, there is a relative paucity of research comparing these methods directly to one another. Here, we compare these various approaches to estimating a causal effect in a longitudinal treatment setting using both simulated data and using the IeDEA-EA cohort.

Six distinct estimators are considered here. They are (i) a simple substitution estimator, based on estimating the iterated conditional expectation (ICE) representation of the longitudinal g-computation formula [75, 3], (ii) an inverse propensity weighted (IPW) method [35, 74, 32], (iii) an augmented IPW (AIPW) estimator that directly solves the estimating equation corresponding to the efficient influence function (EIF) [79, 81, 83, 93, 80, 75], (iv) a double robust iterated conditional expectation (DRICE) estimator as presented by Scharfstein et. al. [93], Robins [75], and Bang and Robins [3], (v) a modified version of DRICE presented by Robins et. al. [82], in which we apply the estimated inverse probabilities as observational weights, and (vi) a targeted minimum loss-based estimator as presented by van der Laan and Gruber [46]. AIPW, DRICE, and TMLE all solve the estimating equation corresponding to the EIF and are therefore efficient estimators. However, due to the different manners in which they solve the estimating equation, finite sample performance may differ.

This chapter builds on the existing comparative methods literature in several important ways. First, much prior work has focused on the point treatment setting, in which both the statistical estimands and corresponding estimators are substantially simpler than their longitudinal counterparts [36, 64, 6, 26, 65, 14, 118, 51, 30, 98, 24]. For example, Decker et. al. [14] compared the ICE and TMLE methods in estimating the causal effect of physical activity and diet on body mass index. Petersen et. al. [64] compared the IPW and TMLE estimators using simulated data in a point treatment setting with positivity violations. Second, most of the prior work comparing the performance of estimators of longitudinal treatment effects has generally been limited to comparison of a single double robust efficient estimator (such as TMLE or AIPW) to a simpler alternative such as IPW or g-computation. To the best of our knowledge, this is the first longitudinal study to directly compare the performance of a range of proposed double robust estimators in a head to head comparison and incorporating machine learning. Furthermore, this study directly compares the performance of the estimators under gradually increasing levels of positivity violations [10, 114, 64]. Lastly, for the applied setting, it considers estimates from both the generalized linear modelling approach and by incorporating Super Learning [47], an ensemble data adaptive machine learning algorithm.

3.1 Data

To reduce the complexity of the analyses in this chapter, we condition on LREC availability. That is, we assume that the LREC program has already been initiated in each of the 15 clinics and analyze the effect of program enrollment. Consider, as a working example, right censored survival data in which subjects are followed from a baseline time point $t = 0$ up to some final time point $K+1$. At each time point t, subjects may enroll into a treatment program. Regardless of whether they enroll, each subject is followed until the first of either (i) some terminal event of interest is observed, (ii) they are right censored due to transfers or administrative censoring, or (iii) they

reach the end of follow-up ($t = K+1$). Time-varying covariates are measured at each time point t which may affect subsequent treatment covariates and outcome. Additionally, baseline covariates are measured at $t = 0$.

More formally, we consider an independent and identically distributed (iid) statistical data structure. Let $Y(t)$ be a failure indicator, a counting process which takes value 0 until the outcome event of interest is observed and subsequently switches to and remains at 1 for all remaining time points. We assume that $Y(0) = 0$ for everyone, i.e. that no subjects have experienced the event at the beginning of follow-up. Let $L1(t)$ be the time-varying covariate values measured at each time point t. We define $L1(0)$ to additionally include all baseline covariates measured at $t = 0$. Similar to Chapter 1.2, we refer to the joint outcome and covariate variables at time t as $L(t) = (Y(t), L1(t)) : t = 0, 1, \ldots, K+1$. Let $A2(t)$ be the indicator of enrollment into the treatment program and $C1(t)$ be the indicator of censoring due to patient transfers and $C2(t)$ be the indicator of censoring due to data base closure. Each of these variables take value 0 until an enrollment or censoring event is observed, respectively, at which point they become fixed at 1. We collectively refer to the treatment and censoring processes as $A(t) = (A2(t), C1(t), C2(t)) : t = 0, 1, \ldots, K$. For notational convenience, we define all variables after failure or right censoring occurs as deterministically equal to their last observed value. Our updated data structure can, similar to Chapter 1.2, be represented as n independent and identically distributed (iid) copies of the longitudinal data structure

$$O = (L(0), A(0), L(1), A(1), \ldots, A(K), L(K+1)) \overset{iid}{\sim} P_0. \tag{3.1}$$

3.2 Estimators for $\mathbb{E}Y_d(t^*)$

We present each of the five estimators considered below for estimating $\mathbb{E}Y_d(t^*)$.

Iterated conditional expectation estimation (ICE)

As stated in Equation (2.12), our parameter of interest can be represented as a series of iterated conditional expectations. From this representation, as described by Scharfstein et. al. [94] and Robins [75], we can form an estimator which starts by estimating the inner most conditional expectation and iterating outward until reaching the outermost conditional expectation. The parameter estimate $\hat{\psi}_n^{ICE}$ is then just the empirical mean over all the observations. The specific algorithm proceeds as follows:

1. Let T denote the failure time. Estimate the innermost conditional expectation $\bar{Q}^d_{0,L(t^*)} = \mathbb{E}_0[Y(t^*)|\bar{L}(t^*-1), \bar{A}(t^*-1) = d(\bar{l}(t^*-1))]$, where the expectation is known to be equal to 1 if $T < t^*$. We denote this estimate as $\bar{Q}^d_{n,L(t^*)}$.

2. Given $\bar{Q}^d_{n,L(t^*)}$, we can recursively iterate outwards for $t^* = t^*-1, t^*-2, \ldots, 1$, estimating $\bar{Q}^d_{0,L(t)} = \mathbb{E}_0[\bar{Q}^d_{n,L(t^+)}|\bar{L}(t^-), \bar{A}(t^-) = d(\bar{l}(t^-))]$, acknowledging our slight abuse of notation,

where again the expectation is known to be equal to 1 if $T < t$. We denote these estimates as $\bar{Q}^d_{n,L(t)}$.

3. At $t = 1$, we have the estimate $\bar{Q}^d_{n,L(1)}$, which now is a function of only $L(0)$. As indicated in Equation (2.12), our parameter estimate is simply the empirical expectation over $L(0)$, i.e. $\hat{\psi}^{ICE}_n = \mathbb{E}_n \bar{Q}^d_{n,L(1)}$.

As with the established parametric g-computation estimator (e.g. Taubman et. al. [102]), this estimator only relies upon the Q_0 portion of the likelihood in estimating our target parameter. Unlike the parametric g-computation estimator, however, which relies on estimating each of the conditional probability distributions given in the g-computational formula above (Equation 2.4), this estimator relies on only the conditional expectations $\bar{Q}_0$. Consequently, it provides a substantial advantage over the parametric g-computation approach in that it avoids the need to estimate all of the conditional densities.

Inverse propensity weighted estimation (IPW)

Our parameter can also be estimated by up-weighting subjects from strata of $L1(t)$ that are underrepresented compared to the representation they would have had under a randomized treatment assignment [35, 74]. This approach can be understood as creating a *pseudo*-population in which the measured covariates are balanced between treatment groups [31]. More formally, we implement the following estimator [74]:

$$\hat{\psi}^{HT}_n = \frac{\frac{1}{n}\sum_{i=1}^{n} \frac{\mathbb{I}(\bar{A}_i(t^*-1)=d(\bar{l}(t^*-1)))}{g_{n,0:t^*-1,i}(d(\bar{l}(t^*-1)))} Y_i(t^*)}{\frac{1}{n}\sum_{i=1}^{n} \frac{\mathbb{I}(\bar{A}_i(t^*-1)=d(\bar{l}(t^*-1)))}{g_{n,0:t^*-1,i}(d(\bar{l}(t^*-1)))}} \tag{3.2}$$

where $g_{n,0:t^*-1,i}(d(\bar{l}(t^*-1))) = \prod_{k=0}^{t^*-1} P_n(A_i(k) = d(\bar{l}(k))|\bar{L}_i(k), \bar{A}_i(k^-) = d(\bar{l}(k^-)))$. By using this inverse weighting, this approach relies upon the consistent estimation of g_0 for proper inference.

Augmented inverse probability weighted estimation (AIPW)

Realizing that the EIF in Equation (2.7) has mean zero and is a function of our target parameter [25, 107], we can straight forwardly form an estimating equation and solve for our parameter [83, 93]. This naturally results in the estimating equation estimator

$$\begin{aligned}\hat{\psi}^{AIPW}_n &= \mathbb{E}_n D^*(Q_n, g_n)(O_i) + \Psi(\bar{Q}^d_n) \\ &= \mathbb{E}_n \left[\sum_{t=1}^{t^*} \frac{\mathbb{I}(\bar{A}_i(t^-) = d(\bar{l}(t^-)))}{g_{n,0:t^-,i}(d(\bar{l}(t^-)))} \left(\bar{Q}^d_{n,L(t^+),i} - \bar{Q}^d_{n,L(t),i} \right) + \bar{Q}^d_{n,L(1),i} \right]\end{aligned} \tag{3.3}$$

where again, for notational convenience, we use $\bar{Q}^d_{n,L(K+2),i}$ to denote $Y_i(t^*)$.

Double robust iterated conditional expectation (DRICE)

As shown by Bang and Robins [3], we can form a DR estimator as a sequential regression estimator quite similar to the ICE approach from Chapter 3.2. This approach, however, additionally uses the inverse propensity estimate $g_{n,0:t^-}(d(\bar{l}(t^-)))^{-1}$ times the indicator of following treatment $\mathbb{I}(\bar{A}_i(t^-) = d(\bar{l}(t^-)))$ as a covariate in estimating $\bar{Q}^d_{0,L(t)}$. Thus, our outcome regression for each time point t can instead be represented as a function of the observed past $(L(t^-),A(t^-))$, the indicator of following the regime of interest $\mathbb{I}(\bar{A}_i(t^-) = d(\bar{l}(t^-)))$, and our probability of treatment $g_{0,0:t^-}(d(\bar{l}(t^-)))$. Our iterated conditional expectation algorithm is therefore updated as follows:

1. Estimate $g_{0,0:t^-}(d(\bar{l}(t^-))) : t = t^*, t^*-1, \ldots, 0$. We denote the estimates as $g_{n,0:t^-}(d(\bar{l}(t^-)))$.

2. Let T denote the failure time. Acknowledging our slight abuse of notation, we estimate the innermost conditional expectation $\bar{Q}^{d,g}_{0,L(t^*)} = \mathbb{E}_0[Y(t^*)|\bar{L}(t^*-1), \bar{A}(t^*-1) = d(\bar{l}(t^*-1)), \mathbb{I}(\bar{A}_i(t^*-1) = d(\bar{l}(t^*-1))) \times g_{n,0:t^*-1}(d(\bar{l}(t^*-1)))^{-1}]$, where the expectation is known to be equal to 1 if $T < t^*$. We denote this estimate as $\bar{Q}^{d,g}_{n,L(t^*)}$.

3. Given $\bar{Q}^{d,g}_{n,L(t^*)}$, we can recursively iterate outwards for $t = t^*, t^*-1, \ldots, 1$, estimating $\bar{Q}^{d,g}_{0,L(t)} = \mathbb{E}_0[\bar{Q}^{d,g}_{n,L(t^+)}|\bar{L}(t^-), \bar{A}(t^-) = d(\bar{l}(t^-)), \mathbb{I}(\bar{A}_i(t^-) = d(\bar{l}(t^-))) \times g_{n,0:t^-}(d(\bar{l}(t^-)))^{-1}]$. The expectation is also known to be equal to 1 if $T < t$.

4. At $t = 1$, we have the estimate $\bar{Q}^{d,g}_{n,L(1)}$, which now is a function of only $L(0)$. As indicated in Equation (2.12), our parameter estimate is simply the empirical expectation over $L(0)$, i.e. $\hat{\psi}^{DRICE}_n = \mathbb{E}_n \bar{Q}^{d,g}_{n,L(1)}$.

Bang and Robins [3] have shown that this estimator is algebraically the same as the AIPW estimator, since

$$\sum_i g_{n,0:t^*-1}(d(\bar{l}(t^*-1)))^{-1}\mathbb{I}(\bar{A}_i(t^*-1) = d(\bar{l}(t^*-1)))(Y_i - \bar{Q}^d_{n,L(1)}(L_i(1)|Pa(L_i(1)))) = 0 \quad (3.4)$$

where $\bar{Q}^d_{n,L(1)}(L_i(1)|Pa(L_i(1)))$ is instead modelled as a link function applied to the sum of $\bar{Q}^d_0$ and $g_0(d(\bar{l}))$ times a parameter ϕ. While this holds in the case where the outcome is continuous and the link function is the identity link, we note that applying this in a binary outcome setting where the link function is the logit link creates a separate bounded estimator [82] which enforces an extra restriction to the estimation procedure not present for the AIPW estimator. The supremum and infimum of transformed values under this link can never fall outside of $[0,1]$ and, consequently, the approach is converted into a substitution based estimation procedure. This extra restriction is beneficial for this setting in that it ensures that our resultant parameter estimates will always respect the parameter space boundaries. We compare this approach to the AIPW approach in our simulations and applications.

The modified DRICE

Rather than using the inverse propensity estimates $g_{n,0:t^-}^{d(\bar{l}(t^-))}$ as part of a covariate as presented by Bang and Robins [3], an alternative approach could be to instead apply the inverse of $g_{n,0:t^-}^{d(\bar{l}(t^-))}$ as observational weights resulting in a *pseudo*-population and estimating $\bar{Q}^d_{0,L(t)}$ in the same manner as the approach above but instead using $\mathbb{I}(\bar{A}_i(t^*-1) = d(\bar{l}(t^*-1)))$ as a covariate [82]. This approach is also double robust in the sense that consistent estimation of the target parameter will be achieved if either the $g_{0,0:t^-}(d(\bar{l}(t^-)))$ or $\bar{Q}^d_{0,L(t)}$ for all t are estimated consistently.

The approach of applying the inverse propensity estimates $g_{n,0:t^-}(d(\bar{l}(t^-)))$ as weights can potentially aid us in the presence of positivity violations, since small values of $g_{n,0:t^-}(d(\bar{l}(t^-)))$ can potentially have large effects on the estimates of $\bar{Q}^d_{0,L(t)}$ by acting as an outlier if used as a covariate. Indeed, Robins et. al. [82] found that this approach resulted in lower mean squared error for their simulations. Therefore, we also considered this version of the estimator for this study and refer to it as mDRICE.

It is important to note here that the DRICE and modified DRICE are efficient only in linear models as a consequence of the condition that the normal equations be solved for the DRICE and AIPTW estimators to be equivalent [3]. Conversely, the TMLE and AIPTW estimators have no such restriction and are therefore efficient under the larger semi-parametric statistical model $\mathscr{M}$.

Targeted minimum loss-based estimation (TMLE)

van der Laan and Gruber [46] present an alternative ICE substitution based approach to solving the EIF, thereby retaining both double robustness and asymptotic efficiency while simultaneously respecting the parameter space and guaranteeing unique solutions. The targeted minimum loss-based estimator requires a number of ingredients, including (i) the EIF $D^*(Q,g)(O)$ defined above, (ii) a generalized loss function possibly indexed by a nuisance parameter $\mathscr{L}_{t,\bar{Q}^d_{L(t^+)}}(\bar{Q}^d_{L(t)})$, (iii) a least favorable parametric submodel $\bar{Q}^d_{L(t)}(\varepsilon_t)$ chosen such that the linear span of the generalized score at zero fluctuation spans the EIF, and (iv) an updating algorithm which iteratively minimizes the generalized loss-based empirical risk over the parameters of the least favorable parametric submodel. Once these ingredients are collected, the algorithm for estimation is as follows:

1. Estimate $g_{0,0:t^-}(d(\bar{l}(t^-))) : t = t^*, t^*-1, \ldots, 1$. We denote the estimates as $g_{n,0:t^-}(d(\bar{l}(t^-)))$.
2. Let T denote the failure time. Estimate $\bar{Q}^d_{0,L(t^*)} = \mathbb{E}_0[Y(t^*)|\bar{L}(t^*-1), \bar{A}(t^*-1) = d(\bar{l}(t^*-1))]$, where the expectation is known to be equal to 1 if $T < t^*$. We denote this estimate as $\bar{Q}^d_{n,L(t^*)}$.
3. Update the initial fit $\bar{Q}^d_{n,L(t^*)}$ based on the t^*-th loss function $\mathscr{L}_{t^*,\bar{Q}^d_{L(t^*+1)}}(\bar{Q}^d_{L(t^*)}(\varepsilon_{t^*}))$ and using the submodel $\bar{Q}^d_{L(t^*)}(\varepsilon_{t^*})$. By setting $\hat{\varepsilon}_{n,t^*} = \arg\min_{\varepsilon} P_n \mathscr{L}_{t^*,\bar{Q}^d_{L(t^*+1)}}(\bar{Q}^d_{L(t^*)}(\varepsilon))$, an updated fit is formed $\bar{Q}^{d,*}_{n,L(t^*)} = \bar{Q}^d_{L(t^*)}(\hat{\varepsilon}_{n,t^*})$ that is targeted at the parameter $\Psi(P_0)$.

4. Given $\bar{Q}^{d,*}_{n,L(t^*)}$, we can recursively for $t = t^*, t^*-1, \ldots, 1$:

   a) Estimate the conditional expectation $\bar{Q}^{d}_{0,L(t)} = \mathbb{E}_0[\bar{Q}^{d,*}_{n,L(t^+)} | \bar{L}(t^-), \bar{A}(t^-) = d(\bar{l}(t^-))]$, where again the expectation is known to be equal to 1 if $T < t$. We denote this estimate as $\bar{Q}^{d}_{n,L(t)}$.

   b) Similar to Step 3, update $\bar{Q}^{d}_{n,L(t)}$ using the loss function $\mathscr{L}_{t,\bar{Q}^{d}_{L(t^+)}}(\bar{Q}^{d}_{L(t)})$ with the parametric submodel $\bar{Q}^{d}_{L(t)}(\varepsilon_t)$. Again, minimizing the empirical loss function $\hat{\varepsilon}_{n,t} = \arg\min_{\varepsilon} P_n \mathscr{L}_{t,\bar{Q}^{d}_{L(t^+)}}(\bar{Q}^{d}_{L(t)}(\varepsilon))$ results in the updated fit $\bar{Q}^{d,*}_{n,L(t)} = \bar{Q}^{d}_{L(t)}(\hat{\varepsilon}_{n,t})$ for time t.

5. At $t = 1$, we have the estimate $\bar{Q}^{d,*}_{n,L(1)}$, which now is a function of only $L(0)$. Our parameter estimate is simply the empirical expectation over $L(0)$, i.e. $\hat{\psi}^{TMLE}_n = \mathbb{E}_n \bar{Q}^{d,*}_{n,L(1)}$.

Similar to the DRICE estimator, this approach solves the EIF in a substitution based setting. The primary difference in the two approaches is that TMLE forms the initial fit and subsequently uses the inverse propensity estimates to update this initial fit, resulting in a 2-step approach. Conversely, DRICE solves for the EIF by including the inverse propensity estimates within the initial fit.

Conditioning the nuisance parameter estimation on $d(\bar{l}(t))$

One approach at estimating the nuisance parameters $g_{0,0:t^-}(d(\bar{l}(t^-)))$ and $\bar{Q}^{d}_{0,L(t)}$, as presented by Bang and Robins [3] and van der Laan and Gruber [46], conditions on having followed the treatment regime of interest $d(\bar{l}(t))$ and uses only these subjects to estimate g_0 and $\bar{Q}^{d}_0$.

An alternative approach is to instead pool across all subjects regardless of treatment history and uses these subjects in estimating g_0 and $\bar{Q}^{d}_0$. We refer to this latter approach as the pooled estimator and the approach that conditions on following $d(\bar{l}(t))$ as the stratified estimator. The pooled approach provides potentially more efficiency and stability due to the increased sample size used to estimate each nuisance parameter. We therefore considered the pooled approach for the ICE estimator and both stratified and pooled approaches for the three DR estimators. We further note that if, as in our worked example, $A(t)$ is a counting process that jumps to one once and remains there deterministically, the stratified and pooled approaches to estimating g_0 will be identical (n.b. the deterministic nature of $A(t)$ naturally implies conditioning on prior $A(t-1) = 0$ when estimating g_0). We therefore only considered the stratified approach for the IPW estimator and and compare stratified vs. pooled approaches to estimating $\bar{Q}^{d}_0$ only.

Data adaptive estimation

The above estimators each require initial estimates of either g_0, $\bar{Q}^{d}_0$, or both. Similar to the analyses in Chapter 2, we use Super Learning [47] for this estimation. This algorithm uses V-fold cross

validation to select the best convex combination of conditional density or probability fits within a user specified library of potential candidates. If none of the candidates in the library is a correctly specified parametric model, then it has been proven to perform at least as well asymptotically as an oracle selector that selects the best candidate from the library based on the (unknown in practice) true distribution P_0; otherwise, the Super Learner achieves the almost parametric rate of convergence of $\log(n)/n$.

Most of the estimators considered, with the exception of DRICE and mDRICE, can incorporate machine learning while retaining valid theory based inference. The exceptions are due to the previously stated requirement that linear statistical models be used. This restriction extends to data adaptive estimation as well, such as stepwise regression, as the covariate incorporating the inverse propensity estimate could be eliminated from the analysis resulting in a loss of the DR property. We therefore did not consider the data adaptive approaches for DRICE and mDRICE.

3.3 Simulations

To compare the various approaches presented at estimating our target parameter in a controlled empirical and finite sample setting, we undertook a simulation using a fixed parametric distribution.

Data generating distribution P_0

Our data was generated for times $t = 0, 1, 2, \ldots, 6$ using the data structure described in Chapter 3.1 under a sample size of $n = 500$ as follows:

$$\begin{aligned}
W_1, W_3 &\sim N(0,1) \\
W_2 &\sim \text{Ber}(\text{logit}^{-1}(-1)) \\
Y(t)|Y(t^-)=0 &\sim \text{Ber}(\text{logit}^{-1}(-1.9 + 1.2W_1 - 2.4W_2 - 1.8L1_1(t^-) \\
&\quad - 1.6L1_2(t^-) + L1_1(t^-)L1_2(t^-) - A1(t^-))) \\
L1_1(t)|Y(t^-)=0 &\sim N(0.1 + 0.4W_1 + 0.6L1_1(t^-) - 0.7L1_2(t^-) \\
&\quad - 0.45A1(t^-), 0.5) \\
L1_2(t)|Y(t^-)=0 &\sim N(-0.55 + 0.5W_1 + 0.75W_2 + 0.1L1_1(t^-) \\
&\quad + 0.3L1_2(t^-) - 0.75A1(t^-), 0.5) \\
A1(t)|Y(t^-)=0, A1(t^-)=0 &\sim \text{Ber}(\text{logit}^{-1}(-1 - 1.5W_1 + 1.75W_2 \\
&\quad + 1.2L1_1(t) - 1.8L1_2(t) + 0.8L1_1(t)L1_2(t))))
\end{aligned} \tag{3.5}$$

Thus, we have that $L(t) = (Y(t), L1_1(t), L1_2(t))$ where at time $t = 0$, $L(0) = (W_1, W_2, W_3, Y(0), L1_1(0), L1_2(0))$ and $A(t) = (A1(t))$ for all t. For ease of notation, we defined $L1_1(-1) = L1_2(-1) = A1(-1) = 0$, fixed $Y(0) = 0$, i.e. assumed that everyone was alive at baseline $t = 0$, and discarded treatment and confounder data gathered at $t = 6$ as our last time

point of interest for the outcome is at $t = 6$ and that extra data would only affect survival $Y(t)$ at $t > 6$. We also enforced the restrictions of our model listed in Chapter 2.1. Thus, we have that $Y(t)$ and $A1(t)$ are counting processes. Once a subject experiences a failure, i.e. $Y(t) = 1$, all the remaining values remain at 1. Subjects enrolling, i.e. having $A1(t) = 1$, stayed enrolled for the remainder of follow-up. For simplicity, we did not include censoring in our data generating distribution for this simulation.

Target parameter $\Psi(P_0)$

For the simulation, we considered the intervention of interest to be the static intervention *"never enroll"*, i.e. $d(\bar{l}(t)) = \bar{a}(t) = 0 : t = 0, 1, \ldots, t^*$. Because the needed identifying assumptions hold by design in this simulation, the target causal and statistical parameter values are identical. The target parameter under this data generating distribution P_0 is defined as the statistical estimand corresponding to $\mathbb{E}Y_{\bar{a}}(t^*)$ for $t^* = 1, 2, \ldots, 6$. This parameter can be thought of as the cumulative probability of failure at each time point t^* up to time 6 within a counterfactual population of subjects where no one ever enrolls into the treatment program of interest, i.e. $\overline{A1}(t^*) = 0$.

We determined the true parameter value $\psi_0(t^*) : t^* = 1, 2, \ldots, 6$ by drawing observations under the post-intervention data generating distribution P_0^d with a sample size of $8x10^6$, where data were generated according to Equation (3.5) but setting $\overline{A1}(t^*) = 0$. Defining our true parameter as $\psi_0 \equiv (\psi_0(1), \psi_0(2), \ldots, \psi_0(6))$, this resulted in the true parameter value

$$\psi_0 \approx (0.232, 0.335, 0.390, 0.428, 0.460, 0.489). \tag{3.6}$$

Time dependent confounding was defined such that failing to adjust for any confounders would result in an underestimation of $\psi_0(t^*)$ for low values of t^* and overestimation for high values of t^*.

Positivity

Under the specified distribution P_0, the degree of practical positivity violations increases with t. Figure 3.1 shows the marginal densities of $g_{0,0:t^*-1}(\bar{a}(t^* - 1) = 0) : t^* = 1, 2, \ldots, 6$ for each final time point t^* under $\bar{a}(t^* - 1) = 0$, allowing us to obtain a sense of the severity of the violations. Because $g_{0,0:t^*-1}(\bar{a}(t^* - 1)) : t^* = 1, 2, \ldots, 6$ are functions of $\bar{L}(t^* - 1)$, the marginal densities were derived by taking the conditional probabilities marginally over the distributions of $\bar{L}(t^* - 1)$ for each final time point t^*. The density plot for time $t^* = 1$ shows a somewhat uniform distribution for $g_{0,0:0}(\bar{a} = 0)$, with only 4% of the marginal distribution below 0.01. As t^* increases, however, the cumulative probability of remaining unenrolled, i.e. having $\bar{a}(t^* - 1) = 0$, decreases significantly. For example, 40% of the marginal distribution $g_{0,0:5}(\bar{a}(5) = 0)$ for time $t^* = 6$ was below 0.01. The resulting distributions become increasingly concentrated close to 0, indicating that the probability of remaining unenrolled is near 0 at later time points. Indeed, we saw in the realized simulations that on average, only 6% of observations were still unenrolled and at risk of failure at $t = 5$.

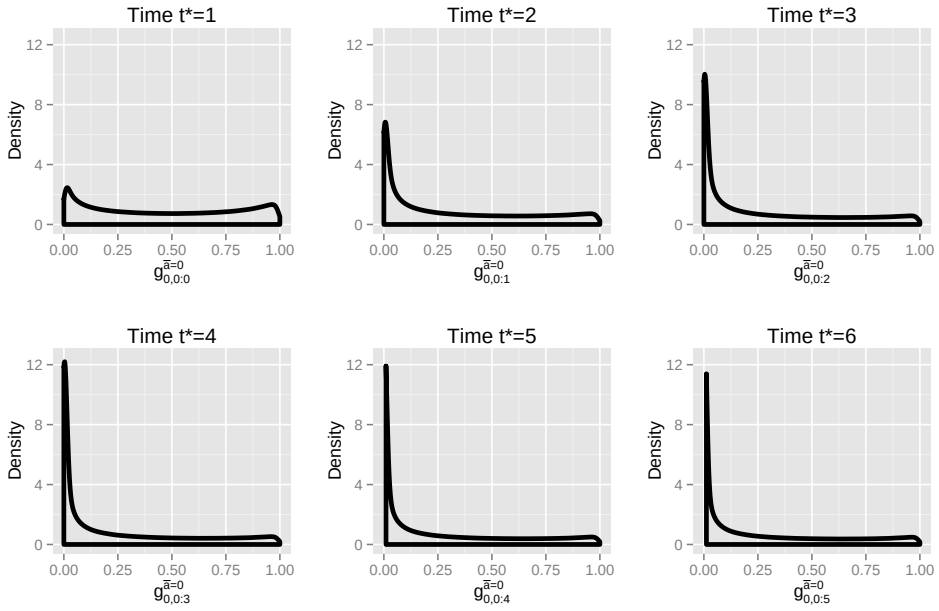

Figure 3.1: Marginal densities of $g_{0,0:t^*-1}(\bar{a}(t^*-1)=0)$ for each time point t^*.

3.4 Estimator practical implementations

We provide summaries below of how each estimator presented above is implemented in our simulation. For ease of presentation, we only cover estimation of our target parameter for the final time point of interest, i.e. $t^* = K+1 = 6$, and note that the same approaches are taken for the earlier time points as well.

ICE

We first condition our data on survival up to the penultimate time point t^*-1, i.e. $t=5$, as well as optionally conditioning on the subset of subjects having $\bar{A}(t^*-1) = d(\bar{l}(t^*-1))$ (depending on whether we are using the stratified versus pooled approach described in Chapter 3.2), where as stated in Chapter 3.3, $\bar{a}(t^*-1)=0$. With this subset, we carry out a logistic regression, regressing $Y(t^*)$ onto $\bar{A}(t^*-1)$ and $\bar{L}(t^*-1)$. Specifically, for this simulation we used the following correctly specified logistic model $\bar{Q}_{L(t^*)}$

$$\begin{aligned}\mathbb{E}[Y(t^*)|\bar{L}(t^*-1),\bar{A}1(t^*-1)] = \text{logit}^{-1}[&\beta_0+\beta_1 W_1+\beta_2 W_2+\beta_3 L1_1(t^*-1)+\beta_4 L1_2(t^*-1)\\ &+\beta_5 L1_1(t^*-1)L1_2(t^*-1)+\beta_6 A1(t^*-1)]\end{aligned} \tag{3.7}$$

This logistic regression provides us with estimates $\hat{\beta} = (\hat{\beta}_0, \hat{\beta}_1, \ldots, \hat{\beta}_6)$, resulting in a fitted object $\bar{Q}_{n,L(t^*)}$ for time $t^*=6$. Under the stratified approach the coefficient on the $A1(t)$ covariate in

the logistic model is absorbed into the intercept. With this fit, we estimate $\bar{Q}^d_{0,L(t^*)}$ for each subject i by setting $\overline{A1}(t^*-1)=0$ for all subjects and calculating estimates under the fitted object $\bar{Q}_{n,L(t^*)}$ for everyone, where the conditional expectation is known to be equal to 1 if $T_i < t^*$. We iterate this by now conditioning on survival up to time $t=4$ and (optionally) on the subset of subjects having $\bar{A1}(t^*-2) = d(\bar{l}(t^*-2))$, using the estimates $\bar{Q}^d_{n,L(t^*),i}$ as the outcome, and regressing it onto $\bar{A}(t^*-2)$ and $\bar{L}(t^*-2)$ using the logistic model from Equation (3.7) but replacing the predictors for time t^*-1 with those of time t^*-2. This gives us a fitted object $\bar{Q}_{n,L(t)}$ for time $t=5$, which we use to estimate $\bar{Q}^d_{0,L(K)}$ for each subject i by setting $\overline{A1}(t^*-2)=0$ for everyone and again calculating estimates under the fitted object, again setting it equal to 1 if $T_i < 5$. We continue iterating these steps backwards over t until we reach time $t=1$, at which we will have the estimate $\bar{Q}^d_{n,L(1)}$ as a function of only $L(0)$. The target parameter is then estimated by taking the empirical mean over the sample, $\hat{\psi}^{ICE}_n = \frac{1}{n}\sum_{i=1}^n \bar{Q}^d_{n,L(1),i}$.

We also considered the performance of this estimator using the same algorithm above, but under the following mis-specified logistic model (acknowledging a slight abuse of notation):

$$\mathbb{E}[Y(t^*)|\bar{L}(t^*-1),\bar{A}1(t^*-1)] = \text{logit}^-[\beta_0+\beta_1 W_2+\beta_2 L1_1(t^*-1)+\beta_3 A1(t^*-1)] \tag{3.8}$$

AIPW

In implementing this estimator, g_0 was estimated as for the IPW estimator, using the same correctly specified and mis-sepcifed parametric models. In contrast to the IPW estimator which only requires an estimate of each subject's cumulative probability of never enrolling through time t^*-1 (i.e. $g_{n,0:t^*-1}(d(\bar{l}=0))$), for the AIPW estimator, we compute $g_{n,0:t}(d(\bar{l}(t)))$ for $t=0,1,\ldots,t^*$, i.e. each subject i's predicted cumulative probability of not enrolling up to each time point t.

We first condition our data on survival up to the penultimate time point t^*-1, i.e. $t=5$, as well as optionally conditioning on the subset of subjects having $\bar{A}(t^*-1)=d(\bar{l}(t^*-1))$ (depending on whether we are using the stratified versus pooled approach described in Chapter 3.2), where as stated in Chapter 3.3, $\bar{a}(t^*-1)=0$. We use this subset to estimate $\bar{Q}^d_{0,L(t^*)}$ by carrying out a logistic regression of $Y(t^*)$ on $\bar{L}(t^*-1)$. Both the correctly specified model from Equation (3.7) and the mis-specified model from Equation (3.8) were considered (where under the stratified approach the coefficient on the $A1(t)$ covariate in the logistic model is absorbed into the intercept). The resulting logistic regression fit $\bar{Q}^d_{n,L(t^*)}$ is then evaluated for each subject i in the study setting the predicted value equal to 1 if $T_i < t^*$. This allows us to evaluate the functional

$$D^*_{t^*}(Q_n,g_n)(O_i) = \frac{\mathbb{I}(\bar{A}_i(t^*-1)=d(\bar{l}(t^*-1)))}{g_{n,0:t^*-1,i}(d(\bar{l}(t^*-1)))}\left(Y_i(t^*)-\bar{Q}^d_{n,L(t^*),i}\right) : i=1,2,\ldots,n.$$

We iterate this by now conditioning on survival up to time $t=4$ and (optionally) on the subset of subjects having $\bar{A}(t^*-2)=d(\bar{l}(t^*-2))$, using the estimates $\bar{Q}^d_{n,L(t^*),i}$ as the outcome, and regressing this outcome onto $\bar{L}(t^*-2)$ using the logistic model from Equation (3.7) or (3.8) but replacing

the predictors for time t^*-1 with those of time t^*-2. This gives us a fitted object $\bar{Q}^d_{n,L(t)}$ for time $t=5$, which we evaluate for each subject i setting the predicted value equal to 1 if $T_i < 5$. The fit is then used to evaluate

$$D^*_{t^*-1}(Q_n,g_n)(O_i) = \frac{\mathbb{I}(\bar{A}_i(t^*-2) = d(\bar{l}(t^*-2)))}{g_{n,0:t^*-2,i}(d(\bar{l}(t^*-2)))}\left(\bar{Q}^d_{n,L(t^*),i} - \bar{Q}^d_{n,L(t^*-1),i}\right) : i = 1,2,\ldots,n.$$

This procedure is iterated until all of $D^*_t(Q_n,g_n)(O_i) : t = 1,\ldots,t^*$ are evaluated and $\bar{Q}^d_{n,L(1),i} : i = 1,2,\ldots,n$ is estimated. The target parameter is then estimated by simply taking the empirical mean of the sum of $\bar{Q}^d_{n,L(1),i}$ and $D^*_t(Q_n,g_n)(O_i)$ over t, i.e.

$$\hat{\psi}^{AIPW}_n = \frac{1}{n}\sum_{i=1}^{n}\left(\bar{Q}^d_{n,L(1),i} + \sum_{t=1}^{t^*} D^*_t(Q_n,g_n)(O_i)\right).$$

DRICE

This estimator is very similar to the ICE approach, with the added step of using the inverse estimates $g_{n,0:t^-}(d(\bar{l}(t^-)))^{-1}$ times the indicator $\mathbb{I}(\bar{A}_i(t^-) = d(\bar{l}(t^-)))$ as an additional predictor in estimating $\bar{Q}^d_{0,L(t)} : t = 1,2,\ldots,t^*$. Similar to Chapter 3.4, we first compute $g_{n,0:t^-}(d(\bar{l}(t^-)))$ for $t = 1,2,\ldots,t^*$.

We then take the subset of subjects conditioned on survival up to the penultimate time point t^*-1; if using the stratified approach we further conditioned on having $\bar{A}(t^*-1) = d(\bar{l}(t^*-1))$. With this subset, we carry out a logistic regression, regressing $Y(t^*)$ on $\bar{A}(t^*-1)$, $\bar{L}(t^*-1)$, and $\mathbb{I}(\bar{A}(t^*-1) = d(\bar{l}(t^*-1)))/g_{n,0:t^*-1}(d(\bar{a}(t^*-1)))$. Thus, rather than just using Equations (3.7) and (3.8), we additionally regressed on $\mathbb{I}(\bar{A}(t^*-1) = d(\bar{l}(t^*-1)))/g_{n,0:t^*-1}(d(\bar{l}(t^*-1)))$, such that the correctly specified and mis-specified models respectively are

$$\begin{aligned}
\mathbb{E}[Y(t^*)|\bar{L}(t^*-1),\bar{A}1(t^*-1)] &= \text{logit}^{-1}[\beta_0 + \beta_1 W_1 + \beta_2 W_2 + \beta_3 L1_1(t^*-1) + \beta_4 L1_2(t^*-1) \\
&\quad + \beta_5 L1_1(t^*-1)L1_2(t^*-1) + \beta_6 A1(t^*-1) \\
&\quad + \beta_7 \frac{\mathbb{I}(\bar{A}(t^*-1) = d(\bar{l}(t^*-1)))}{g_{n,0:t^*-1}(d(\bar{l}(t^*-1)))}] \\
\mathbb{E}[Y(t^*)|\bar{L}(t^*-1),\bar{A}1(t^*-1)] &= \text{logit}^{-1}[\beta_0 + \beta_1 W_2 + \beta_2 L1_1(t^*-1) + \beta_3 A1(t^*-1) \\
&\quad + \beta_4 \frac{\mathbb{I}(\bar{A}(t^*-1) = d(\bar{l}(t^*-1)))}{g_{n,0:t^*-1}(\bar{a}(t^*-1))}],
\end{aligned} \tag{3.9}$$

noting that in the stratified approach the coefficient on $A1(t^*-1)$ is simply absorbed into the intercept. Following this approach of fitting $\bar{Q}_0$, the remaining steps are implemented identically to the ICE approach in Chapter 3.4.

mDRICE

As discussed in Chapter 3.2, it is also possible to form a DR estimator by instead using the inverse propensity estimates $g_{n,0:t^-}(d(\bar{l}(t^-)))^{-1}$ as observational weights [82]. This approach is similar to the DRICE approach above, with the lone exception that the inverse propensity estimates $g_{n,0:t^-}(d(\bar{l}(t^-)))^{-1}$ instead be used as observational weights in the logistic regressions, leaving the indicator of treatment $\mathbb{I}(\bar{A}_i(t^-) = d(\bar{l}(t^-)))$ as the model covariate. We denote estimates under this approach as $\hat{\psi}_n^{mDRICE}$.

While generally the use of logistic regression solved Equation (3.4), we noticed a small proportion of the time that convergence was not truly achieved, i.e. the equation corresponding to the EIF was not truly solved. This occurrence increased with the presence of practical positivity issues. To ensure that resulting estimates would not be biased, we implemented a customized optimization function which would directly solve the EIF in this setting.

TMLE

This estimator differs from the ICE approach in that at each sequential regression step, $\bar{Q}^d_{0,L(t)}$ is first estimated (without use of the estimated inverse propensity score) and then this initial fit is updated in a second regression step, using the initial fit as offset, and including the identical term $\mathbb{I}(\bar{A}_i(t) = d(\bar{l}(t)))$ as a single covariate. As noted in Section 3.2, this modification facilitates the use of machine learning approaches for $\bar{Q}_0$.

As with the DRICE approach, we first take the subset of the data conditioned on survival up to the penultimate time point $t^* - 1$ (as well as optionally, $\bar{A}(t^* - 1) = d(\bar{l}(t^* - 1))$). With this subset, we carry out a logistic regression, regressing $Y(t^*)$ onto $\bar{A}(t^* - 1)$, $\bar{L}(t^* - 1)$. Both models from Equations (3.7) and (3.8) were considered. The transformed fit $\operatorname{logit} \bar{Q}^d_{n,L(t^*)}$ is then used as an offset in a univariate logistic regression with no intercept, the covariate $h = 1$, and observational weight $\mathbb{I}(\bar{A}(t^* - 1) = d(\bar{l}(t^* - 1)))/g_{n,0:t^*-1}(d(\bar{l}(t^* - 1)))$ to form the following parametric submodel [46]

$$\operatorname{logit} \bar{Q}^{d,s}_{n,L(t)}(\varepsilon_t) = \operatorname{logit} \bar{Q}^{d,s}_{n,L(t)} + \varepsilon_t h \tag{3.10}$$

where $t = t^*$. The sole parameter ε_{t^*} is estimated using maximum likelihood estimation, i.e. the negative likelihood loss function, resulting in the estimate $\hat{\varepsilon}_{n,t^*}$. This allows the initial fit $\bar{Q}^d_{n,L(t^*)}$ to be updated to $\bar{Q}^d_{n,L(t^*)}(\hat{\varepsilon}_{n,t^*})$, which we denote as $\bar{Q}^{d,*}_{n,L(t^*)}$. This updated fit is evaluated for each subject i in the study, setting equal to 1 if $T_i < t^*$. We iterate by now conditioning on survival up to time 4 and (optionally) on the subset of subjects having $\bar{A}(t^* - 2) = d(\bar{l}(t^* - 2))$, using the estimates $\bar{Q}^{d,*}_{n,L(t^*),i}$ as the outcome, and regressing it onto $\bar{A}(t^* - 2)$ and $\bar{L}(t^* - 2)$ using the logistic model from Equation (3.7) or (3.8) but replacing the predictors for time $t^* - 1$ with those of time $t^* - 2$. This gives us an initial fitted object $\bar{Q}^d_{n,L(t)}$ for time $t = 5$, which we again update using the parametric submodel specified in Equation 3.10, with $t = 5$ and observational weight $\mathbb{I}(\bar{A}(t^* - 2) = d(\bar{l}(t^* - 2)))/g_{n,0:t^*-2}(d(\bar{l}(t^* - 2)))$. The updated fit is evaluated for each subject i in the study setting equal to 1 if $T_i < 5$, and the procedure is iterated backwards over t until we reach time

$t = 1$, at which we will have the estimate $\bar{Q}^{d,*}_{n,L(1)}$ as a function of only $L(0)$. The target parameter is then estimated by taking the empirical mean over the sample, $\hat{\psi}^{TMLE,s}_n = \frac{1}{n}\sum_{i=1}^{n} \bar{Q}^{d,*}_{n,L(1),i}$.

For completeness, we also implemented a version of TMLE which uses the covariate $h_t = \mathbb{I}(\bar{A}(t) = d(\bar{l}(t)))/g_{n,0:t}(d(\bar{l}(t)))$ and observational weight 1 which we refer to as the covariate submodel. Similar to DRICE, we also noticed some convergence issues with this estimator as practical positivity issues increased and therefore implemented our own customized optimization function to directly solve the EIF in this setting.

R-packages

While TMLE is potentially complex in its execution, we note that an `ltmle` R-package has been developed for this estimator [65, 96] and uploaded to The Comprehensive R Archive Network (CRAN). Additionally, estimates using the ICE and IPW estimators can be obtained using this package. A built in option allows the user to either stratify or pool the subjects when estimating $\bar{Q}^{d}_{0,L(t)}$ and $g_{n,0:t^-}(d(\bar{l}(t^-)))$ for $t = 1, 2, \ldots, t^*$. Furthermore, it can conduct the estimations using either a parametric generalized linear model or the Super Learner estimation approach discussed in Section 3.2. We used version `0.9-6` of this package for this study.

We additionally developed a new `lrecCompare` R-package which was designed specifically for all of the remaining analyses for the current study. This package contains all of the functions to generate the simulation data, as well as code to perform the AIPW and DRICE computations. Similar to the `ltmle` R-package, a built in option allows the user to either stratify or pool the subjects and computations can be conducted using either a parametric generalized linear model or Super Learner. A further option allows the user to use the modified version of the DRICE estimator. Corresponding code for the package is presented in Appendix C.

Performance

Estimator performance was evaluated by comparing the bias and mean squared error of each estimator across 1000 iterations. While the sample variance of the empirical IF/n can provide a straightforward variance estimator for most of the estimators considered (with the exception of the ICE estimator), IF based variance estimation has been shown in past work to result in anti-conservative confidence intervals in settings, such as the one deliberately studied here, with practical positivity violations (e.g. van der Laan and Gruber [46], Petersen et. al. [65]). While the comparative evaluation of variance estimators is an exciting area in its own right, we focused here on performance of estimators of the target parameter $\mathbb{E}Y_d(t^*)$ and leave the evaluation of variance estimators for Chapter 4. To evaluate the extent to which an estimator's bias/se ratio threatened valid inference, we therefore calculated a modified version of 95% confidence interval "coverage", in which we used standard error estimators based on the empirical variance of an estimator across the 1000 repetitions.

3.5 Results

The mean squared error of each estimator from the simulations are presented in Figure 3.2 with further results presented in Tables 3.1 and 3.2. To limit the impact of the positivity violations and in accordance with theory [48, 49], estimates of $g_{0,0:t^*-1}(d(\bar{l}(t^*)))$ were truncated at 0.001.

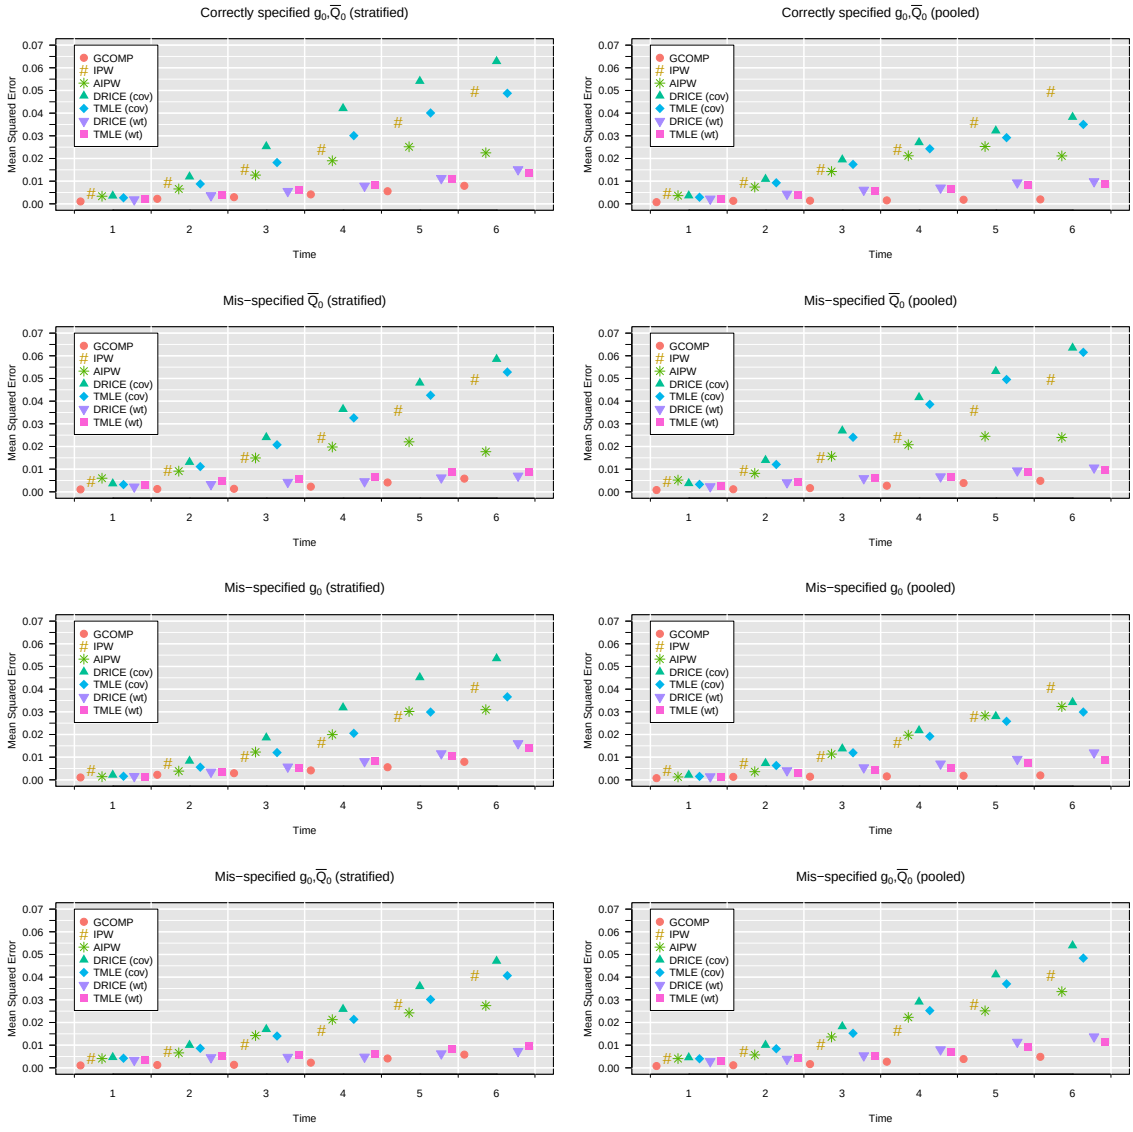

Figure 3.2: Mean squared error of each estimator under correct and mis-specification of g_0 and $\bar{Q}_0$.

ICE

Under correctly specified models, the ICE estimator had the lowest MSE and bias among all the estimators studied. As time increased and practical positivity issues arose, estimator variance increased. However, the increase was minimal and estimates remained stable. Under a mis-specification of $\bar{Q}_0$, the bias and MSE were slightly higher, though also remained stable across time. Coverage for this estimator remained valid over all time points under the correctly specified model and (expectedly) fell to as low as 0.69 at $t^* = 6$ under the mis-specified stratified model.

IPW

As predicted by underlying theory, the IPW estimator generally had among the highest bias and MSE within the estimators studied. For example, under the correctly specified model at $t^* = 1$, a bias of 0.009 was observed. Additionally, as expected given the complete reliance of this estimator on estimation of the propensity score, the bias increased by a factor of almost 6 when the g_0 model was mis-specified. Positivity issues noticeably affected this estimator, with a steady increase in the MSE at each time point. Within this setting, the high bias in mis-specifying g_0 was reversed for later time points. Due to positivity, the MSE tended to remain lower for this estimator when g_0 was mis-specified. At $t^* = 6$ the MSE was 0.050 for the correctly specified model and 0.041 for the mis-specified model. Even under a correctly specified model, coverage reduced with positivity. While at $t^* = 1$, a coverage of 0.96 was observed, this fell to 0.63 at $t = 6$. The mis-specified model also resulted in reduced coverage though at later time points where positivity issues were present, the coverage was higher than those from correctly specified models.

AIPW

At $t^* = 1$, the AIPW had minimal bias when at least one of the two nuisance parameters were estimated consistently. With both the g_0 and $\bar{Q}_0$ models specified correctly, a bias close to 0 was observed. As time progressed and the extent of practical positivity violations increased, the estimates for the target parameter became biased and MSE increased. Coverage, however, remained valid at approximately 95% over all time points. In settings of mis-specification, lower bias and MSE was seen when $\bar{Q}$ was correctly specified (and g was mis-specified) at the early time points, while the relationship was reversed at later time points. In comparison to IPW (which is also an estimating equation), the bias and MSE in AIPW was considerably improved, even under model mis-specification and positivity issues. Compared to the ICE estimator, the bias and MSE were higher at early time points. At later time points, the MSE remained higher although the ICE bias was higher. A proportion of estimates was noticed to reside outside the parameter boundary of $[0,1]$, which increased with time. For example, at $t = 1$ between 0.1% to 0.05% of the estimates were outside $[0,1]$, while this increased to 2.1% at $t = 6$.

DRICE

The DRICE estimator resulted in minimal bias when at least one of g_0 or $\bar{Q}_0$ was specified correctly. Under model mis-specification, lower bias and MSE was seen when only g_0 was mis-specified. Positivity issues also appeared to affect this estimator, as both the bias and MSE increased considerably over time. Similar to the AIPW, under large levels of practical positivity issues a lower MSE and bias was seen from the mis-specified models. Coverage remained valid over all time points. When compared to AIPW, performance was worse over all settings with higher bias and MSE.

Modifying DRICE had a large impact on the estimator's performance, such that the MSE was lower than those from AIPW. For example, under a correctly specified model at $t^* = 6$ the MSE was 0.015 compared to 0.022 for AIPW. This improvement was also seen under a fully mis-specified setting, with a MSE of 0.007 and 0.027 for the mDRICE and AIPW estimators respectively. Similar to the unmodified DRICE, under low positivity issues a low bias was observed when either of g_0 or $\bar{Q}$ was correctly specified.

TMLE

The TMLE performance was similar to that of DRICE. Under the covariate submodel, positivity issues resulted in an increase in the bias and MSE. For example, at $t^* = 1$ a bias and MSE of 0.003 were observed while $t^* = 6$ they were 0.111 and 0.049 respectively. Mis-specifying g_0 or $\bar{Q}_0$ appeared to result in lower bias and MSE with increases in positivity issues. Coverage also remained mostly at 95% when at least one of the two models were correctly specified.

Under the weighted submodel, TMLE tended to perform the best among the DR estimators, with the lowest bias and MSE observed. Indeed, it's performance was near that of the ICE estimator for correct specification and with minimal positivity issues. Similar performance was seen under a mis-specification of either g_0 or $\bar{Q}_0$. As with the weighted DRICE, mis-specifying $\bar{Q}_0$ under high levels of positivity resulted in lower bias and MSE. Coverage also remained mostly at 95%.

Pooling vs stratifying

Pooling observations in estimating $\bar{Q}_0$ increased estimator performance under most scenarios (Figure 3.2). The benefit of pooling was seen both with and without model mis-specification. Once significant levels of positivity violations were present, all estimators benefited greatly from pooling across different strata of $\bar{a}$ when at least one of the two nuisance parameters were specified correctly. The benefit was highest for DRICE, with the primary gains being seen when $\bar{Q}_0$ was correctly specified. Whereas at high levels of positivity a lower bias and MSE were seen from mis-specification of the nuisance parameters under the stratified approach for the DR estimators, the opposite was true for pooling observations. In other words, lower bias and MSE were seen for correct specification when pooling observations, but not when stratifying by treatment followed.

Table 3.1: Simulation results for estimating $\psi_0(t^*) = \mathbb{E}_{P_0^{\bar{a}}} Y(t^*) : t^* = 1,2,\ldots,6$ stratifying data by treatment followed.

Time (t^*)		1	2	3	4	5	6
	$\psi_0(t^*)$	0.232	0.335	0.390	0.428	0.460	0.489
		Bias (MSE) [Cov.]	Bias (MSE) [Cov.]	Bias (MSE) [Cov.]	Bias (MSE) [Cov.]	Bias (MSE) [Cov.]	Bias (MSE) [Cov.]
	Unadjusted	-0.043 (0.003) [0.60]	-0.040 (0.003) [0.79]	-0.017 (0.002) [0.93]	0.012 (0.002) [0.95]	0.034 (0.003) [0.89]	0.052 (0.005) [0.83]
ICE	True P_0	0.001 (0.001) [0.95]	0.004 (0.002) [0.95]	0.000 (0.003) [0.96]	0.004 (0.004) [0.95]	0.006 (0.006) [0.95]	0.015 (0.008) [0.94]
	Mis-$\bar{Q}_0$	0.018 (0.001) [0.89]	0.010 (0.001) [0.93]	0.009 (0.001) [0.95]	0.030 (0.002) [0.87]	0.049 (0.004) [0.79]	0.063 (0.006) [0.69]
IPW	True P_0	0.009 (0.005) [0.96]	0.003 (0.009) [0.94]	0.049 (0.015) [0.91]	0.099 (0.024) [0.89]	0.139 (0.036) [0.84]	0.187 (0.050) [0.63]
	Mis-g_0	0.052 (0.004) [0.69]	0.048 (0.007) [0.90]	0.009 (0.010) [0.95]	0.044 (0.016) [0.93]	0.086 (0.028) [0.92]	0.143 (0.041) [0.83]
AIPW	True P_0	0.000 (0.003) [0.97]	0.002 (0.007) [0.96]	0.001 (0.013) [0.96]	0.007 (0.019) [0.96]	0.015 (0.025) [0.96]	0.024 (0.022) [0.97]
	Mis-$\bar{Q}_0$	0.003 (0.006) [0.98]	0.008 (0.009) [0.96]	0.007 (0.015) [0.96]	0.005 (0.020) [0.97]	0.016 (0.022) [0.95]	0.021 (0.018) [0.96]
	Mis-g_0	0.001 (0.001) [0.95]	0.001 (0.004) [0.95]	0.007 (0.012) [0.96]	0.017 (0.020) [0.96]	0.026 (0.030) [0.97]	0.029 (0.031) [0.96]
	Mis-$\bar{Q}_0, g_0$	0.049 (0.004) [0.82]	0.049 (0.007) [0.92]	0.038 (0.014) [0.96]	0.017 (0.021) [0.96]	0.001 (0.024) [0.96]	0.012 (0.027) [0.95]
DRICE (covariate)	True P_0	0.003 (0.004) [0.98]	0.016 (0.012) [0.96]	0.052 (0.025) [0.94]	0.095 (0.042) [0.94]	0.121 (0.054) [0.96]	0.127 (0.063) [0.99]
	Mis-$\bar{Q}_0$	0.003 (0.004) [0.98]	0.011 (0.013) [0.98]	0.022 (0.024) [0.95]	0.043 (0.036) [0.97]	0.056 (0.048) [0.99]	0.066 (0.059) [1.00]
	Mis-g_0	0.003 (0.002) [0.96]	0.002 (0.008) [0.96]	0.029 (0.019) [0.95]	0.062 (0.032) [0.95]	0.099 (0.045) [0.97]	0.098 (0.054) [0.99]
	Mis-$\bar{Q}_0, g_0$	0.058 (0.005) [0.60]	0.058 (0.010) [0.95]	0.037 (0.017) [0.96]	0.006 (0.026) [0.97]	0.010 (0.036) [0.98]	0.027 (0.047) [1.00]
TMLE (covariate)	True P_0	0.003 (0.003) [0.97]	0.008 (0.009) [0.95]	0.035 (0.018) [0.94]	0.075 (0.030) [0.92]	0.092 (0.040) [0.94]	0.111 (0.049) [0.94]
	Mis-$\bar{Q}_0$	0.000 (0.003) [0.98]	0.007 (0.011) [0.97]	0.022 (0.021) [0.95]	0.045 (0.033) [0.96]	0.052 (0.043) [0.98]	0.070 (0.053) [1.00]
	Mis-g_0	0.001 (0.002) [0.96]	0.003 (0.006) [0.95]	0.025 (0.012) [0.94]	0.051 (0.021) [0.92]	0.079 (0.030) [0.93]	0.085 (0.037) [0.94]
	Mis-$\bar{Q}_0, g_0$	0.055 (0.004) [0.62]	0.055 (0.009) [0.94]	0.033 (0.014) [0.97]	0.005 (0.021) [0.97]	0.011 (0.030) [0.98]	0.033 (0.041) [0.99]
DRICE (weighted)	True P_0	0.001 (0.002) [0.95]	0.004 (0.004) [0.96]	0.002 (0.006) [0.96]	0.008 (0.008) [0.95]	0.013 (0.011) [0.95]	0.024 (0.015) [0.94]
	Mis-$\bar{Q}_0$	0.006 (0.002) [0.96]	0.011 (0.003) [0.95]	0.005 (0.004) [0.94]	0.004 (0.005) [0.95]	0.013 (0.006) [0.94]	0.022 (0.007) [0.94]
	Mis-g_0	0.002 (0.002) [0.95]	0.005 (0.003) [0.95]	0.000 (0.006) [0.96]	0.007 (0.008) [0.95]	0.015 (0.012) [0.94]	0.026 (0.016) [0.94]
	Mis-$\bar{Q}_0, g_0$	0.047 (0.003) [0.72]	0.043 (0.005) [0.87]	0.028 (0.005) [0.93]	0.008 (0.005) [0.94]	0.007 (0.006) [0.95]	0.020 (0.007) [0.94]
TMLE (weighted)	True P_0	0.000 (0.002) [0.95]	0.003 (0.004) [0.95]	0.002 (0.006) [0.94]	0.008 (0.008) [0.95]	0.011 (0.011) [0.94]	0.023 (0.014) [0.94]
	Mis-$\bar{Q}_0$	0.005 (0.003) [0.96]	0.010 (0.005) [0.95]	0.006 (0.006) [0.94]	0.003 (0.006) [0.94]	0.014 (0.009) [0.95]	0.021 (0.009) [0.94]
	Mis-g_0	0.001 (0.001) [0.95]	0.000 (0.003) [0.95]	0.006 (0.005) [0.95]	0.016 (0.008) [0.94]	0.020 (0.011) [0.94]	0.029 (0.014) [0.94]
	Mis-$\bar{Q}_0, g_0$	0.048 (0.004) [0.72]	0.048 (0.005) [0.87]	0.033 (0.006) [0.93]	0.014 (0.006) [0.95]	0.003 (0.008) [0.95]	0.015 (0.010) [0.95]

Estimates are based on 1000 simulations with a fixed sample size of n=500. $\psi_0(t^*)$ represents the true parameter value at time t^* under the intervention $\bar{a}(K) = 0$.
Mis-specified model for g_0 was $g_{0,A(t)} = \text{logit}^{-1}(\beta_0 + \beta_1 W_1 + \beta_2 W3 + \beta_3 L1_1(t^-))$. Mis-specified model for $\bar{Q}_0$ was $\bar{Q}_{L(t)} = \text{logit}^{-1}(\beta_0 + \beta_1 W_2 + \beta_2 L1_1(t^-) + \beta_3 A(t^-))$.
Confidence intervals were computed using the empirical standard error estimates across the $1,000$ simulations. Stratified and pooled models are identical for IPW.

Table 3.2: Simulation results for estimating $\psi_0(t^*) = \mathbb{E}_{P_0^{\bar{a}}} Y(t^*) : t^* = 1,2,\ldots,6$ pooling data over all treatment regimes.

Time (t^*)		1	2	3	4	5	6
	$\psi_0(t^*)$	0.232	0.335	0.390	0.428	0.460	0.489
		Bias (MSE) [Cov.]	Bias (MSE) [Cov.]	Bias (MSE) [Cov.]	Bias (MSE) [Cov.]	Bias (MSE) [Cov.]	Bias (MSE) [Cov.]
	Unadjusted	-0.043 (0.003) [0.60]	-0.040 (0.003) [0.79]	-0.017 (0.002) [0.93]	0.012 (0.002) [0.95]	0.034 (0.003) [0.89]	0.052 (0.005) [0.83]
ICE	True P_0	0.001 (0.001) [0.95]	0.003 (0.001) [0.94]	0.007 (0.001) [0.94]	0.001 (0.002) [0.95]	0.005 (0.002) [0.94]	0.007 (0.002) [0.94]
	Mis-$\bar{Q}_0$	0.009 (0.001) [0.94]	0.003 (0.001) [0.95]	0.020 (0.002) [0.92]	0.036 (0.003) [0.83]	0.048 (0.004) [0.79]	0.055 (0.005) [0.75]
IPW	True P_0	0.009 (0.005) [0.96]	0.003 (0.009) [0.94]	0.049 (0.015) [0.91]	0.099 (0.024) [0.89]	0.139 (0.036) [0.84]	0.187 (0.050) [0.63]
	Mis-g_0	0.052 (0.004) [0.69]	0.048 (0.007) [0.90]	0.009 (0.010) [0.95]	0.044 (0.016) [0.93]	0.086 (0.028) [0.92]	0.143 (0.041) [0.86]
AIPW	True P_0	0.000 (0.004) [0.96]	0.001 (0.007) [0.96]	0.003 (0.014) [0.96]	0.001 (0.021) [0.97]	0.005 (0.025) [0.96]	0.000 (0.021) [0.97]
	Mis-$\bar{Q}_0$	0.002 (0.005) [0.97]	0.004 (0.008) [0.96]	0.003 (0.016) [0.96]	0.006 (0.021) [0.97]	0.011 (0.024) [0.95]	0.011 (0.024) [0.96]
	Mis-g_0	0.000 (0.001) [0.96]	0.016 (0.004) [0.95]	0.016 (0.011) [0.96]	0.003 (0.020) [0.97]	0.006 (0.028) [0.97]	0.004 (0.032) [0.96]
	Mis-$\bar{Q}_0, g_0$	0.046 (0.004) [0.89]	0.037 (0.006) [0.94]	0.021 (0.014) [0.96]	0.004 (0.022) [0.97]	0.020 (0.025) [0.96]	0.027 (0.034) [0.96]
DRICE (covariate)	True P_0	0.001 (0.004) [0.98]	0.020 (0.011) [0.94]	0.040 (0.019) [0.92]	0.053 (0.027) [0.92]	0.051 (0.032) [0.93]	0.056 (0.038) [0.92]
	Mis-$\bar{Q}_0$	0.001 (0.004) [0.98]	0.013 (0.014) [0.97]	0.038 (0.027) [0.97]	0.067 (0.042) [0.98]	0.082 (0.053) [1.00]	0.092 (0.063) [1.00]
	Mis-g_0	0.001 (0.002) [0.96]	0.000 (0.007) [0.95]	0.019 (0.014) [0.94]	0.045 (0.022) [0.93]	0.054 (0.028) [0.92]	0.070 (0.034) [0.92]
	Mis-$\bar{Q}_0, g_0$	0.058 (0.005) [0.59]	0.054 (0.010) [0.96]	0.022 (0.018) [0.97]	0.017 (0.029) [0.97]	0.045 (0.041) [0.99]	0.069 (0.054) [1.00]
TMLE (covariate)	True P_0	0.002 (0.003) [0.98]	0.017 (0.009) [0.93]	0.040 (0.017) [0.92]	0.054 (0.024) [0.91]	0.057 (0.029) [0.91]	0.069 (0.035) [0.91]
	Mis-$\bar{Q}_0$	0.001 (0.003) [0.98]	0.010 (0.012) [0.97]	0.038 (0.024) [0.96]	0.068 (0.039) [0.97]	0.085 (0.050) [0.98]	0.110 (0.062) [1.00]
	Mis-g_0	0.000 (0.002) [0.96]	0.004 (0.006) [0.95]	0.017 (0.012) [0.94]	0.043 (0.019) [0.92]	0.056 (0.026) [0.91]	0.067 (0.030) [0.92]
	Mis-$\bar{Q}_0, g_0$	0.054 (0.004) [0.63]	0.050 (0.008) [0.95]	0.018 (0.015) [0.97]	0.021 (0.025) [0.97]	0.052 (0.037) [0.98]	0.077 (0.048) [1.00]
DRICE (weighted)	True P_0	0.001 (0.002) [0.95]	0.003 (0.004) [0.96]	0.000 (0.006) [0.95]	0.001 (0.007) [0.94]	0.006 (0.009) [0.94]	0.006 (0.010) [0.94]
	Mis-$\bar{Q}_0$	0.001 (0.002) [0.95]	0.002 (0.004) [0.95]	0.004 (0.006) [0.95]	0.010 (0.007) [0.94]	0.018 (0.009) [0.95]	0.022 (0.011) [0.94]
	Mis-g_0	0.002 (0.001) [0.95]	0.028 (0.004) [0.94]	0.025 (0.005) [0.95]	0.013 (0.007) [0.95]	0.004 (0.009) [0.95]	0.001 (0.012) [0.95]
	Mis-$\bar{Q}_0, g_0$	0.040 (0.003) [0.80]	0.028 (0.004) [0.93]	0.005 (0.005) [0.95]	0.019 (0.008) [0.95]	0.035 (0.011) [0.93]	0.044 (0.014) [0.93]
TMLE (weighted)	True P_0	0.001 (0.002) [0.95]	0.002 (0.004) [0.95]	0.003 (0.006) [0.95]	0.001 (0.007) [0.95]	0.002 (0.008) [0.95]	0.003 (0.009) [0.94]
	Mis-$\bar{Q}_0$	0.003 (0.003) [0.95]	0.006 (0.004) [0.95]	0.001 (0.006) [0.95]	0.006 (0.007) [0.94]	0.013 (0.009) [0.95]	0.017 (0.010) [0.94]
	Mis-g_0	0.000 (0.001) [0.95]	0.015 (0.003) [0.95]	0.015 (0.004) [0.94]	0.005 (0.005) [0.95]	0.001 (0.007) [0.95]	0.003 (0.009) [0.95]
	Mis-$\bar{Q}_0, g_0$	0.045 (0.003) [0.75]	0.037 (0.004) [0.91]	0.018 (0.005) [0.95]	0.003 (0.007) [0.96]	0.018 (0.009) [0.94]	0.028 (0.012) [0.94]

Estimates are based on 1000 simulations with a fixed sample size of n=500. $\psi_0(t^*)$ represents the true parameter value at time t^* under the intervention $\bar{a}(K) = 0$.
Mis-specified model for g_0 was $g_{0,A(t)} = \text{logit}^{-1}(\beta_0 + \beta_1 W_1 + \beta_2 W3 + \beta_3 L1_1(t^-))$. Mis-specified model for $\bar{Q}_0$ was $\bar{Q}_{L(t)} = \text{logit}^{-1}(\beta_0 + \beta_1 W_2 + \beta_2 L1_1(t^-) + \beta_3 A(t^-))$.
Confidence intervals were computed using the empirical standard error estimates across the $1,000$ simulations. Stratified and pooled models are identical for IPW.

3.6 Estimating $\mathbb{E}Y_d(t^*)$ in the IeDEA-EA cohort

Recall that, as opposed to using the entire IeDEA-EA cohort in Chapter 2, we conditioned on LREC availability for this analysis. Our baseline time point, $t = 0$ is therefore the first point of patient eligibility once the program has been initiated at each of the 15 clinics. With this restriction, $15,225$ subjects were found eligible for the LREC program. Of them, $2,011$ immediately enrolled into the program at the start of follow-up (i.e. within the first 90 day interval). During follow-up, subjects continued to enroll into the LREC program. By 21 months from initial eligibility, $1,819$ subjects were still alive and remained un-enrolled. The patients cumulatively contributed a total of $229,941$ person-months (Interquartile range: $6,18$) of follow-up to the analyses. In regards to their final observed outcomes, $1,440$ experienced either loss to follow-up ($1,362$) or death (78) by 24 months, while 140 were censored due to clinic transfers. Subjects for whom the database closed less than 24 months after eligibility were administratively censored. Some of the variables (i.e. cd4 count, clinic visited, WHO stage, pregnancy, tuberculosis treatment) had missing values. Rather than omitting the subjects with missing values completely, we imputed values by using either their last observation carried forward (for time-varying covariates with prior measures) or by using the median of the value across all subjects. This did not noticeably change any of the overall summary statistics for the covariates when compared to a complete case analysis.

Target parameter

We are interested in comparing 1 minus the cumulative probability of dying or dropping out of care over time under the intervention to prevent enrollment at all time points ($\bar{a1} = 0$) and enforce no censoring ($\bar{c1} = \bar{c2} = 0$). Under the counterfactual intervention we estimate 1 minus the counterfactual probability of failure by time t^*, $\psi_{\bar{a}}(t^*) = \mathbb{E}_{P_d^0}[Y_d(t^*)]$ for each time point $t^* = 1,2,\ldots,7$, where $t^* = 7$ corresponds to the 24^{th} month of follow-up.

Estimation results

We first used parametric generalized linear models (GLMs) to estimate the g_0 and $\bar{Q}_0$ portions of the likelihood. Fits for g_0 were formed by pooling observations over the 7 time points, as opposed to stratifying by each time point and fitting separate models (as was done in the simulations). This aided us in estimating the clinic transfer mechanism, as the number of subjects transferring clinics was low ($n = 140$). Figure 3.3 shows the estimated marginal densities of $g_{n,0:t^*-1}(\bar{a}(t^*)) : t^* = 1,2,\ldots,7$ at each time point t^* taken marginally over the $15,225$ subjects. While the trend of shifting to lower probabilities over time is similar to that seen in the simulated data (Figure 3.1), the degree of positivity violations is smaller.

Figure 3.4 shows the results of applying each estimator to our data using the GLM fits. 95% confidence intervals were calculated using standard error estimates from each estimator's influence function. For the ICE estimator, the empirical EIF was also used. At early time points, all approaches yielded very similar estimates. As time progressed all the parameter estimates decreased and only slightly deviated from one another (with the lone exception of IPW). Confidence

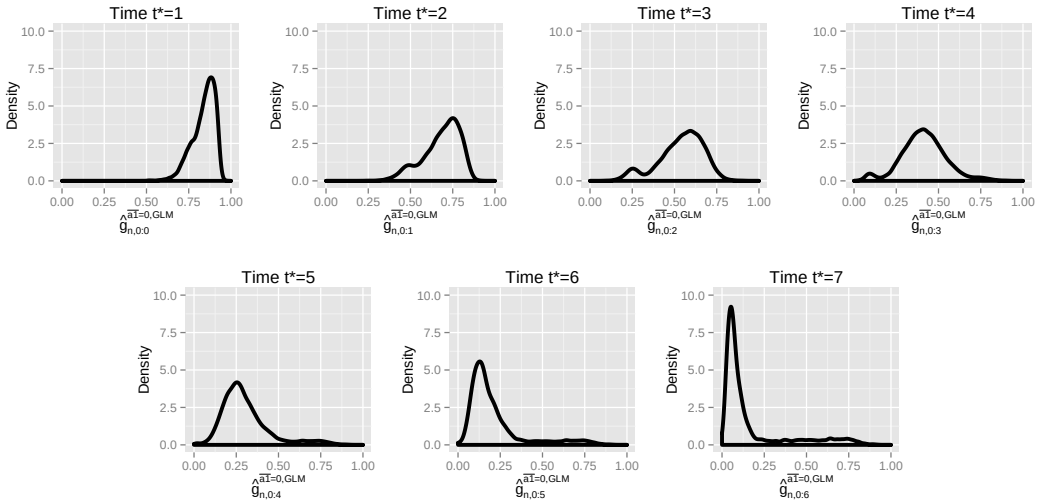

Figure 3.3: Marginal densities of $g_{n,0:t^*-1}(\bar{a}(t^*))$ for each time point t^* taken over the $15,225$ subjects using GLMs, where $\bar{a1} = 0$.

intervals also increased with time. Pooling observations across treatment regimes did not affect the estimates noticeably. The estimators all (within rounding errors) respected the monotonicity of the cumulative failure distribution over time.

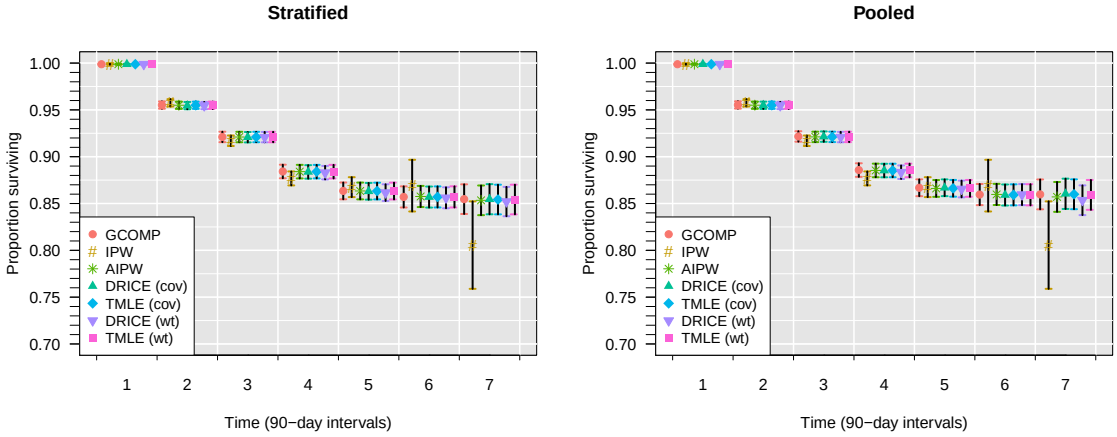

Figure 3.4: Applied GLM estimates and 95% confidence intervals for estimating $\psi_{\bar{a}=0}(t^*) = \mathbb{E}Y_{\bar{a}=0}(t^*) : t^* = 1,2,\ldots,7$.

As stated in Section 3.2, Super Learner was also used to estimate the conditional probability g_0 and conditional expectations $\bar{Q}_0$. The library of potential candidates used here consisted of: a generalized linear model, Bayesian genearlized linear model, multivariate adaptive regression spline,

gradient boosting machine, support vector machine, neural network, LASSO, ridge regression, and a stepwise selected model using the Akaike information criterion. We took the linear combination which minimized the cross-validated non-negative binomial likelihood risk.

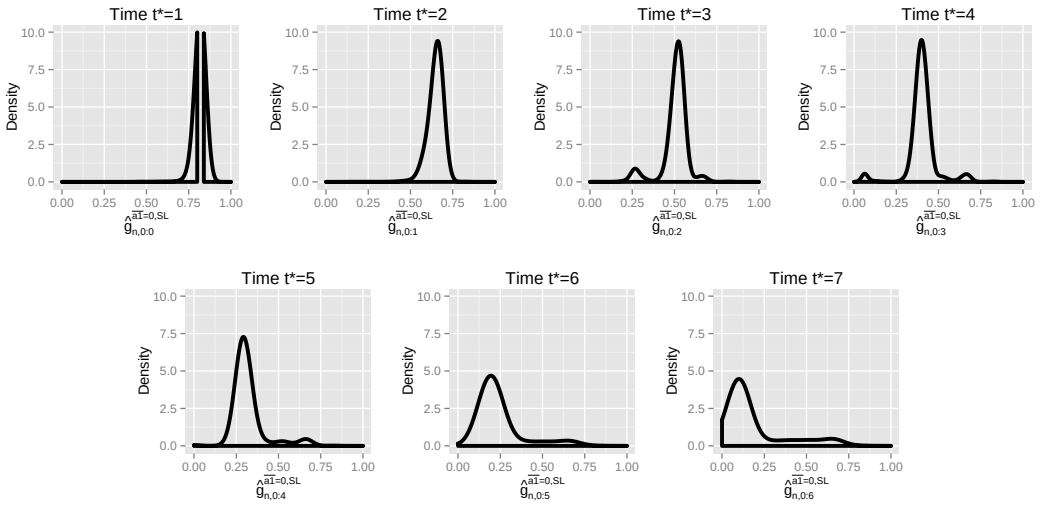

Figure 3.5: Marginal densities of $g_{n,0:t^*-1}(\bar{a})$ for each time point t^* taken over the $15,225$ subjects using Super Learner, where $\bar{a1} = 0$.

Figure 3.5 shows the estimated density for $g_{n,0:t^*-1}(\bar{a}(t^*-1))$ for time point $t^* = 1,2,\ldots,7$ using Super Learner. A comparison of the GLM and Super Learner plots implies that the fits between the two approaches are very similar, with Super Learner generally resulting in more concentrated distributions. Correlations in the estimated cumulative probabilities at each time point t^* resulting from the two approaches varied from 0.56 to 0.97 with higher correlations toward the later time points. Estimates of probabilities tended to be higher for $g_{n,0:t^*-1}(\bar{a}(t^*-1)=0)$ fits using Super Learner. Figure 3.6 shows the estimates of our target parameter, with influence curve based 95% confidence intervals. We excluded DRICE here, as a number of the candidates we used in the Super Learner library were non-linear. While the estimates are similar to the approach using GLMs, a number of differences are seen. Firstly, the similarity of the estimates across the set of estimators (within each time point) is reduced. A noticeable reduction is seen in estimates from TMLE with inverse propensity score as covariate at the later time points. Monotonicity is still preserved (within rounding errors). Confidence intervals were generally larger than those estimated using GLM fits.

3.7 Discussion

Numerous approaches have been proposed for estimating causal parameters in a longitudinal treatment setting, several of which have the attractive theoretical properties of double robustness and semi-parameteric efficiency. However, there are few published direct comparisons of the relative

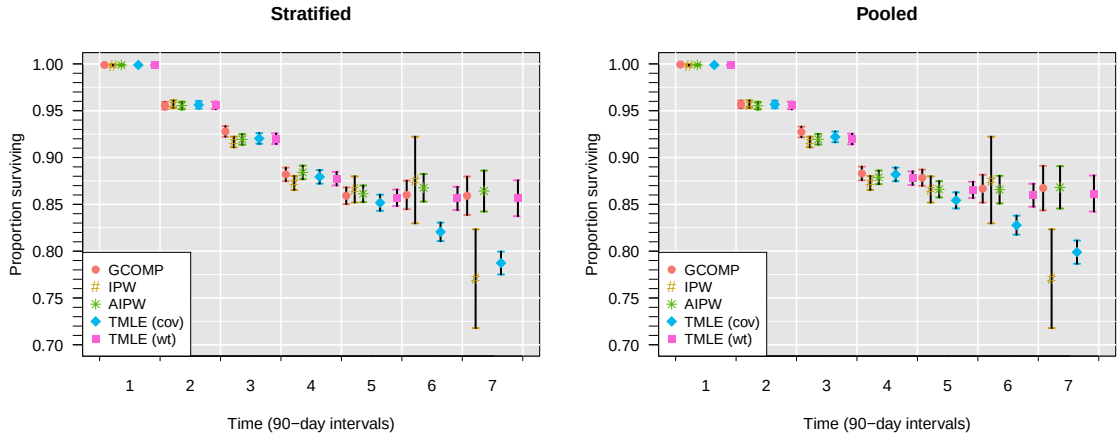

Figure 3.6: Applied Super Learner estimates and 95% confidence intervals for estimating $\psi_{\bar{a}=0}(t^*) = \mathbb{E}Y_{\bar{a}=0}(t^*) : t^* = 1,2,\ldots,7$.

performance of the DR estimators in a longitudinal setting or with machine learning approaches for estimating the nuisance parameters. In this article, we presented and analyzed seven specific estimators, including 4 double robust estimators, in a longitudinal treatment setting using both simulated and real world data. The DR estimators we considered further included possible modifications to estimators [3, 46], in that we considered estimators that (a) integrated the inverse propensity score estimate as a weight rather than a covariate [82] and (b) pooled rather than stratified on treatment history when estimating nuisance parameters. We evaluated performance under model mis-specifications and steadily increasing levels of positivity violations for our simulated data. For the applied setting, we considered both a parametric and data adaptive approach at estimating g_0 and $\bar{Q}_0$. Simulation results showed decreasing estimator performance with increasing positivity violations and with model mis-specifications. Pooling observations tended to result in lower MSE, as well as reduced bias when at least one of the two nuisance parameters were estimated consistently. The positivity violations affected the IPW estimator the most and the ICE estimator the least, while the DR estimators showed an affect intermediate to the two. Without positivity violations, as predicted by theory, the DR estimators retained valid inference, unbiasedness, and low MSE when at least one of the two models were specified correctly. When both nuisance parameters were specified incorrectly, the ICE approach out performed the DR estimators in terms of bias and MSE. Among the DR estimators, DRICE and TMLE using the inverse propensity score estimate as covariate performed the worse in simulations with the highest bias and MSE. Integrating the inverse propensity score as observational weights in DRICE and TMLE resulted in considerably improved performance, outperforming AIPW in terms of MSE. The magnitude of difference increased over time with the magnitude of positivity issues.

In the applied setting, similar estimates were observed across estimators when using GLM fits of the two nuisance parameters. The use of Super Learner to estimate g_0 and $\bar{Q}_0$ resulted estimates

similar to the GLM approach, though with more variability between estimators and slightly wider confidence intervals. Integrating the inverse propensity scores as observational weights in the two DR substitution based estimators resulted in improved performance in simulations when compared with the covariate form.

In simulations, pooling observations generally resulted in lower bias and MSE for all estimators when at least one of the two relevant portions of the likelihood were estimated consistently. This became even more apparent as time progressed and practical positivity violations set in. In our applied analyses, however, no large difference was seen from pooling observations.

A number of debates have arisen regarding which of the estimators considered here are preferred in practice as well as theory. For example, these estimators can be stratified into an estimating equation approach (IPW, AIPW) or a substitution based estimator (ICE, DRICE, TMLE). Our preference is for substitution based estimators, as previous research has noted that estimating equations may result in no or multiple solutions [46]. Furthermore, as seen in the simulations, these methods do not always obey the constraints of the parameter space. Consequently, the use of these methods can, as seen here, lead to estimates and confidence intervals that are outside of the $[0, 1]$ range for our parameter. In our simulations, we saw that up to 2% of the estimates from AIPW estimator were outside that range. Substitution estimators are defined as mappings applied to probability distributions and consequently will never result in values beyond the bounds of the parameter space.

There is also a debate regarding whether DR estimators really have an advantage, as in practice, both the outcome and treatment models could be mis-specified. Our simulations here, those from Kang and Schafer [36], and those from Robins et al [82] in the point treatment setting show that when both are wrongly specified, the ICE approach can outperform the DR estimators in terms of bias and MSE. However, in using machine learning approaches to estimate both nuisance parameters, we largely side step this concern and gain several important advantages. By using machine learning approaches, which respect that our true statistical model is non-parametric, we can ensure consistency of our nuisance parameters estimators, providing an essential condition for accurate inference as well as estimator efficiency. This further translates into finite sample gains in both bias and variance. Furthermore, in settings such as randomized controlled trials where g_0 is known, the DR estimators are guaranteed to be consistent and their use will only improve efficiency.

Our intention in comparing the various estimators was to provide a sense of the performance of these estimators in practice. Given its efficiency, ability to respect the parameter space, and observed performance, we recommend the pooled and weighted TMLE approach as the preferred estimator. The weighted DRICE estimator also performed well as a substitution based estimator, though is not generalizable to machine learning which can make use of non-parametric models. While the ICE approach generally resulted in the lowest MSE, the simulation was conducted in a parametric setting which is unlikely to occur in practice.

We end by stating that the approaches presented here can easily be generalized to a number of various other estimation problems. An example includes the estimation of causal effects of dynamic treatment regimes, where treatment decisions are functions of the time-dependent covariate process. Extensions involving marginal structural models may also be applied in summarizing the treatment effect over multiple time points, levels of treatments, or treatment effects that are also

functions baseline covariates [77, 65].

Chapter 4

Robust variance estimation and inference for causal effect estimation

4.1 Introduction

As shown in Chapter 3, a number of different estimators are available for estimating the treatment specific mean outcome parameter (and the corresponding causal contrasts). Variance estimation for these estimators are conventionally achieved by using their corresponding influence functions (IF) or by re-sampling methods such as the bootstrap. However, a number of shortcomings exists with these variance estimation approaches. In particular, no theory for exists for the non-parametric bootstrap when using data adaptive methods for estimation nuisance parameters, and both IF-based and bootstrap based confidence intervals can become anti-conservative with increasing levels of practical positivity violations. For example, van der Laan and Gruber [46] found IF-based variance estimates for the intervention specific mean outcome that were overly-conservative when compared with the Monte-Carlo variance of the TMLE, leading to anti-conservative confidence intervals. Petersen et. al. [65] found poor coverage for influence function-based confidence intervals, owing to both a result of practical positivity violations and relatively rare outcomes. This behaviour is especially true under sparsity in finite samples, even when the assumptions for asymptotic validity of these estimators hold [64]. As a consequence, statistical inference based on these estimators and these variance estimates becomes unreliable when the treatment mechanism (i.e. the missingness mechanism) practically or theoretically violates the underlying positivity assumption. Despite this, to the best of our knowledge, very little research has conducted into variance estimation.

Additionally, under sparsity issues, the estimated variance may also fail to raise a red flag for unreliable statistical inference [64]. For example, these estimates of the asymptotic variance are not sensitive to theoretical violations of the positivity assumptions under which the asymptotic variance would be infinity, i.e. when positivity fails. Consequently, it is less likely that the analyst will be able to determine if the data at hand provides insufficient information to estimate the desired causal parameter with any reasonable degree of accuracy.

Previous work [64, 65] proposed estimating the asymptotic variance of the estimator with a

parametric bootstrap-based on a fit of the density of the data generating distribution, involving estimation of the treatment mechanism and the G-computation factor of the likelihood. This proposal corresponds with evaluation of the variance of a given estimator using the data at hand as a given data generating experiment. The consistency of this estimator relies on correct specification of both the treatment mechanism and the G-computation factor. As a consequence, this parametric bootstrap-based variance estimate was only proposed as a measure to raise a red flag for unreliable statistical inference. In addition, in the context of sparsity, one needs to sample many bootstrap samples and refit the likelihood in each iteration in order to obtain a valid evaluation of the estimator variance, in order to capture the rare observations that nonetheless heavily contribute to this variance. Thus, this semi-parametric bootstrap method is extremely computer intensive, making it an intractable method for complex estimators and complex data generating distributions.

In this chapter, we use analytic expressions to compute the variance of the efficient influence function (EIF) [25, 79] which provide the asymptotic variance of estimators solving the estimating equation corresponding to the EIF, such as the AIPW or TMLE. These analytic expressions naturally integrate over the rare observations, and thereby avoid the finite sample bias in variance estimation using standard influence curve or non-parametric bootstrap based methods due to rare observations mentioned above. With this, we construct plug-in type estimators of these asymptotic variances that are consistent if the treatment mechanism is consistently estimated. These estimators only require estimation of the treatment mechanism and several treatment specific means of specified outcomes (defined as a function of the observed data structure, indexed by the estimator of the treatment mechanism), which can thereby be estimated with either an estimating equation type IPW estimator or an efficient double robust method such as a targeted minimum loss-based estimator. The resulting variance estimator, unlike current alternatives based on taking the empirical variance of the estimated influence function, or using a non-parametric bootstrap, will become very large whenever the estimated treatment mechanism reflects practical or theoretical violations of the positivity assumption and, consequently, a lack of identifiability.

While this newly presented approach performs well in estimating the asymptotic estimator variance, a lower finite sample variance is expected for substitution based estimators such as TMLE, due to the guaranteed parameter boundaries provided by the estimator. To address this, we additionally present a bootstrap based approach of estimating the finite sample variance which does not require re-estimation of the treatment mechanisms and the Q-factor of the likelihood. The approach is shown to, under certain conditions, be asymptotically linear. The resulting reduction in the computational load (compared to a fully non-parametric bootstrap approach which refits the likelihood for each iteration) allows for a more tractable approach at estimating the variance.

To summarize, we start by reviewing the current approach of influence function (IF) based estimator variance estimation. We then present the approach for robust estimation of the variance of the EIF, which performs well under sparsity. The expression for the variance of the efficient influence function is presented along with both an IPW and TMLE based estimation approach for this parameter. To help illustrate, an example is given for a point treatment setting under a static treatment regime and advantages of this new approach are covered. We additionally introduce a bootstrap approach to estimate the estimator variance. To reduce the computational intensity required, a TMLE for estimating the intervention specific mean outcome is presented. The TMLE

is shown to be asymptotically linear with an influence function equal to the EIF. This bootstrap approach also solves the EIF and therefore has a normal limiting distribution, implying a consistent approach at estimating the variance.

We illustrate the results of our work by applying the presented variance estimation approaches to both a single time-point static treatment setting and longitudinal setting with three time points and time-dependent confounding. We show that the robust approach at estimating variance performs well for the AIPW variance. The robust approach over-estimates the TMLE variance, while the bootstrap approach results in estimates close to the observed Monte-Carlo variance. The resulting confidence intervals are shown to be valid, while the bootstrap approach is shown to retain higher statistical power. Lastly, we apply the new and current variance estimation approaches in evaluating the effect of enrollment into the LREC program in our IeDEA-EA cohort. Point estimates, standard errors, and confidence intervals for the additive treatment effect are presented for seven distinct time points. We conclude with a discussion, which reviews the benefits of this new approach, potential limitations, and future directions.

4.2 Semi-targeted estimation of the EIF variance

Recall that an estimator $\hat{\Psi}(P_n)$ is considered to be asymptotically linear if and only if

$$\hat{\Psi}(P_n) - \Psi(P_0) = \frac{1}{n}\sum_{i=1}^{n} D(P_0)(O_i) + o_p(n^{-1/2})$$

for some mean 0 finite variance influence function $D(P_0)(O)$ [25]. Any estimators solving the estimating equation corresponding to an influence function will therefore, by the central limit theorem, be asymptotically normally distributed with mean zero and variance equal to the variance of the influence function σ^2 divided by n. The asymptotic variance of the estimator can therefore be consistently estimated with the empirical variance of the estimated influence function $D(P_n)(O)$, i.e. $\hat{\text{var}}[\hat{\Psi}(P_n)] = \text{var}[D(P_n)(O)]/n$, which implies an asymptotically valid confidence interval.

Alternatively, we can directly target the variance of $D(P_0)(O)$ as an expectation. The following describes how to obtain a TMLE of the variance of each component of the EIF σ_t^2 in the setting of a scalar parameter. We provide a proof for the more general working MSM setting in Appendix B.

Expression for variance of the EIF for $\mathbb{E}Y_d$

Under regimens $d(\bar{l}(K))$, we have

$$\sigma_0^2 \equiv \mathbb{E}_0[D^*(P_0)(O)]^2 = \sum_{t=0}^{K+1} \mathbb{E}_0[H_t^2(g_0)(\bar{Q}^d_{0,t^+} - \bar{Q}^d_{0,t})^2].$$

Using the expression for $H_t(g)$ from Equation (2.8), and first taking the conditional expectation w.r.t. $\bar{A}(t^-)$ given $X = (\bar{L}^d : d)$, it follows that this can be written as:

$$\sigma_0^2 = \sum_{t=0}^{K+1} \mathbb{E}_{P_0^d}\left[\frac{(\bar{Q}^d_{0,t^+} - \bar{Q}^d_{0,t})^2(\bar{L}^d(t))}{g_0(d(\bar{l}(t^-)), \bar{L}^d(t^-))}\right], \tag{4.1}$$

where we define $g_0(d(\bar{l}(-1)),\bar{L}^d(-1)) = 1$ so that the term at $t = 0$ equals $\mathbb{E}_{L(0)}[\bar{Q}^d_{0,1}(L(0)) - \mathbb{E}_0 Y^d]^2$. This is simply a sum of expectations over $t \in \{0,1,\ldots,K+1\}$. We can write

$$\sigma_0^2 = \sum_{t=0}^{K+1} \sigma_t^{2,d} = \sum_{t=0}^{K+1} \mathbb{E}_{P_0^d}\left[S_t^d(\bar{Q}_0,g_0)(\bar{L}^d(t))\right] \tag{4.2}$$

for the specified function

$$S_t^d(\bar{Q}_0,g_0)(\bar{L}^d(t)) \equiv \frac{(\bar{Q}^d_{0,t^+} - \bar{Q}^d_{0,t})^2(\bar{L}^d(t))}{g_0(d(\bar{l}(t^-)),\bar{L}^d(t^-))} : t = 0,1,\ldots,K+1$$

Note that, given $(\bar{Q}_0,g_0)$, $\mathbb{E}_{P_0^d} S_t^d(\bar{Q}_0,g_0)$ is the mean of a counterfactual $S_t^d(\bar{Q}_0,g_0)(\bar{L}^d(t))$, i.e., the mean of a real valued function (indexed by $d(\bar{l})$ itself) of $\bar{L}^d(j)$, which needs to be estimated based on the longitudinal data structure $L(0),A(0),\ldots,A(t-1),L(t)$. Given $\bar{Q}_0,g_0$, we observe the outcome $S_t^d(\bar{Q}_0,g_0)(\bar{L}_i(t))$, $i = 1,2,\ldots,n$, so that we can represent the observed data structure as $L(0),A(0),\ldots,A(t-1),S_t^d(\bar{Q}_0,g_0)(\bar{L}(t))$, and we wish to estimate the statistical target parameter

$$\mathbb{E}_{P_0^d} S_t^d(\bar{Q}_0,g_0) = \sum_{\bar{l}(t)} S_t^d(\bar{Q}_0,g_0)(\bar{l}(t)) P_0^d(\bar{L}^d(t) = \bar{l}(t)) : t = 0,1,\ldots,K+1, \tag{4.3}$$

where again we assume $l(t)$ is discrete for sake of presentation.

Estimation of variance of the EIF

With the expression for the variance of the efficient influence function in hand (Equation 4.2), we can now form estimators which target this parameter. $\bar{Q}_0$ and g_0 are not known in practice, though estimates $\bar{Q}_n^*$ and g_n will be readily available if estimating $\mathbb{E}Y_d$ using a double robust estimator such as TMLE, thus providing us with the observed outcome $S_t^d(\bar{Q}_n^*,g_n)(\bar{L}(t))$. Treating this variable as our new time point specific outcome, our goal is to estimate the mean of this variable over the post-intervention of $\bar{L}^d(t)$. For notational convenience, let $Z^d(t) \equiv S_t^d(\bar{Q}_0,g_0)(\bar{L}(t))$, and represent the observed data structure as $(L(0),A(0),\ldots,A(t-1),Z^d(t))$.

One possible approach to estimating each of the components (Equation (4.3)) is to use a simple IPW estimator [35]

$$\hat{\sigma}^{2,d}_{t,n,IPW} = \frac{1}{n}\sum_{i=1}^{n} \frac{\mathbb{I}(\bar{A}_i(t^-) = d(\bar{l}(t^-)))}{g_{0:t^-,n}(\bar{A}_i(t^-),\bar{L}_i(t^-))} Z_n^d(t)$$

where $Z_n^d(t) = S_t^d(\bar{Q}_n,g_n)(\bar{L}(t))$. However, such an estimator would still be subject to underestimation of the variance by ignoring the contribution of observations that selected a likely treatment $\bar{A}_i$, even though their probability of following $d(\bar{l})$ is very small. In other words, subjects i with small probabilities of following $d(\bar{l})$ would be unlikely to be observed with $\bar{A}_i = d(\bar{l})$ resulting in an indicator value of 0 for the numerator and, consequently, a contribution of 0 to the IPW estimator. Therefore, we stress that it is important to use a plug-in estimator such as the TMLE [46] to estimate this parameter. A plug-in estimator will integrate over all $\bar{l}(t)$ in the support of $P^d_{t,n}$ and

thus contribute many large values of $S^d_{t,n}(\bar{Q}^*_n, g_n)$ when there are practical or theoretical positivity assumption violations. In addition, the TMLE is a double robust estimator so that it will yield a consistent estimator of this variance if g_n is consistent for the true g_0.

Given $\bar{Q}_0, g_0$, we will now provide a succinct summary of the TMLE of $\sigma^{2,d}_{0,t} = \mathbb{E}_{P^d_0} Z^d(t)$ that is based on iterative sequential regression. Note that this iterative sequential regression approach is similar to the one presented by van der Laan and Gruber [46] for the intervention specific mean outcome parameter. Denote the counterfactual of $Z^d(t)$ under treatment d' with $Z^{d,d'}(t)$, and let $P^{d'}_0$ be the G-computation formula [71] corresponding with this intervention $\bar{A}(t^-) = d'(\bar{l}(t^-))$. We wish to estimate $\sigma^{2,d}_{0,t} = \mathbb{E}_{P^d_0} Z^{d,d}(t)$, which can be represented as a series of iterated conditional expectations

$$\sigma^{2,d}_{0,t} = \mathbb{E}[\mathbb{E}[\cdots \mathbb{E}[\mathbb{E}[Z^d(t)|\bar{L}(t-1), \bar{A}(t-1) = d(\bar{l}(t-1))]|\bar{L}(t-2), \bar{A}(t-2) = d(\bar{l}(t-2))] \cdots |\bar{L}(0)]].$$

The EIF for this target parameter $\sigma^{2,d}_t$ is given by

$$D^*_{\sigma^{2,d}_t}(P)(O) = \sum_{m=0}^{t} H^{d,t}_m(g)(\bar{Q}^{d,\sigma^2_t}_{m^+} - \bar{Q}^{d,\sigma^2_t}_m),$$

where we define

$$\begin{aligned} \bar{Q}^{d,\sigma^2_t}_{t+1} &= Z^d(t) \\ H^{d,t}_m(g) &= \frac{\mathbb{I}(\bar{A}(m^-) = d(\bar{l}(m^-)))}{g_{0:m^-}(\bar{A}(m^-), \bar{L}(m^-))} : m = 1, 2, \ldots, t \\ H^{d,t}_0 &= 1. \end{aligned}$$

Therefore, the EIF for $\sigma^2 = \sum_t \sigma^{2,d}_t$ is simply $D^*_{\sigma^2} = \sum_t D^*_{\sigma^{2,d}_t}$.

With the EIF established, the TMLE of $\sigma^{2,d}_t$ is now defined as follows.

1. Estimates $g_{0:m^-,n} : m = 1, 2, \ldots, t$ are readily available if estimating $\mathbb{E}Y_d$ using a double robust estimator such as TMLE.

2. Set $\bar{Q}^{d,\sigma^2_t}_{t,n} = Z^d_i(t)$. Determine the range (a,b) for $Z^d_i(t)$, $i = 1, \ldots, n$ and target this initial fit using a parametric submodel respecting this range (a,b) by adding the clever covariate $H^{d,t}_t$ on the logistic scale (or by using $H^{d,t}_t$ as observational weight with clever covariate 1), using the initial fit as off-set. The resulting updated fit is denoted with $\bar{Q}^{d,\sigma^2_t,*}_{t,n}$.

3. Given $\bar{Q}^{d,\sigma^2_t,*}_{t,n}$, we can recursively for $m = t-1, t-2, \ldots, 1$:

   a) Regress the targeted fit $\bar{Q}^{d,\sigma^2_t,*}_{m^+,n}$ onto $\bar{A}(m^-) = d(\bar{l}(m^-)), \bar{L}(m^-)$, using logistic regression to respect the range (a,b). Denote the fit $\bar{Q}^{d,\sigma^2_t}_{m,n}$.

b) Target this initial fit respecting the range (a,b) with clever covariate $H^{d,t}_{m^-}$ on the logistic scale (or by using $H^{d,t}_{m^-}$ as observational weight with clever covariate 1), and denote this targeted fit of $\bar{Q}^{d,\sigma_t^2}_m$ with $\bar{Q}^{d,\sigma_t^2,*}_{m,n}$.

4. At $m=1$, we have the estimate $\bar{Q}^{d,\sigma_t^2,*}_{1,n}$, which now is a function of only $L(0)$. Finally, we take the average of $\bar{Q}^{d,\sigma_t^2,*}_{1,n}$ w.r.t. the empirical distribution of $L_i(0)$. The resulting $\hat{\sigma}^{2,d}_{t,n,TMLE} = \bar{Q}^{d,\sigma_t^2,*}_{0,n}$ is the desired TMLE of $\sigma_t^{2,d}$.

Application to single time-point treatment setting

For the sake of illustration, let us consider the method presented above for estimation of the variance of the EIF for the case that $O=(L(0),A(0),Y=L(1))$ and the target parameter is $\mathbb{E}Y^a$ for a static point treatment a.

In this case, the variance of the efficient influence curve is represented as

$$\begin{aligned}
\sigma_0^2 &= \mathbb{E}_0[D^*(P_0)(O)]^2 \\
&= \mathbb{E}_0\left[\frac{\mathbb{I}(A=a)}{g_0(a\mid L(0))}(Y-\bar{Q}^a_0(L(0)))+\bar{Q}^a_0(L(0))-\mathbb{E}Y^a\right]^2 \\
&= \mathbb{E}_0\left[\frac{\mathbb{I}(A=a)}{g_0(a\mid L(0))}(Y-\bar{Q}^a_0(L(0)))\right]^2+\mathbb{E}_0[\bar{Q}^a_0(L(0))-\mathbb{E}Y^a]^2 \\
&= \mathbb{E}_{P_0^a}\left[\frac{(Y^a-\bar{Q}^a_0(L(0)))^2}{g_0(a\mid L(0))}\right]+\mathbb{E}_0[\bar{Q}^a_0(L(0))-\mathbb{E}Y^a]^2.
\end{aligned} \tag{4.4}$$

If using a double robust estimator for the estimation of $\mathbb{E}Y^a$ such as TMLE, we are provided with estimators g_n and $\bar{Q}^*_n$ of $g_0(A\mid L(0))$ and $\bar{Q}^a_0(L(0))=\mathbb{E}[Y^a\mid L(0)]=\mathbb{E}_0[Y\mid A=a,L(0)]$. The second term in the final expression of Equation (4.4) is easily estimated with the empirical distribution. Given g_0 and $\bar{Q}_0$, the first term can be represented as the mean of a counterfactual $S^a(L^a(0))\equiv(Y^a-\bar{Q}^a_0(L(0)))^2/g_0(a\mid L(0))$ which needs to be estimated based on $(L(0),A,S^a(L(0),Y))$, where $S^a(L(0),Y)=(Y-\bar{Q}_0(a,L(0)))^2/g_0(a\mid L(0))$ represents the observed outcome. For example, we can use a TMLE estimator $\mathbb{E}^*_{a,n}S^a(L(0),Y^a)$ of $\mathbb{E}_0S^a(L(0),Y^a)=\mathbb{E}_{L(0),0}[\mathbb{E}_0[S^a\mid A=a,L(0)]]$. The TMLE estimate $\mathbb{E}^*_n[S^a\mid A=a,L(0)]$ of $\mathbb{E}_0[S^a\mid A=a,L(0)]$ is defined by determining the range (a,b) of $S^a(L_i(0),Y_i)$, obtaining an initial regression fit of $\mathbb{E}_0[S^a\mid L(0),A]$ that respects this range, representing it as a logistic regression fit bounded by (a,b), and updating the latter by fitting a univariate logistic regression with clever covariate $\mathbb{I}(A=a)/g_0(a\mid L(0))$, using the initial fit as an off-set. Regarding the initial fit $\mathbb{E}_n[S^a\mid A=a,L(0)]$, recall from above that S^a is a function of $L(0)$ which results in the initial fit being exactly $(Y-\bar{Q}_0(a,L(0)))^2/g_0(a\mid L(0))$ such that regression is unneeded. Following the update step, the TMLE of $\mathbb{E}_0S^a(L(0),Y^a)$ is now given by $\frac{1}{n}\sum_{i=1}^n\mathbb{E}^*_n[S^a\mid L_i(0),A=a]$, so that

$$\hat{\sigma}^{2,*}_n=\frac{1}{n}\sum_{i=1}^n\mathbb{E}^*_n[S^a(\bar{Q}^*_n,g_n)\mid A=a,L_i(0)]+\frac{1}{n}\sum_{i=1}^n(\bar{Q}^*_n(L_i(0),a)-\hat{\psi}^*_n)^2.$$

Advantages of the plug-in estimator of σ_0^2

Since σ_0^2/n equals the asymptotic variance of an asymptotically efficient estimator, it provides a good measure of the amount of information in the data for the target parameter of interest. Therefore, it is sensible to view σ_0^2/n as a measure of sparsity for the target parameter of interest. If g_n is a good estimator of g_0, then our proposed plug-in estimator $\hat{\sigma}_n^2$ is much less subject to underestimation due to sparsity than currently available estimators such as the sample variance of the estimated influence function, and the bootstrap-based estimate of the variance of an efficient estimator. Indeed, the non-parametric bootstrap generally is not valid, except when using a parametric model to estimate g_0 and $\bar{Q}_0$ which will never capture a true model in practice. This plug-in estimate $\hat{\sigma}_n^2$ represents a variance of the estimate of the EIF which involves the integration of rare combinations of treatment and covariates that are unlikely to occur in the actual sample.

In particular, if there are theoretical violations of the positivity assumption, then this true variance σ^2 equals infinity, and, if g_n approximates g_0 well, then also the estimate $\hat{\sigma}_n^2$ will generate very large values, demonstrating the lack of identifiability and thereby raising a red flag for finite sample sparsity bias in the estimators (beyond the large confidence intervals generated by $\hat{\sigma}_n^2$). We also note that if there are serious practical violations of the positivity assumption, then the estimate of this variance should be imprecise, since it is itself a highly variable estimator of a weakly identifiable parameter.

4.3 Variance estimation for substitution based estimators

The plug-in estimator of the asymptotic variance of the EIF presented above is superior to the more common approach of taking the empirical EIF variance over the sample (i.e., var$[D^*(P_n)(O)]$), in that there is a much stronger contribution of combinations of treatment and covariates that are unlikely to occur in the actual sample. In finite samples, however, the use of substitution based estimators such as TMLE (which are guaranteed to solve the EIF within a bounded range) will expectedly provide lower estimator variance due to the mere fact that they are guaranteed to respect the global constraints of the statistical model and target parameter mapping. That is, as opposed to estimating equations that tend to result in estimates outside parameter boundaries as the EIF variance increases, the use of substitution estimators in finite samples will retain an estimator variance that is smaller than the EIF variance divided by the sample size, n. Thus, using the newly presented robust EIF variance method will likely result in over-estimation of the estimator variance for these types of estimators.

One alternative approach at estimating the variance for substitution based estimators is to conduct a nonparametric bootstrap. The n observations are sampled with replacement and used to form an estimate of the parameter over B iterations. However, as stated above, the non-parametric bootstrap is generally invalid. Additionally, this is a very computer intensive method that usually requires estimating the full likelihood (i.e., P_0^d) of the longitudinal data structure within each sampled iteration and is therefore normally infeasible in practice unless conducted within an a prior selected smaller parametric statistical model such as logistic regression. In this section we thus

present an alternative bootstrap based approach that, unlike the standard non-parametric bootstrap, is both computationally feasible and theoretically valid. That is, this bootstrap approach allows us to estimate the variance of the estimator while avoiding re-estimation of g_0 and $\bar{Q}_0$. To facilitate this, we propose a modification of the usual TMLE such that the targeting step is separated.

Modified TMLE for $\mathbb{E}Y^d$

To reduce the computational burden that bootstrapping requires, we first present the modified TMLE approach at estimating the parameter $\mathbb{E}Y^d$. This parameter can be estimated by

1. Estimate $g_{0:t^-}(\bar{A},\bar{L}) : t = 1,2,\ldots,K+1$ and denote the fits $g_{0:t^-,n}$.

2. Determine the range (a,b) for $\mathbb{E}Y^d$. Recursively for $t = K+1,K,\ldots,1$, estimate the conditional expectation $\bar{Q}^d_t = \mathbb{E}[\bar{Q}^d_{t^+}|\bar{L}(t^-),\bar{A}(t^-) = d(\bar{l}(t^-))]$ respecting this range. Denote the fits $\bar{Q}^d_{t,n}$.

3. For time $t = K+1$, target the initial fit $\bar{Q}^d_{K+1,n}$ by using a parametric submodel respecting the range (a,b) by adding the covariate $H_{K+1}(g_n)$ (on the logistic scale), using the initial fits as off-set, and setting Y as the dependent variable. Denote this updated fit as $\bar{Q}^{d,*}_{K+1,n}$.

4. Given $\bar{Q}^{d,*}_{K+1,n}$, we can recursively for $t = K,K,K-1,\ldots,1$ target the initial fits $\bar{Q}^d_{t,n}$ by using parametric submodels respecting the range (a,b), adding the covariates $H_t(g_n)$ (on the logistic scale), using the initial fits as off-set, and setting $\bar{Q}^{d,*}_{t^+,n}$ as the dependent variable. Denote the updated fits as $\bar{Q}^{d,*}_{t,n}$.

5. At $t = 1$, we have the estimate $\bar{Q}^{d,*}_{1,n}$, which now is a function of only $L(0)$. Taking the average of $\bar{Q}^{d,*}_{1,n}$ w.r.t. the empirical distribution of $L_i(0)$ gives us the desired TMLE estimate of $\mathbb{E}Y^d$.

This estimator also solves the EIF and is therefore also asymptotically linear and efficient. To see this, first note that we are given an initial fit $P_n = (g_{0:t^-,n}, \bar{Q}^d_{t,n} : t = 1,2,\ldots,K+1)$. This fit is updated using a submodel such that $P^*_n = P_n(\varepsilon)$. Looking at this estimator minus the truth, we see that

$$\Psi(P^*_n) - \Psi(P_0) = -P_0 D^*(P_n)(O) + R_2(P_n - P_0)$$

for some remainder term $R_2(P_n - P_0)$. Assuming that P_n is a good enough fit, then it will converge to P_0 at a rate fast enough such that the remainder term is $o_p(1/\sqrt{n})$. Consequently, we have that (under empirical process conditions)

$$\begin{aligned}\Psi(P^*_n) - \Psi(P_0) &= (P_n - P_0)D^*(P^*_n)(O) \\ &= (P_n - P_0)D^*(P_0)(O) + o_p(1/\sqrt{n}).\end{aligned}$$

We emphasize that this estimator is proposed for the sake of the bootstrap method. It is recursive, in that each fit $\bar{Q}^d_{t,n}$ is dependent upon the fit at t^+. As opposed to the TMLE presented by van

der Laan and Gruber [46], the recursive nature of this TMLE is self contained within each step. In other words, each estimation step in this TMLE can be performed independently of the other steps. This allows the analyst to form all of the initial fits P_n prior to performing any of the targeted updates.

Bootstrapping the modified TMLE

If desired, the new TMLE approach presented above can be bootstrapped in a fully non-parametric manner, such that observations are drawn with replacement prior to fitting the full likelihood P_0^d and used to form an estimates of the parameter, leading to an estimator variance estimate. For the sake of computational ease, our general recommendation is to only bootstrap the targeting step. More specifically, once the fits $g_{0:t^-,n}$ and $\bar{Q}_{t,n}^d$ are formed for $t = 1,2,\ldots,K+1$, steps 3-5 above are carried out in the bootstrap such that for $b = 1,2,\ldots,B$ we have

$$P_{n,b}^* = P_n(\varepsilon_b)$$

for a user selected submodel $P(\varepsilon)$. The estimator variance is then estimated by taking the variance over the bootstrapped estimates, i.e., $\text{var}(\Psi(\hat{P}_n)) = \text{var}[\Psi(P_{n,b}^*)]$.

It is easy to see that the EIF is also solved for each bootstrapped sample in the updating step, i.e. $P_{n,b}^* D^*(P_n(\varepsilon_b))(O) = 0$. More precisely, if $P_n \to P_0$, then this bootstrap estimates $\sigma_0^2 = P_0 D^*(P_0)^2$ and thus consistently estimates the asymptotic variance of our TMLE. If, for example, only g_n is consistent, then this TMLE will be asymptotically linear with IF $D^*(\bar{Q}, g_0)$, which is a conservative influence function, since the influence function of the TMLE is $D^*(\bar{Q}, g_0)$ minus its projection onto the tangent space [48, 46]. We note that, just like the empirical variance of $D^*(P_n)(O)$, this inference is not double robust in sense that if, for example, g_n is inconsistent but $\bar{Q}_n$ is consistent, then it will not be consistent.

4.4 Simulations

The variance estimators presented in this paper provide novel estimators of the asymptotic variance of estimators solving the estimating equations corresponding to the EIF. Simulation studies presented in this section illustrate their applications in two settings: the estimation of the effect of treatment in a simpler point treatment setting, and the estimation of the effect of treatment on survival in a longitudinal observational study setting with three time points (i.e. $K+1 = 3$) under time-dependent confounding. To analyze performance, we first compare the variance estimation approaches covered above in estimating the estimator variance. Both the AIPW and TMLE estimators are considered in order to demonstrate the difference in estimating equations and substitution based estimators, respectively. The mean of the variance estimates are compared to the Monte-Carlo variance of each estimator. The Monte-Carlo variance of each variance approach is also reported. Additionally, we present the empirical coverage, Type I, and Type II errors resulting from each variance estimation approach.

Data generating distribution P_0

Point treatment setting

Consider a point treatment setting, such as patient enrollment into a care program, in which the treatment $A(0)$ is only assigned at a single time point. We are interested in determining whether the treatment of interest has a significant effect on the outcome on an additive scale. Our target parameter is therefore the difference of the mean outcomes under treatment and control, i.e., $\psi_0 = \mathbb{E}Y_1 - \mathbb{E}Y_0$. Under this setting, the simulated data were generated as follows:

$$\begin{aligned}
W_1, W_3 &\sim N(0,1), \text{ bounded at [-2,2]} \\
W_2 &\sim Ber(\text{logit}^{-1}(-1)) \\
L_1(0) &\sim N(0.1 + 0.4W_1, 0.5^2) \\
L_2(0) &\sim N(-0.55 + 0.5W_1 + 0.75W_2, 0.5^2) \\
\bar{g}_{0,0}(A(0)|Pa(A(0))) &= \text{logit}^{-1}(\beta_p - (\beta_p + 2.5)W_1 + 1.75W_2 \\
&\quad + (\beta_p + 3.2)L_1(0) - 1.8L_2(0) + 0.8L_1(0)L_2(0))) \\
\bar{Q}_{0,1}(Y|Pa(Y)) &= \text{logit}^{-1}(-0.5 + 1.2W_1 - 2.4W_2 - 1.8L_1(0) - 1.6L_2(0) \\
&\quad + L_1(0)L_2(0) - \beta_{\psi_0}A(0))
\end{aligned}$$

with a positivity associated parameter β_p ranging from -2 (minor positivity violations) to 1 (strong practical positivity violations) and the treatment effect associated parameter β_{ψ_0} ranging from 0 (no treatment effect) to 1 (strong treatment effect). Here, $L_1(0)$ and $L_2(0)$ are not time-dependent confounders and are therefore considered baseline covariates along with (W_1, W_2), which affect both the treatment and the outcome.

Longitudinal treatment setting

For the longitudinal setting, we considered a treatment $A(t)$ which was allowed to vary over time as a counting process. That is, if $A(t) = 1$ then we have that $\underline{A}(t) = 1$. Similar to the point treatment setting, we are interested in whether the treatment of interest has a significant effect on the outcome at the final time point $t^* = 3$ on an additive scale. Thus, our target parameter is the difference of the mean outcomes under treatment and control at this final time point, i.e., $\psi_0 = \mathbb{E}Y_1(t^*) - \mathbb{E}Y_0(t^*)$ where $Y(t^*) = L_3(3)$. Under this setting, data for the first time point was generated in the same manner as the point treatment setting above. For the remaining two time points, the data were generated conditional on survival (i.e. $L_3(t^-) = 0$) as follows:

$$
\begin{aligned}
L_1(t) &\sim N(0.1+0.4W_1+0.6L_1(t^-)-0.7L_2(t^-)+0.45\beta_{\psi_0}A(t^-),0.5^2)\\
L_2(t) &\sim N(-0.55+0.5W_1+0.75W_2+0.1L_1(t^-)+0.3L_2(t^-)\\
&\quad +0.75\beta_{\psi_0}A(t^-),0.5^2)\\
\bar{g}_{0,t}(A(t)|Pa(A(t))) &= \text{logit}^{-1}(\beta_p-(\beta_p+2.5)W_1+1.75W_2\\
&\quad +(\beta_p+3.2)L_1(t)-1.8L_2(t)+0.8L_1(t)L_2(t)))\\
\bar{Q}_{0,t}(L_3(t)|Pa(L_3(t))) &= \text{logit}^{-1}(-0.5+1.2W_1-2.4W_2-1.8L_1(t^-)-1.6L_2(t^-)\\
&\quad +L_1(t^-)L_2(t^-)-\beta_{\psi_0}A(t^-))
\end{aligned}
$$

Similar to the point treatment setting, the treatment effect associated parameter β_{ψ_0} also ranged from 0 to 1. We note, however, that the positivity issues faced in this scenario will be even more severe due to taking the cumulative probabilities of treatment over time, which resulted in smaller probabilities. We therefore considered only β_p values from -2 to 0 and imposed a truncation level of 0.001 to the estimates of $g_{0:t}$. Figure 4.1 shows the proportion of observations with truncated $g_{0:t}$ as a function of β_p at a null effect, i.e. $\beta_{\psi_0}=0$

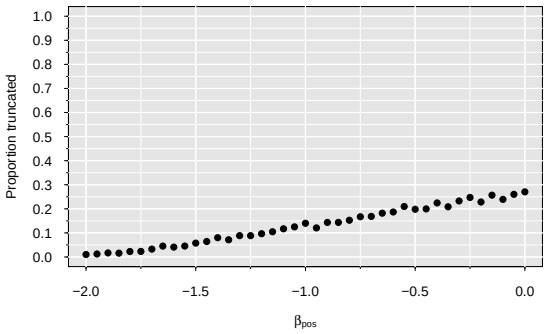

Figure 4.1: Proportion of observations with $g_{0:t}$ truncated.

Under these settings, the true parameter values ψ_0 were achieved by generating 8×10^7 observations under the counterfactual distribution for each β_{ψ_0} considered. Simulation results were obtained for 500 simulations of size $n=500$. Within each simulation, the bootstrap estimates of variance were formed from $B=1000$ iterations.

Submodels used

Any submodel and loss function for which its loss-function specific score

$$\frac{\partial}{\partial\varepsilon}L(P(\varepsilon))\Big|_{\varepsilon=0}$$

spans $D^*(P_0)(O)$ can be chosen in TMLE for both estimation of the mean outcome $\mathbb{E}Y_d$ and the variance of the EIF σ^2. As these submodels solve the equation corresponding to the EIF, they will all be asymptotically equivalent and thus, all asymptotically efficient. That is, no difference will be seen between the use of various submodels as the sample size grows to infinity. However, we note that the use of various submodels in finite samples can have varying performance. For example, under increasing levels of positivity violations the use of linear submodels which use $H_t(g)$ as a covariate can have higher variance due to observations with low probabilities of treatment acting as outliers which result in highly influential points for the estimation of the submodel parameter ε. Alternatively, using $H_t(g)$ as weights for the submodel avoids this issue.

Recall that the catalyst for this work was the anti-conservative estimates of estimator variance resulting from the use of the empirical EIF variance. We therefore wish to establish a robust estimator of the variance of estimators which solve the EIF, particularly under violations or near violations of positivity. In other words, we desire a variance estimator which will asymptotically converge to the true variance of the estimator, but also simultaneously act on the conservative side in finite samples. Keeping this in mind, we used two submodel and loss function combinations for our simulations. For the estimation of the target parameter and the robust estimator of the EIF variance, we used submodels which define $H_t(g)$ and $H_m^{d,t}(g)$ to be observational weights such that

$$\text{logit}\bar{Q}(\varepsilon) = \text{logit}\bar{Q} + \varepsilon,$$

acknowledging our slight abuse of notation. Alternatively, in our bootstrap approach at estimating the TMLE variance, we define the clever covariate using $H_t(g)$ such that

$$\text{logit}\bar{Q}(\varepsilon) = \text{logit}\bar{Q} + \varepsilon H_t(g).$$

Both submodels use, as loss function, the negative log-likelihood loss. As stated previously, both of the submodels presented solve the equation corresponding to the EIF and are therefore asymptotically equivalent.

Simulation results

Point treatment results

Figure 4.2 shows the Monte-Carlo variance under no treatment effect ($\beta_{\psi_0} = 0$) for both the AIPW and TMLE estimators, along with the mean of the variance estimates from each estimation approach. To keep the differences in perspective, we plotted results only for the positivity associated parameter $\beta_p \leq 0$. At the lower end of β_p where positivity violations are minor, the observed estimator variance is low for both the AIPW and TMLE estimators, with TMLE showing lower variance between the two. For example, at $\beta_p = -2$ the Monte-Carlo variance was 0.0027 and 0.0025 for the AIPW and TMLE estimators, respectively. Equivalently, the TMLE estimator resulted in a variance that is 6% lower than that of the AIPW estimator, despite solving the same estimating equation corresponding to the EIF. As β_p increased, introducing higher levels of positivity violations, the estimator variance increased for both estimators. Additionally, this occurred

at a noticeably higher rate for the AIPW estimator than for TMLE, resulting in an increase in the magnitude of difference between the two estimators. For example, at $\beta_p = 0$ the simulations resulted in a Monte-Carlo variance of 0.0167 and 0.0077 for the AIPW and TMLE estimators, respectively.

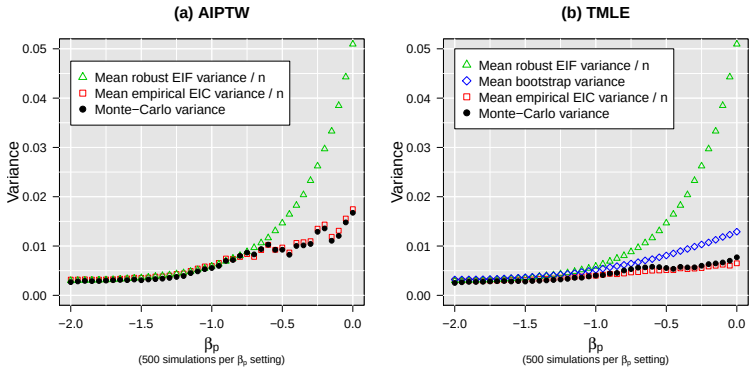

Figure 4.2: Mean of variance estimates for each estimator under no treatment effect ($\beta_{\psi_0} = 0$) at each positivity (β_p) value under the point treatment setting, overlaid with the estimator's Monte-Carlo variance. Robust variance estimates are identical for the two estimators.

For the AIPW estimator, the empirical EIF based approach of estimating variance performed well at estimating the variance of the estimator. For example, at $\beta_p = -2$ the mean of the EIF approach was 0.0031 compared with the Monte-Carlo variance of 0.0027. At $\beta_p = 0$, they were 0.0167 and 0.0175, respectively. Alternatively, the robust approach of estimating variance appeared to overestimate the variance at higher instances of practical positivity violations (e.g.$\beta_p = 0$).

In the TMLE estimator, all three approaches to variance estimation performed almost identically at low values of β_p. At $\beta_p = -2$, the mean of the estimates was 0.0030, 0.0030, and 0.0032 for the empirical EIF, robust, and bootstrapped based approaches, respectively, compared to the estimator's Monte-Carlo variance of 0.0025. As β_p increased and practical positivity issues arose, the empirical EIF approach tended to result in slightly anti-conservative estimates of variance, while the bootstrap approach resulted in slightly conservative estimates. As opposed to the AIPW estimator, the robust EIF approach severely overestimated the TMLE estimator variance.

Figure 4.3 shows the Monte-Carlo variance for each approach taken at estimating the estimator's variance. Lower values in this figure can be interpreted as coming from a variance estimator with more precision. In the AIPW estimator, the empirical EIF approach has noticeably higher variance than the robust approach, with a variance of 2.93 at $\beta_p = 0$ compared with 1.50 for the robust approach. This implies that the empirical EIF approach to estimating the AIPW estimator variance is less reliable than the robust EIF approach. For the TMLE estimator (Figure 4.3b), the empirical EIF approach to estimating variance showed much lower Monte-Carlo variance. The bootstrap approach also resulted in very low variance, implying a high reliability of this approach at estimating the variance.

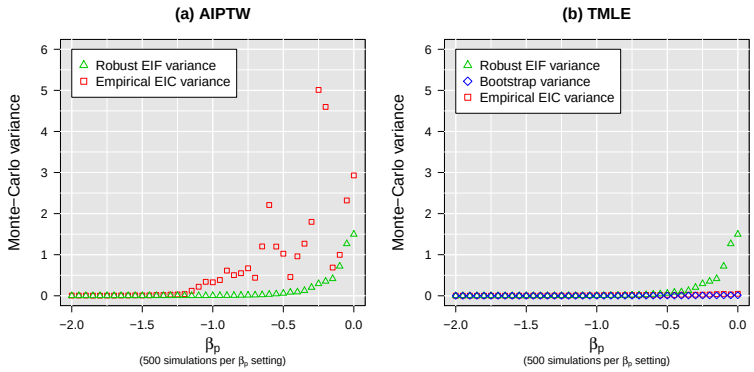

Figure 4.3: Monte-Carlo variance of variance estimators for each mean outcome estimator under no treatment effect ($\beta_{\psi_0} = 0$) at each positivity (β_p) value under the point treatment setting. Robust variance estimates are identical for the two mean outcome estimators.

We evaluated 95% confidence interval coverage for the TMLE estimator of $\mathbb{E}Y_d$ under the three approaches to variance estimation. Due to the lower variance seen in Figures 4.2 and 4.3, we focused only on the TMLE estimator here. Figure 4.4 shows a heat map overlaid with a contour plot of the resulting coverage estimates (i.e. the observed proportion of times the true parameters were captured by the confidence intervals) over the different combinations of β_{ψ_0} and β_p. Additionally, we estimated the power to reject the null hypothesis (at a level of 0.05) corresponding to each variance estimation approach under the range of treatment effect sizes and degrees of positivity violation considered above. Figure 4.5 shows a heat map overlaid with a contour plot of the resulting power estimates. Results at $\beta_{\psi_0} = 0$ can be interpreted as Type I errors, as they inform us of the times that the null hypothesis of no treatment effect is incorrectly rejected.

At low instances of positivity issues, coverage appears valid for all three variance estimation approaches with the proportion of time the true parameter was captured consistently at 0.95 or larger (Figure 4.4). Where positivity issues were low ($\beta_{\psi_0} < -0.5$), the empirical EIF approach maintained nominal to conservative coverage. Where severe positivity violations were present, coverage dropped substantially below 0.95. For example, at $\beta_p = 1$ coverage for this approach varied from 0.41 to 0.85. In contrast, the robust EIF approach consistently resulted in coverage at around $0.95 - 0.96$ at low values of β_p and *increased* with β_p, consistent with prior results showing overestimation of the variance under increasing positivity by this approach. For example, at $\beta_p = -.7$, coverage remained at 0.98 at all values of β_{ψ_0}. At $\beta_p \geq -0.1$, the observed coverage was almost always greater than or equal to 0.99 at all values of β_{ψ_0}. The bootstrap based coverage shown in Figure 4.4c varied the least, with coverage consistently between $0.95 - 0.97$ irrespective of the treatment effect (β_{ψ_0}) and positivity severity (β_p) considered.

Regarding the observed power (Figure 4.5), the empirical-EIF based variance approach resulted in the highest power among all three variance estimation approaches when an effect was present. For example, at $\beta_{\psi_0} = 1$ and $\beta_p = -1$, the observed power was 0.71, 0.51, and 0.51 for the empirical-EIF, robust-EIF, and bootstrap approaches respectively. While this result implies a more

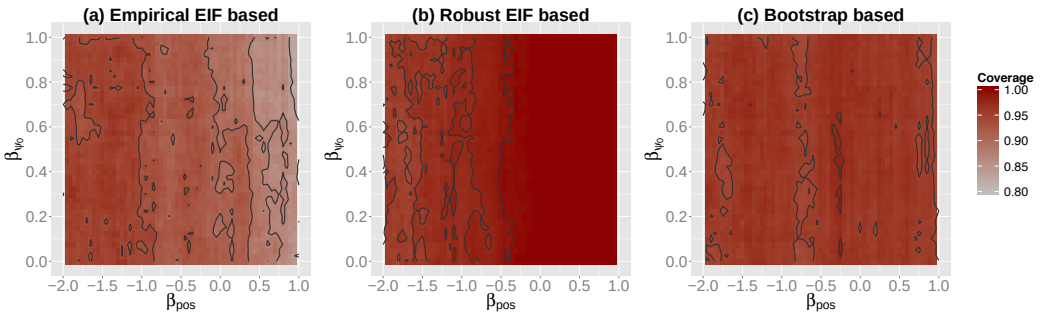

Figure 4.4: Simulated coverage for each variance estimation approach for the TMLE estimator under various treatment (β_{ψ_0}) and positivity (β_p) values under the point treatment setting.

efficient approach, it expectedly came at a cost of higher Type I error which became uncontrolled with an increase in β_p. For example, at $\beta_p = -2$ an observed 4.2% of the simulations incorrectly rejected the null hypothesis. This proportion increased to as high as 15% at $\beta_p = 1$. Alternatively, the robust EIF estimation approach resulted in low Type I errors (i.e. between $0 - 5.8\%$) with none of the simulations incorrectly rejecting the null beyond $\beta_p = -0.1$. The bootstrap approach resulted in an intermediate performance, with higher power than the robust EIF approach when an effect was present while simultaneously retaining appropriate control of the Type I error at all levels of β_p when no effect was present. For example, at $\beta_p = 1$ only 4.8% of the simulations incorrectly rejected the null hypothesis.

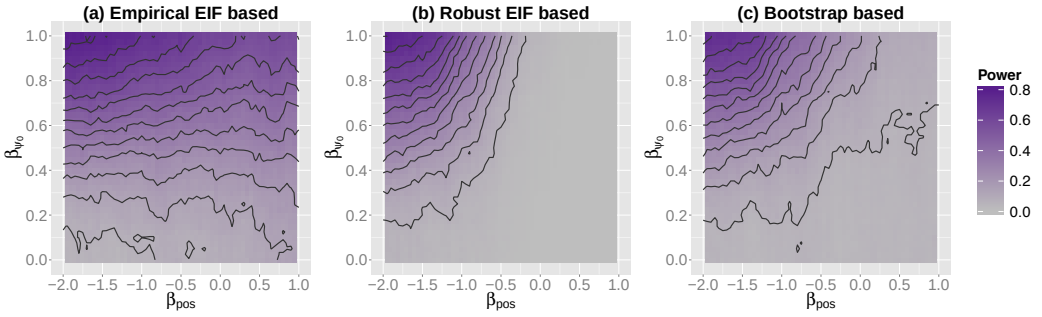

Figure 4.5: Simulated power for each variance estimation approach for the TMLE estimator under various treatment (β_{ψ_0}) and positivity (β_p) values under the point treatment setting.

Longitudinal treatment results

Results for the longitudinal setting were less stable, though still similar to the point treatment setting. Figure 4.6 shows the mean of the variance estimates under each approach, overlaid with the Monte-Carlo variance of the estimators. The same trend over the different levels of positivity was seen as in Figure 4.2, with the variance increasing with the magnitude of positivity issues. The

empirical EIF approach also performed well here at low levels of β_p for both the AIPW and TMLE estimators. At high values of β_p, the approach tended to slightly underestimate the variance both both intervention specific mean outcome estimators. Conversely, the robust EIF approach consistently over estimated the variance for both estimators. The bootstrap approach resulted in slightly conservative variance, though were still very similar to the Monte-Carlo variance esitmates.

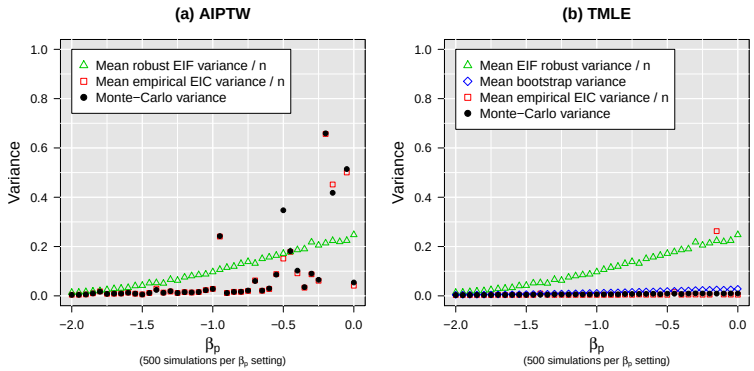

Figure 4.6: Mean of variance estimates for each estimator under no treatment effect ($\beta_{\psi_0} = 0$) at each positivity (β_p) value under the longitudinal treatment setting, overlaid with the estimator's Monte-Carlo variance. Robust variance estimates are identical for the two estimators.

Figure 4.7 shows the coverage corresponding to each variance estimation approach for the TMLE estimator of the intervention specific mean outcome. Coverage for the empirical EIF approach dropped considerably with an increase in positivity issues. For example, at a null effect (i.e. β_{ψ_0}) the observed coverage was 0.93 at $\beta_p = -2$ and 0.78 at $\beta_p = 0$. For the robust EIF approach, coverage increased with positivity. This became as high as 1.00 (i.e. all simulated confidence intervals captured the true parameter value) at higher levels of positivity issues. For the bootstrap approach, a higher level of coverage was also seen. For example, under a null effect, a coverage of 0.95 was observed at $\beta_p = -2$ and 0.98 at $\beta_p = 0$.

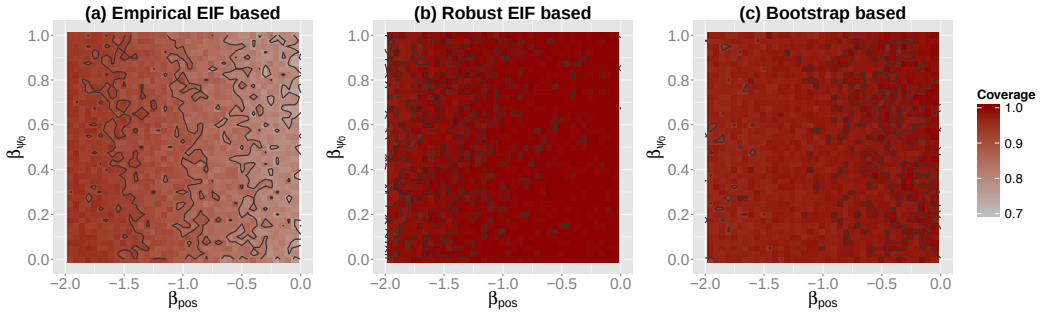

Figure 4.7: Simulated coverage for each variance estimation approach for the TMLE estimator under various treatment (β_{ψ_0}) and positivity (β_p) values under the longitudinal treatment setting.

Results for the Type I error and power were also similar to the point treatment setting. When there was an effect, the empirical EIF approach resulted in the highest power. At $\beta_{\psi_0} = 1$ and $\beta_p = -2$, we observed a power of 0.99. However, the Type I error was also uncontrolled here, becoming as high as 0.22 at $\beta_p = 0$. While the robust EIF approach maintained valid Type I error rates, the power for this approach when an effect was present was the lowest. For example, for an treatment effect size of $\beta_{\psi_0} = 1$ we observed a power ranging from 0.996 at $\beta_p = -2$ to 0.14 at $\beta_p = 0$. The bootstrap approach also resulted in controlled Type I error rates, with observed values below 0.05 over all values of β_p considered. Power was higher than the robust EIF approach across all values of β_{ψ_0} and β_p. For a treatment effect size of $\beta_{\psi_0} = 1$, we observed a power ranging from 0.988 at $\beta_p = -2$ to 0.40 at $\beta_p = 0$ for the bootstrap approach. Compared with the robust EIF approach, this is almost a 3-fold increase in power!

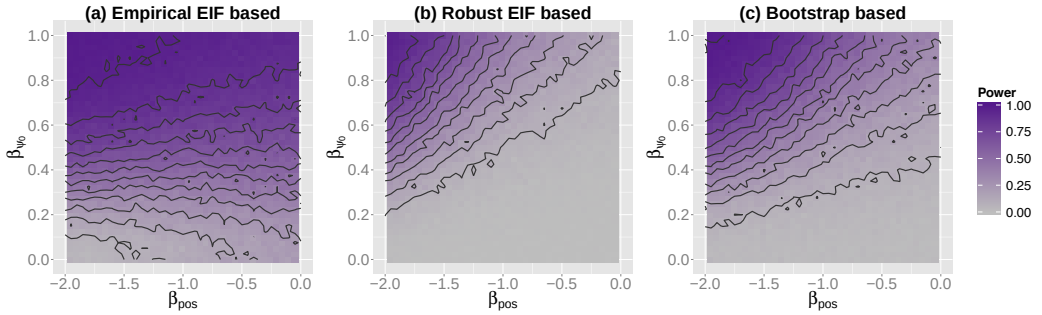

Figure 4.8: Simulated power for each variance estimation approach for the TMLE estimator under various treatment (β_{ψ_0}) and positivity (β_p) values under the longitudinal treatment setting.

4.5 Variance estimates for the impact of LREC enrollment in the IeDEA-EA cohort

Similar to the approach taken in Chapter 3.6, we condition on LREC availability for these analyses. To avoid potential theoretical positivity issues from subjects losing eligibility during follow up, we define our target parameter as the difference in the mean outcomes between the intervention to enforce immediate enrollment for all subjects (i.e. $d(\bar{l}(t^*)) = 1$) and the intervention to prevent subjects from ever enrolling into the program (i.e. $d(\bar{l}(t^*)) = 0$). Under both interventions, subjects are also intervened to never be censored. Our target parameter is therefore $\psi_0(t^*) = \mathbb{E}_0 Y_{d=1}(t^*) - \mathbb{E}_0 Y_{d=0}(t^*) : t^* = 1, 2, \ldots, 7$.

Nuisance parameters were estimated using Super Learning [47]. Candidates considered in the library included (i) an intercept only model (i.e. the mean), (ii) logistic regression [113], (iii) Bayesian logistic regression [20], and (iv) step-wise logistic regression using the Akaike Information Criterion [111]. A truncation level of 0.001 was applied to the probabilities of treatment.

For the robust estimation of the EIF variance, we bounded S_t^d and used a quasi-binomial logistic regression approach to estimate $\bar{Q}_m^{d,\sigma_t^2} : t = 1,2,\ldots,K+1, m = 0,1,\ldots,t$.

As the TMLE estimator out-performed the AIPW estimator in the simulations, we also considered only that estimator for this analysis. While our primary interest for this analysis is the additive treatment effect, positivity issues can have an effect on our variance estimates, potentially resulting in different inference or conclusions. We therefore considered all three variance estimation approaches.

Estimation results

Table 4.1: Point and standard error estimates for the additive treatment effect of LREC program enrollment. Bootstrap estimates were based on 1000 repetitions.

Time	Point	Standard Errors		
point	Estimate	Empirical EIF	Robust EIF	Bootstrap
1	-0.0001	0.0008	0.0010	0.0010
2	-0.0348	0.0032	0.0036	0.0035
3	-0.0523	0.0048	0.0070	0.0070
4	-0.0698	0.0056	0.0091	0.0077
5	-0.0681	0.0062	0.0107	0.0077
6	-0.0671	0.0065	0.0147	0.0088
7	-0.0660	0.0067	0.0229	0.0089

Table 4.1 shows the results of the applied analyses. The treatment effect estimates increased up to time $t^* = 4$ and plateaued indicating a (mostly) monotonic effect as time passed. Similar to the results in the simulations, the empirical EIF approach resulted in the lowest standard errors, while the robust EIF approach resulted in the highest estimates. This difference increased with time, with a factor of over three times the magnitude of the empirical EIF approach at $t^* = 7$. The bootstrap approach resulted in standard error estimates that were in between the two previous approaches, with estimates closer to those from the empirical EIF approach.

Figure 4.9 shows the corresponding 95% confidence intervals for the point estimates using each of the variance estimation approaches. As expected from the standard error estimates, the empirical EIF approach resulted in the tightest confidence intervals, while the robust EIF approach resulted in confidence interval lengths that were up to 3 times those from the empirical EIF approach. Confidence intervals created using the bootstrap approach were intermediate to the previous two, with lengths closer to those of the empirical EIF.

4.6 Discussion

The goal of the current study was to establish a consistent approach of estimating the variance of estimators solving the EIF which, in contrast to the common approach based on the empirI-

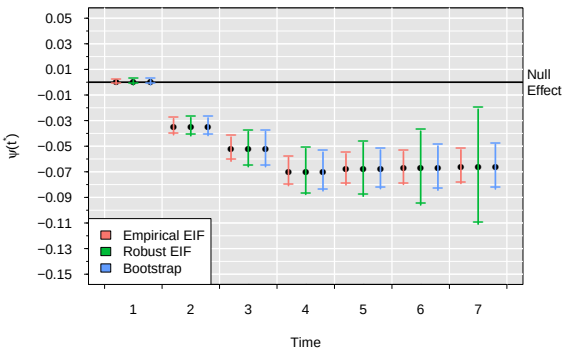

Figure 4.9: Point estimates and 95% confidence intervals for the additive treatment effect of LREC program enrollment. Bootstrap estimates were based on 1000 repetitions.

cal variance of the estimated EIF, do not act anti-conservatively when confronted with posItivity violations. We have presented two such novel approaches at estimating this variance: 1) a robust approach that directly targets the asymptotic variance of the EIF, and 2) a bootstrap approach based on fitting the initial density of the data once, followed by a non-parametric bootstrap of the targeting step. We noticed in simulations that the variance of AIPW increases with the variance of the EIF as positivity increases. In contrast, the variance of the TMLE was constrained in the face of increasing positivity violations by being a substitution estimator, and as a result, while the empirical EIF approach underestimated variance, the robust EIF approach increasingly over-estimated the variance as the degree of positivity violations increased. In contrast, the bootstrap based approach provided less conservative variance estimation, while maintaining Type I error control in the face of extreme positivity violations, both in the point treatment and longitudinal setting. In our applied setting, the empirical EIF approach also resulted in the lowest standard errors, while the robust EIF approach resulted in noticeably higher estimates. This resulted in noticeable 95% confidence interval length differences between the three approaches.

We emphasize that, as the robust EIF approach directly targets the asymptotic variance of the EIF, extremely large values of estimates from this can be used to raise a red flag for unreliable statistical inference due to sparsity, thereby declaring that the target parameter is practically not identifiable from the data, and that the reported variance estimates (though large) will themselves be imprecise. As such, we recommend that this approach be used if there is concern regarding identifiability of the data for the target parameter of interest.

While the EIF can raise a red flag for lack of identifiability, for substitution estimators such as TMLE we suggest that it is overly conservative for constructing valid confidence intervals and tests in finite sample in the face of substantial positivity violations. In previous work [64], we suggested the parametric bootstrap as a diagnostic tool for sparsity-bias in the point treatment setting. The approach can become cumbersome, as the analyst would need to refit the whole likelihood for each iteration of the bootstrap. Our robust EIF approach is able to avoid estimating the whole

likelihood by targeting the required means under the post intervention distribution defined by the longitudinal g-computation formula directly. Even if we use an actual TMLE of P_0^d, our analytic estimate of the variance is still much less computer intensive than the parametric bootstrap method, in particular, in view that one would need to run many replicate samples of the data set in order to pick up the observations that correspond with a rare treatment and thus contribute large numbers to the variance expression. Therefore, we believe that the proposed analytic method will be (at least, practically) superior to the earlier proposed parametric bootstrap method. Our presented bootstrap approach, while more computationally intensive than the robust EIF approach, is also superior to the earlier proposed approach, in that we do not have to refit the entire likelihood under each iteration. This also significantly reduces the computational resources required to obtain estimates of the target parameter, particularly if computationally intensive non-parametric machine learning algorithms are used to estimate these densities.

Further refinements can be applied in an attempt to obtain variance estimates with an even smaller finite sample bias. One such approach is a convex combination of the variance estimators considered above. For example, we noticed in supplementary analyses that taking

$$\hat{\alpha}_n \hat{\sigma}^2_{eEIF,n} + (1-\hat{\alpha}_n)\hat{\sigma}^2_{rEIF,n}$$

had good performance, where $\hat{\sigma}^2_{eEIF,n}$ is the variance estimate using the empirical EIF approach, $\hat{\sigma}^2_{rEIF,n}$ is the variance estimate using the robust EIF approach, and $\hat{\alpha}_n = |\hat{\sigma}^2_{rEIF,n} - \hat{\sigma}^2_{eEIF,n}|/(\hat{\sigma}^2_{rEIF,n} + \hat{\sigma}^2_{eEIF,n})$. We note, however, that such an approach is ad-hoc and may lead to varying results in other simulations or distributions. We therefore chose not to present the results here.

A potential limitation of the robust approach at estimating the variance involves the conditions for asymptotic linearity to be met. Specifically, we have the requirement that $\bar{Q}_t^d : t = 0, 1, \ldots, K+1$ and g_0 be estimated consistently. Without this, we do not have a limiting distribution so that we are not estimating a well defined σ_0^2 and inference is impossible. Furthermore, it is also required that $\bar{Q}_0^{d,\sigma_t^2}$ be estimated both consistently and at a fast enough rate. In our simulations, we used a simple parametric model to estimate $\bar{Q}_0^{d,\sigma_t^2}$, though a more non-parametric approach such as Super Learning could have been applied. Such an approach is extremely computationally expensive and infeasible in practice, as it would have to be applied to estimate the variance of the EIF at every single time point, σ_t^2. In this regard, the bootstrap approach is superior as it does not require the additional estimation of $\bar{Q}_0^{d,\sigma_t^2}$.

It would be of interest to further evaluate not only the practical performance of these variance estimation approaches in future studies, but also the application of the approaches to other parameters. Appendix B derives the general approach for working marginal structural models. Further research into the practical performance of this approach is needed for this setting. These variance estimation approaches can also be used for more complex parameters, such as mean outcomes under dynamic regimes, stochastic interventions, or marginal structural working models. Future research could also develop a collaborative TMLE [45] or cross-validated [117] based approach at robustly estimating the EIF variance.

Bibliography

[1] Ns Altman. "An introduction to kernel and nearest-neighbor nonparametric regression". In: *The American Statistician* 46.3 (1992), pp. 175–185. ISSN: 0003-1305. DOI: 10.1080/00031305.1992.10475879.

[2] Suzanna Attia et al. "Sexual transmission of HIV according to viral load and antiretroviral therapy: systematic review and meta-analysis." In: *AIDS (London, England)* 23.11 (July 2009), pp. 1397–404. ISSN: 1473-5571. DOI: 10.1097/QAD.0b013e32832b7dca. URL: http://www.ncbi.nlm.nih.gov/pubmed/19381076.

[3] Heejung Bang and James M Robins. "Doubly robust estimation in missing data and causal inference models." In: *Biometrics* 61.4 (Dec. 2005), pp. 962–73. ISSN: 0006-341X. DOI: 10.1111/j.1541-0420.2005.00377.x. URL: http://www.ncbi.nlm.nih.gov/pubmed/16401269.

[4] L. M. Bodnar. "Marginal Structural Models for Analyzing Causal Effects of Time-dependent Treatments: An Application in Perinatal Epidemiology". In: *American Journal of Epidemiology* 159.10 (May 2004), pp. 926–934. ISSN: 0002-9262. DOI: 10.1093/aje/kwh131. URL: http://aje.oupjournals.org/cgi/doi/10.1093/aje/kwh131.

[5] Bernhard E. Boser, Isabelle M. Guyon, and Vladimir N. Vapnik. "A Training Algorithm for Optimal Margin Classifiers". In: *Proceedings of the 5th Annual ACM Workshop on Computational Learning Theory* (1992), pp. 144–152. ISSN: 0-89791-497-X. DOI: 10.1.1.21.3818. URL: http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.21.3818.

[6] Jordan C. Brooks et al. "Targeted Minimum Loss-Based Estimation of Causal Effects in Right-Censored Survival Data with Time-Dependent Covariates: Warfarin, Stroke, and Death in Atrial Fibrillation". In: *Journal of Causal Inference* 1.2 (Jan. 2013), pp. 235–254. ISSN: 2193-3677. DOI: 10.1515/jci-2013-0001. URL: http://www.degruyter.com/view/j/jci.2013.1.issue-2/jci-2013-0001/jci-2013-0001.xml.

[7] J. Bryan. "Analysis of longitudinal marginal structural models". In: *Biostatistics* 5.3 (July 2004), pp. 361–380. ISSN: 1465-4644. DOI: 10.1093/biostatistics/kxg041. URL: http://biostatistics.oupjournals.org/cgi/doi/10.1093/biostatistics/kxg041.

[8] CM Cassel, Carl-Erik Sarndal, and Jan Wretman. "Some uses of statistical models in connection with the nonresponse problem". In: *Incomplete Data in Sample Surveys III. Symposium on Incomplete Data, Proceedings*. New York: Academic Press, 1983.

[9] MS Cohen et al. "Prevention of HIV-1 infection with early antiretroviral therapy". In: *The New England journal of medicine* 365.6 (2011), pp. 493–505. URL: http://www.nejm.org/doi/full/10.1056/nejmoa1105243.

[10] Stephen R Cole and Miguel A Hernan. "Constructing Inverse Probability Weights for Marginal Structural Models". In: *American Journal of Epidemiology* 168.6 (2008), pp. 656–664. DOI: 10.1093/aje/kwn164.

[11] Corinna Cortes and Vladimir N. Vapnik. "Support-Vector Networks". In: *Machine Learning* 20 (1995), pp. 273–297. ISSN: 1747-0285.

[12] Christine Danel et al. "Conference on Retroviruses and Opportunistic Infections". In: *Early ART and IPT in HIV-Infected African Adults With High CD4 Count (Temprano Trial)*. Vol. 17. 5. May 2015. DOI: 10.1016/S0737-0806(97)80016-0.

[13] Moupali Das et al. "Decreases in community viral load are accompanied by reductions in new HIV infections in San Francisco." In: *PloS one* 5.6 (Jan. 2010), e11068. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0011068. URL: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2883572%7B%5C&%7Dtool=pmcentrez%7B%5C&%7Drendertype=abstract.

[14] Anna L. Decker et al. "Semiparametric Estimation of the Impacts of Longitudinal Interventions on Adolescent Obesity using Targeted Maximum-Likelihood: Accessible Estimation with the ltmle Package". In: *Journal of Causal Inference* 2.1 (Jan. 2014), pp. 95–108. ISSN: 2193-3677. DOI: 10.1515/jci-2013-0025. URL: http://www.degruyter.com/view/j/jci.2014.2.issue-1/jci-2013-0025/jci-2013-0025.xml.

[15] CW Dieffenbach and AS Fauci. "Thirty years of HIV and AIDS: future challenges and opportunities". In: *Annals of internal medicine* 154.11 (2011), pp. 766–771. URL: http://annals.org/article.aspx?articleid=746972.

[16] David A. Freedman. *Statistical Models: Theory and Practice*. New York: Cambridge University Press, 2009, p. 458. ISBN: 0521743850.

[17] Jerome Friedman. "Greedy Function Approximation: A Gradient Boosting Machine". In: *Annals of statistics* 29.5 (2001), pp. 1189–1232.

[18] Jerome H Friedman. "Multivariate Adaptive Regression Splines". In: 19.1 (1991), pp. 1–67.

[19] Jerome H Friedman. "Stochastic Gradient Boosting". In: *Computational Statistics and Data Analysis* 38.4 (2002), pp. 367–378.

[20] Andrew Gelman et al. "A weakly informative default prior distribution for logistic and other regression models". In: *The Annals of Applied Statistics* 2.4 (2008), pp. 1360–1383. DOI: 10.1214/08-AOAS191.

[21] T. P. Giordano et al. “Retention in Care: A Challenge to Survival with HIV Infection”. In: *Clinical Infectious Diseases* 44.11 (June 2007), pp. 1493–1499. ISSN: 1058-4838. DOI: 10.1086/516778. URL: http://cid.oxfordjournals.org/lookup/doi/10.1086/516778.

[22] Thomas P Giordano et al. “Patients referred to an urban HIV clinic frequently fail to establish care: factors predicting failure.” In: *AIDS care* 17.6 (Aug. 2005), pp. 773–83. ISSN: 0954-0121. DOI: 10.1080/09540120412331336652. URL: http://www.ncbi.nlm.nih.gov/pubmed/16036264.

[23] RM Gulick et al. “Treatment with indinavir, zidovudine, and lamivudine in adults with human immunodeficiency virus infection and prior antiretroviral therapy”. In: *The New England journal of medicine* 337 (1997), pp. 734–739. URL: http://www.nejm.org/doi/full/10.1056/NEJM199709113371102.

[24] Roee Gutman and Donald B Rubin. “Estimation of causal effects of binary treatments in unconfounded studies.” In: *Statistics in medicine* April (May 2015). ISSN: 1097-0258. DOI: 10.1002/sim.6532. URL: http://www.ncbi.nlm.nih.gov/pubmed/26013308.

[25] Frank R Hampel. “The Influence Curve and its Role in Robust Estimation”. In: *Journal of the American Statistical Association* 69.346 (1974), pp. 383–393. ISSN: 0162-1459. DOI: 10.1080/01621459.1974.10482962. URL: http://www.tandfonline.com/doi/abs/10.1080/01621459.1974.10482962.

[26] P. Han and L. Wang. “Estimation with missing data: beyond double robustness”. In: *Biometrika* 100.2 (Apr. 2013), pp. 417–430. ISSN: 0006-3444. DOI: 10.1093/biomet/ass087. URL: http://biomet.oxfordjournals.org/cgi/doi/10.1093/biomet/ass087.

[27] Trevor Hastie and Robert Tibshirani. *Generalized Additive Models*. 1990, p. 352. DOI: 978-0-412-34390-2.

[28] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *Elements of Statistical Learning*. 2nd ed. Stanford, CA: Springer, 2008, p. 745. DOI: 0387848576.

[29] Trevor Hastie et al. “The Entire Regularization Path for the Support Vector Machine”. In: *Test* 5.2 (2004), pp. 1391–1415. ISSN: 15324435. DOI: 10.1145/347090.347165. URL: http://portal.acm.org/citation.cfm?id=1005332.1044706.

[30] Satoshi Hattori and Masayuki Henmi. “Stratified doubly robust estimators for the average causal effect.” In: *Biometrics* 70.2 (June 2014), pp. 270–7. ISSN: 1541-0420. DOI: 10.1111/biom.12157. URL: http://www.ncbi.nlm.nih.gov/pubmed/24571129.

[31] M a Hernán, B Brumback, and J M Robins. “Marginal structural models to estimate the causal effect of zidovudine on the survival of HIV-positive men.” In: *Epidemiology (Cambridge, Mass.)* 11.5 (2000), pp. 561–570. ISSN: 1044-3983.

[32] Miguel a Hernán and James M Robins. "Estimating causal effects from epidemiological data." In: *Journal of epidemiology and community health* 60.7 (July 2006), pp. 578–86. ISSN: 0143-005X. DOI: 10.1136/jech.2004.029496. URL: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2652882%7B%5C&%7Dtool=pmcentrez%7B%5C&%7Drendertype=abstract.

[33] Arthur E Hoerl and Robert W Kennard. "Ridge Regression: Biased Estimation for Nonorthogonal Problems". In: *Technometrics* 12.1 (1970), pp. 55–67. ISSN: 0040-1706. DOI: 10.1080/00401706.1970.10488634.

[34] Elizabeth Horstmann et al. "Retaining HIVfffdfffdfffdInfected Patients in Care: Where Are We? Where Do We Go from Here?" In: *Clinical Infectious Diseases* 50 (Feb. 2010), pp. 100201102709029–000. ISSN: 1058-4838. DOI: 10.1086/649933. URL: http://cid.oxfordjournals.org/lookup/doi/10.1086/649933.

[35] DG Horvitz and DJ Thompson. "A Generalization of Sampling Without Replacement From a Finite Universe". In: *Journal of the American Statistical Association* 47.260 (1952), pp. 663–685.

[36] Joseph DY Kang and Joseph L Schafer. "Demystifying Double Robustness: A Comparison of Alternative Strategies for Estimating a Population Mean from Incomplete Data". In: *Statistical Science'* 22.4 (Jan. 2007), pp. 523–239. ISSN: 0883-4237. DOI: 10.1214/07-STS227. arXiv: arXiv:0804.2958v1. URL: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2397555%7B%5C&%7Dtool=pmcentrez%7B%5C&%7Drendertype=abstract.

[37] Olivia Keiser, P Taffé, and Marcel Zwahlen. "All cause mortality in the Swiss HIV Cohort Study from 1990 to 2001 in comparison with the Swiss population". In: *Aids* March (2004), pp. 1835–1843. URL: http://journals.lww.com/aidsonline/Abstract/2004/09030/All%7B%5C_%7Dcause%7B%5C_%7Dmortality%7B%5C_%7Din%7B%5C_%7Dthe%7B%5C_%7DSwiss%7B%5C_%7DHIV%7B%5C_%7Dcohort%7B%5C_%7Dstudy.13.aspx.

[38] Patricia Kissinger et al. "Compliance with public sector HIV medical care." In: *Journal of the National Medical Association* 87.1 (1995), pp. 19–24. URL: http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2607741/.

[39] Rex B. Kline. *Principles and Practice of Structural Equation Modeling*. 3rd. New York, 2011, p. 428. ISBN: 978-1-60623-877-6.

[40] Olivier Koole and Robert Colebunders. "ART in low-resource settings: how to do more with less." In: *Lancet* 376.9739 (Aug. 2010), pp. 396–8. ISSN: 1474-547X. DOI: 10.1016/S0140-6736(10)61020-3. URL: http://www.ncbi.nlm.nih.gov/pubmed/20638119.

[41] H Koul, V Susarla, and J Van Ryzin. "Regression analysis with randomly right-censored data". In: *The Annals of Statistics* 9.6 (1981), pp. 1276–1288. URL: http://www.jstor.org/stable/2240417.

[42] T Kredo et al. *Decentralising HIV treatment in lower- and middle-income countries (Review)*. Tech. rep. 6. Tygerberg: The Cochrane Collaboration, 2013, p. 78. DOI: 10.1002/14651858.CD009987.pub2..

[43] Mark J van der Laan. "Targeted Learning : Application to Optimal Dynamic Treatments". In: (2013), pp. 1–48.

[44] Mark J van der Laan and Sandrine Dudoit. "Unified Cross-Validation Methodology For Selection Among Estimators and a General Cross-Validated Adaptive Epsilon-Net Estimator : Finite Sample Oracle Inequalities and Examples". In: *The Berkeley Electronic Press* 130 (2003).

[45] Mark J van der Laan and Susan Gruber. "Collaborative Double Robust Targeted Maximum Likelihood estimation". In: *The international journal of biostatistics* 6.1 (2010).

[46] Mark J van der Laan and Susan Gruber. "Targeted Minimum Loss Based Estimation of an Intervention Specific Mean Outcome". In: *The Berkeley Electronic Press* 290 (2011).

[47] Mark J van der Laan, Eric C Polley, and Alan E Hubbard. "Super Learner". In: *U.C. Berkeley Division of Biostatistics Working Paper Series* 222 (2007), pp. 1–20.

[48] Mark J van der Laan and James M Robins. *Unified methods for censored longitudinal data and causality*. New York: Springer, 2003. DOI: 0387955569.

[49] Mark J van der Laan and Sherri Rose. *Targeted Learning*. New York: Springer, 2011. ISBN: 9781441997814.

[50] Mark J. van der Laan, Sherri Rose, and Susan Gruber. *Readings on targeted maximum likelihood estimation*. Tech. rep. Bepress, 2009. URL: http://www.bepress.com/ucbbiostat/paper254.

[51] L Li, K Kleinman, and M W Gillman. "A comparison of confounding adjustment methods with an application to early life determinants of childhood obesity." In: *Journal of developmental origins of health and disease* 5.6 (Dec. 2014), pp. 435–47. ISSN: 2040-1752. DOI: 10.1017/S2040174414000415. URL: http://www.ncbi.nlm.nih.gov/pubmed/25171142.

[52] RJ Little and Donald B Rubin. *Statistical Analysis with Missing Data*. 2nd. New York: Wiley, 2002, p. 381. ISBN: 0-471-18386-5.

[53] Roderick Little and Hyonggin An. "Robust Likelihood-Based Analysis of Multivariate Data with Missing Values". In: *Statistica Sinica* 14 (2004), pp. 949–968.

[54] Jens Lundgren et al. "Why START? Reflections that led to the conduct of this large long-term strategic HIV trial." In: *HIV medicine* 16 Suppl 1 (Apr. 2015), pp. 1–9. ISSN: 1468-1293. DOI: 10.1111/hiv.12227. URL: http://www.ncbi.nlm.nih.gov/pubmed/25711317.

[55] Margaret May et al. “Impact of late diagnosis and treatment on life expectancy in people with HIV-1: UK Collaborative HIV Cohort (UK CHIC) Study.” In: *BMJ (Clinical research ed.)* 343 (Jan. 2011), p. d6016. ISSN: 1756-1833. DOI: 10.1136/bmj.d6016. URL: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3191202%7B%5C&%7Dtool=pmcentrez%7B%5C&%7Drendertype=abstract.

[56] Warren S. McCulloch and Walter Pitts. “A logical calculus of the ideas immanent in nervous activity”. In: *The Bulletin of Mathematical Biophysics* 5.4 (1943), pp. 115–133. ISSN: 00074985. DOI: 10.1007/BF02478259.

[57] Romain Neugebauer and Mark J Van Der Laan. “G-computation Estimation of Nonparametric Causal Effects on Time-Dependent Mean Outcomes in Longitudinal Studies G-computation Estimation of Nonparametric Causal Effects on Time-Dependent Mean Outcomes in Longitudinal Studies”. In: *The Berkeley Electronic Press* 183 (2005).

[58] Romain Neugebauer and Mark J Van Der Laan. “Locally Efficient Estimation of Nonparametric Causal Effects on Mean Outcomes in Longitudinal Studies Locally Efficient Estimation of Nonparametric Causal Effects on Mean Outcomes in Longitudinal Studies”. In: *The Berkeley Electronic Press* 134 (2003).

[59] Frank J Palella et al. “Declining morbidity and mortality among patients with advanced human immunodeficiency virus infection”. In: *The New England journal of medicine* 338 (1998), pp. 853–860. URL: http://www.nejm.org/doi/full/10.1056/NEJM199803263381301.

[60] Judea Pearl. “Causal diagrams for empirical research”. In: *Biometrika* 82.4 (1995), pp. 669–688. ISSN: 00063444. DOI: 10.1093/biomet/82.4.669.

[61] Judea Pearl. *Causality*. 2nd ed. New York: Cambridge University Press, 2009. ISBN: 978-0-521-89560-6.

[62] Maya L Petersen. “Commentary: Applying a causal road map in settings with time-dependent confounding.” In: *Epidemiology (Cambridge, Mass.)* 25.6 (Nov. 2014), pp. 898–901. ISSN: 1531-5487. DOI: 10.1097/EDE.0000000000000178. URL: http://www.ncbi.nlm.nih.gov/pubmed/25265135.

[63] Maya L Petersen et al. “Delayed switch of antiretroviral therapy after virologic failure associated with elevated mortality among HIV-infected adults in Africa”. In: *AIDS* 28 (2014), pp. 2097–2107. DOI: 10.1097/QAD.0000000000000349.

[64] Maya L Petersen et al. “Diagnosing and responding to violations in the positivity assumption.” In: *Statistical methods in medical research* 21.1 (Feb. 2012), pp. 31–54. ISSN: 1477-0334. DOI: 10.1177/0962280210386207. URL: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=4107929%7B%5C&%7Dtool=pmcentrez%7B%5C&%7Drendertype=abstract.

[65] Maya L Petersen et al. “Targeted Maximum Likelihood Estimation for Dynamic and Static Longitudinal Marginal Structural Working Models Targeted Maximum Likelihood Estimation for Dynamic and Static Longitudinal Marginal Structural Working Models”. In: *The Berkeley Electronic Press* 312 (2013).

[66] Maya Petersen and Mark J van der Laan. "Causal Models and Learning from Data: Integrating Causal Modeling and Statistical Estimation". In: *Epidemiology* 25.3 (2014), pp. 418–426. DOI: 10.1097/EDE.0000000000000078.Causal.

[67] Eric Polley and Mark J van der Laan. *SuperLearner: Super Learner Prediction*. 2014. URL: https://github.com/ecpolley/SuperLearner.

[68] Thomas C Quinn et al. "Viral load and heterosexual transmission of human immunodeficiency virus type 1". In: *The New England journal of medicine* 342 (2000), pp. 921–929. URL: http://www.nejm.org/doi/full/10.1056/NEJM200003303421303.

[69] C Roberts and SA Roberts. "Design and analysis of clinical trials with clustering effects due to treatment". In: *Clinical Trials* 2 (2005), pp. 152–162. URL: http://ctj.sagepub.com/content/2/2/152.short.

[70] James M Robins. "A graphical approach to the identification and estimation of causal parameters in mortality studies with sustained exposure periods". In: *Journal of Chronic Disease* 40.2 (1987), 139S–161S.

[71] James M Robins. "A New Approach to Causal Inference in Mortality Studies with a Sustained Exposure Period - Application to Control of the Healthy Worker Survivor Effect". In: *Mathematical Modelling* 7 (1986), pp. 1393–1512.

[72] James M. Robins. "Association, Causation, and Marginal Structural Models". In: *Synthese* 121 (1999), pp. 151–179.

[73] James M Robins. "Marginal Structural Models". In: *Computer Music Journal* 7.2 (1983), pp. 43–55.

[74] James M Robins. "Marginal Structural Models". In: *1997 Proceedings of the American Statistical Association, Section on Bayesian Statistical Science* (1998), pp. 1–10. URL: http://link.springer.com/chapter/10.1007/978-1-4419-9782-1%7B%5C_%7D9.

[75] James M Robins. "Robust Estimation in Sequentially Ignorable Missing Data and Causal Inference Models". In: *Proceedings of the American Statistical Association Section on Bayesian Statistical Science* (2000), pp. 6–10.

[76] James M Robins and Miguel A Hernán. "Estimation of the causal effects of time-varying exposures". In: *Longitudinal Data Analysis*. Ed. by Garrett M Fitzmaurice et al. i. CRC Press, 2009. Chap. 1.

[77] James M. Robins, Miguel Ángel Hernán, and Babette Brumback. "Marginal Structural Models and Causal Inference in Epidemiology". In: *Epidemiology* 11.5 (Sept. 2000), pp. 550–560. ISSN: 1044-3983. DOI: 10.1097/00001648-200009000-00011. URL: http://content.wkhealth.com/linkback/openurl?sid=WKPTLP:landingpage%7B%5C&%7Dan=00001648-200009000-00011.

[78] James M. Robins and Andrea Rotnitzky. "Comment on the Bickel and Kwon article, fffdfffdfffdInference for semiparametric models: some questions and an answer.fffdfffdfffd". In: *Statistica Sinica* 11.4 (2001), pp. 920–936.

[79] James M Robins and Andrea Rotnitzky. "Recovery of Information and Adjustment for Dependent Censoring Using Surrogate Markers". In: *AIDS Epidemiology*. Ed. by Nicholas P Jewell, Klaus Dietz, and Vernon T Farewell. Boston: Birkhäuser, 1992. Chap. 3, pp. 297–331. DOI: 10.1007/978-1-4757-1229-2{_}14.

[80] James M Robins, Andrea Rotnitzky, and Mark J van der Laan. "Discussion of fffdfffdfffdOn profile likelihoodfffdfffdfffd by Murphy and van der Vaart". In: *Journal of the American Statistical Association* 95.450 (2000), pp. 477–482.

[81] James M. Robins, Andrea Rotnitzky, and Lue Ping Zhao. "Estimation of Regression Coefficients When Some Regressors Are Not Always Observed". In: *Journal of the American Statistical Association* 89.427 (1994), pp. 846–867.

[82] James Robins et al. "Comment : Performance of Double-Robust Estimators When fffdfffdfffd Inverse Probability fffdfffdfffd Weights Are Highly Variable". In: *Statistical Science* 22.4 (2007), pp. 544–559. DOI: 10.1214/07-STS227.

[83] JM Robins. "Marginal structural models versus structural nested models as tools for causal inference". In: *Statistical models in epidemiology, the environment, ...* (1999), pp. 1–30. URL: http://link.springer.com/chapter/10.1007/978-1-4612-1284-3%7B%5C_%7D2.

[84] JM Robins, S Greenland, and FC Hu. "Estimation of the causal effect of a time-varying exposure on the marginal mean of a repeated binary outcome". In: *Journal of the American ...* 94.447 (1999), pp. 687–700. URL: http://amstat.tandfonline.com/doi/abs/10.1080/01621459.1999.10474168.

[85] Alison J Rodger et al. "Mortality in well controlled HIV in the continuous antiretroviral therapy arms of the SMART and ESPRIT trials compared with the general population." In: *AIDS (London, England)* 27.6 (Mar. 2013), pp. 973–9. ISSN: 1473-5571. DOI: 10.1097/QAD.0b013e32835cae9c. URL: http://www.ncbi.nlm.nih.gov/pubmed/23698063.

[86] P R Rosenbaum and D B Rubin. "a. fffdfffdfffdThe Central Role of the Propensity Score in Observational Studies for Causal Effectsfffdfffdfffd Biometrika ." In: 70 SRC -.1 (1983), pp. 41–55.

[87] Andrea Rotnitzky and James Robins. *Inverse probability weighted estimation in survival analysis* . 2005.

[88] Daniel Rubin. "Estimating causal effects of treatments in randomized and nonrandomized studies." In: *Journal of educational Psychology* 66.5 (1974), pp. 688–701. URL: http://psycnet.apa.org/journals/edu/66/5/688/.

[89] Jeffrey H. Samet et al. "Discontinuation From HIV Medical Care: Squandering Treatment Opportunities". In: *Journal of Health Care for the Poor and Underserved* 14.2 (2003), pp. 244–255. ISSN: 1548-6869. DOI: 10.1353/hpu.2010.0798. URL: http://muse.jhu.edu/content/crossref/journals/journal%7B%5C_%7Dof%7B%5C_%7Dhealth%7B%5C_%7Dcare%7B%5C_%7Dfor%7B%5C_%7Dthe%7B%5C_%7Dpoor%7B%5C_%7Dand%7B%5C_%7Dunderserved/v014/14.2.samet.html.

[90] Carl-Erik Sarndal, Bengt Swensson, and Jan Wretman. *Model Assisted Survey Sampling*. Springer, 2003, p. 695. ISBN: 0387406204.

[91] Glen A Satten and Somnath Datta. "Marginal analysis of multistage data". In: *Handbook of Statistics: Advances in Survival Analysis*. Ed. by N Balakrishnan and CR Rao. 23rd ed. Elsevier, 2004. Chap. 32, pp. 559–574. DOI: 0444500790.

[92] Glen A Satten and Somnath Datta. "The Kaplan-Meier Estimator as an Weighted Average". In: *The American Statistician* 55.3 (2001), pp. 207–210. URL: http://www.jstor.org/stable/2685801.

[93] Daniel O Scharfstein, Andrea Rotnitzky, and James M Robins. "Adjusting for Nonignorable Drop-Out Using Semiparametric Nonresponse Models". In: *Journal of the American Statistical Association* 94.448 (1999), pp. 1096–1120.

[94] Daniel O Scharfstein, Andrea Rotnitzky, and James M Robins. "Adjusting for Nonignorable Drop-Out Using Semiparametric Nonresponse Models: Rejoinder". In: *Journal of the American Statistical Association* 94.448 (1999), pp. 1135–1146.

[95] Mireille E Schnitzer et al. "Modeling the impact of hepatitis C viral clearance on end-stage liver disease in an HIV co-infected cohort with targeted maximum likelihood estimation." In: *Biometrics* 70.1 (Mar. 2014), pp. 144–52. ISSN: 1541-0420. DOI: 10.1111/biom.12105. URL: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3954273%7B%5C&%7Dtool=pmcentrez%7B%5C&%7Drendertype=abstract.

[96] Joshua Schwab et al. *ltmle: Longitudinal Targeted Maximum Likelihood Estimation*. 2014.

[97] Jasjeet S Sekhon. "The Neyman-Rubin Model of Causal Inference and Estimation via Matching Methods". In: *The Oxford Handbook of Political Methodology* (2007).

[98] Michelle Shardell, Gregory E Hicks, and Luigi Ferrucci. "Doubly robust estimation and causal inference in longitudinal studies with dropout and truncation by death." In: *Biostatistics (Oxford, England)* 16.1 (Jan. 2015), pp. 155–68. ISSN: 1468-4357. DOI: 10.1093/biostatistics/kxu032. URL: http://www.ncbi.nlm.nih.gov/pubmed/24997309.

[99] Ori M Stitelman et al. "Targeted maximum likelihood estimation of effect modification parameters in survival analysis." In: *The international journal of biostatistics* 7.1 (Jan. 2011), p. 19. ISSN: 1557-4679. DOI: 10.2202/1557-4679.1307. URL: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3083138%7B%5C&%7Dtool=pmcentrez%7B%5C&%7Drendertype=abstract.

[100] JSA Stringer et al. "Rapid scale-up of antiretroviral therapy at primary care sites in Zambia: feasibility and early outcomes". In: *JAMA* 296.7 (2006), pp. 782–793. URL: http://jama.jamanetwork.com/article.aspx?articleid=203173.

[101] Amitabh B Suthar et al. "Improving antiretroviral therapy scale-up and effectiveness through service integration and decentralization." In: *AIDS (London, England)* 28 Suppl 2 (Mar. 2014), S175–85. ISSN: 1473-5571. DOI: 10.1097/QAD.0000000000000259. URL: http://www.ncbi.nlm.nih.gov/pubmed/24849478.

[102] Sarah L Taubman et al. "Intervening on risk factors for coronary heart disease: an application of the parametric g-formula." In: *International journal of epidemiology* 38.6 (Dec. 2009), pp. 1599–611. ISSN: 1464-3685. DOI: 10.1093/ije/dyp192. URL: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2786249%7B%5C&%7Dtool=pmcentrez%7B%5C&%7Drendertype=abstract.

[103] Eric J Tchetgen Tchetgen and Tyler J VanderWeele. "On causal inference in the presence of interference." In: *Statistical methods in medical research* 21.1 (Feb. 2012), pp. 55–75. ISSN: 1477-0334. DOI: 10.1177/0962280210386779. URL: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=4216807%7B%5C&%7Dtool=pmcentrez%7B%5C&%7Drendertype=abstract.

[104] R Core Team. *R: A language and environment for statistical computing*. Vienna, Austria, 2014.

[105] Robert Tibshirani. "Regression Shrinkage and Selection via the Lasso". In: *Journal of the Royal Statistical Society* 58.1 (1996), pp. 267–288.

[106] Linh Tran et al. "Evaluating the Impact of a HIV Low-Risk Express Care Task-Shifting Program: a case study of the targeted learning roadmap". In: *U.C. Berkeley Division of Biostatistics Working Paper Series* (2016).

[107] Anastasios Tsiatis. *Semiparametric theory and missing data*. New York: Springer, 2006, p. 388. DOI: 1441921850.

[108] UNAIDS. *Global Report: UNAIDS report on the global AIDS epidemic 2013*. Tech. rep. 2013.

[109] Aad W. van der Vaart, Sandrine Dudoit, and Mark J. van der Laan. "Oracle inequalities for multi-fold cross validation". In: *Statistics & Decisions* 24.3 (Jan. 2006), pp. 351–371. ISSN: 0721-2631. DOI: 10.1524/stnd.2006.24.3.351. URL: http://www.degruyter.com/view/j/stnd.2006.24.issue-3/stnd.2006.24.3.351/stnd.2006.24.3.351.xml.

[110] Wim Van Damme, Katharina Kober, and Marie Laga. "The real challenges for scaling up ART in sub-Saharan Africa". In: *AIDS* 20.December 2005 (2006), pp. 653–656.

[111] Bill Venables and Brian Ripley. *Modern Applied Statistics with S*. Fourth. Springer, 2002.

[112] Karel Vermeulen. *Semiparametric Efficiency*. Gent, Belgium, 2011, p. 257.

[113] Strother H Walker and David B Duncan. "Estimation of the Probability of an Event as a Function of Several Independent Variables". In: *Biometrika* 54.1 (1967), pp. 167–179.

[114] Daniel Westreich and Stephen R Cole. "Invited commentary: positivity in practice." In: *American journal of epidemiology* 171.6 (Mar. 2010), pp. 674–7, 674–7. ISSN: 1476-6256. DOI: 10.1093/aje/kwp436. URL: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2877454%7B%5C&%7Dtool=pmcentrez%7B%5C&%7Drendertype=abstract.

[115] World Health Organization. *Consolidated Guidelines on the Use of Antiretroviral Drugs for Treating and preventing HIV Infection*. June. London, 2013, p. 272. ISBN: 978 92 4 150572 7. URL: http://apps.who.int/iris/bitstream/10665/85321/1/9789241505727%7B%5C_%7Deng.pdf.

[116] World Health Organization. *Global Update on HIV Treatment 2013: Results, Impact, and Opportunities*. Tech. rep. June. Geneva, 2013, p. 126.

[117] Wenjing Zheng and MJ van der Laan. "Asymptotic theory for cross-validated targeted maximum likelihood estimation". In: *U.C. Berkeley Division of Biostatistics Working Paper Series* 273 (2010). URL: http://biostats.bepress.com/ucbbiostat/paper273/.

[118] Jie Zhou et al. "Coarsened Propensity Scores and Hybrid Estimators for Missing Data and Causal Inference". In: *International Statistical Review* (Oct. 2014), n/a–n/a. ISSN: 03067734. DOI: 10.1111/insr.12082. URL: http://doi.wiley.com/10.1111/insr.12082.

[119] Marcel Zwahlen et al. "Mortality of HIV-infected patients starting potent antiretroviral therapy: comparison with the general population in nine industrialized countries." In: *International journal of epidemiology* 38.6 (Dec. 2009), pp. 1624–33. ISSN: 1464-3685. DOI: 10.1093/ije/dyp306. URL: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3119390%7B%5C&%7Dtool=pmcentrez%7B%5C&%7Drendertype=abstract.

Appendix A

Proof that O and O_r have equivalent influence functions

We provide a proof here showing that the influence function for the full data O is equivalent to the influence function we solve for in our analysis based on the reduced data $O_r = O/L1(K+1)$.

Firstly, we know the efficient influence function $D(P)$ for the reduced data O_r. Our goal is to compute the efficient influence function for the full data O for this same parameter (but now as a function on a model on O instead of O_r).

Recall that one can compute an efficient influence function $D^*(P)$ by first deriving an influence function of any estimator, and then projecting this influence function on the tangent space $T(P)$ of the model. $D(P)$ is one such influence function since it is the influence function of the TMLE ignoring $L1(t^*)$. Thus, the desired efficient influence function is $D^*(P) = \pi(D(P)|T(P))$ where π is the projection operator acting onto $T(P)$.

Now, we note that the likelihood of O can be factorized as $P(L1(t^*)|Y(t^*))P(O_r)$. Thus, the tangent space is the orthogonal sum of the tangent space $T_1(P)$ of the first factor $P(L1(t^*)|Y(t^*))$ and the tangent space $T_r(P)$ of $P(O_r)$. Consequently, $\pi(D(P)|T(P)) = \pi(D(P)|T_1(P)) + \pi(D(P)|T_r(P))$. However, the target parameter is only a function of $P(O_r)$, such that $P(L1(t^*)|Y(t^*))$ is actually a nuisance parameter. Because the efficient influence function is always orthogonal to a nuisance tangent space (i.e. the space of scores one gets by varying the nuisance parameters), it is also orthogonal to the tangent space of $P(L1(t^*)|Y(t^*))$. It follows that $D^*(P) = \pi(D(P)|T_r(P))$ (i.e. the component in $T_1(P)$ is zero).

However, $D(P)$ is the efficient influence function of the target parameter based on O_r and is therefore an element of the tangent space of $P(O_r)$, so that $D(P)$ is in $T_r(P)$. This results in $\pi(D(P)|T_r(P)) = D(P)$, which completes the proof that $D^*(P) = D(P)$. That is, the efficient influence function of our parameter based on O_r is the same as the efficient influence function of our parameter based on O.

Appendix B

Efficient influence function variance for working marginal structural models

B.1 TMLE of σ^2_{K+1} for marginal structural working models

For the general working marginal structural model (MSM) $\Theta \equiv \{m_\beta : \beta\}$ from Petersen2013, we have that the component corresponding with the last time point $K+1$ equals

$$\begin{aligned} D^*_{K+1}(P) &= \sum_{d \in D} h_1(d, K+1) \frac{\mathbb{I}(\bar{A}(K) = d(\bar{L}(K)))}{g_{0:K}(\bar{A}(K), \bar{L}(K))} (Y - \bar{Q}_{K+1}(\bar{A}(K), \bar{L}(K))) \\ &= C_{K+1}(P)(\bar{A}, \bar{L})(Y - \bar{Q}_{K+1}), \end{aligned}$$

where, for some user defined weight function $h(d, K+1)$,

$$\begin{aligned} C_{K+1}(P)(\bar{A}, \bar{L}) &= \sum_{d \in D} h_1(d, K+1) \frac{\mathbb{I}(\bar{A}(K) = d(\bar{L}(K)))}{g_{0:K}(\bar{A}, \bar{L})}. \\ h_1(d, K+1) &= h(d, K+1) \frac{\frac{\partial}{\partial \beta} m_\beta(d, K+1)}{m_\beta(1 - m_\beta)} \end{aligned}$$

We want to obtain a representation of the variance of this component D^*_{K+1} so that we can use a semi-substitution estimator of this part of the variance of the EIF, hopefully, thereby obtaining a variance estimator that is more accurate under violations of practical positivity, and a variance estimator that can be used as a red flag for lack of identifiability. This variance can thus be written

as

$$\begin{aligned}
\sigma^2_{K+1} &= \mathbb{E}[C^2(Y-\bar{Q}_{K+1})^2] \\
&= \mathbb{E}[C^2\bar{Q}_{K+1}(1-\bar{Q}_{K+1})] \\
&= \mathbb{E}\left[\left(\sum_{d\in D} h_1(d,K+1)\mathbb{I}(\bar{A}=d(\bar{L}))\right)^2 \frac{\bar{Q}_{K+1}(1-\bar{Q}_{K+1})}{g^2_{0:K}}(O)\right] \\
&= \mathbb{E}\left[\left(\sum_{d_1,d_2\in D} h_1(d_1,K+1)h_1(d_2,K+1)\mathbb{I}(\bar{A}=d_1(\bar{L}))\mathbb{I}(\bar{A}=d_2(\bar{L}))\right)\frac{\bar{Q}(1-\bar{Q})}{g^2_{0:K}}(O)\right] \\
&= \sum_{d_1,d_2} h_1(d_1,K+1)h_1(d_2,K+1)\mathbb{E}\left[\mathbb{I}(\bar{A}=d_1(\bar{L}))\mathbb{I}(\bar{A}=d_2(\bar{L}))\frac{\bar{Q}(1-\bar{Q})}{g^2_{0:K}}(O)\right].
\end{aligned}$$

The latter expectation equals:

$$\begin{aligned}
&\textstyle\int_{\bar{L}} \mathbb{I}(d_1(\bar{L})=d_2(\bar{L}))\prod_{t=0}^{K+1} Q(L_t \mid \bar{A}(t^-)=d_1(\bar{L}(t^-)),\bar{L}(t^-))\frac{\bar{Q}(1-\bar{Q})}{g_{0:t}(d_1(\bar{L}),\bar{L})} \\
&= \mathbb{E}_{P_0^{d_1}}\left[\mathbb{I}(d_1(\bar{L}_{d_1})=d_2(\bar{L}_{d_1}))\frac{\bar{Q}(1-\bar{Q})}{g_{0:K}}(d_1(\bar{L}_{d_1}),\bar{L}_{d_1})\right].
\end{aligned}$$

This yields the following expression:

$$\begin{aligned}
\sigma^2_{K+1} &= \sum_{d_1,d_2\in D} h_1(d_1,K+1)h_1(d_2,K+1) \\
&\qquad \mathbb{E}\left[\mathbb{I}(d_1(\bar{L}_{d_1})=d_2(\bar{L}_{d_1}))\frac{\bar{Q}(1-\bar{Q})}{g_{0:K}}(d_1(\bar{L}_{d_1}),\bar{L}_{d_1})\right] \\
&= \sum_{d_1\in D} h_1(d_1,K+1) \\
&\qquad \mathbb{E}\left[\left(\sum_{d_2\in D} h_1(d_2,K+1)\mathbb{I}(d_1(\bar{L}_{d_1})=d_2(\bar{L}_{d_1}))\right)\frac{\bar{Q}(1-\bar{Q})}{g_{0:K}}(d_1(\bar{L}_{d_1}),\bar{L}_{d_1})\right] \\
&= \sum_{d_1\in D} h_1(d_1,K+1)\mathbb{E}Z_{d_1}(d_1,K+1)
\end{aligned}$$

where

$$Z(d_1,K+1) = \left(\sum_{d_2\in D} h_1(d_2,K+1)\mathbb{I}(d_1(\bar{L}(K))=d_2(\bar{L}(K)))\right)\frac{\bar{Q}(1-\bar{Q})}{g_{0:K}(d(\bar{L}(K)),\bar{L}(K))},$$

so that the counterfactual of $Z(d_1,K+1)$ under intervention d_1 is given by

$$\begin{aligned}
Z_{d_1}(d_1,K+1) = &\left(\sum_{d_2\in D} h_1(d_2,K+1)\mathbb{I}(d_1(\bar{L}_{d_1}(K))=d_2(\bar{L}_{d_1}(K)))\right) \\
&\frac{\bar{Q}(1-\bar{Q})}{g_{0:K}}(d_1(L_{d_1}(K)),L_{d_1}(K)).
\end{aligned}$$

Static regimens

In the special case that the class of dynamic regimens D consists only of static regimens $\bar{a}(K)$ so that there is only one and exactly one treatment such that $\bar{A}(K) = d(\bar{L}(K))$, then we have

$$Z(K+1) = h_1(\bar{A}, K+1)\frac{\bar{Q}(1-\bar{Q})}{g_{0:K}}(\bar{A},\bar{L}),$$

so that

$$Z_d(K+1) = h_1(d, K+1)\frac{\bar{Q}(1-\bar{Q})}{g_{0:K}}(d(\bar{L}_d),\bar{L}_d).$$

In that case, we have

$$\sigma^2_{K+1} \quad = \quad \sum_{d\in D} h_1(d, K+1)^2 \mathbb{E}Z_{1d}(K+1)$$

where $Z_1(K+1) = \bar{Q}(1-\bar{Q})/g_{0:K}(\bar{A},\bar{L})$ and $Z_{1d}(K+1) = \bar{Q}(1-\bar{Q})/g_{0:K}(d(\bar{L}_d),\bar{L}_d)$.

It is important to note that in expressing our variance this way, we integrate out the indicator of treatment over $\bar{A}$, i.e. $\mathbb{I}(\bar{A} = d(\bar{L}))$. By getting rid of this indicator, we no longer rely as heavily on observations from subjects following treatment in estimating the variance of D^*_{K+1}. This particularly helps us when there is a lack of positivity, since subjects with low probabilities of desired treatment simply are not observed.

We have now shown that

$$\sigma^2_{K+1} = \sum_{d\in D} h_1(d, K+1)\mathbb{E}Z_d(d, K+1),$$

where

$$Z(d_1, K+1) = \left\{ \sum_{d_2} h_1(d_2, K+1)\mathbb{I}(d_1(\bar{L}) = d_2(\bar{L})) \right\} \frac{\bar{Q}(1-\bar{Q})}{g_{0:K}}(d_1(\bar{L}),\bar{L}).$$

We can now define $Z(K+1)(\bar{A},\bar{L}) = \sum_{d\in D} h_1(d, K+1)\mathbb{I}(\bar{A} = d(\bar{L}))\frac{\bar{Q}(1-\bar{Q})}{g_{0:K}}(\bar{A},\bar{L})$ (as function of $\bar{A},\bar{L}$) as a new outcome for our longitudinal data structure such that $Z_d(d, K+1) = Z(K+1)(d(\bar{L}_d),\bar{L}_d)$. Our variance σ^2_{K+1} is then represented as $\sum_{d\in D} h_1(d, K+1)\mathbb{E}Z_d(d, K+1)$. Thus, if we redefine the longitudinal data as $(\bar{A},\bar{L})$ with the final outcome of interest as $Z(K+1) = Z(K+1)(\bar{A},\bar{L})$, and use the working MSM parameter $\mathbb{E}Z_d(K+1) = \beta_0$ with $\beta_0 = \arg\min_\beta \sum_{d\in D} h_1(d, K+1)(\mathbb{E}Z_d(K+1) - \beta)^2$, then we have that

$$\beta_0 = \sum_{d\in D} h_1(d, K+1)\mathbb{E}Z_d(K+1) / \sum_{d\in D} h_1(d, K+1).$$

This demonstrates that we can obtain σ^2_{K+1} by simply multiplying β_0 by $\sum_{d\in D} h_1(d, K+1)$, i.e.

$$\sigma^2_{K+1} = \beta_0 \sum_d h_1(d, K+1).$$

We can therefore also estimate this variance component σ^2_{K+1} by computing the TMLE of the intercept β_0 in the working MSM for our newly defined outcome $Z(K+1)$ using weights $h_1(d, K+1)$, and then multiplying it against $\sum_{d\in D} h_1(d, K+1)$.

B.2 TMLE of σ_t^2 for marginal structural working models

We now present the how to obtain a TMLE of the variance of the t-th component of the EIF, σ_t^2. For the general working MSM from Petersen2013, we have that the component corresponding with the t-th time point equals

$$\begin{aligned} D_t^*(P) &= \sum_{d\in D} h_1(d,t)\frac{\mathbb{I}(\bar{A}(t^-)=d(\bar{L}(t^-)))}{g_{0:t^-}(\bar{A}(t^-),\bar{L}(t^-))}(\bar{Q}_{t^+}^d(\bar{A}(t),\bar{L}(t))-\bar{Q}_t^d)(\bar{A}(t^-),\bar{L}(t^-))) \\ &= \sum_{d\in D} C_t(P,d)(\bar{Q}_{t^+}^d-\bar{Q}_t^d). \end{aligned}$$

Similar to above, we want to obtain a representation of the variance of this component so that we can use a semi-substitution estimator of this part of the variance of the EIF, hopefully, thereby obtaining a variance estimator that is more accurate under violations of practical positivity, and a variance estimator that can be used as a red flag for lack of identifiability. This variance σ_t^2 can thus be written as

$$\sigma_t^2 = \sum_{d_1,d_2} h_1(d_1,t)h_1(d_2,t)\mathbb{E}\left[\mathbb{I}(\bar{A}(t^-)=d_1)\mathbb{I}(\bar{A}(t^-)=d_2)\frac{\Sigma_t(d_1,d_2)}{g_{0:t^-}^2}(\bar{A}(t^-),\bar{L}(t^-))\right]$$

where

$$\Sigma_t(d_1,d_2)(\bar{A}(t^-),\bar{L}(t^-)) = \mathbb{E}\left[(\bar{Q}_{t^+}^{d_1}-\bar{Q}_t^{d_1})(\bar{Q}_{t^+}^{d_2}-\bar{Q}_t^{d_2}) \,\middle|\, \bar{A}(t^-),\bar{L}(t^-)\right]$$

is the conditional covariance of $\bar{Q}_{t^+}^{d_1}$ and $\bar{Q}_{t^+}^{d_2}$, given $(\bar{A}(t^-),\bar{L}(t^-))$. Note that this can be obtained by regressing this cross-product on $(\bar{A}(t^-),\bar{L}(t^-))$. The latter sum can be further worked out giving us

$$\sigma_t^2 = \sum_{d_1\in D} h_1(d_1,t)\mathbb{E}Z_{d_1}(d_1,t),$$

where

$$Z(d_1,t) = \left(\sum_{d_2\in D} h_1(d_2,t)\mathbb{I}(d_1(\bar{L}(t^-))=d_2(\bar{L}(t^-)))\right)\frac{\Sigma_t(d_1,d_2)}{g_{0:t^-}}(d_{1,t^-}(\bar{L}(t^-)),\bar{L}(t^-)).$$

so that the counterfactual of Z_t under intervention d_1 is given by

$$Z_{d_1}(d_1,t) = \left(\sum_{d_2\in D} h_1(d_2,t)\mathbb{I}(d_1(\bar{L}_{d_1}(t^-))=d_2(\bar{L}_{d_1}(t^-)))\right)\frac{\Sigma_t(d_1,d_2)}{g_{0:t^-}}(d_{1,t^-}(\bar{L}_{d_1}(t^-)),\bar{L}_{d_1}(t^-)).$$

With this expression, we can now use a TMLE to estimate $\mathbb{E}Z_{d_1}(d_1,t)$ for each $d_1\in D$ by using the longitudinal data structure with final outcome $Z(d_1,t)$, for each d_1 separately. To create the observed outcome $Z(d_1,t)$ we need a fit of the treatment mechanism $g_{A(m)} : m = 0,1,\ldots,t^-$, evaluated at $\bar{A}(t^-) = d_{t^-}(\bar{L}(t^-))$, and for each rule compatible with d_1 (for that subject) we need to have an estimate of $\Sigma_t(d_1,d_2)$. Thus, given a priori estimates of the full treatment mechanism and all $(\Sigma_t(d_1,d_2) : d_1,d_2\in D)$ we can construct this observed outcome $Z(d_1,t)$ and run the TMLE.

B.3 Estimation of the variance of the EIF

The above approach defines for each time point t and each rule d an observed longitudinal outcome $Z(d,t)$, where $Z(d,t)$ is a function of $(\bar{A}(t),\bar{L}(t))$. The TMLE of $\mathbb{E}Z_d(d,t)$ can then be computed based on the longitudinal data structure $(L(0),A(0),\ldots,L(t),A(t),Z(d,t))$ for each d and each $t \in \{0,1,\ldots,K+1\}$. As a result, we have that $\sigma_t^2 = \sum_{d\in D} h_1(d,t)\mathbb{E}Z_d(d,t)$ and

$$\begin{aligned}\sigma^2 &= \sum_{t=0}^{K+1} \sigma_t^2 \\ &= \sum_{d\in D}\left(\sum_{t=0}^{K+1} h_1(d,t)\mathbb{E}Z_d(d,t)\right) \\ &= \sum_{d\in D}\mathbb{E}\left[\sum_{t=0}^{K+1} h_1(d,t)Z_d(d,t)\right].\end{aligned}$$

Let's now define the counterfactual outcome

$$\bar{Z}_d(d) \equiv \sum_{t=0}^{K+1} h_1(d,t)Z_d(d,t),$$

and the corresponding observed outcome

$$\bar{Z}(d) \equiv \sum_{t=0}^{K+1} h_1(d,t)Z(d,t).$$

We could apply the TMLE to estimate $\mathbb{E}\bar{Z}_d(d)$ based on the longitudinal data structure $(L(0),A(0),\ldots,L(K),A(K),\bar{Z}(d,K+1))$, for each $d \in D$, and use that

$$\sigma^2 = \sum_{d\in D}\mathbb{E}\bar{Z}_d(d).$$

In applying TMLE here, we should be using that

$$\mathbb{E}\left[\bar{Z}_d \,\middle|\, \bar{A}(m),\bar{L}(m)\right] = \sum_{t\le m} h_1(d,t)Z(d,t) + \mathbb{E}\left[\sum_{t>m} h_1(d,t)Z(d,t) \,\middle|\, \bar{A}(m),\bar{L}(m)\right].$$

To start with, let

$$\begin{aligned}\bar{Q}_d^{Z(K+1)} &= \mathbb{E}[\bar{Z}(d) \mid \bar{A}(K),\bar{L}(K)] \\ &= \sum_{t\le K} h_1(d,t)Z(d,t) + \mathbb{E}[h_1(d,K+1) + Z(d,K+1) \mid \bar{A}(K),\bar{L}(K)].\end{aligned}$$

Denote last conditional expectation with $\bar{Q}_d^{Z(K+1),d}$ so that

$$\bar{Q}_d^{Z(K+1)} = \sum_{t\le K} h_1(d,t)Z(d,t) + \bar{Q}_d^{Z(K+1),d}.$$

Then,

$$\begin{aligned}\bar{Q}_d^{Z(K)} &= \mathbb{E}\left[\bar{Q}_d^{Z(K+1)} \,\middle|\, \bar{A}(K-1), \bar{L}(K-1)\right] \\ &= \sum_{t \leq K-1} h_1(d,t) Z(d,t) + \mathbb{E}\left[h_1(d,K) Z(d,K) + \bar{Q}_d^{Z(K+1),d} \,\middle|\, \bar{A}(K-1), \bar{L}(K-1)\right].\end{aligned}$$

Again, denote the latter conditional expectation by $\bar{Q}_d^{Z(K),d}$ so that

$$\bar{Q}_d^{Z(K)} = \sum_{t \leq K-1} h_1(d,t) Z(d,t) + \bar{Q}_d^{Z(K),d}.$$

Then,

$$\begin{aligned}\bar{Q}_d^{Z(K-1)} &= \mathbb{E}\left[\bar{Q}_d^{Z(K)} \,\middle|\, \bar{A}(K-2), \bar{L}(K-2)\right] \\ &= \sum_{t \leq K-2} h_1(d,t) Z(d,t) + \\ &\qquad \mathbb{E}\left[h_1(d,K-1) Z(d,K-1) + \bar{Q}_d^{Z(K),d} \,\middle|\, \bar{A}(K-2), \bar{L}(K-2)\right].\end{aligned}$$

Again, denote the latter conditional expectation with $\bar{Q}_d^{Z(K-1),d}$ so that

$$\bar{Q}_d^{Z(K-1)} = \sum_{t \leq K-2} h_1(d,t) Z(d,t) + \bar{Q}_d^{Z(K-1),d}.$$

This is then iterated:

$$\bar{Q}_d^{Z(m)} = \sum_{t \leq m-1} h_1(d,m) Z(d,m) + \bar{Q}_d^{Z(m),d},$$

where $\bar{Q}_d^{Z(m),d} = \mathbb{E}\left[h_1(d,m) Z(d,m) + \bar{Q}_d^{Z(m+1),d} \,\middle|\, \bar{A}(m-1), \bar{L}(m-1)\right]$.

Before we go to the next conditional expectation we need to target with a parametric submodel, such as

$$\operatorname{Logit} \bar{Q}_d^m(\varepsilon) = \operatorname{Logit} \bar{Q}_d^m + \varepsilon \frac{\mathbb{I}(\bar{A}(m-1) = d(\bar{L}(m-1)))}{g_{0:m-1}}.$$

In this way, we will only have to run one TMLE for each rule d, which still utilizes that the outcome is a sum of outcomes that are known for histories including that outcome.

Appendix C

R code for intervention specific mean outcome estimators

C.1 Data generation

```
############################
## GENERATE SIMULATED DATA ##
############################
#' @export
generateData = function(n, time.pt, abar=NULL) {

       #Time ordering: W, Y(t), L(t), A(t) : W=(W1,W2) and L(t) = (L2(t),L1(t))
       #n.b. Within L(t) there is no implied time-ordering...i.e. either of L2(t)
            or L1(t) can go first
       rexpit = function(x) rbinom(n=length(x), size=1, prob=x)
       QW1    = function(n) rnorm(n, mean=0, sd=1)
       QW2    = function(n) rep(plogis(-1), n)
       QW3    = function(n) rnorm(n, mean=0, sd=1)
       QY.t   = function(prev_y, w1, w2, w3, prev_l1, prev_l2, prev_a) ifelse(
           prev_y==1, 1, plogis(-1.9 + 1.2*w1 - 2.4*w2 - 1.8*prev_l1 - 1.6*prev_l2
            + 1*prev_l1*prev_l2 - 1*prev_a))
       QL1.t  = function(y, w1, prev_l1, prev_l2, prev_a) ifelse(y==1, prev_l1,
           0.1 + 0.4*w1 + 0.6*prev_l1 - 0.7*prev_l2 - 0.45*prev_a - rnorm(length(
           w1), sd=0.5))
       QL2.t  = function(y, w1, w2, prev_l1, prev_l2, prev_a) ifelse(y==1,
           prev_l2, -0.55 + 0.5*w1 + 0.75*w2 + 0.1*prev_l1 + 0.3*prev_l2 - 0.75*
           prev_a - rnorm(length(w1), sd=0.5))
       gA.t   = function(y, w1, w2, l1, l2, prev_a) ifelse(y==1, prev_a, ifelse(
           prev_a==1, 1, plogis(-1 - 1.5*w1 + 1.75*w2 + 1.2*l1 - 1.8*l2 + 0.8*l1*
           l2)))
```

```
# nb. Distribution is set up such that:
#       Y(0)=0 for everyone, ie. Everyone is alive at the beginning of
    follow-up
#       if Y(t)=1, then all remaining covariate last values get carried
    forward
#       if A(t-1)=1 then A(t)=1

g.matrix = matrix(ncol=time.pt, nrow=n, dimnames=list(NULL, paste0("A.",
    0:(time.pt-1))))

## CHECKS ##
if(time.pt==0) stop("time.pt has to be greater than 0")
if(any(cummax(abar)!=abar)) stop("A is a counting process & cannot
    decrease")
if(!is.null(abar) & length(abar) != time.pt) stop("abar has to be either
    NULL or length of time.pt")

## INITIALIZATION ##
o.names = NULL
for(i in 0:time.pt){
        if(i<time.pt) {
                o.names = c(o.names, paste0(c("Y", "L1","L2", "A"), ".", i))
        } else {
                o.names = c(o.names, paste0(c("Y"), ".", i))
        }
}
O = matrix(nrow=n, ncol=length(o.names)+3, dimnames=list(NULL, c("W1", "W2
    ", "W3", o.names)))

## OBSERVED VALUES ##
O[,"W1"] = QW1(n)
O[,"W2"] = rexpit(QW2(n))
O[,"W3"] = QW3(n)
for(i in 0:time.pt){
        #nb. "prev" values are set to 0 for t=0
        if(i==0) {
                #Y(t)
                O[,"Y.0"] = rep(0,n)
                #L1(t)
                O[,"L1.0"] = QL1.t(y=O[,"Y.0"], w1=O[,"W1"], prev_l1=rep(0,
                    n), prev_l2=rep(0, n), prev_a=rep(0, n))
                #L2(t)
                O[,"L2.0"] = QL2.t(y=O[,"Y.0"], w1=O[,"W1"], w2=O[,"W2"],
                    prev_l1=rep(0, n), prev_l2=rep(0, n), prev_a=rep(0, n))
```

```
        #A(t)
        if(is.null(abar)) {
               g.matrix[,"A.0"] = gA.t(y=O[,"Y.0"], w1=O[,"W1"], w2
                   =O[,"W2"], l1=O[,"L1.0"], l2=O[,"L2.0"], prev_a=
                   rep(0, n))
               O[,"A.0"] = rexpit(g.matrix[,"A.0"])
        } else {
               g.matrix[,"A.0"] = gA.t(y=O[,"Y.0"], w1=O[,"W1"], w2
                   =O[,"W2"], l1=O[,"L1.0"], l2=O[,"L2.0"], prev_a=
                   rep(0, n))
               O[,"A.0"] = rep(abar[i+1], n)
        }
} else if (i<time.pt) {
        #Y(t)
        O[,paste0("Y.",i)] = rexpit(QY.t(prev_y=O[,paste0("Y.",i-1)
            ], w1=O[,"W1"], w2=O[,"W2"], w3=O[,"W3"], prev_l1=O[,
            paste0("L1.",i-1)], prev_l2=O[,paste0("L2.",i-1)],
            prev_a=O[,paste0("A.",i-1)]))
        #L1(t)
        O[,paste0("L1.",i)] = QL1.t(y=O[,paste0("Y.",i)], w1=O[,"W1
            "], prev_l1=O[,paste0("L1.",i-1)], prev_l2=O[,paste0("L2
            .",i-1)], prev_a=O[,paste0("A.",i-1)])
        #L2(t)
        O[,paste0("L2.",i)] = QL2.t(y=O[,paste0("Y.",i)], w1=O[,"W1
            "], w2=O[,"W2"], prev_l1=O[,paste0("L1.",i-1)], prev_l2=
            O[,paste0("L2.",i-1)], prev_a=O[,paste0("A.",i-1)])
        #A(t)
        if(is.null(abar)) {
               g.matrix[,paste0("A.",i)] = gA.t(y=O[,paste0("Y.",i)
                   ], w1=O[,"W1"], w2=O[,"W2"], l1=O[,paste0("L1.",i
                   )], l2=O[,paste0("L2.",i)], prev_a=O[,paste0("A
                   .",i-1)])
               O[,paste0("A.",i)] = rexpit(g.matrix[,paste0("A.",i)
                   ])
        } else {
               g.matrix[,paste0("A.",i)] = gA.t(y=O[,paste0("Y.",i)
                   ], w1=O[,"W1"], w2=O[,"W2"], l1=O[,paste0("L1.",i
                   )], l2=O[,paste0("L2.",i)], prev_a=O[,paste0("A
                   .",i-1)])
               O[,paste0("A.",i)] = rep(abar[i+1], n)
               O[O[,paste0("Y.",i)]==1,paste0("A.",i)] = O[O[,
                   paste0("Y.",i)]==1,paste0("A.",i-1)]
        }
} else if (i==time.pt) {
```

```
                    #Y(t)
                    O[,paste0("Y.",i)] = rexpit(QY.t(prev_y=O[,paste0("Y.",i-1)
                        ], w1=O[,"W1"], w2=O[,"W2"], w3=O[,"W3"], prev_l1=O[,
                        paste0("L1.",i-1)], prev_l2=O[,paste0("L2.",i-1)],
                        prev_a=O[,paste0("A.",i-1)]))
                }
        }
        O = data.frame(O)
        O$Y.0 = NULL

        return(list(O=O, g.matrix=g.matrix))

}
```

C.2 Augmented Inverse Probability of Treatment Weighted Estimation

```
####################
## AIPW ESTIMATOR ##
####################
#' Augmented Inverse Probability of Treatment Weighted Estimation
#'
#' \code{aiptw} Estimates the parameter of interest (E[Y_d]) by directly solving
    the efficient influence function as an estimating equation.
#'
#' Please refer to Scharfstein (1999), Robins (2000), or van der Laan and Gruben
(2007) for details regarding derivation of the efficient influence function,
    which this estimator solves.
#'
#' @param data data frame following the time-ordering of the nodes.
#' @param cum.g a matrix of the cumulative probabilities of treatment (and being
    uncensored) given the parents.
#' @param Ynodes column names or indicies in \code{data} of outcome nodes.
#' @param Anodes column names or indicies in \code{data} of treatment nodes.
#' @param Cnodes column names or indicies in \code{data} of censoring nodes.
#' @param abar binary vector (numAnodes x 1) of counterfactual treatment
#' @param Qform character vector of regression formulas for Q.
#' @param SL.library optional character vector of libraries to pass to
    SuperLearner. NULL indicates glm should be called instead of SuperLearner.
#' @param stratify if \code{TRUE} condition on following \code{abar} when
    estimating Q and g. If \code{FALSE}, pool over all subjects.
#'
#' @return \code{BangRobinsDR} returns a list of items as an object of class \
    code{aiptw}, which include
```

```
#' \itemize{
#'     \item {The estimate of the parameter value under the intervention \code{
    abar}}
#'     \item {The empirical influence function for the point estimate}
#'     \item {The time-specific empirical influence functions}
#'     \item {The conditional expectation fits for \code{Qbar}}
#'     \item {The call to the function}
#' }
#'
#' @references
#' Scharfstein, D. O., A. Rotnitzky, and J. M. Robins (1999a): "Adjusting for
Nonignor- able Drop-Out Using Semiparametric Nonresponse Models," Journal of the
    American Statistical Association, 94, 1096-1120.
#' Robins, J. M., A. Rotnitzky, and M. J. van der Laan (2000b): "Discussion of
On profile likelihood by Murphy and van der Vaart," Journal of the American
    Statistical Association, 95, 477-482.
#' van der Laan, M. J. and S. Gruber (2011): "Targeted Minimum Loss Based
    Estimation of an Intervention Specific Mean Outcome," The Berkeley Electronic
    Press.
#' @export
aiptw = function(data, cum.g, Ynodes, Lnodes, Anodes, Cnodes=NULL, abar, Qform,
    SL.library=NULL, family="quasibinomial", stratify=TRUE) {

       ## Initializes ##
       Q.kplus1 = data[,Ynodes[length(Ynodes)]]
       if(is.null(dim(cum.g))) {
              cum.g = matrix(cum.g, ncol=1)
       }
       IC.all = matrix(nrow=nrow(data), ncol=length(Ynodes), dimnames=list(NULL,
           Ynodes))
       Qfits = list()

       ## Recursive regression ##
       for(i in length(Ynodes):1) {
              ## Initializes ##
              if(family %in% c("quasibinomial", "binomial")) {
                     # nb. In binary setting, rather than using previous outcome
                         as predictor, we stratify
                     Q.kplus1[data[,Ynodes[i]]==1] = 1
                     if(i==1) {
                            at.risk = rep(TRUE, nrow(data))
                     } else {
                            at.risk = data[,Ynodes[i-1]]==0
                     }
```

```
} else if(family=="gaussian") {
        at.risk = rep(TRUE, nrow(data))
} else {
        stop("Function will only work with two scenarios (see help
            file).")
}
P_na.data = data
for(a in 1:length(Anodes)) {
        P_na.data[,Anodes[a]] = abar[a]
}
P_n = cbind(Q.kplus1, data[,all.vars(as.formula(Qform[i])[[3]])])
P_na = cbind(Q.kplus1, P_na.data[,all.vars(as.formula(Qform[i])
    [[3]])])

## Followed regime of interest ##
followed.abar = apply(data[,Anodes[1:i], drop=FALSE], 1, function(x
    ) all(x==abar[1:i], na.rm=T))
if(is.null(Cnodes)) {
        uncensored = rep(TRUE, nrow(data))
} else {
        uncensored = apply(data[,Cnodes[1:(i*2)], drop=FALSE], 1,
            function(x) all(x=="uncensored", na.rm=T))
}
followed.abar = followed.abar * uncensored

## Fits Qbar ##
if(stratify) {
        fitData = P_n[at.risk & followed.abar & uncensored,, drop=
            FALSE]
} else {
        fitData = P_n[at.risk & uncensored,, drop=FALSE]
}
if(is.null(SL.library)) {
        Qfit.abar = glm(as.formula(Qform[i]), data=fitData, family=
            family)
} else {
        tmp = complete.cases(P_n[,-1, drop=FALSE])
        if(family=="quasibinomial") family = "binomial"
        if(ncol(fitData[,-1, drop=FALSE])>0) {
                Qfit.abar = mcSuperLearner(Y=fitData$Q.kplus1, X=
                    fitData[,all.vars(as.formula(paste(Qform[i]))
                    [[3]]), drop=FALSE], newX=P_na[tmp,all.vars(as.
                    formula(paste(Qform[i]))[[3]]), drop=FALSE], SL.
                    library=SL.library, family=family, control = list
```

```
                                (trimLogit=.001, saveFitLibrary=FALSE), cvControl
                                =list(V=8), method="method.NNloglik.LT", verbose=
                                TRUE)
                        }
                }

                ## Gets Q.k ##
                if(is.null(SL.library)) {
                        Q.k = suppressWarnings(predict(Qfit.abar, newdata=P_na, type
                            ="response"))
                } else {
                        Q.k = rep(NA, length(tmp))
                        if(ncol(fitData[,-1, drop=FALSE])>0) {
                                library.predict = Qfit.abar$library.predict
                                library.predict[is.na(library.predict)] = 0
                                Qfit.abar$SL.predict = library.predict %*% Qfit.
                                    abar$coef
                                Q.k[tmp] = Qfit.abar$SL.predict
                        } else {
                                Q.k[tmp] = mean(fitData$Q.kplus1)
                        }
                }
                if(family %in% c("quasibinomial", "binomial")) {
                        Q.k[!at.risk] = 1
                }

                ## IC and Updates Q.kplus1 ##
                IC.all[,i] = calcIC(Q.kplus1=Q.kplus1, Q.k=Q.k, h.g.ratio=1/cum.g[,
                    i], uncensored=uncensored, intervention.match=followed.abar)
                Q.kplus1 = Q.k
                Qfits = c(Qfits, list(Qfit.abar))
                names(Qfits)[length(Qfits)] = Ynodes[i]
        }

        ## Temp IC ##
        tmpIC = apply(IC.all, 1, sum) + Q.kplus1
        out = list(estimate=mean(tmpIC), IC=tmpIC-mean(tmpIC), IC.all=IC.all,
            Qfits=Qfits, call=match.call())
        class(out) = "aiptw"

        return(out)

}
#'@export
```

```
print.aiptw = function(x, ...) {
       PrintCall(x$call)
       cat("AIPTW Estimate:\t", x$estimate, "\n")
       invisible(x)
}
```

C.3 Double Robust Sequential Regression Estimation

```
#####################
## DRICE ESTIMATOR ##
#####################
#' Double Robust Sequential Regression Estimation
#'
#' \code{drice} Estimates the parameter of interest (E[Y_d]) by using the
    sequential
#' regression approach as presented by Bang and Robins (2005). This function only
#' works in the setting where the outcome is binary.
#'
#' @param data data frame following the time-ordering of the nodes.
#' @param Anodes column names or indicies in \code{data} of treatment nodes.
#' @param Cnodes column names or indicies in \code{data} of censoring nodes.
#' @param Lnodes column names or indicies in \code{data} of time-dependent
    covariate nodes.
#' @param Ynodes column names or indicies in \code{data} of outcome nodes.
#' @param Qform character vector of regression formulas for \code{Qbar}.
#' @param cum.g a matrix of the cumulative probabilities of treatment (and being
    uncensored) given the parents.
#' @param abar binary vector (numAnodes x 1) of counterfactual treatment
#' @param stratify if \code{TRUE} condition on following \code{abar} when
    estimating \code{Qbar}. If \code{FALSE}, pool over all subjects.
#' @param type if \code{weight}, then use the g fits as weights in the \code{Qbar
    } fit. If \code{covariate}, then use the g fits as a covariate.
#'
#' @return \code{icedr} returns a list of items as an object of class \code{icedr
    }, which include
#' \itemize{
#'     \item {The estimate of the parameter value under the intervention \code{
    abar}}
#'     \item {A matrix of the condition expectation fit \code{Qbar.hat} under \
    code{abar} for each time point.}
#'     \item {The conditional expectation fits for \code{Qbar}}
#'     \item {The call to the function}
#' }
#'
```

```
#' @references
#' Bang, H. and J. M. Robins (2005): "Doubly robust estimation in missing data
    and causal inference models." Biometrics, 61, 962-73, URL http://www.ncbi.nlm.
    nih.gov/pubmed/16401269.
#' @export
drice = function(data, Anodes, Cnodes=NULL, Lnodes, Ynodes, Qform, cum.g, abar,
    stratify=FALSE, type="covariate") {

        ## INITIALIZE ##
        Qbar.hat = IC.all = matrix(nrow=nrow(data), ncol=length(Ynodes), dimnames=
            list(NULL, Ynodes))
        Q.kplus1 = data[,Ynodes[length(Ynodes)]]
        Q = list()
        P_na = data
        for(k in 1:length(Anodes)) P_na[,Anodes[k]] = rep(abar[k], each=nrow(data)
            )

        ## ICE ##
        for(k in length(Ynodes):1) {
                if (k>1) is.deterministic = !is.na(data[,Ynodes[k-1]]) & data[,
                    Ynodes[k-1]]==1 else is.deterministic = rep(FALSE,nrow(data))
                intervention.match = !is.na(data[,Anodes[k]]) & data[,Anodes[k]]==
                    abar[k]
                if(is.null(Cnodes)) {
                        uncensored = rep(TRUE, nrow(data))
                } else {
                        uncensored = apply(data[,Cnodes[1:(k*2)], drop=FALSE], 1,
                            function(x) all(x=="uncensored", na.rm=T))
                }
                intervention.match = intervention.match & uncensored
                Qform.icedr = paste(Qform[k], "+ cleverCov")
                if (stratify) {
                        index = intervention.match & !is.deterministic
                } else index = !is.deterministic
                if(type=="covariate") {
                        QbarData = model.frame(Qform.icedr, cbind(data, Q.kplus1,
                            cleverCov = ifelse(intervention.match, 1/cum.g[,k], 0)),
                             na.action=NULL)
                        tmpData = model.frame(Qform.icedr, cbind(data, Q.kplus1,
                            cleverCov = ifelse(intervention.match, 1/cum.g[,k], 0))[
                            index,])
                        icedrFit = glm(Qform.icedr, data=tmpData, family="
                            quasibinomial", control=glm.control(trace=FALSE, maxit
                            =1000))
```

```
            coef = icedrFit$coef[!is.na(icedrFit$coef)]
            X = model.matrix(as.formula(Qform.icedr), QbarData, na.
                action=NULL)[,names(coef), drop=FALSE]
            X_a = model.matrix(as.formula(Qform.icedr), cbind(P_na,
                cleverCov=1/cum.g[,k]), na.action=NULL)[, names(coef),
                drop=FALSE]
            if (nrow(X)!=nrow(QbarData) | nrow(X_a)!=nrow(P_na)) {
                X = as.matrix(cbind("(Intercept)"=1, QbarData)[,
                    names(coef), drop=FALSE])
                X_a = as.matrix(cbind("(Intercept)"=1, P_na,
                    cleverCov=1/cum.g[,k])[, names(coef), drop=FALSE
                    ])
            }
            score.icedrFit = CalcScore2(coef=coef, X=X, Q.kplus1=Q.
                kplus1, h.g.ratio=1/cum.g[,k], uncensored=uncensored,
                intervention.match=intervention.match, is.deterministic=
                is.deterministic)
            if (score.icedrFit>0.0001) {
                icedrFix = FixScore2(X, Q.kplus1, 1/cum.g[,k], rep(
                    TRUE,nrow(data)), intervention.match, is.
                    deterministic, start=coef)
                ## nb. If couldn't solve, just stuck with glm
                if (icedrFix$solved) {
                    icedrFit = icedrFix
                }
            }
        } else if (type=="weight") {
            QbarData = cbind(model.frame(Qform.icedr, cbind(data, Q.
                kplus1, cleverCov = ifelse(intervention.match, 1, 0)),
                na.action=NULL), weight.vec=1/cum.g[,k])
            tmpData = cbind(model.frame(Qform.icedr, cbind(data, Q.
                kplus1, cleverCov = ifelse(intervention.match, 1, 0)),
                na.action=NULL)[index,], weight.vec=1/cum.g[index,k])
            icedrFit = glm(Qform.icedr, data=tmpData, weight=scale(
                weight.vec, center=FALSE), family="quasibinomial",
                control=glm.control(trace=FALSE, maxit=1000))
            coef = icedrFit$coef[!is.na(icedrFit$coef)]
            X = model.matrix(as.formula(Qform.icedr), QbarData)[,names(
                coef), drop=FALSE]
            X_a = model.matrix(as.formula(Qform.icedr), cbind(P_na,
                cleverCov=1/cum.g[,k]))[, names(coef), drop=FALSE]
            if (nrow(X)!=nrow(QbarData) | nrow(X_a)!=nrow(P_na)) {
                X = as.matrix(cbind("(Intercept)"=1, QbarData)[,
                    names(coef), drop=FALSE])
```

```
                        X_a = as.matrix(cbind("(Intercept)"=1, P_na,
                            cleverCov=1)[,names(coef), drop=FALSE])
                    }
                    score.icedrFit = CalcScore2(coef=coef, X=X, Q.kplus1=Q.
                        kplus1, h.g.ratio=1/cum.g[,k], uncensored=uncensored,
                        intervention.match=intervention.match, is.deterministic=
                        is.deterministic)
                    if (score.icedrFit>0.0001) {
                        icedrFix = FixScore2(X, Q.kplus1, 1/cum.g[,k],
                            uncensored, intervention.match, is.deterministic,
                             start=coef)
                        ## nb. If couldn't solve, just stuck with glm
                        if (icedrFix$solved) {
                            icedrFit = icedrFix
                        }
                    }
                }

                IC.all[,k] = calcIC(Q.kplus1, plogis(X_a %*% coef), 1/cum.g[,k],
                    uncensored, intervention.match)
                Q.kplus1 = plogis(X_a %*% coef)
                if(sum(is.deterministic)>0) Q.kplus1[is.deterministic] = 1
                Qbar.hat[,k] = Q.kplus1
                Q = c(Q, list(icedrFit))
                names(Q)[length(Q)] = Ynodes[k]
        }

        out = list(estimate=mean(Q.kplus1), Qbar.hat=Qbar.hat, fit=list(Q=Q), IC=
            IC.all, call=match.call())
        class(out) = "drice"
        return(out)
}
#'@export
print.drice = function(x, ...) {
        cat("Call:\n", paste(deparse(x$call), sep = "\n", collapse = "\n"), "\n\n
            ", sep = "")
        cat("DRICE Estimate:\t", x$estimate, "\n")
        invisible(x)
}
```

C.4 Targeted minimum loss-based estimation

```
####################
```

```
## TMLE ESTIMATOR ##
####################
#' Targeted minimum loss-based estimation
#'
#' \code{icedr} Estimates the parameter of interest (E[Y_d]) by using the
    sequential regression approach as presented by van der Laan (2011).
#' This function only works in the setting where the outcome is binary.
#'
#' @param data data frame following the time-ordering of the nodes.
#' @param Anodes column names or indicies in \code{data} of treatment nodes.
#' @param Cnodes column names or indicies in \code{data} of censoring nodes.
#' @param Ynodes column names or indicies in \code{data} of time-dependent
    covariate nodes.
#' @param Ynodes column names or indicies in \code{data} of outcome nodes.
#' @param Qform character vector of regression formulas for \code{Qbar}.
#' @param cum.g a matrix of the cumulative probabilities of treatment (and being
    uncensored) given the parents.
#' @param abar binary vector (numAnodes x 1) of counterfactual treatment
#' @param stratify if \code{TRUE} condition on following \code{abar} when
    estimating \code{Qbar}. If \code{FALSE}, pool over all subjects.
#' @param type if \code{weight}, then use the g fits as weights in the submodel
    fit. If \code{covariate}, then use g fits as covariate.
#'
#' @return \code{icedr} returns a list of items as an object of class \code{tmle
    }, which include
#' \itemize{
#'     \item {The estimate of the parameter value under the intervention \code{
    abar}}
#'     \item {A matrix of the condition expectation fit \code{Qbar.hat} under \
    code{abar} for each time point.}
#'     \item {The conditional expectation fits for the initial fit \code{Q} and
    update step \code{Qstar}}
#'     \item {The call to the function}
#' }
#'
#' @references
#' van der Laan, M. J. and S. Gruber (2011): "Targeted Minimum Loss Based
    Estimation of an Intervention Specific Mean Outcome," The Berkeley Electronic
    Press.
#' @export
tmle = function(data, Anodes, Cnodes=NULL, Lnodes, Ynodes, Qform, cum.g, abar,
    stratify=FALSE, type="weight", SL.library=NULL) {

       ## INITIALIZE ##
```

```
Qbar.hat = IC.all = matrix(nrow=nrow(data), ncol=length(Ynodes), dimnames=
    list(NULL,Ynodes))
Q.kplus1 = data[,Ynodes[length(Ynodes)]]
Q = Qstar = list()
P_na = data
for(k in 1:length(Anodes)) P_na[,Anodes[k]] = rep(abar[k], each=nrow(data)
    )

## ICE ##
for(k in length(Ynodes):1) {
       if (k>1) is.deterministic = !is.na(data[,Ynodes[k-1]]) & data[,
           Ynodes[k-1]]==1 else is.deterministic = rep(FALSE, nrow(data))
       intervention.match = !is.na(data[,Anodes[k]]) & data[,Anodes[k]]==
           abar[k]
       if(is.null(Cnodes)) {
              uncensored = rep(TRUE, nrow(data))
       } else {
              uncensored = apply(data[,Cnodes[1:(k*2)], drop=FALSE], 1,
                  function(x) all(x=="uncensored", na.rm=T))
       }
       intervention.match = intervention.match & uncensored
       if (stratify) {
              index = intervention.match & !is.deterministic
       } else index = !is.deterministic & uncensored

       if(is.null(SL.library)) {
              QbarData = model.frame(Qform[k], cbind(data, Q.kplus1)[index
                  ,], na.action=NULL)
              Qinit = glm(Qform[k], data=QbarData, family="quasibinomial",
                   control=glm.control(trace=FALSE, maxit=1000))
              Qbar = predict(Qinit, newdata=P_na)
       } else {
              QbarData = cbind(data, Q.kplus1)[index,]
              tmp = complete.cases(P_na[,-1, drop=FALSE])
              if(ncol(QbarData[,-1, drop=FALSE])>0) {
                     Qbar = rep(NA, length(tmp))
                     vars = all.vars(as.formula(Qform[k])[[3]])
                     Qinit = mcSuperLearner(Y=QbarData$Q.kplus1, X=
                         QbarData[, vars, drop=FALSE], newX=P_na[tmp, vars
                         , drop=FALSE], SL.library=SL.library, family="
                         binomial", control = list(trimLogit=.001,
                         saveFitLibrary=FALSE), cvControl=list(V=8),
                         method="method.NNloglik.LT", verbose=TRUE)
                     library.predict = Qinit$library.predict
```

```
                        library.predict[is.na(library.predict)] = 0
                        Qinit$SL.predict = library.predict %*% Qinit$coef
                        Qbar[tmp] = qlogis(Qinit$SL.predict)
                        Qbar[!tmp] = qlogis(mean(QbarData$Q.kplus1))
                } else {
                        Qbar[tmp] = mean(QbarData$Q.kplus1)
                }
        }
        Qbar[is.deterministic] = qlogis(.9999)
        Qstar.data = data.frame(Q.kplus1, Qbar)
        Qupdate = Qupdate(Q.kplus1=Q.kplus1, Qbar=Qbar, cum.g=cum.g[,k],
            uncensored=uncensored, intervention.match=intervention.match,
            is.deterministic=is.deterministic, type=type)
        IC.all[,k] = calcIC(Q.kplus1, Qupdate$Q.kplus1, 1/cum.g[,k],
            uncensored, intervention.match)
        Q.kplus1 = Qbar.hat[,k] = Qupdate$Q.kplus1
        Q = c(Q, list(Qinit))
        Qstar = c(Qstar, list(Qupdate$Qstar))
        names(Q)[length(Q)] = names(Qstar)[length(Qstar)] = Ynodes[k]

    }

    out = list(estimate=mean(Q.kplus1), Qbar.hat=Qbar.hat, fit=list(Q=Q, Qstar
        =Qstar), IC=IC.all, call=match.call())
    class(out) = "tmle"
    return(out)

}
#'@export
print.tmle = function(x, ...) {
    cat("Call:\n", paste(deparse(x$call), sep = "\n", collapse = "\n"), "\n\n
        ", sep = "")
    cat("TMLE Estimate:\t", x$estimate, "\n")
    invisible(x)
}
```

C.5 Customized optimizers solving the EIF

```
####################
## TMLE OPTIMIZER ##
####################
FixScore = function(Qstar.kplus1, Qlogit, h.g.ratio, uncensored, intervention.
    match, is.deterministic, weight, max.objective = 1e-8) {
```

```
        m <- nlminb(start=0, objective=CalcScore, Qstar.kplus1=Qstar.kplus1,
            Qlogit=Qlogit, h.g.ratio=h.g.ratio, uncensored=uncensored, intervention
            .match=intervention.match, is.deterministic=is.deterministic, weight=
            weight, control=list(abs.tol=0.0001^2, eval.max=500, iter.max=500, x.
            tol=1e-14, rel.tol=1e-14))
        m$minimizer = "nlminb"
        names(m)[which(names(m)=="objective")] = "value"
        if (m$convergence!=0) {
                m = optim(par=0, fn=CalcScore, Qstar.kplus1=Qstar.kplus1, Qlogit=
                    Qlogit, h.g.ratio=h.g.ratio, uncensored=uncensored,
                    intervention.match=intervention.match, is.deterministic=is.
                    deterministic, weight=weight, control=list(abstol=max.objective
                    , reltol=1e-14, maxit=2000))
                m$minimizer = "optim"
                if (m$convergence!=0) {
                        return("Convergence unsuccessful.")
                }
        }
        names(m)[which(names(m)=="par")] = "coef"
        return(m)
}

######################
## DRICE OPTIMIZER ##
######################
FixScore2 = function(X, Q.kplus1, h.g.ratio, uncensored, intervention.match, is.
    deterministic, start) {

        FindMin = function(minimizer, start, objective, X, Q.kplus1, h.g.ratio,
            uncensored, intervention.match, is.deterministic, num.tries = 30 , max.
            objective = 1e-8) {

            for (i in 1:num.tries) {
                        if (minimizer == "nlminb") {
                                m <- nlminb(start=start, objective=CalcScore2, X=X,
                                    Q.kplus1=Q.kplus1, h.g.ratio=h.g.ratio,
                                    uncensored=uncensored, intervention.match=
                                    intervention.match, is.deterministic=is.
                                    deterministic, control=list(abs.tol=0.0001^2,
                                    eval.max=500, iter.max=500, x.tol=1e-14, rel.tol
                                    =1e-14))
                                coef <- m$par
                                obj.val <- m$objective
```

```
                    } else if (minimizer == "optim") {
                            m = optim(par=start, fn=CalcScore2, X=X, Q.kplus1=Q.
                                kplus1, h.g.ratio=h.g.ratio, uncensored=
                                uncensored, intervention.match=intervention.match
                                , is.deterministic=is.deterministic, control=list
                                (abstol=max.objective, reltol=1e-14, maxit=2000))
                            coef <- m$par
                            obj.val <- m$value
                    } else if (minimizer == "nlm") {
                            m <- nlm(f=CalcScore2, p=start, X=X, Q.kplus1=Q.
                                kplus1, h.g.ratio=h.g.ratio, uncensored=
                                uncensored, intervention.match=intervention.match
                                , is.deterministic=is.deterministic)
                            coef <- m$estimate
                            obj.val <- m$minimum
                    } else {
                            stop("bad minimizer")
                    }
                    if (obj.val < max.objective) {
                            return(list(coef=coef, solved=TRUE, m=m))
                    }
                    init.e = rnorm(length(start)) #if the first try didn't work,
                         try a random initial estimate of epsilon
            }
            return(list(coef=numeric(length(start)), solved=FALSE, m="score
                equation not solved!"))
        }

        ## Minimizes ##
        l <- FindMin("nlminb", start, objective, X, Q.kplus1, h.g.ratio,
            uncensored, intervention.match, is.deterministic, max.objective = 1e-8)
        if (!l$solved) l <- FindMin("optim", start, objective, X, Q.kplus1, h.g.
            ratio, uncensored, intervention.match, is.deterministic, max.objective
            = 1e-8)
        if (!l$solved) l <- FindMin("nlm", start, objective, X, Q.kplus1, h.g.
            ratio, uncensored, intervention.match, is.deterministic, max.objective
            = 1e-8)
        if (l$solved) return(l)
        if (!l$solved) return(l)     #stop("all minimizers failed")

}
```

C.6 Utility functions

```
########################
## INFLUENCE FUNCTION ##
########################
calcIC = function(Q.kplus1, Q.k, h.g.ratio, uncensored, intervention.match) {

       IC = rep(0, length(Q.k))
       index = uncensored & intervention.match
       if (any(h.g.ratio[index] != 0)) {
              IC[index] = (Q.kplus1[index] - Q.k[index]) * h.g.ratio[index]
       }
       return(IC)

}

####################
## SCORE FOR TMLE ##
####################
CalcScore <- function(e, Qstar.kplus1, Qlogit, h.g.ratio, uncensored,
    intervention.match, is.deterministic, weight=TRUE) {
        if (weight) {
               Qstar <- QstarFromE(e, Qlogit, rep(1,length(Qlogit)), is.
                   deterministic)
        } else Qstar <- QstarFromE(e, Qlogit, h.g.ratio, is.deterministic)
        ICtemp <- calcIC(Qstar.kplus1, Qstar, h.g.ratio, uncensored, intervention.
            match)
        return(sum(ICtemp)^2)
}

######################
## SCORE FOR DRICE ##
######################
CalcScore2 = function(coef, X, Q.kplus1, h.g.ratio, uncensored, intervention.
    match, is.deterministic) {
        coef[is.na(coef)] = 0
        Q.k = plogis(X %*% coef)
        Q.k[is.deterministic] = rep(1, sum(is.deterministic))
        ICtemp <- calcIC(Q.kplus1, Q.k, h.g.ratio, uncensored, intervention.match)
        return(sum(ICtemp)^2)
}

########################
## TMLE PERTURBATION ##
########################
```

```
QstarFromE <- function(e, off, X, is.deterministic) {
       Qstar <- plogis(off + e*X)
       Qstar[!is.na(is.deterministic) & is.deterministic] = rep(1, sum(is.
          deterministic, na.rm=TRUE))
       return(Qstar)
}

##########################
## TMLE UPDATE (WEIGHT) ##
##########################
Qupdate.wt = function(Q.kplus1, Qbar, cum.g, uncensored, intervention.match, is.
    deterministic) {
       cleverCov = ifelse(intervention.match, 1/cum.g, 0)
       Qstar.data = data.frame(Q.kplus1, Qbar)
       Qstar = glm(Q.kplus1 ~ offset(Qbar), weight=scale(cleverCov, center=FALSE)
          , data=Qstar.data, subset = !is.deterministic & (cleverCov > 0), family
          ="quasibinomial", control=glm.control(trace=FALSE, maxit=1000))
       Qstar$cleverCov = cleverCov
       score.tmle.wt = CalcScore(e=Qstar$coef, Qstar.kplus1=Q.kplus1, Qlogit=Qbar
          , h.g.ratio=1/cum.g, uncensored=uncensored, intervention.match=
          intervention.match, is.deterministic=is.deterministic, weight=TRUE)
       if (score.tmle.wt>0.0001) {
            Qstar = FixScore(Qstar.kplus1=Q.kplus1, Qlogit=Qbar, h.g.ratio=1/
               cum.g, uncensored=uncensored, intervention.match=intervention.
               match, is.deterministic=is.deterministic, weight=TRUE)
       }
       Q.kplus1 = plogis(Qbar + Qstar$coef)
       Q.kplus1[is.deterministic] = 1
       return(list(Qstar=Qstar, Q.kplus1=Q.kplus1))
}

#############################
## TMLE UPDATE (COVARIATE) ##
#############################
Qupdate.cov = function(Q.kplus1, Qbar, cum.g, uncensored, intervention.match, is.
    deterministic) {
       cleverCov = ifelse(intervention.match, 1/cum.g, 0)
       Qstar.data = data.frame(Q.kplus1, Qbar, cleverCov)
       Qstar = glm(Q.kplus1 ~ -1 + offset(Qbar) + cleverCov, data=subset(Qstar.
          data,!is.deterministic), family="quasibinomial", control=glm.control(
          trace=FALSE, maxit=1000))
       score.tmle.cov = CalcScore(e=Qstar$coef, Qstar.kplus1=Q.kplus1, Qlogit=
          Qbar, h.g.ratio=1/cum.g, uncensored=uncensored, intervention.match=
          intervention.match, is.deterministic=is.deterministic, weight=FALSE)
```

```
        if (score.tmle.cov>0.0001) {
                Qstar = FixScore(Qstar.kplus1=Q.kplus1, Qlogit=Qbar, h.g.ratio=1/
                    cum.g, uncensored=uncensored, intervention.match=intervention.
                    match, is.deterministic=is.deterministic, weight=FALSE)
        }
        Q.kplus1 = plogis(Qbar + Qstar$coef/cum.g)
        Q.kplus1[is.deterministic] = 1
        return(list(Qstar=Qstar, Q.kplus1=Q.kplus1))
}

#################
## TMLE UPDATE ##
#################
Qupdate = function(Q.kplus1, Qbar, cum.g, uncensored, intervention.match, is.
    deterministic, type = c("covariate", "weight")) {
        type <- match.arg(type)
        switch(type,
                        covariate = Qupdate.cov(Q.kplus1=Q.kplus1, Qbar=Qbar, cum.g=
                            cum.g, uncensored, intervention.match=intervention.match
                            , is.deterministic=is.deterministic),
                        weight = Qupdate.wt(Q.kplus1=Q.kplus1, Qbar=Qbar, cum.g=cum.
                            g, uncensored, intervention.match=intervention.match, is
                            .deterministic=is.deterministic))
}
```