Title
Semiparametric and Robust Methods for Complex Parameters in Causal Inference

Author
Zheng, Wenjing

Publication Date
2014

Peer reviewed|Thesis/dissertation

Semiparametric and Robust Methods for Complex Parameters in Causal Inference

by

Wenjing Zheng

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

in

Biostatistics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Mark van der Laan, Co-chair
Professor Antoine Chambaz, Co-chair
Professor Maya Petersen
Professor Jennifer Ahern

Fall 2014

Semiparametric and Robust Methods for Complex Parameters in Causal Inference

Abstract

Semiparametric and Robust Methods for Complex Parameters in Causal Inference

by

Wenjing Zheng

Doctor of Philosophy in Biostatistics

University of California, Berkeley

Professor Mark van der Laan, Co-chair

Professor Antoine Chambaz, Co-chair

This dissertation focuses on developing robust semiparametric methods for complex parameters that emerge at the interface of causal inference and biostatistics, with applications to epidemiological and medical research. Specifically, it address three important topics: Part I (chapter 1) presents a framework to construct and analyze group sequential *covariate-adjusted response-adaptive (CARA) randomized controlled trials (RCTs)* that admits the use of data-adaptive approaches in constructing the randomization schemes and in estimating the conditional response model. This framework adds to the existing literature on CARA RCTs by allowing flexible options in both their design and analysis. Part II (chapters 2 and 3) concerns two parameters that arise in longitudinal causal effect analysis using marginal structural models (MSMs). Chapter 2 presents a targeted maximum likelihood estimator (TMLE) for the the dynamic MSM for the hazard function. This estimator improves upon the existing inverse probability weighted (IPW) estimators by providing efficiency gain and robustness protection against model misspecification. Chapter 3 addresses the issue of effect modification (in a MSM) by an effect modifier that is post exposure. This parameter is particularly relevant if an effect modifier of interest is missing at random; or if one wishes to evaluate the effect modification of a second-line-treatment by a post first-line-treatment variable, where assignment of the first-line-treatment shares common determinants with the outcome of interest. We also present a TMLE for this parameter. Part III (chapters 4 and 5) addresses semiparametric inference for mediation analysis. Chapter 4 presents a TMLE estimator for the natural direct and indirect effects in a one-time point setting; it improves upon existing estimators by offering robustness, weakened sensitivity to near positivity violations, and potential applications to situations with high-dimensional mediators. Chapter 5 studies longitudinal mediation analysis with time-varying exposure and mediators. In it, we propose a reformulation of the mediation

problem in terms of stochastic interventions, establish an identification formula for the mediation functional, and present a TMLE for this parameter. This chapter contributes to existing literature by presenting a nonparametrically defined parameter of interest in longitudinal mediation and a multiply robust and efficient estimator for it.

Chapter 1: An adaptive trial design allows pre-specified modifications to some aspects of the on-going trial based on analysis of the accruing data, while preserving the validity and integrity of the trial. This flexibility potentially translates into more efficient studies (e.g. shorter duration, fewer subjects) or greater chance of answering clinical questions of interest (e.g. detecting a treatment effect if one exists, broader does-response information, etc). In an adaptive CARA RCT, the treatment randomization schemes are allowed to depend on the patient's pre-treatment covariates, and the investigators have the opportunity to adjust these schemes during the course of the trial based on accruing information, including previous responses, in order to meet some pre-specified objectives. In a *group-sequential* CARA RCT, such adjustments take place at interim time points given by sequential inclusion of blocks of c patients, where $c \geq 1$ is a pre-specified integer. In this chapter, we present a novel group-sequential CARA RCT design and corresponding analytical procedure that admits the use of flexible approaches in constructing randomization schemes and a wide range of data-adaptive techniques in estimating the conditional response model. Under the proposed framework, the sequence of randomization schemes is group-sequentially determined, using the accruing data, by targeting a formal, user-specified optimal randomization design. The parameter of interest is nonparametrically defined and is estimated using the paradigm of targeted minimum loss estimation. We establish that under appropriate empirical process conditions, the resulting sequence of randomization schemes converges to a fixed design, and the proposed estimator is consistent and asymptotically Gaussian, with an asymptotic variance that is estimable from data, thus giving rise to valid confidence intervals of given asymptotic levels. To illustrate the proposed framework, we consider LASSO regression in estimating the conditional outcome given treatment and baseline covariates. The asymptotic results ensue under minimal condition on the growth of the dimension of the regression coefficients and mild conditions on the complexity of the classes of randomization schemes.

Chapter 2: In many applications, one is often interested in the effect of a longitudinal exposure on a time-to-event process. In particular, consider a study where subjects are followed over time; in addition to their baseline covariates, at various time points we also record their time-varying exposure of interest, time-varying covariates, and indicators for the event of interest (say death). Time varying confounding is ubiquitous in these situations: the exposure of interest depends on past covariates that confound the effect of the exposure on the outcome of interest, in turn exposure affects future confounders;

right censoring may also be present in a study of this nature, often in response to past covariates and exposure. One way to assess the comparative effect of different regimens of interest is to study the hazard as a function of such regimens. The features of this hazard are often encoded in a marginal structural model. This chapter builds upon the work of Petersen, Schwab, Gruber, Blaser, Schomaker, and van der Laan (2014) to present a targeted maximum likelihood estimator for the marginal structural model for the hazard function under longitudinal dynamic interventions. The proposed estimator is efficient and doubly robust, hence offers an improvement over the incumbent IPW estimator.

Chapter 3: A crucial component of comparative effectiveness research is evaluating the modification of an exposure's effect by a given set of baseline covariates (effect modifiers). In complex longitudinal settings where time-varying confounding exists, this effect modification analysis is often performed using a marginal structural model. Generally, the conditioning effect modifiers in a MSM are cast as variables of the observed past. Yet, in some applications the effect modifiers of interest are in fact *counterfactual*. For instance, for a specific value of the first-line treatment, one may wish to evaluate the effect modification of a second-line-treatment by a post first-line-treatment variable, wherein the first-line-treatment assignment shares common determinants with the outcome of interest. In this case a simple stratification on the first-line treatment will only yield effect modification over a subpopulation given by said determinants. Hence, the wished parameter of interest should be formulated in terms of randomization on first-line treatment as well. In another example, the effect modifiers may be subject to missingness, which may depend on other baseline confounders; a simple complete-case analysis may introduce selection bias due to the high correlation of these confounders with the missingness of the effect modifier. In this case, one would formulate the wish parameter of interest in terms of an intervention on missingness. We call these *counterfactual effect modifiers*. In such situations, analysis by stratification alone may harbor selection bias. In this chapter, we investigate MSM defined by counterfactual effect modifiers. Firstly, we determine the identification of the causal dose-response curve and MSM parameters in this setting. Secondly, we establish the semiparametric efficiency theory for these statistical parameters, and present a substitution-based, semiparametric efficient and doubly robust estimator using the targeted maximum likelihood estimation methodology. However, as we shall see, due to the form of the efficient influence curve, the implementation of this estimator may prove arduous in applications where the effect modifier is high dimensional. To address this problem, our third contribution is a projected influence curve (and the corresponding TMLE estimator), which retains most of the robustness of its efficient peer and can be easily implemented in applications where the use of the efficient influence curve becomes taxing. In addition to these two robust estimators, we also present an IPW estimator, and a non-targeted G-computation estimator.

Chapter 4: In many causal inference problems, one is interested in the direct causal effect of an exposure on an outcome of interest that is not mediated by certain intermediate variables. Robins and Greenland (1992) and Pearl (2001) formalized the definition of two types of direct effects (natural and controlled) under the counterfactual framework. The efficient influence curves (under a nonparametric model) for the various natural effect parameters and their general robustness conditions, as well as an estimating equation based estimator using the efficient influence curve, are provided in Tchetgen Tchetgen and Shpitser (2011a). In this chapter, we apply the targeted maximum likelihood framework to construct a semiparametric efficient, multiply robust, substitution estimator for the natural direct effect which satisfies the efficient influence curve equation derived in Tchetgen Tchetgen and Shpitser (2011a). We note that the robustness conditions in Tchetgen Tchetgen and Shpitser (2011a) may be weakened, thereby placing less reliance on the estimation of the mediator density. More precisely, the proposed estimator is asymptotically unbiased if either one of the following holds: i) the conditional mean outcome given exposure, mediator, and confounders, and the mediated mean outcome difference are consistently estimated; (ii) the exposure mechanism given confounders, and the conditional mean outcome are consistently estimated; or (iii) the exposure mechanism and the mediator density, or the exposure mechanism and the conditional distribution of the exposure given confounders and mediator, are consistently estimated. If all three conditions hold, then the effect estimate is asymptotically efficient. Extensions to the natural indirect effect are also discussed.

Chapter 5: In this chapter, we study the effect of a time-varying exposure mediated by a time-varying intermediate variable. More specifically, consider a study where baseline covariates, time-varying treatment, time-varying mediator, time-varying covariates, and an outcome process are observed on subjects that are followed over time. The treatment of interest is influenced by past covariates and mediator, and affects future covariates and mediator. Right censoring, if present, occurs in response to past covariates and treatment. We also allow the outcome to be a time-to-event (say survival) process, in which case, at each time we record whether death has occurred. Due to subtleties that are unique to time-varying exposures and mediators, we reformulate the mediation problem in terms of stochastic interventions, as proposed by Didelez, Dawid, and Geneletti (2006) in the one-time point setting. Upon establishing the estimands of interest, we derive the efficient influence curves and establish their robustness properties. Applying the targeted maximum likelihood methodology, we use these efficient influence curves to construct multiply robust and efficient estimators. We also present an IPW estimator and a non-targeted substitution estimator for these parameters.

To my mother Xiaoling, my husband Karén and our son Ruben.

Contents

Acknowledgments

The completion of this doctoral work marks a milestone on a journey that has been nothing short of exhilarating, rewarding and inspiring since day one, and it was only possible thanks to the help and support of my mentors, colleagues, friends and family. I owe them my deepest gratitude. But the most heartfelt sentiments are often the most inarticulate ones; hard as I try to express them here, my gratitude towards these people and their contribution to my growth, both as a professional and as an individual, cannot be gauged by the limitations of my words.

I am eternally grateful to Mark van der Laan for being a passionate, dedicated, patient, exciting, challenging, generous and inspiring mentor. It is a privilege to learn statistics and a whole new way of thinking from him, benefits of which I will reap for the rest of my career. In addition to fostering my intellectual growth, his immaculate work ethics, drive, and resilience constantly inspire me to hold myself to the same standard. More than a nurturing educator and an inspiring role model, I am also deeply thankful to Mark for what he has helped me recognize in myself. I first came to Mark a lost soul, fresh out of undergraduate and not knowing my place in the vast scientific world. Mark welcomed me with open arms, and in his mentorship I found a home. His encouragement and exemplary tutelage helped me envision a career path and brought out my inner determination to achieve those goals. For all these and more — my most heartfelt thanks to Mark.

I am equally grateful to Maya Petersen for her patience, her contagious enthusiasm about our work, and her tireless support and guidance, for being the best complement to Mark's mentorship, and for showing me a whole new world in application problems and in HIV research. I am also most grateful to Maya for her generous guidance in helping me navigate a career path post graduation. Through learning causal inference and solving application problems from Maya, I gained a whole new perspective and saw how, as a statistician, one can have direct immediate contribution to, and inherent responsibility in, solving real world problems. Working with her (and the group of equally passionate and dedicated investigators) on HIV research, is an inspiring, humbling, exciting and rewarding experience, and is what motivates me to get up every morning.

My deepest thanks also goes to Antoine Chambaz, for his patience and encouragement, for being one of the most fun, exciting and inspiring collaborators and mentors I have had the honor to work with, for nurturing not only my intellectual growth but also my work-life balance. I am also thankful to Antoine for the precious opportunity to work with him in Paris, which is one of the most wonderful memories of my graduate career. Along with Antoine, I would like to extend my deep gratitude to Julie, Lou, Fausto and Claire, for welcoming me into their home during my stay in France; that epic Parisian winter was no match to the warmth of their beautiful family.

I am also thankful to Alan Hubbard, Sherri Rose, and Paul Chaffee for their encouragements to join Mark's group, which turned out to be a wonderful family and marked a turning point in my professional and personal development. I am also thankful to Alan for his continuing encouragement and moral support. Sherri, with her dedication and work ethics, has been a role model for me since day one; I am also thankful to her for her many generous advice in planning for a scientific career post graduation. I am grateful to Susan Gruber for our endless intellectually stimulating conversations, for always finding time to answer my questions no matter how busy she was, for her infinitely warm moral support, and for sharing her wisdom on everything from statistics to life and to the world at large. I thank my fellow students, Ivan Diaz, Sam Lendle, Molly Davies, Anna Decker, Linh Tran, Laura Balzer, to name a few, for all the fun we had taking our first TMLE classes and learning about influence curves together, as well for all the laughs and wonderful memories in Havilland.

The R code used in three out of five chapters in this dissertation are adapted from a beautifully crafted software by Josh Schwab. I am grateful to Josh for taking on the tremendous challenge that was implementing LTMLE and for his patience in helping me understanding his work. His bravery and skill has made my work a much more smooth sailing.

I also owe many thanks to Sharon Norris and Burke the-computing-meister Bundy — the unsung heroes of the Division of Biostatistics. Sharon has made my life so much easier and has pulled me out of bureaucratic hot water in more times than I can count. If it weren't for Burke's diligent upkeep of a smooth running linux cluster, my research would have easily taken twice as long.

Finally, the long trek of this doctoral training was fueled by the untiring and steadfast support of my family. The final days of completing this dissertation and the start of the next stage of my career would not haven been possible without my mother lovingly caring for my young son. My wonderful husband, Karén, has been my rock in good and bad times; it is thanks to his support, understanding and belief in me, that I have the freedom and confidence to pursue my dreams and jump into a new chapter of my profession life.

Part I

Covariate-Adjusted Response-Adaptive (CARA) Randomized Controlled (RCT) Trial Designs

Chapter 1

When Adaptive CARA RCT meets Data-Adaptive Estimation: Targeted Maximum Likelihood Estimation for Adaptive Randomized Controlled Trial Designs

1.1 Introduction

Adaptive clinical trial design methods have garnered growing attention in the recent years, in large part due to their greater flexibility over their traditional counterparts. In a traditional trial design, key aspects of the trials are set before the start of the data collection, usually based on assumptions about certain parameters of the study that are yet unsure at the design stage. The success of the trial, therefore, depends on the accuracy of these original assumptions. By contrast, an adaptive trial design allows pre-specified modifications to some aspects of the on-going trial based on analysis of the accruing data, while preserving the validity and integrity of the trial. This flexibility potentially translates into more efficient studies (e.g. shorter duration, fewer subjects) or greater chance of answering clinical questions of interest (e.g. detecting a treatment effect if one exists, broader does-response information, etc).

We focus here on the study of the so-called adaptive group-sequential *covariate-adjusted response-adaptive* (CARA) *randomized controlled trials* (RCT). In an adaptive CARA RCT, the treatment randomization schemes are allowed to depend on the patient's pre-treatment covariates, and the investigators have the opportunity to adjust these schemes

during the course of the trial based on accruing information, including previous responses, in order to meet some pre-specified objectives. In a *group-sequential* CARA RCT, such adjustments take place at interim time points given by sequential inclusion of blocks of c patients, where $c \geq 1$ is a pre-specified integer. We consider the case of $c = 1$ for simplicity of exposition, though the discussions generalize to any $c > 1$.

The trial protocol pre-specifies the observed data structure, scientific parameters of interest, analysis methods, and a criterion characterizing an optimal randomization scheme, which ought to reflect the goals of the adaptation and can be approximated using the accruing data. Here, baseline covariates and a primary outcome of interest are measured on each patient. We choose the marginal treatment effect of a binary treatment as the scientific parameter of interest, ψ_0. The analysis employs targeted minimum loss estimation (TMLE). The TMLE methodology was first introduced by van der Laan and Rubin (2006) in the independent identically distributed (i.i.d.) setting; its extension to adaptive RCTs was considered in van der Laan (2008) and Chambaz and van der Laan (2013), upon which a large part of this chapter relies. For concreteness sake, we choose the so-called Neyman design as our optimal randomization scheme. The Neyman design minimizes the Cramér-Rao lower bound on the asymptotic variances of a large class of estimators of ψ_0. The resulting Neyman allocation probabilities are evaluated conditionally on the baseline covariates. By targeting the Neyman design, we aim at improving the efficiency of the study, *i.e.*, at reaching a valid result using as few blocks of patients as possible. We emphasize that the results and procedures presented here are generally applicable to other parameters and optimal randomization schemes, after corresponding adjustments.

Since the randomization is response-adaptive, a consistent estimator of the conditional response model can more effectively steer the the adaptation towards the optimal randomization scheme. Moreover, since a patient's primary outcome is often correlated with many individual characteristics, greater latitude in adjusting for these baseline covariates, in both treatment allocation and outcome estimation, allows the investigators to better account for heterogeneity in the patient population. With the complexity of modern trials, information is often available on vast number of covariates; traditional parametric regression techniques are often too restrictive in such a high-dimensional scenario. While the use of data-adaptive techniques is very common in i.i.d. context (e.g. traditional RCT), its applicability in an adaptive RCT remains rather uncharted.

In this chapter, we present a general framework to construct and analyze group sequential CARA RCTs that admits the use of flexible approaches in constructing randomization schemes and in estimating the conditional response model. The proposed framework is *targeted* in the sense that: *(i)* the sequence of randomization schemes is group-sequentially determined, using the accruing data, by targeting a formal, user-specified optimal randomization design; *(ii)* the paradigm of targeted minimum loss estimation aims to optimize the bias-variance tradeoff of the estimates of the nuisance parameters towards the nonparamet-

rically defined parameter of interest. This framework injects flexibility into randomization adaptation and response prediction through loss-based estimation over classes of functions that may change with sample size. We establish that under appropriate empirical process conditions, the resulting sequence of randomization schemes converges to a fixed design, and the proposed estimator is consistent (regardless of the consistency of the response model) and asymptotically Gaussian, with an asymptotic variance that is estimable from data, hence giving rise to valid confidence intervals of given asymptotic levels. Moreover, the limiting design equals the target Neyman design if the response model is consistent and if the Neyman design can be approximated by the user-supplied classes of randomization schemes.

To illustrate the proposed framework, we consider LASSO regression in estimating the conditional outcome given treatment and baseline covariates. This example encompasses the parametric approach of Chambaz and van der Laan (2013) as a special case. The asymptotic results ensue under minimal condition on the growth of the dimension of the regression coefficients and mild conditions on the complexity of the classes of randomization schemes. The performance of the procedure is evaluated in a simulation study.

Before we delve into the main contents, let us motivate our discussion with a bird's eye view of the landscape of CARA designs.

Literature Review

Adaptive randomization has a long history that can be traced back to the 1930s. We refer to Rosenberger (1996), Rosenberger, Sverdlov, and Hu (2012), (Hu and Rosenberger, 2006, Section 1.2) and (Jennison and Turnbull, 2000, Section 17.4) for a comprehensive historical perspective. Many chapters are devoted to the study of response-adaptive randomizations, which select current treatment probabilities based on responses of previous patients, but not on the covariates of the current patients. We refer to Hu and Rosenberger (2006), Chambaz and van der Laan (2011a), Rosenberger et al. (2012) for a bibliography on that topic. In a heterogeneous population, however, it is often sensible to take into account the patients' characteristics for treatment assignment. CARA randomization tackles heterogeneity by dynamically calculating the allocation probabilities based on previous responses and current and past values of certain covariates. Compared to the broader literature on response-adaptive randomization, the advances in CARA randomization are relatively recent, but growing steadily. Among the first approaches, Rosenberger, Vidyashankar, and Agarwal (2001), Bandyopadhyay and Biswas (2001) considered randomization procedures defined as explicit functions of the conditional responses, which are modeled by generalized linear models. Though these procedures are not defined based on formal optimality criteria, their general goal is to allocate more patients to their corresponding "better" treatment arm. Atkinson and Biswas (2005) presented a biased-coin

design with skewed allocation, which is determined by sequentially maximizing a function that combines the variance of the parameter estimate, based on a Gaussian linear model for the conditional response, and the conditional treatment effect given covariates. Up till here, very little work had been devoted to asymptotic properties of CARA designs. Subsequently, Zhang, Hu, Cheung, and Chan (2007), Zhang and Hu (2009) established the efficiency theory for CARA designs converging to a given target design when the responses follow a generalized linear model, and proposed a covariate-adjusted doubly-adaptive biased coin design whose asymptotic variance achieves the efficiency bound. Chang and Park (2013) proposed a sequential estimation of CARA designs under generalized linear models for the response. This procedure allocates treatment based on the patients' baseline covariates, accruing information and sequential estimates of the treatment effect and uses a stopping rule that depends on the observed Fisher information. With regard to hypothesis testing, Shao, Yu, and Zhong (2010), Shao and Yu (2013) provided asymptotic results for valid tests under generalized linear models for the responses. Most recently, progress has also been made in CARA designs in the longitudinal settings, see for example Biswas, Bhattacharya, and Park (2014), Huang, Liu, and Hu (2013), Sverdlov, Rosenberger, and Ryeznik (2013). In the bulk of literature listed here, the analytic strategy fundamentally relies on defining the parameter of interest as regression coefficients on a generalized linear model. While this approach yields accessible estimators, the validity of its inference is at the mercy of the model specification. In particular, if the model is misspecified, the parameter estimates will be biased under the adaptive RCT sampling; in these cases, a standard RCT would be preferable.

Chambaz and van der Laan (2013) proposed a targeted CARA design where the treatment allocation is conditional on a summary measure of the covariates that takes only finitely many values. This framework defines the parameter of interest nonparametrically, and applies TMLE methodology in the analysis, thus leading to consistent and asymptotic normal parameter estimates that are robust to misspecification of the parametric working model for the response. However, assigning treatment based on finitely valued summary measures is perhaps too restrictive in real-life RCTs where response to treatment may be correlated with a large number of a patient's baseline characteristics, some of which continuous. Moreover, as mentioned before, although a misspecified parametric working model for the response does not hinder the consistency of the effect estimate, it may still hamper the estimate's efficiency and the convergence of the CARA design to the targeted optimal design.

We generalize the results of Chambaz and van der Laan (2013) to address the two issues mentioned above. We adopt a loss-based approach to the construction of more flexible CARA randomization schemes while exploiting data-adaptive estimators for the estimation of the response model, in search for more effective targeting of the optimal scheme, greater efficiency of the parameter estimate through better variable adjustments,

and more accurate estimation of the variance of the estimator.

Organization

The remainder of this chapter is organized as follows. In section 1.2, we introduce a general framework for constructing and analyzing an adaptive group-sequential CARA RCT design using data-adaptive loss-based estimation and TMLE procedure. In section 1.3, we present the theoretical results on the convergence of this targeted CARA design and the asymptotics, consistency and central limit theorem of the TMLE estimator. We illustrate this framework using the LASSO methodology to model the conditional response given baseline covariates and treatment in section 1.4. The performance of the LASSO-based group-sequential CARA RCT is assessed in a simulation study in section 3.5. The chapter closes on a summary in Section 1.6.

1.2 Targeted CARA RCT using Data-adaptive Loss-based Estimation

In the introduction, we have outlined the incentives to use flexible procedures to estimate the conditional response given treatment and covariates and to construct the randomization schemes. Such procedures we consider here stem from data-adaptive loss-based estimators for the nuisance parameters (van der Laan and Robins (2003a)); in particular, both the conditional response estimator and the adaptive randomization scheme are defined as minimizers of a user-specified weighted loss, over (possibly) changing classes of functions.

We begin by establishing the key features of the trial, namely, the parameter of interest, analysis method, and the optimal randomization scheme. Then, we describe the data generating process and the targeted minimum loss estimation procedure.

Observed Data Structure, Parameter of Interest and Optimal Design

Prior to data collection, the trial protocol specifies the observed data structure, parameter of interest and the target optimal randomization design, the latter two expressed in terms of features of the true, unknown data-generating process in the population of interest. In this chapter we consider a basic data structure and a simple parameter of interest. The range of application of the methods presented here extends beyond this limited yet instructive case.

Sections 2.2, 2.2 and 1.2 are respectively devoted to the presentation and discussion of the observed data structure, parameter of interest, and optimal randomization design.

Observed Data Structure

The data structure O writes as $O \equiv (W,A,Y)$, where $W \in \mathscr{W}$ consists of the baseline covariates (some of which may be continuous), $A \in \mathscr{A} \equiv \{0,1\}$ is the binary treatment of interest, and $Y \in \mathscr{Y}$ is the primary outcome of interest. We assume that the outcome space $\mathscr{O} \equiv \mathscr{W} \times \mathscr{A} \times \mathscr{Y}$ is bounded. Without loss of generality, we may then assume that $Y \in \mathscr{Y} \equiv (0,1)$ is bounded away from 0 and 1.

Every distribution of O consists of three components. On one hand, the marginal distribution of W and the conditional distribution of Y given (A,W) form a couple which is given by nature. On the other hand, the conditional distribution of A given W, also known as (a.k.a.) randomization scheme, is controlled by the investigators of the RCT. To reflect this dichotomy, we denote the distribution of O as $P_{Q,g}$, where Q equals the couple formed by the marginal distribution of W and the conditional distribution of Y given (A,W), and g equals the randomization scheme. We shall use $\mathscr{G}$ to denote the set of all randomization schemes. For a given Q, we denote Q_W the related marginal distribution of W and Q_Y the related conditional expectation of Y given (A,W). Moreover, we denote Q_0 the true couple in our population of interest, which is unknown to us, and we assume that this Q_0 does not vary during the whole duration of the RCT. Thus, for any Q and g, $P_{Q_0,g}$ is the true, partially unknown distribution of O when treatment is drawn using g, and $E_{P_{Q,g}}(Y|A,W) = Q_Y(A,W)$, $P_{Q,g}(A=1|W) = g(1|W) = 1 - g(0|W)$ $P_{Q,g}$-almost surely.

Parameter of Interest

In this chapter, the parameter of interest we consider is the marginal treatment effect on an additive scale:

$$\psi_0 \equiv E_{P_{Q_0,g}}\{Q_{Y,0}(1,W) - Q_{Y,0}(0,W)\} = \int (Q_{Y,0}(1,w) - Q_{Y,0}(0,w)) dQ_{W,0}(w),$$

which evidently depends on $P_{Q_0,g}$ only through Q_0. Of particular interest in medical, epidemiological and social sciences research, this parameter can be interpreted causally under assumptions on the data-generating process (Rosenbaum and Rubin (1983) and Pearl (1995)). Let $\mathscr{M}$ denote the set of all possible distributions of O. Central to our approach is formulating ψ_0 as the value at any $P_{Q_0,g}$ of the mapping $\Psi : \mathscr{M} \to [-1,1]$ characterized by

$$\Psi(P_{Q,g}) \equiv E_{P_{Q,g}}\{Q_Y(1,W) - Q_Y(0,W)\} = \int (Q_Y(1,w) - Q_Y(0,w)) dQ_W(w).$$

Since Ψ only depends on $P_{Q,g}$ through Q, we may sometimes write $\Psi(Q)$ in place of $\Psi(P_{Q,g})$.

The mapping Ψ is pathwise differentiable; its efficient influence curve sheds light on the asymptotic properties of all regular and asymptotically linear estimators of $\Psi(P_{Q_0,g})$.

The latter statement is formalized in the following lemma —we refer the reader to Bickel, Klaassen, Ritov, and Wellner (1998), van der Laan and Robins (2003a), van der Vaart (1998b) for definitions and proofs.

Lemma 1.1. *The mapping $\Psi : \mathscr{M} \to [-1,1]$ is pathwise differentiable at every $P_{Q,g} \in \mathscr{M}$ with respect to (wrt) the maximal tangent space. Its efficient influence curve at $P_{Q,g}$, denoted as $D^*(P_{Q,g})$, orthogonally decomposes as*

$$D^*(P_{Q,g})(O) = D_W(Q,g)(W) + D_Y(Q,g)(O)$$

with

$$\begin{aligned} D_W(Q)(W) &\equiv Q_Y(1,W) - Q_Y(0,W) - \Psi(Q), \\ D_Y(Q,g)(O) &\equiv \frac{2A-1}{g(A|W)}\left(Y - Q_Y(A,W)\right). \end{aligned}$$

The variance $\mathrm{Var}_{P_{Q,g}} D^*(P)(O)$ is a generalized Cramér-Rao lower bound for the asymptotic variance of any regular and asymptotically linear estimator of $\Psi(P_{Q,g})$ when sampling independently from $P_{Q,g}$.

Moreover, if either $Q_Y = Q'_Y$ or $g = g'$ then $E_{P_{Q,g}} D^*(P_{Q',g'})(O) = 0$ implies $\Psi(P_{Q,g}) = \Psi(P_{Q',g'})$.

The last statement of Lemma 2.1, often referred to as a "double-robustness" property, assures that D^* can be deployed to safeguard against model misspecifications when estimating ψ_0. This is especially relevant in an RCT setting, since the randomization scheme g is known whenever one samples an observation from $P_{Q,g}$.

Optimal Design

Suppose our goal of adaptation is to reach a randomization scheme of higher efficiency, *i.e.*, to obtain a valid estimate of ψ_0 using as few blocks of patients as possible. By Lemma 2.1, the asymptotic variance of a regular, asymptotically linear estimator is lower-bounded by $\min_{g \in \mathscr{G}} \mathrm{Var}_{P_{Q_0,g}} D^*(P_{Q_0,g})$. In this light, the *Neyman design* (see e.g. Hu and Rosenberger (2006))

$$g_0 \equiv \arg\min_{g \in \mathscr{G}} \mathrm{Var}_{P_{Q_0,g}} D^*(P_{Q_0,g}) = \arg\min_{g \in \mathscr{G}} E_{P_{Q_0,g}} \frac{\left(Y - Q_{Y,0}(A,W)\right)^2}{g^2(A|W)} \tag{1.1}$$

can be considered as an optimal randomization design ("optimal design" for short). Since its definition involves the unknown Q_0, the optimal design g_0 is unknown too. It is readily seen that g_0 is characterized by $g_0(1|W) = \sigma_0(1,W)/(\sigma_0(1,W) + \sigma_0(0,W))$, where $\sigma_0^2(A,W)$ is the conditional variance of Y given (A,W) under Q_0. It therefore appears that,

under this randomization scheme, the treatment arm with higher probability for a patient with baseline covariates W is the one for which the conditional variance of the outcome is higher.

In a situation where we knew the optimal design, we could undertake the covariate-adjusted trial consisting of drawing independently observations from P_{Q_0,g_0}. The next task would be to build a regular, asymptotically linear estimator with asymptotic variance $\mathrm{Var}_{P_{Q_0,g_0}} D^*(P_{Q_0,g_0})$ based on the resulting data. In the present situation, we are going to "target" g_0 at some pre-determined interim steps. By targeting g_0 we mean estimating g_0 based on past observations and relying on the resulting estimator to collect the next block of data. In addition to targeting g_0, each interim analysis will also consist of building an adaptive, targeted, regular and asymptotically linear estimator of ψ_0. The details of this procedure are presented in Section 1.2.

Data-Generating Mechanism and Estimation Procedures

Describing the data-generating mechanism amounts to presenting how we target the optimal design g_0 at each interim step, which involves the estimation of the conditional expectation $Q_{Y,0}$. We initiate the data-generating process in Section 5.2, describe a data-adaptive loss-based estimation procedure of $Q_{Y,0}$ in Section 1.2 and the related targeting procedure of g_0 in Section 1.2. By then, the data-generating mechanism is fully characterized by recursion.

Initiating the Data-Generating Mechanism

In what follows, we denote $O_i \equiv (W_i, A_i, Y_i)$ the i-th observation that we sample. The indexing indicates the time ordering of the data collection: $j < i$ means that O_j was collected before or at the same time as O_i. For convenience, we let $\mathbf{O}_n \equiv (O_1, \ldots, O_n)$ be the ordered vector of first n observations, with convention $O_0 \equiv \emptyset$. In the adaptive trial, the treatment A_i is drawn conditionally on W_i from the Bernoulli law with parameter $g_i(1|W_i)$, where the randomization scheme $g_i : \mathscr{A} \to [0,1]$ depends on past observations $\mathbf{O}_{i-1}$. We set $\mathbf{g}_n \equiv (g_1, \ldots, g_n)$, the ordered vector of first n randomization schemes. The data-generating distribution of $\mathbf{O}_n$ is denoted $\mathbf{P}_{Q_0,\mathbf{g}_n}$. It is formally characterized by the following factorization of the density of $\mathbf{O}_n$ wrt the product of the dominating measures: for any $g \in \mathscr{G}$,

$$\mathbf{P}_{Q_0,\mathbf{g}_n}(\mathbf{O}_n) = \prod_{i=1}^{n} P_{Q_0,g_i}(O_i) = \prod_{i=1}^{n} Q_{W,0}(W_i) \times g_i(A_i|W_i) \times P_{Q_0,g}(Y_i|A_i,W_i).$$

Let g^b be the balanced randomization scheme wherein each arm is assigned with probability $1/2$ regardless of baseline covariates. For a pre-specified n_0, we first draw n_0 independent observations $O_1, \ldots, O_{n_0}$ from P_{Q_0,g^b}. At an interim point, suppose one has thus

far drawn n observations $\mathbf{O}_n \sim \mathbf{P}_{Q_0,\mathbf{g}_n}$. An estimator of $Q_{Y,0}$ is obtained based on $\mathbf{O}_n$; the next randomization scheme g_{n+1} is defined using said estimator and $(\mathbf{O}_n, \mathbf{g}_n)$, then the $(n+1)$-th observation O_{n+1} is drawn from $P_{Q_0,g_{n+1}}$. We will describe the estimation of $Q_{Y,0}$ and construction of g_{n+1} in the two following sections.

Data-Adaptive Loss-based Estimation of the Conditional Outcome Expectation

At an interim point, we have drawn n observations $\mathbf{O}_n \sim \mathbf{P}_{Q_0,\mathbf{g}_n}$. The construction of our estimator of $Q_{Y,0}$ relies on a working model $\mathscr{Q}_{1,n}$ and on a loss function L for $Q_{Y,0}$, both specified by the investigators. Specifically, $Q_{Y,0}$ is the minimizer of $Q_Y \mapsto P_{Q_0,g}L(Q_Y)$ over $\mathscr{Q}_Y$, the set of all conditional expectations of Y given (A,W), and $\mathscr{Q}_{1,n}$ is a user-specified subset of $\mathscr{Q}_Y$. Note that the value of $g \in \mathscr{G}$ plays no role in this characterization. We can represent $\mathscr{Q}_{1,n} \equiv \{Q_{Y,\beta} : \beta \in B_n\}$, where B_n is an indexing set for $\mathscr{Q}_{1,n}$. Recall from section 1.2 that Y is bounded away from 0 and 1; therefore, we also require that $\mathscr{Q}_{1,n}$ be uniformly bounded away from 0 and 1 across all n. That is, there exists constants m_Y and M_Y satisfying $0 < m_Y < \inf_{Q_Y \in \bigcup_n \mathscr{Q}_{1,n}} \inf_O Q_Y(A,W)$ and $\sup_{Q_Y \in \bigcup_n \mathscr{Q}_{1,n}} \sup_O Q_Y(A,W) < M_Y < 1$.

Common among the possible loss functions are the least-square loss function

$$L(Q_{Y,\beta})(O) \equiv (Y - Q_{Y,\beta}(A,W))^2, \tag{1.2}$$

and, since Y is assumed bounded from 0 and 1, the quasi negative-log-likelihood loss function

$$-L(Q_{Y,\beta})(O) \equiv Y \log(Q_{Y,\beta}(A,W)) + (1-Y)\log\left(1 - Q_{Y,\beta}(A,W)\right). \tag{1.3}$$

Given a user-specified reference $g^r \in \mathscr{G}$ that is bounded away from 0 and 1, we estimate $Q_{Y,0}$ with the minimizer of the weighted empirical risk:

$$Q_{Y,\beta_n} \in \underset{Q_{Y,\beta} \in \mathscr{Q}_{1,n}}{\arg\min} \frac{1}{n}\sum_{i=1}^{n}\left(L(Q_{Y,\beta})(O_i)\frac{g^r(A_i|W_i)}{g_i(A_i|W_i)}\right). \tag{1.4}$$

Though different loss functions can yield risks that are minimized at $Q_{Y,0}$ over $\mathscr{Q}_Y$, the choice of loss function, however, distinctly affects the behavior of the estimator and conveys differential interpretation for its performance. Therefore, it is an important decision left to be addressed by the investigator. The class $\mathscr{Q}_{1,n}$ may depend on n, in which case its complexity grows with sample size; for inference sake, such growth should remain tethered. In section 1.3, we will learn sufficient conditions on the complexity of $\mathscr{Q}_{1,n}$ under which the empirical risk converges to the true risk over $\mathscr{Q}_{1,n}$. The choice of $\mathscr{Q}_{1,n}$, together with the loss function L, determines the technique used to estimate $Q_{Y,0}$. For instance, in the traditional parametric approach, $\mathscr{Q}_{1,n} = \mathscr{Q}_1$ for all n, and this class is indexed by a finite dimensional parameter set. In a spline regression model, $\mathscr{Q}_{1,n}$ is the

the set of smooth piecewise polynomial functions over the bounded $\mathscr{A} \times \mathscr{W}$ with a given degree and knot sequence. In partitioning estimation, for a given partition Π_n of $\mathscr{A} \times \mathscr{W}$, $\mathscr{Q}_{1,n}$ is the set of all piecewise constant functions wrt Π_n, and can be represented as the span of indicator functions on the partition cells. Under the LASSO methodology, which we depict in section 1.4, $\mathscr{Q}_{1,n}$ is the linear span of a given basis, with B_n being the set of coefficient vectors with a bounded L_1-norm and dimension that may depend on n.

Here, the reference randomization scheme g^r delivers the option to differentially weight each observation when fitting of the estimator Q_{Y,β_n}. But more than that, as we shall see in section 1.3, g^r will also serve to define an L^2 norm on $\mathscr{Q}_Y$.

Adapting Towards the Optimal Design

We now turn to the construction of the next randomization scheme g_{n+1}.

Our optimal design minimizes $g \mapsto \mathrm{Var}_{P_{Q_0,g}} D^*(P_{Q_0,g})$ over the class $\mathscr{G}$ of all randomization schemes, see (1.1). We adopt a loss-based approach, by defining g_{n+1} as the minimizer in g of an estimator of $\mathrm{Var}_{P_{Q_0,g}} D^*(P_{Q_0,g})$ over a user-specified class of randomization schemes. This approach is applicable in the largest generality. In the case that W is discrete, or if one is willing to assign treatment only based on a discrete summary measure V of W, g_{n+1} can be defined explicitly as an estimator of the Neyman design based on Q_{Y,β_n} and observations $\mathbf{O}_n$; we refer the readers to Chambaz and van der Laan (2013) for details.

To proceed, we first note that, for all $g' \in \mathscr{G}$,

$$g_0 = \arg\min_{g \in \mathscr{G}} E_{P_{Q_0,g'}} \frac{(Y - Q_{Y,0}(A,W))^2}{g(A|W)g'(A|W)}.$$

This equality teaches us that for the sake of estimating g_0 using observations drawn from $P_{Q_0,g'}$ we may consider the loss function L_{Q_Y} characterized over $\mathscr{G}$ by

$$L_{Q_Y}(g)(O) \equiv \frac{(Y - Q_Y(A,W))^2}{g(A|W)},$$

provided it is weighted by $1/g'(A|W)$. Note that this loss function is indexed by a given Q_Y.

Recall that we have already drawn n observations $\mathbf{O}_n \sim \mathbf{P}_{Q_0,\mathbf{g}_n}$ and estimated $Q_{Y,0}$ with Q_{Y,β_n}. Now, consider a class of randomization schemes, $\mathscr{G}_{1,n} \subset \mathscr{G}$ that may depend on n and are uniformly bounded away from 0 and 1 across all n. In other words, for $\mathscr{G}_1 \equiv \bigcup_{n \geq 1} \mathscr{G}_{1,n}$, $0 < m_g \leq \inf_{g \in \mathscr{G}_1} \inf_{O \in \mathscr{O}} g(A \mid W)$ and $\sup_{g \in \mathscr{G}_1} \sup_{O \in \mathscr{O}} g(A \mid W) \leq M_g < 1$, for some fixed constants m_g and M_g. We define the next randomization scheme as

$$g_{n+1} \in \arg\min_{g \in \mathscr{G}_{1,n+1}} \frac{1}{n} \sum_{i=1}^{n} \frac{L_{Q_{Y,\beta_n}}(g)(O_i)}{g_i(A_i|W_i)} = \arg\min_{g \in \mathscr{G}_{1,n+1}} \frac{1}{n} \sum_{i=1}^{n} \frac{\left(Y - Q_{Y,\beta_n}(O_i)\right)^2}{g(A_i|W_i)g_i(A_i|W_i)} \tag{1.5}$$

We emphasize the importance of the uniform boundedness of the union $\mathscr{G}_1$ when choosing the classes of estimators $\mathscr{G}_{1,n}$. In particular, this uniform boundedness implies the following property: given any $g^r \in \mathscr{G}$ that is bounded away from 0 and 1, there exists some constant $\kappa > 0$, such that $\left\| \frac{g}{g^r} \right\|_\infty \leq \kappa$, for all $g \in \mathscr{G}_1$. This property is pivotal in obtaining asymptotic results in the next section; for ease of reference, we shall call it the *dominated ratio property* of $\mathscr{G}_1$.

In many applications, it suffices to use a fixed class $\mathscr{G}_{1,n} = \mathscr{G}_1$ for all n. However, in some situations, e.g if the population is very heterogeneous, a suitable but fixed $\mathscr{G}_1$ may be too large for the adaptation to begin at a reasonable point (as a sizable sample may be required), and thereby depriving the trial of many of the advantages of an adaptive design. By allowing the class of randomization schemes to depend on n, one gains the flexibility to modify such classes according to the modesty or generosity of the sample size.

This completes the description of our data-generating mechanism.

Targeted Minimum Loss Estimation

Given n observations $\mathbf{O}_n \sim P_{Q_0,\mathbf{g}_n}$ and the estimator Q_{Y,β_n} of $Q_{Y,0}$ defined in (1.4), we may carry out the estimation of the parameter of interest ψ_0. We choose the targeted minimum loss estimation methodology. In the setting of a covariate-adjusted RCT with fixed design, a TMLE estimator is unbiased and asymptotically Gaussian regardless of the specification of the working model used for the estimation of $Q_{Y,0}$. It is known that unbiasedness and asymptotic normality still hold in the context of this chapter (CARA RCT for the estimation of ψ_0 based on copies of O), if the randomization schemes depend on W only through a summary measure taking finitely many values and the working model used for the estimation of $Q_{Y,0}$ is a simple linear model (this basically amounts to taking $d_n = d$ constant and $b_n = M$) in Section 1.4), see Chambaz and van der Laan (2013). Yet by relying on more flexible randomization schemes and on more adaptive estimators of $Q_{Y,0}$ we may achieve a greater efficiency through better estimation of the optimality criteria that may facilitate adaptation toward the optimal design, better adjustment of the variables that may directly improve the estimation of the parameter of interest, and more accurate estimation of the variance of the estimator.

In a glimpse, the proposed strategy consists of targetedly updating the initial estimator Q_{Y,β_n} by minimizing a pre-specified loss along a least favorable (wrt ψ_0) submodel through Q_{Y,β_n}, and then evaluating Ψ at the resulting updated estimator of Q_0. Formally, consider the negative-log-likelihood loss, see (1.3), which we denote L^*. This is a valid loss function for Q_Y upon our assumption that Y takes values within the unit interval. Correspondingly, consider the following one-dimensional parametric working model through Q_{Y,β_n}: for a

given closed, bounded interval $\mathscr{E} \subset \mathbb{R}$ containing 0 in its interior,

$$\left\{ Q_{Y,\beta_n}(\varepsilon) \equiv \text{expit}\left(\text{logit}(Q_{Y,\beta_n}) + \varepsilon H(g_n)\right) : \varepsilon \in \mathscr{E} \right\}, \tag{1.6}$$

with notation $H(g)(O) \equiv \frac{2A-1}{g(A|W)}$. This model passes through Q_{Y,β_n} at $\varepsilon = 0$ and satisfies the score condition $\frac{\partial}{\partial \varepsilon} L^*(Q_{Y,\beta_n}(\varepsilon))|_{\varepsilon=0} = D_Y(Q_{Y,\beta_n}, g_n)$. The optimal fluctuation parameter ε_n minimizes the weighted empirical risk along the working model:

$$\varepsilon_n \in \arg\min_{\varepsilon \in \mathscr{E}} \frac{1}{n} \sum_{i=1}^{n} L^*(Q_{Y,\beta_n}(\varepsilon))(O_i) \frac{g_n(A_i|W_i)}{g_i(A_i|W_i)}. \tag{1.7}$$

Set $Q^*_{Y,\beta_n} \equiv Q_{Y,\beta_n}(\varepsilon_n)$ and $Q^*_{\beta_n} \equiv (Q_{W,n}, Q^*_{Y,\beta_n})$, where $Q_{W,n}$ is the empirical marginal distribution of the W. The TMLE estimator of ψ_0 is defined as

$$\psi^*_n \equiv \frac{1}{n} \sum_{i=1}^{n} Q^*_{Y,\beta_n}(1, W_i) - Q^*_{Y,\beta_n}(0, W_i).$$

It satisfies $\psi^*_n = \Psi(Q^*_{\beta_n})$.

1.3 Asymptotics

In this section, we assay the theoretical properties of the targeted CARA design and its corresponding estimator depicted in section 1.2. Throughout this section, we should bear in mind that $\mathscr{G}_1 = \bigcup_n \mathscr{G}_{1,n}$ satisfies the dominated ratio property outlined in section 1.2. We first introduce further notation in Section 1.3, then we investigate the convergence of the targeted CARA design in Section 1.3 and the asymptotic behavior of the TMLE estimator in Section 1.3.

Notation

In general, given a known $g \in \mathscr{G}$ and an observation O drawn from $P_{Q_0,g}$, $Z \equiv g(A|W)$ is a deterministic function of g and O. Note that Z should be interpreted as a weight associated with O and will be used as such. Therefore, we can augment O with Z, *i.e.*, substitute (O,Z) for O, while still denoting $(O,Z) \sim P_{Q_0,g}$. In particular, during the course of our trial, conditionally on $\mathbf{O}_{i-1}$, the randomization scheme g_i is known and we can substitute $(O_i, Z_i) = (O_i, g_i(A_i|W_i)) \sim P_{Q_0,g_i}$ for O_i drawn from P_{Q_0,g_i}. By uniform boundedness of $\mathscr{G}_1$, the inverse weights $1/g_i(A_i|W_i)$ are bounded.

The empirical distribution of $\mathbf{O}_n$ is denoted P_n. For a function $f : \mathscr{O} \times [0,1] \to \mathbb{R}^d$, we will use the notation $P_n f \equiv n^{-1} \sum_{i=1}^{n} f(O_i, Z_i)$. Likewise, for any fixed $P_{Q,g} \in \mathscr{M}$,

$P_{Q,g}f \equiv E_{P_{Q,g}}f(O,Z)$ and, for each $i = 1,\ldots,n$, $P_{Q_0,g_i}f \equiv E_{Q_0,g_i}[f(O_i,Z_i)|\mathbf{O}_{i-1}]$, $\mathbf{P}_{Q_0,\mathbf{g}_n}f \equiv n^{-1}\sum_{i=1}^n E_{Q_0,g_i}[f(O_i,Z_i)|\mathbf{O}_{i-1}]$.

We endow $\mathscr{Q}_Y$ with the norm $\|\cdot\|_{2,P_{Q_0,g^r}}$ given by

$$\|Q_Y - Q'_Y\|^2_{2,P_{Q_0,g^r}} \equiv E_{P_{Q_0,g^r}}\left(Q_Y(A,W) - Q'_Y(A,W)\right)^2.$$

Similarly, we endow the set $\mathscr{G}$ with the norm $\|\cdot\|_{2,Q_{W,0}}$ given by

$$\|g - g'\|^2_{2,Q_{W,0}} \equiv E_{Q_{W,0}}\left(g(1|W) - g'(1|W)\right)^2.$$

More generally, for a function f on $\mathscr{O}\times[0,1]$, we shall use $\|\cdot\|_\infty$ to denote the sup norm $\|f\|_\infty \equiv \sup_{\mathscr{O}\times[0,1]}|f(O,Z)|$, and for an given $P \in \mathscr{M}$ and r, and we shall use $\|\cdot\|_{r,P}$ to denote the $L^r(P)$ norm $\|f\|_{r,P} \equiv (Pf^r)^{1/r}$.

For any class of functions $\mathscr{F}$, equipped with a norm $\|\cdot\|$ and $\delta > 0$, $N(\delta,\mathscr{F},\|\cdot\|)$ is the δ-bracketing number of $\mathscr{F}$ wrt $\|\cdot\|$ and $J(1,\mathscr{F},\|\cdot\|) \equiv \int_0^1 \sqrt{\log N(\delta,\mathscr{F},\|\cdot\|)}d\delta$ is the corresponding bracketing entropy (evaluated at 1).

Convergence of the Targeted CARA Design

Our first concern is the convergence of the estimators Q_{Y,β_n}, see (1.4). The equivalent of this result in the i.i.d. setting is well established (e.g. van der Vaart (1998b), Pollard (1984)). The following proposition revises those results for the current adaptive RCT setting.

Proposition 1.1 (Convergence of Q_{Y,β_n}).
Consider the following assumptions:

A1. *The conditional density under Q_0 of Y given (A,W) wrt some dominating measure is bounded away from 0.*

A2. *There exists a Q_{Y,β_0}, bounded away from 0 and 1, such that for all $\delta > 0$,*

$$P_{Q_0,g^r}L(Q_{Y,\beta_0}) < \inf_{\left\{\|Q_Y - Q_{Y,\beta_0}\|_{2,P_{Q_0,g^r}} \geq \delta : Q_Y \in \mathscr{Q}_Y\right\}} P_{Q_0,g^r}L(Q_Y).$$

A3. *For each n, there exists some $Q_{Y,\beta_{n,0}} \in \mathscr{Q}_{1,n}$ satisfying*

$$P_{Q_0,g^r}L(Q_{Y,\beta_{n,0}}) = \inf_{Q_{Y,\beta}\in\mathscr{Q}_{1,n}} P_{Q_0,g^r}L(Q_{Y,\beta}).$$

Moreover, $P_{Q_0,g^r}L(Q_{Y,\beta_{n,0}}) \to P_{Q_0,g^r}L(Q_{Y,\beta_0})$ as $n \to \infty$.

A4. *The classes* $\mathscr{Q}_{1,n}$ *satisfy* $J(1, L(\mathscr{Q}_{1,n}), \|\cdot\|_{2,P_{Q_0,g^r}}) = o(\sqrt{n})$.

Under ***A1–A4****,* $\left\|Q_{Y,\beta_n} - Q_{Y,\beta_0}\right\|_{2,P_{Q_0,g^r}} \to 0$ *in probability.*

Proof. See appendix. □

This proposition stipulates conditions under which Q_{Y,β_n} converges to some limiting Q_{Y,β_0}, which may depend on the user-supplied reference design g^r. If the true response model $Q_{Y,0}$ can be approximated by $\mathscr{Q}_{1,n}$, then $Q_{Y,\beta_0} = Q_{Y,0}$ by virtue of the loss function selection. Assumption **A2** requires that approximately minimizing the risk $P_{Q_0,g^r}L(Q_Y)$ should specify Q_{Y,β_0}. Assumption **A3** requires that the approximation error be tending to 0 as $n \to \infty$. The main constraint is assumption **A4**, which concerns the "size" or complexity of the classes $L(\mathscr{Q}_{1,n})$ (and hence of $\mathscr{Q}_{1,n}$). Understandably, the class of candidate estimators $\mathscr{Q}_{1,n}$ may grow with sample size; if this rate of increasing complexity (expressed in terms of bracketing integral) is controlled at $o(\sqrt{n})$, then we can still achieve convergence of the estimators Q_{Y,β_n}.

We now turn to the convergence of the targeted CARA design $\{g_n\}_{n\geq 1}$, see (1.5), toward a fixed, limiting design $g_0^* \in \mathscr{G}$.

Proposition 1.2 (Convergence of the targeted CARA Design).
Consider the setup of Proposition 1.1 and the following additional assumptions:

A5. *There exists* $g_0^* \in \mathscr{G}$*, bounded away from* 0 *and* 1*, such that for all* $\delta > 0$

$$P_{Q_0,g^r}\frac{L_{Q_{Y,\beta_0}}(g_0^*)}{g^r} < \inf_{\left\{\|g-g_0^*\|_{2,Q_{W,0}}\geq\delta : g\in\mathscr{G}\right\}} P_{Q_0,g^r}\frac{L_{Q_{Y,\beta_0}}(g)}{g^r}. \tag{1.8}$$

Similarly, for each n, there exists some $g_{n,0} \in \mathscr{G}_{1,n}$ *satisfying*

$$P_{Q_0,g^r}\frac{L_{Q_{Y,\beta_0}}(g_{n,0})}{g^r} = \inf_{g\in\mathscr{G}_{1,n}} P_{Q_0,g^r}\frac{L_{Q_{Y,\beta_0}}(g)}{g^r}.$$

Moreover, $P_{Q_0,g^r}\frac{L_{Q_{Y,\beta_0}}(g_{n,0})}{g^r} \to P_{Q_0,g^r}\frac{L_{Q_{Y,\beta_0}}(g_0^*)}{g^r}$*, as* $n \to \infty$.

A6. *The classes* $1/\mathscr{G}_{1,n} \equiv \{1/g : g \in \mathscr{G}_{1,n}\}$ *satisfy* $J(1, 1/\mathscr{G}_{1,n}, \|\cdot\|_{2,P_{Q_0,g^r}}) = o(\sqrt{n})$.

A7. *Let* $h_1(Q_Y)(O,Z) \equiv \left|Q_Y(O) - Q_{Y,\beta_0}(O)\right|$.
The classes $h_1(\mathscr{Q}_{1,n})$ *satisfy* $J(1, h_1(\mathscr{Q}_{1,n}), \|\cdot\|_{2,P_{Q_0,g^r}}) = o(\sqrt{n})$.

Under ***A1–A7****,* $\|g_n(1|W) - g_0^*(1|W)\|_{2,Q_{W,0}} \to 0$ *in probability.*

Proof. See appendix. □

We have already emphasized that through the choice of $\mathscr{G}_{1,n}$, the investigators of the RCT benefit from a great flexibility in treatment allocation. The main constraint on the classes $\mathscr{G}_{1,n}$ is **A6**, a condition on the complexity/richness of the class. We refer the reader to (van der Vaart, 1998b, Examples 19.7-19.11, Lemma 19.15) for typical examples. They notably include "well-behaved" parametric classes and VC classes. In particular, $\mathscr{G}_{1,n}$ can be a fixed class (not dependent on n) that consists of randomization schemes such that the allocation probabilities only depend on W through a discrete summary measure of it, as considered in Chambaz and van der Laan (2013). Under our proposed framework, since each adaptation of the randomization scheme depends on a loss function indexed by the estimator of $Q_{Y,0}$, convergence of the design will also depend on the complexity of the classes of outcome estimators $\mathscr{Q}_{1,n}$ — analogous to **A4**, assumption **A7** controls the rate of growth in complexity of $\mathscr{Q}_{1,n}$, as characterized by the bracketing entropy.

The limiting randomization scheme g_0^* depends on the user-supplied reference design g^r only through Q_{Y,β_0}: replacing g^r with any $g \in \mathscr{G}$ in (1.8) does not alter the definition of g_0^*. Furthermore, g_0^* can be interpreted as the most optimal design in some $\mathscr{G}_1' \supset \mathscr{G}_1$ given the limiting conditional outcome model Q_{Y,β_0}:

$$g_0^* \in \underset{g\in\mathscr{G}_1'}{\arg\min}\, \mathrm{Var}_{P_{Q_0,g}} D_Y^*(Q_{Y,\beta_0},g) = \underset{g\in\mathscr{G}_1'}{\arg\min} \left\{ \mathrm{Var}_{P_{Q_0,g}} D_Y^*(Q_0,g) + P_{Q_0,g} \frac{\left(Q_{Y,0}-Q_{Y,\beta_0}\right)^2}{g^2} \right\}.$$

Comparing the above equality with (1.1) yields that $g_0^* = g_0$, the Neyman design, whenever $Q_{Y,\beta_0} = Q_{Y,0}$ and $g_0 \in \mathscr{G}_1'$. In general, g_0^* minimizes an objective function that is the sum of the Cramér-Rao lower bound and a second-order residual. This underscores the motivation for using a flexible estimator in estimating $Q_{Y,0}$: by minimizing this second-order residual of the limiting conditional outcome model, we are closer to adapting toward the desired optimal design in $\mathscr{G}_1'$.

The convergence in probability of g_n also imply the following convergences that we shall use later.

Corollary 1.1 (Convergence of $\frac{1}{g_n}$, $\frac{1}{n}\sum_{i=1}^n g_i$, $\frac{1}{n}\sum_{i=1}^n \frac{1}{g_i}$).
$\|g_n - g_0^*\|_{2,Q_{W,0}}$ converges to 0 in probability, implies the following useful convergences:

1. *$\|g_n - g_0^*\|_{2,Q_{W,0}}$ converges to 0 in L^1.*
2. *$\left\| \frac{1}{g_n} - \frac{1}{g_0^*} \right\|_{2,Q_{W,0}}$ converges to 0 in probability and in L^1.*
3. *$\left\| \frac{1}{n}\sum_{i=1}^n g_i - g_0^* \right\|_{2,Q_{W,0}}$ converges to 0 in probability and in L^1.*

4. $\left\| \frac{1}{n}\sum_{i=1}^{n} \frac{1}{g_i} - \frac{1}{g_0^*} \right\|_{2,Q_{W,0}}$ *converges to* 0 *in probability and in* L^1.

Proof. See appendix. □

Consistency and Central Limit Theorem

Having secured the convergence of the data-adaptive initial conditional outcome estimator and of the targeted CARA RCT, we are now ready to obtain the consistency and a central limit theorem for the targeted parameter estimate ψ_n^* of ψ_0. As with the initial LASSO estimators of the conditional outcome, we are firstly concerned with the convergence of the updated estimators Q^*_{Y,β_n}, as this shall lay the stones for the consistency of ψ_n^* and a central limit theorem.

Proposition 1.3 (Consistency of ψ_n^*).
Consider the setups of Propositions 1.1 and 1.2 and the following additional assumption:

A8. *There exists a unique* $\varepsilon_0 \in \mathscr{E}$ *such that*

$$\varepsilon_0 \in \arg\min_{\varepsilon \in \mathscr{E}} P_{Q_0,g_0^*} L^*(Q_{Y,\beta_0}(\varepsilon)).$$

*Assume that **A1– A8** are met and define* $Q^*_{Y,\beta_0} \equiv Q_{Y,\beta_0}(\varepsilon_0)$. *It holds that* $\|Q^*_{Y,\beta_n} - Q^*_{Y,\beta_0}\|_{2,P_{Q_0,g^r}} \to 0$ *in probability. Moreover,* ψ_n^* *consistently estimates* ψ_0.

Proof. See appendix. □

If $Q_{Y,\beta_0} = Q_{Y,0}$ then $\varepsilon_0 = 0$; therefore, the updating procedure in TMLE will preserve the consistency of the initial estimator $\Psi(Q_{Y,\beta_n})$. More importantly, proposition 1.3 guarantees that even if $Q_{Y,\beta_0} \neq Q_{Y,0}$, ψ_n^* still consistently estimates ψ_0, by double-robustness of the methodology. As we shall see in proposition 1.4, the convergence of the updated outcome estimators Q^*_{Y,β_n} (to the truth or otherwise) is crucial for studying the asymptotic behavior of the parameter estimate ψ_n^*.

The following exact linear expansion of the TMLE estimate ψ_n^* is useful in proving the central limit theorem for ψ_n^*.

Lemma 1.2 (Exact Linear Expression of ψ_n^*).
For both $\beta = \beta_0$ *and* $\beta = \beta_n$, *introduce* $d^*_{Y,\beta}$ *and* $q^*_{Y,\beta}$ *given by*

$$\begin{aligned} d^*_{Y,\beta}(O,Z) &\equiv \frac{2A-1}{Z}\left(Y - Q^*_{Y,\beta}(A,W)\right), \\ q^*_{Y,\beta}(W) &\equiv Q^*_{Y,\beta}(1,W) - Q^*_{Y,\beta}(0,W), \end{aligned}$$

It follows from the definition of $\psi_n^* \equiv \Psi(Q_{\beta_n}^*)$ *that*

$$\begin{aligned}\psi_n^* - \psi_0 &= -P_{Q_0,g_0^*} D^*(P_{Q_{\beta_n}^*,g_0^*}) \quad (1.9)\\ &= (P_n - \mathbf{P}_{Q_0,\mathbf{g}_n})\left\{d_{Y,\beta_0}^* + D_W^*(Q_{\beta_0}^*)\right\}\\ &+ (P_n - \mathbf{P}_{Q_0,\mathbf{g}_n})\left\{\left(d_{Y,\beta_n}^* - d_{Y,\beta_0}^*\right) + \left(q_{Y,\beta_n}^* - q_{Y,\beta_0}^*\right)\right\} \quad (1.10)\end{aligned}$$

Proof. See appendix. □

In specifying the working model $\mathscr{Q}_{1,n}$ for the conditional outcome expectation, we have allowed the class of estimators to change with n. In general, given a sequence of classes $\mathscr{F}_n$, the nature of the functions may be very different between the classes, but in order to establish a central limit theorem, it is required that the sequence of envelope functions F_n satisfy the so-called Lindeberg condition:

$$\begin{aligned}P_{Q_0,g^r} F_n^2 &= O(1),\\ P_{Q_0,g^r} F_n^2 \left\{F_n > \delta\sqrt{n}\right\} &\to 0, \text{ for every } \delta > 0. \quad (1.11)\end{aligned}$$

This condition is fulfilled by the envelopes of the classes

$$\mathscr{Q}_{1,n}^* \equiv \left\{\text{expit}\left(\text{logit}(Q_{Y,\beta}) + \varepsilon H(g)\right) : Q_{Y,\beta} \in \mathscr{Q}_{1,n}, g \in \mathscr{G}_{1,n}, \varepsilon \in \mathscr{E}\right\},$$

due to their uniform boundedness. As we shall see in the next proposition, this Lindeberg condition and assumptions **A9** and **A10**, which concerns the asymptotic complexity of the classes $\mathscr{Q}_{1,n}$ and $\mathscr{G}_{1,n}$, pave the way to the convergence of random variables of the form $\sqrt{n}\left(P_n - \mathbf{P}_{Q_0,\mathbf{g}_n}\right)\theta(Q_{Y,\beta_n}^*)$, for a function $Q_Y \mapsto \theta(Q_Y)$.

Proposition 1.4 (Asymptotic Linearity and Central Limit Theorem for ψ_n^*).
Using the notations in lemma 1.2, define

$$\Sigma_n \equiv \frac{1}{n}\sum_{i=1}^{n}\left(d_{Y,\beta_n}^*(O_i,Z_i) + D_W^*(Q_{\beta_n}^*)(W_i)\right)^2. \quad (1.12)$$

Consider the setups of Propositions 1.1, 1.2 and 1.3 and the following additional assumption:

A9. *The sequence of entropies* $J\left(\delta_n, \mathscr{Q}_{1,n}, \|\cdot\|_{2,P_{Q_0,g^r}}\right) \to 0$ *for every* $\delta_n \downarrow 0$.

A10. *The sequence of entropies* $J\left(\delta_n, 1/\mathscr{G}_{1,n}, \|\cdot\|_{2,P_{Q_0,g^r}}\right) \to 0$ *for every* $\delta_n \downarrow 0$.

A11. *For any deterministic function* F, $F(O) = 0$ P_{Q_0,g_0^*}*-almost surely implies that* $F = 0$

*Assume that **A1–A11** are met. Then*

$$\begin{aligned}\psi_n^* - \psi_0 &= (P_n - \mathbf{P}_{Q_0,\mathbf{g}_n})\left\{d^*_{Y,\beta_0} + D^*_W(Q^*_{\beta_0})\right\} + o_P(1/\sqrt{n}) \\ &= \frac{1}{n}\sum_{i=1}^{n}\left\{d^*_{Y,\beta_0}(O_i,Z_i) + D^*_W(Q^*_{\beta_0})(W_i) - P_{Q_0,g_0^*}D^*(P_{Q^*_{\beta_0},g_0^*})\right\} + o_P(1/\sqrt{n}). \qquad (1.13)\end{aligned}$$

Moreover, $(\Sigma_n/n)^{-1/2}(\psi_n^* - \psi_0)$ *converges in distribution to the standard normal distribution.*

Proof. See appendix. □

The expression (1.13) amounts to the version asymptotic linearity under the current adaptive RCT setting. The last statement in proposition 1.8 underpin the statistical analysis of the proposed targeted CARA RCT. In particular, denoting $\xi_{1-\alpha/2}$ the $(1-\alpha/2)$-quantile of the standard normal distribution, $\left[\psi_n^* \pm \xi_{1-\alpha/2}(\Sigma_n/n)^{1/2}\right]$ is a confidence interval of asymptotic level $(1-\alpha)$.

1.4 Example: Targeted LASSO-based CARA RCT

In the previous two sections we have presented a general framework for constructing and analyzing CARA RCTs using data-adaptive loss-based estimators of the conditional outcome expectation, coupled with the TMLE methodology for parameter estimation. As described in the introduction, high-dimensional settings are increasingly common in clinical trials working with heterogenous populations. A popular device in high-dimensional statistics, due to its computational feasibility and amenability to theoretical study, is the LASSO methodology (Tibshirani (1996)) — a shrinkage and selection method for generalized regression models that optimizes a loss function of the regression coefficients subject to constraint on the L^1 norm. In this section, we illustrate the application of the proposed framework using a LASSO estimator for the conditional outcome expectation; the parametric estimators considered in Chambaz and van der Laan (2013) are a special case of a LASSO estimator.

LASSO Estimation of the Outcome's Conditional Expectation

Consider $\{b_n\}_{n\geq 1}$ and $\{d_n\}_{d\geq 1}$ two non-decreasing, possibly unbounded sequences over $\mathbb{R}_+$ and, for some $M > 0$ and every $n \geq 1$, introduce the subset

$$B_{M,n} \equiv \left\{\beta \in \ell^1 : \|\beta\|_1 \leq \min(b_n, M) \text{ and } \forall j \geq d_n,\ \beta^j = 0\right\} \qquad (1.14)$$

of $\ell^1 \equiv \left\{\beta \in \mathbb{R}^{\mathbb{N}} : \sum_{j\in\mathbb{N}} |\beta^j| < \infty\right\}$. Let $\left\{\phi_j : j \in \mathbb{N}\right\}$ be a uniformly bounded set of functions from $\mathscr{A} \times \mathscr{W}$ to $\mathbb{R}$. Without loss of generality, we may assume that $\|\phi_j\|_\infty = 1$ for all $j \in \mathbb{N}$. For all $\beta \in \ell^1$, we denote $\Phi_\beta : \mathscr{A} \times \mathscr{W} \to \mathbb{R}$ the function $\Phi_\beta(A,W) \equiv \sum_{j\in\mathbb{N}} \beta^j \phi_j(A,W)$.

The construction of our LASSO estimators for $Q_{Y,0}$ relies on user-specified working model $\mathscr{Q}_{1,n}$ and loss function L for $Q_{Y,0}$. For instance, we can take $\mathscr{Q}_{1,n} \equiv \{Q_{Y,\beta} \equiv \Phi_\beta : \beta \in B_{M,n}\}$ with $M = 1$, and the least-square loss function L in (1.2). We can also take $\mathscr{Q}_{1,n} \equiv \{Q_{Y,\beta} \equiv \text{expit}(\Phi_\beta) : \beta \in B_{M,n}\}$ with M a deterministic upper-bound on $|\text{logit}(Y)|$ (recall that Y is assumed bounded away from 0 and 1), and the quasi negative-log-likelihood loss function L in (1.3). Note that in both cases, for all $\beta \in B_{M,n}$, $\|Q_{Y,\beta}\|_\infty$ is upper-bounded by a deterministic upper-bound on $|Y|$.

Recall that we have already drawn n observations $\mathbf{O}_n \sim \mathbf{P}_{Q_0,\mathbf{g}_n}$. Given a user-specified reference $g^r \in \mathscr{G}$ that is bounded away from 0 and 1, we estimate $Q_{Y,0}$ with Q_{Y,β_n}, where

$$\beta_n \in \underset{\beta\in B_{M,n}}{\arg\min} \frac{1}{n}\sum_{i=1}^{n}\left(L(Q_{Y,\beta})(O_i)\frac{g^r(A_i|W_i)}{g_i(A_i|W_i)}\right). \tag{1.15}$$

The above minimization with the constraint $\|\beta\|_1 \leq \min(b_n, M)$, see (1.14), can be rewritten as a minimization free of the latter constraint by adding a term of the form $\lambda_n\|\beta\|_1$ to the empirical criterion, where λ_n depends on b_n. This is the so-called LASSO procedure introduced by Tibshirani (1996) for the sake of obtaining estimators with fewer nonzero parameter values, thus effectively reducing the number of variables upon which the given solution is dependent. Note that when $d_n = d$ is held constant by choice, (1.15) should be interpreted as a standard parametric procedure rather than as a LASSO.

Asymptotics for the Targeted LASSO-based CARA RCT

Using the theoretical results procured in the section 1.3, we show that the LASSO-based target CARA RCT design is indeed convergent, and the corresponding TMLE estimator is consistent and satisfies a central limit theorem.

Proposition 1.5 (Convergence of LASSO Q_{Y,β_n}).

Consider either the working model $\mathscr{Q}_{1,n}$ given by $Q_{Y,\beta} \equiv \Phi_\beta$ and the squared error loss L in (1.2), or the logistic model given by $Q_{Y,\beta} \equiv \text{expit}\left(\Phi_\beta\right)$ and the negative log-likelihood loss in (1.3).

Consider the following assumptions:

B1. *The conditional density under Q_0 of Y given (A,W) wrt some dominating measure is bounded away from* 0.

B2. *There exists a unique* $\beta_0 \in \bigcup_{n\geq 1} B_{M,n}$ *such that*

$$\beta_0 \in \underset{\beta\in\bigcup_{n\geq 1} B_{M,n}}{\arg\min}\ P_{Q_0,g^r} L(Q_{Y,\beta}).$$

B3. *For each n, there exists* $\beta_{n,0} \in B_{M,n}$ *such that*

$$\beta_{n,0} \in \underset{\beta\in B_{M,n}}{\arg\min} P_{Q_0,g^r} L(Q_{Y,\beta}).$$

B4. *It holds that* $d_n = O(n^r)$ *for some* $0 < r < 1$.

Under ***B1–B4****,* $\left\|Q_{Y,\beta_n} - Q_{Y,\beta_0}\right\|_{2,P_{Q_0,g^r}} \to 0$ *in probability.*

Proof. See appendix. □

The sequence of dimensions $\{d_n\}_n$ for the lasso coefficients are allowed to grow with n. However, to ensure that the resulting class $\mathscr{Q}_{1,n}$ has a manageable complexity (in terms of the entropy condition **A4**), it is required that the speed of $\{d_n\}_n$ be controlled at $O(n^r)$ for $0 < r < 1$ (assumption **B4**).

Proposition 1.6 (**Convergence of the targeted LASSO-based CARA Design**)**.**
Consider the setup of Proposition 1.5 and the following additional assumptions:

B5. *There exists* $g_0^* \in \mathscr{G}$*, bounded away from* 0 *and* 1*, such that for all* $\delta > 0$

$$P_{Q_0,g^r}\frac{L_{Q_{Y,\beta_0}}(g_0^*)}{g^r} < \inf_{\left\{\left\|g-g_0^*\right\|_{2,Q_{W,0}}\geq\delta : g\in\mathscr{G}\right\}} P_{Q_0,g^r}\frac{L_{Q_{Y,\beta_0}}(g)}{g^r}.$$

Similarly, for each n, there exists some $g_{n,0} \in \mathscr{G}_{1,n}$ *satisfying*

$$P_{Q_0,g^r}\frac{L_{Q_{Y,\beta_0}}(g_{n,0})}{g^r} = \inf_{g\in\mathscr{G}_{1,n}} P_{Q_0,g^r}\frac{L_{Q_{Y,\beta_0}}(g)}{g^r}.$$

Moreover, $P_{Q_0,g^r}\frac{L_{Q_{Y,\beta_0}}(g_{n,0})}{g^r} \to P_{Q_0,g^r}\frac{L_{Q_{Y,\beta_0}}(g_0^*)}{g^r}$.

B6. *The classes* $1/\mathscr{G}_{1,n} \equiv \{1/g : g \in \mathscr{G}_{1,n}\}$ *satisfy* $J(1, 1/\mathscr{G}_{1,n}, \|\cdot\|_{2,P_{Q_0,g^r}}) = o(\sqrt{n})$.

Under ***B1–B6****,* $\|g_n(1|W) - g_0^*(1|W)\|_{2,Q_{W,0}} \to 0$ *in probability.*

Proof. See appendix. □

As commented following proposition 1.2, assumption **B6** concerns the complexity of the class $\mathscr{G}_1$. The condition can be satisfied for well-behaved classes such as parametric classes or VC classes.

The assumptions **B1 – B6** readily guarantee convergence of the updated conditional outcome estimators and consistency of the TMLE parameter estimate.

Proposition 1.7 (Consistency of LASSO-based ψ_n^*).

Consider the setups of Propositions 1.5 and 1.6, and the following additional assumption:

B7. *There exists a unique $\varepsilon_0 \in \mathscr{E}$ such that*

$$\varepsilon_0 \in \underset{\varepsilon \in \mathscr{E}}{\arg\min} P_{Q_0,g_0^*} L^*(Q_{Y,\beta_0}(\varepsilon)).$$

*Assume that **B1– B7** are met and define $Q^*_{Y,\beta_0} \equiv Q_{Y,\beta_0}(\varepsilon_0)$.*

It holds that $\|Q^*_{Y,\beta_n} - Q^*_{Y,\beta_0}\|_{2,P_{Q_0,g^r}} \to 0$ in probability. Moreover, ψ_n^* consistently estimates ψ_0.

Proposition 1.8 (Central Limit Theorem for LASSO-based ψ_n^*).

Using the notations in lemma 1.2, define

$$\Sigma_n \equiv \frac{1}{n}\sum_{i=1}^{n}\left(d^*_{Y,\beta_n}(O_i,Z_i) + D^*_W(Q^*_{\beta_n})(W_i)\right)^2.$$

Consider the setups of Propositions 1.5, 1.6 and 1.7 and the following additional assumption:

B8. *The sequence of entropies $J\left(\delta_n, 1/\mathscr{G}_{1,n}, \|\cdot\|_{2,P_{Q_0,g^r}}\right) \to 0$ for every $\delta_n \downarrow 0$.*

B9. *For any deterministic function F, $F(O) = 0$ P_{Q_0,g_0^*}-almost surely implies that $F = 0$*

*Assume that **B1–B9** are met. Then*

$$\begin{aligned}\psi_n^* - \psi_0 &= (P_n - \mathbf{P}_{Q_0,\mathbf{g}_n})\left\{d^*_{Y,\beta_0} + D^*_W(Q^*_{\beta_0})\right\} + o_P(1/\sqrt{n}) \\ &= \frac{1}{n}\sum_{i=1}^{n}\left\{d^*_{Y,\beta_0}(O_i,Z_i) + D^*_W(Q^*_{\beta_0})(W_i) - P_{Q_0,g_0^*}D^*(P_{Q^*_{\beta_0},g_0^*})\right\} + o_P(1/\sqrt{n}).\end{aligned}$$

Moreover, $(\Sigma_n/n)^{-1/2}(\psi_n^* - \psi_0)$ converges in distribution to the standard normal distribution.

Proof. See appendix. □

1.5 Simulation Study

We present here the results of a simulation study of the performances of the targeted CARA RCT using LASSO estimators for the conditional outcome expectation

Simulation Scheme

We rely on the same simulation scheme as in Chambaz and van der Laan (2013). For completeness, let us recall that Q_0 is such that:

- the baseline covariate W equals (U,V), where U and V are independently drawn with U uniformly distributed on $[0,1]$ and $Q_{W,0}(V=1)=1/2$, $Q_{W,0}(V=2)=1/3$, $Q_{W,0}(V=3)=1/6$;
- the conditional distribution of Y given (A,W) is the Gamma distribution with conditional mean
$$Q_{Y,0}(A,W)=2U^2+2U+1+\left(AV+\frac{1-A}{1+V}\right)$$
and conditional variance
$$\sigma_0^2(Y|A,W)=\left(U+A(1+V)+\frac{1-A}{1+V}\right)^2.$$

The marginal treatment effect on an additive scale satisfies $\psi_0=\frac{91}{72}\simeq 1.264$.

In this study, we consider the simple case of a fixed class, as opposed to changing with n, of randomization schemes. Specifically, we target the optimal designs corresponding to eight parametric working models $\mathscr{G}_{11},\ldots,\mathscr{G}_{18}$ that we present in Table 1.1.

In addition to the latter parametric working models, we consider eight statistical procedures for the estimation of the conditional expectation $Q_{Y,0}$. Four of them consist of parametric estimation on small-dimensional models $\mathscr{Q}_{11},\ldots,\mathscr{Q}_{14}$. In contrast, the four others rely on moderate-dimensional parametric models, ℓ^1-penalization and cross-validation to select the best regularization parameter. We denote $\mathscr{Q}_{15},\ldots,\mathscr{Q}_{18}$ these "machine-learning", as opposed to "parametric", procedures/models, which embody the LASSO estimating procedure of Section 1.4. We summarize in Table 1.2 what are $\mathscr{Q}_{11},\ldots,\mathscr{Q}_{18}$. All procedures involve the logistic loss, even though the support of the marginal distribution of Y under P_0 is $\mathbb{R}_+$, not $[0,1]$. In fact, given a sample $O_1,\ldots,O_n$, we first scale $Y_1,\ldots,Y_n$ to $[0,1]$, then regress the scaled outcomes on (A,W) based on the logistic loss and one procedure among $\mathscr{Q}_{11},\ldots,\mathscr{Q}_{18}$, then scale back the resulting conditional expectation to the original range of the observed outcomes.

working model	parametric form	dimension	optimal variance
$\mathscr{G}_{11}$	θ_0	1	18.50
$\mathscr{G}_{12}$	$\sum_{v=1}^{3} \theta_v \mathbf{1}\{V=v\}$	3	18.18
$\mathscr{G}_{13}$	$\theta_0 + \theta_1 U$	2	18.37
$\mathscr{G}_{14}$	$\sum_{v=1}^{3} \theta_v \mathbf{1}\{V=v\} + \theta_4 U$	4	18.05
$\mathscr{G}_{15}$	$\theta_0 + \sum_{v=1}^{3} \theta_v \mathbf{1}\{V=v\} U$	4	18.12
$\mathscr{G}_{16}$	$\sum_{v=1}^{3} \theta_v \mathbf{1}\{V=v\} + \theta_4 U + \sum_{v=2}^{3} \theta_{3+v} \mathbf{1}\{V=v\} U$	6	18.01
$\mathscr{G}_{17}$	$\theta_0 + \sum_{v=1}^{3} \theta_v \mathbf{1}\{V=v\} U + \sum_{v=1}^{3} \theta_{4+v} \mathbf{1}\{V=v\} U^2$	7	18.36
$\mathscr{G}_{18}$	$\sum_{v=1}^{3} \theta_v \mathbf{1}\{V=v\} + \theta_4 U + \theta_5 U^2 + \sum_{v=2}^{3} \theta_{4+v} \mathbf{1}\{V=v\} U$ $+ \sum_{v=2}^{3} \theta_{6+v} \mathbf{1}\{V=v\} U^2$	9	18.03

Table 1.1: **Parametric working models $\mathscr{G}_{1k}$ ($k = 1, \ldots, 8$).** In the second column, we report the parametric forms of $\text{logit}\left((g_\theta(W) - \delta)/(1-2\delta)\right)$ for generic elements $g_\theta \in \mathscr{G}_{1k}$ ($k = 1, \ldots, 8$). We set $\delta = 10^{-2}$. In the third column, we give the dimensions of the models. In the fourth column, we report the numerical values of $\arg\min_{g \in \mathscr{G}_{1k}} \text{Var}_{P_{Q_0,g}} D^*(P_{Q_0,g})(O)$ ($k = 1, \ldots, 8$), with precision 10^{-2}. Recall that $\text{Var}_{P_{Q_0,g^b}} D^*(P_{Q_0,g})(O) = 23.87$, with precision 10^{-2}.

Set $B = 1000$ and let $n = (250, 500, 750, 1000, 1250, 1500, 1750, 2000, 2250, 2500)$ be a sequence of sample sizes. For each combination $(k,l) \in \{1, \ldots, 8\}^2$, we repeatedly simulate $B = 1000$ times a targeted CARA RCT based on $\mathscr{G}_{1k}$ and $\mathscr{Q}_{1l}$, performing an update of the randomization scheme and the computation of the TMLE of ψ_0 at every intermediate sample size n_i ($1 \leq i \leq 10$), which we denote $\psi^*_{n_i,klb}$. The simulations are mutually independent. The associated 95%-confidence intervals $\mathscr{I}_{n_i,klb}$ rely on estimated variances of the TMLE as given in (1.12). For each combination (k,l) and intermediate sample size n_i, we compute the empirical variance of the corresponding TMLE

$$\hat{S}_{n_i,kl} = \frac{1}{B} \sum_{b=1}^{B} \psi^{*2}_{n_i,klb} - \left(\frac{1}{B} \sum_{b=1}^{B} \psi_{n_i,klb} \right)^2$$

and the empirical coverage of the corresponding confidence interval

$$\hat{C}_{n_i,kl} = \frac{1}{B} \sum_{b=1}^{B} \mathbf{1}\{\psi_0 \in \mathscr{I}_{n_i,klb}\}.$$

The simulation study is conducted using `R` R Core Team (2014) and the package `glmnet` Friedman, Hastie, and Tibshirani (2010).

	working model	parametric form	dimension
parametric	$\mathscr{Q}_{11}$	$\sum_{v=1}^{3}\theta_v\mathbf{1}\{V=v\}+\theta_4U+\theta_5A$	5
	$\mathscr{Q}_{12}$	$\theta_0+A\left(\theta_1U+\sum_{v=2}^{3}\theta_v\mathbf{1}\{V=v\}\right)+(1-A)\left(\theta_4U+\sum_{v=2}^{3}\theta_{3+v}\mathbf{1}\{V=v\}\right)$	7
	$\mathscr{Q}_{13}$	$A\left(\sum_{v=1}^{3}\theta_v\mathbf{1}\{V=v\}+\theta_4U\right)+(1-A)\left(\sum_{v=1}^{3}\theta_{4+v}\mathbf{1}\{V=v\}+\theta_8U\right)$	8
	$\mathscr{Q}_{14}$	$A\left(\sum_{v=1}^{3}\theta_v\mathbf{1}\{V=v\}+\theta_4U+\theta_5U^2\right)+(1-A)\left(\sum_{v=1}^{3}\theta_{5+v}\mathbf{1}\{V=v\}+\theta_9U+\theta_{10}U^2\right)$	10
LASSO	$\mathscr{Q}_{15}$	$A\left(\sum_{v=1}^{3}\theta_v\mathbf{1}\{V=v\}+\theta_4U+\theta_5U^2\right)+(1-A)\left(\sum_{v=1}^{3}\theta_{5+v}\mathbf{1}\{V=v\}+\theta_9U+\theta_{10}U^2\right)$	10
	$\mathscr{Q}_{16}$	$A\left(\sum_{v=1}^{3}\theta_v\mathbf{1}\{V=v\}+\sum_{l=1}^{5}\theta_{3+l}U^l\right)+(1-A)\left(\sum_{v=1}^{3}\theta_{8+v}\mathbf{1}\{V=v\}+\sum_{l=1}^{5}\theta_{11+l}U^l\right)$	16
	$\mathscr{Q}_{17}$	$A\left(\sum_{v=1}^{3}\theta_v\mathbf{1}\{V=v\}+\sum_{l=1}^{10}\theta_{3+l}U^l\right)+(1-A)\left(\sum_{v=1}^{3}\theta_{13+v}\mathbf{1}\{V=v\}+\sum_{l=1}^{10}\theta_{16+l}U^l\right)$	26
	$\mathscr{Q}_{18}$	$A\left(\sum_{v=1}^{3}\theta_v\mathbf{1}\{V=v\}+\sum_{l=1}^{20}\theta_{3+l}U^l\right)+(1-A)\left(\sum_{v=1}^{3}\theta_{23+v}\mathbf{1}\{V=v\}+\sum_{l=1}^{20}\theta_{26+l}U^l\right)$	46

Table 1.2: **Working models $\mathscr{Q}_{1k}$(k=1, ..., 8) for the conditional expectation $Q_{Y,0}$.** In the second column, we report the parametric form of $\text{logit}\left((q_\theta(A,W)-\delta)/(1-2\delta)\right)$ for generic elements $q_\theta\in\mathscr{Q}_{1k}$ $(k=1,\ldots,8)$. We set $\delta=10^{-2}$. In the third column, we give the dimensions of the models. All working models are exploited in combination with the quasi negative-log-likelihood loss function (1.3). Models $\mathscr{Q}_{11},\mathscr{Q}_{12},\mathscr{Q}_{13},\mathscr{Q}_{14}$ are straightforwardly fitted by relying on the `R` function `glm`. Models $\mathscr{Q}_{15},\mathscr{Q}_{16},\mathscr{Q}_{17},\mathscr{Q}_{18}$ are LASSO-fitted by relying on the `R` function `glmnet`.

Discussion of the Results

Coverage

We propose an evaluation of the coverage performances based on testing. For every $(k,l)\in\{1,\ldots,8\}^2$ and n_i $(1\leq i\leq 10)$, the statistic $B\times\hat{C}_{n_i,kl}$ follows a Binomial distribution with parameter $(B,\pi_{n_i,kl})$ for some $\pi_{n_i,kl}\in[0,1]$. Denote $\hat{p}^{95}_{n_i,kl}$ the exact p-value of the one-sided binomial test of $H^{95}_{n_i,kl}$: "$\pi_{n_i,kl}\geq 95\%$" against "$\pi_{n_i,kl}<95\%$". Under $H^{95}_{n_i,kl}$, $\hat{p}^{95}_{n_i,kl}$ is drawn from the uniform distribution on $[0,1]$.

For every n_i $(1\leq i\leq 10)$, we carry out one-sample Kolmogorov-Smirnov tests of the null stating that the common law of $\{\hat{p}^{95}_{n_i,kl}:1\leq k\leq 8,l\in\mathscr{L}\}$ $(\mathscr{L}\subset\{1,\ldots,8\})$ is the uniform distribution on $[0,1]$ against the alternative that the common law is stochastically smaller than the uniform distribution on $[0,1]$. Rejecting the null for its alternative indicates a defective coverage. The p-values of four such Kolmogorov-Smirnov tests are

n_i		250	500	750	1000	1250	1500	1750	2000	2250	2500
$\bigcap_{1\leq k\leq 8, 1\leq l\leq 8}$	$H^{95}_{n_i,kl}$	< 0.001	< 0.001	0.011	0.003	0.011	0.006	0.110	0.362	0.003	0.059
$\bigcap_{1\leq k\leq 8, 1\leq l\leq 4}$	$H^{95}_{n_i,kl}$	< 0.001	0.015	0.023	< 0.001	0.151	0.034	0.025	0.080	0.281	0.414
$\bigcap_{1\leq k\leq 8, 5\leq l\leq 8}$	$H^{95}_{n_i,kl}$	< 0.001	< 0.001	0.175	0.567	0.004	0.037	0.785	0.804	0.004	0.072
$\bigcap_{1\leq k\leq 8, 1\leq l\leq 8}$	$H^{94}_{n_i,kl}$	0.028	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999

Table 1.3: **Evaluating the coverage performances based on testing.** The first row gives p-values of Kolmogorov-Smirnov tests of the null consisting of the intersection of all $H^{95}_{n_i,kl}$. The second and third rows give p-values of Kolmogorov-Smirnov tests of the nulls consisting of the intersections of all $H^{95}_{n_i,kl}$ based on parametric procedures (second row) and of all $H^{95}_{n_i,kl}$ based on LASSO procedures (third row). The fourth row gives p-values of Kolmogorov-Smirnov tests of the null consisting of the intersection of all $H^{94}_{n_i,kl}$.

reported in Table 1.3. The first row corresponds to the choice $\mathscr{L} = \{1,\ldots,8\}$. It teaches us that the expected 95%-coverage is generally not guaranteed. One may wonder if the same conclusion holds when focusing in turn on the parametric procedures (set $\mathscr{L} = \{1,\ldots,4\}$) or on the LASSO procedures (set $\mathscr{L} = \{5,\ldots,8\}$). Inspecting the second and third rows of Table 1.3 does not reveal an interesting pattern. One may now wonder to what extent the 95%-coverage is deficient. To answer this question, we proceed similarly. We denote $\hat{p}^{94}_{n_i,kl}$ the exact p-value of the one-sided binomial test of $H^{94}_{n_i,kl}$: "$\pi_{n_i,kl} \geq 94\%$" against "$\pi_{n_i,kl} < 94\%$". Under $H^{94}_{n_i,kl}$, $\hat{p}^{94}_{n_i,kl}$ is drawn from the uniform distribution on $[0,1]$. For every n_i ($1 \leq i \leq 10$), we carry out a one-sample Kolmogorov-Smirnov test of the null stating that the common law of $\{\hat{p}^{94}_{n_i,kl} : 1 \leq k \leq 8, 1 \leq l \leq 8\}$ is the uniform distribution on $[0,1]$ against the alternative that the common law is stochastically smaller than the uniform distribution on $[0,1]$. The p-values of these tests are reported in the fourth row of Table 1.3. The conclusion is clear and satisfactory: even if the 95%-confidence intervals fail to guarantee the wished coverage, one can safely consider them as valid 94%-confidence intervals.

Standard Deviation

Here we investigate how the targeted CARA RCT behaves in terms of standard deviation of the produced estimators. As in the previous subsection, the investigation relies on testing. For every $(k,l) \in \{1,\ldots,8\}^2$ and n_i $(1 \leq i \leq 10)$, we first compute the statistic

$$T_{n_i,kl} = \frac{\frac{1}{B}\sum_{b=1}^{B}(\Sigma_{n_i,klb})^{1/2} - \left(\hat{S}_{n_i,kl}\right)^{1/2}}{\left(\frac{1}{B}\sum_{b=1}^{B}\Sigma_{n_i,klb} - \left(\frac{1}{B}\sum_{b=1}^{B}(\Sigma_{n_i,klb})^{1/2}\right)^2\right)^{1/2}},$$

where $\Sigma_{n_i,klb}$ is the estimated variance of the TMLE produced at intermediate sample size n_i by the b-th simulated targeted CARA RCT based on $\mathscr{G}_{1k}$ and $\mathscr{Q}_{1l}$, see (1.12). Thus, $T_{n_i,kl}$ sheds some light on the estimation of the standard deviation of the TMLE $\psi^*_{n_i}$ at sample size n_i by $(\Sigma_{n_i}/n)^{1/2}$ for the targeted CARA RCT based on $\mathscr{G}_{1k}$ and $\mathscr{Q}_{1l}$.

For every n_i $(1 \leq i \leq 10)$, we perform a Lilliefors test of normality based on the sample $\{T_{n_i,kl} : 1 \leq k \leq 8, l \in \mathscr{L}\}$ with $\mathscr{L} = \{1,\ldots,8\}$. The p-values of these tests are reported in Table 1.4. They teach us that there is no stark evidence of non-normality across the ten intermediate sample sizes. This first conclusion justifies the next step: for every n_i $(1 \leq i \leq 10)$, we perform a one-sided Student test of "$\mu_{n_i} \geq 0$" against "$\mu_{n_i} < 0$", where μ_{n_i} denotes the mean of the common distribution of $\{T_{n_i,kl} : 1 \leq k \leq 8, l \in \mathscr{L}\}$ with $\mathscr{L} = \{1,\ldots,8\}$. The p-values of these tests are reported in the two first rows of Table 1.4. Adjusting for multiple testing in terms of the Benjamini and Yekutieli procedure for controlling the false discovery rate at the 5% level, we conclude that estimating the variance as in (1.12) is over-optimistic at least for intermediate sample sizes smaller than or equal to $n_3 = 750$. One may wonder if the same conclusions hold when focusing in turn on the parametric procedures (set $\mathscr{L} = \{1,\ldots,4\}$) or on the LASSO procedures (set $\mathscr{L} = \{5,\ldots,8\}$). Inspecting separately the third and fourth rows of Table 1.4 on one hand then the fifth and sixth rows on the other hand leads to the conclusion that estimating the variance as in (1.12) is over-optimistic only for intermediate sample sizes smaller than or equal to $n_2 = 500$, still adjusting for multiple testing in terms of the Benjamini and Yekutieli procedure for controlling the false discovery rate at the 5% level.

The gap between the conclusions reached when considering all procedures or the parametric and LASSO ones separately may be simply explained by a loss of power due to the reduction of sample size (64 versus 32), or by subtle differences induced by the nature of $\mathscr{Q}_{1l}$. In any case, in light of Section 1.5, the under-estimation of the true variance based on (1.12) is necessarily slight at most.

n_i	250	500	750	1000	1250	1500	1750	2000	2250	2500
Lilliefors	0.670	0.330	0.866	0.033	0.538	0.837	0.133	0.528	0.466	0.022
Student	< 0.001	< 0.001	0.002	0.006	0.008	0.012	0.007	0.008	0.044	0.038
Lilliefors	0.755	0.043	0.270	0.021	0.543	0.620	0.206	0.172	0.685	0.206
Student	< 0.001	< 0.001	0.013	0.026	0.025	0.026	0.021	0.036	0.226	0.420
Lilliefors	0.561	0.894	0.864	0.517	0.500	0.314	0.251	0.783	0.971	0.283
Student	< 0.001	< 0.001	0.044	0.059	0.087	0.116	0.084	0.063	0.050	0.011

Table 1.4: **Investigating the targeted CARA RCT in terms of standard deviation of the produced estimators.** In the first row we report the p-values of the Lilliefors tests of normality of the sample $\{T_{n_i,kl} : 1 \leq k,l \leq 8\}$ $(1 \leq i \leq 10)$. In the second row, we report the p-values of the Student tests of "$\mu_{n_i} \geq 0$" against "$\mu_{n_i} < 0$", where μ_{n_i} denotes the mean of the common distribution of $\{T_{n_i,kl} : 1 \leq k,l \leq 8\}$. In the third and fourth rows (fifth and sixth rows, respectively), we report the p-values of the same Lilliefors and Student tests based on the samples $\{T_{n_i,kl} : 1 \leq k \leq 8, 1 \leq l \leq 4\}$ corresponding to parametric procedures (on the samples $\{T_{n_i,kl} : 1 \leq k \leq 8, 5 \leq l \leq 8\}$ corresponding to LASSO procedures, respectively).

1.6 Summary

We have presented in this chapter a new group-sequential CARA RCT design and corresponding analytical procedure that admits the use of flexible data-adaptive techniques. The procedure is *targeted* in the sense that *(i)* the sequence of randomization schemes is group-sequentially determined by targeting a user-specified optimal randomization design based on accruing data and, *(ii)* the paradigm of targeted minimum loss estimation aims to optimize the bias-variance tradeoff of the estimates of the nuisance parameters towards the nonparametrically defined parameter of interest. For clarity sake, we focused on the marginal effect of a binary treatment as the parameter of interest and the Neyman allocation as the targeted optimal design, in an effort to produce an estimator with smaller asymptotic variance, but our methodology extends beyond this instructive framework.

Targeted minimum loss estimation is doubly robust, as it yields a consistent parameter estimate in the RCT setting regardless of the specification of the conditional response model. Nonetheless, when the randomization is response-adaptive, a consistent estimator of the conditional response model may lead to a more effective adaptation towards the optimal randomization scheme. Moreover, as a patient's primary outcome is often correlated with many individual characteristics, greater latitude in adjusting for these baseline covariates, in both treatment allocation and outcome estimation, allows the investigators to better account for heterogeneity in the patient population. These two observations motivate the use of flexible data-adaptive techniques in estimating the conditional outcome expectation, as well as in constructing the randomization schemes. The proposed framework incorporates such techniques through loss-based estimation over classes of estimators that may

change with n. Under assumptions on the rate of growth of these classes, the resulting sequence of randomization schemes converges to a limiting design, and the TMLE estimator is consistent and asymptotically Gaussian, with an asymptotic variance that we can estimate too. Consequently, we can build valid confidence intervals of given asymptotic levels.

To illustrate the proposed framework, we considered the case that the estimator of the conditional outcome given treatment and baseline covariates, a key element of the procedure, is obtained by LASSO regression. The asymptotic results can be achieved under minimal condition on the growth of the dimension of the regression coefficients and mild conditions on the complexity of the classes of randomization schemes. A simulation study confirms our theoretical results. Across 64 different choices of pairs of working models and 10 intermediate sample sizes ranging from 250 to 2500, there is no empirical evidence that the 95%-confidence intervals do not provide at least 94%-coverage, based on 1000 independent replications. In addition, in the same framework, there is no empirical evidence that the estimators of the variances of the TMLE estimators are over-optimistic for sample sizes larger than 500, adjusting for multiple testing in terms of the Benjamini and Yekutieli procedure for controlling the false discovery rate at the 5% level. For smaller sample sizes, the under-estimation is slight at most.

We will soon make available a `R` package to allow interested readers to test the procedure. In the future, we will also consider alternative strategies to randomly assign successive patients to the treatment arms in such a way that the overall empirical conditional distribution of treatment given baseline covariates be as close as possible to the current best estimator of the targeted optimal design. This will require both new theoretical developments and simulation studies.

1.7 Acknowledgements

The simulation study in section 3.5 is performed by Antoine Chambaz.

1.8 Chapter Appendix

A.1 Proofs of results in main content

Proof of proposition 1.1.

Define $\mathbf{M}(Q_Y) = P_{Q_0,g^r} L(Q_Y)$ and $\mathbf{M}_n(Q_Y) = \frac{1}{n}\sum_{i=1}^n \frac{g^r(A_i|W_i)}{Z_i} L(Q_Y)(O_i)$. To apply lemma 1.3 below, it suffices to show $\sup_{Q_{Y,\beta}\in\mathscr{Q}_{1,n}} \left|\mathbf{M}_n(Q_{Y,\beta}) - \mathbf{M}(Q_{Y,\beta})\right| = o_P(1)$.

Define $\ell(Q_Y)(O,Z) \equiv \frac{g^r(A|W)}{Z} L(Q_Y)(O)$. Then,

$$\sup_{Q_{Y,\beta} \in \mathscr{Q}_{1,n}} \left|\mathbf{M}_n(Q_{Y,\beta}) - \mathbf{M}(Q_{Y,\beta})\right| = \sup_{Q_{Y,\beta} \in \mathscr{Q}_{1,n}} \left| \frac{1}{n} \sum_{i=1}^{n} \frac{g^r(A_i \mid W_i)}{g_i(A_i \mid W_i)} L(Q_{Y,\beta})(O_i) - P_{Q_0,g^r} L(Q_{Y,\beta}) \right|$$

$$= \sup_{Q_{Y,\beta} \in \mathscr{Q}_{1,n}} \left| \frac{1}{n} \sum_{i=1}^{n} \ell(Q_{Y,\beta})(O_i, Z_i) - P_{Q_0,g_i} \ell(Q_{Y,\beta}) \right| = \sup_{f \in \ell(\mathscr{Q}_{1,n})} \left| (P_n - \mathbf{P}_{Q_0,\mathbf{g}_n}) f \right|.$$

To see that the last expression equals $o_P(1)$, we first note that due to uniform boundedness of $\mathscr{G}_1$ and boundedness of g^r, a δ-bracket for $L(\mathscr{Q}_{1,n})$ corresponds to a $k\delta$-bracket for $\ell(\mathscr{Q}_{1,n})$, for a fixed constant k. Therefore, $J\left(1, \ell(\mathscr{Q}_{1,n}), \|\cdot\|_{2,P_{Q_0,g^r}}\right)$ equals, up to a universal constant, to $J\left(1, L(\mathscr{Q}_{1,n}), \|\cdot\|_{2,P_{Q_0,g^r}}\right)$. Hence, we can apply lemma 1.7 with assumption **A4** to conclude that $\sup_{f \in \ell(\mathscr{Q}_{1,n})} \left|\left(P_n - \mathbf{P}_{Q_0,\mathbf{g}_n}\right) f\right| = o_P(1)$.

□

Proof of proposition 1.2.

Let $\ell_{Q_Y}(g)(O,Z) = \frac{L_{Q_Y}(g)(O)}{Z}$. We will apply lemma 1.3 with $\mathbf{M}(g) = P_{Q_0,g^r} \frac{L_{Q_{Y,\beta_0}}(g)}{g^r}$, and $\mathbf{M}_n(g) = \frac{1}{n} \sum_{i=1}^{n} \frac{L_{Q_{Y,\beta_n}}(g)(O_i)}{Z_i}$. It suffices to show that $\sup_{\mathscr{G}_{1,n}} |\mathbf{M}_n(g) - \mathbf{M}(g)| = o_P(1)$. Indeed,

$$\sup_{g \in \mathscr{G}_{1,n}} |\mathbf{M}_n(g) - \mathbf{M}(g)| = \sup_{g \in \mathscr{G}_{1,n}} \left| \frac{1}{n} \sum_{i=1}^{n} \frac{L_{Q_{Y,\beta_n}}(g)(O_i)}{Z_i} - P_{Q_0,g^r} \frac{L_{Q_{Y,\beta_0}}(g)}{g^r} \right|$$

$$\leq \sup_{g \in \mathscr{G}_{1,n}} \left| \frac{1}{n} \sum_{i=1}^{n} \frac{L_{Q_{Y,\beta_0}}(g)(O_i)}{Z_i} - P_{Q_0,g^r} \frac{L_{Q_{Y,\beta_0}}(g)}{g^r} \right| + \sup_{\mathscr{G}_{1,n}} \left| \frac{1}{n} \sum_{i=1}^{n} \frac{L_{Q_{Y,\beta_n}}(g)(O_i) - L_{Q_{Y,\beta_0}}(g)(O_i)}{Z_i} \right|$$

The first term on the right hand side of the inequality is $\sup_{f \in \ell_{Q_{Y,\beta_0}}(\mathscr{G}_{1,n})} \left|\left(P_n - P_{Q_0,\mathbf{g}_n}\right) f\right|$. From boundedness of Q_{Y,β_0} and uniform boundedness of $\mathscr{G}_1$, the bracketing numbers of $\ell_{Q_{Y,\beta_0}}(\mathscr{G}_{1,n})$ and $1/\mathscr{G}_{1,n}$ are proportional by a fixed constant. Therefore, applying lemma 1.7 with assumption **A6** yields $\sup_{f \in \ell_{Q_{Y,\beta_0}}(\mathscr{G}_{1,n})} \left|\left(P_n - P_{Q_0,\mathbf{g}_n}\right) f\right| = o_P(1)$.

We now study the second term:

$$\begin{aligned}
&\sup_{g\in\mathscr{G}_{1,n}}\left|\frac{1}{n}\sum_{i=1}^{n}\frac{L_{Q_{Y,\beta_n}}(g)(O_i)-L_{Q_{Y,\beta_0}}(g)(O_i)}{Z_i}\right| \\
&\lesssim \frac{1}{n}\sum_{i=1}^{n}\frac{\left|\left(Y_i-Q_{Y,\beta_n}(O_i)\right)^2-\left(Y_i-Q_{Y,\beta_0}(O_i)\right)^2\right|}{Z_i} \lesssim \frac{1}{n}\sum_{i=1}^{n}\frac{\left|Q_{Y,\beta_n}(O_i)-Q_{Y,\beta_0}(O_i)\right|}{Z_i} \\
&= P_{Q_0,\mathbf{g}_n}\frac{\left|Q_{Y,\beta_n}-Q_{Y,\beta_0}\right|}{Z}+(P_n-P_{Q_0,\mathbf{g}_n})\frac{\left|Q_{Y,\beta_n}-Q_{Y,\beta_0}\right|}{Z} \\
&= P_{Q_0,g^r}\left|Q_{Y,\beta_n}-Q_{Y,\beta_0}\right|+(P_n-P_{Q_0,\mathbf{g}_n})\frac{\left|Q_{Y,\beta_n}-Q_{Y,\beta_0}\right|}{Z} \\
&= o_P(1).
\end{aligned}$$

The first inequality follows from the uniform boundedness of $\mathscr{G}_{1,n}$. The second inequality follows boundedness of Q_{Y,β_n}, Q_{Y,β_0}, Y, and an application of the mean value theorem, which states that if a function f is continuous on $[a,b]$ and differentiable on (a,b), then $f(b)-f(a)=f'(c)(b-a)$, with $c\in(a,b)$. For the last equality, we first note that first term on the left hand side converges to 0 in probability. Indeed, from the Cauchy-Schwartz inequality we obtain $P_{Q_0,g^r}\left|Q_{Y,\beta_n}-Q_{Y,\beta_0}\right| \leq \left\{P_{Q_0,g^r}\left(Q_{Y,\beta_n}-Q_{Y,\beta_0}\right)^2\right\}^{1/2}$, where the larger term converges to 0 in probability by proposition 1.1. The second term on the right hand side is upper bounded by $\sup_{f\in h_1(\mathscr{Q}_{1,n})}\left|\left(P_n-P_{Q_0,\mathbf{g}_n}\right)\frac{f}{Z}\right|$. To see that this upper bound converges in probability to 0, it suffices note that since Z is uniformly bounded away from 0 and 1, the bracketing numbers for the classes $h_1(\mathscr{Q}_{1,n})$ and $\{f/Z : f\in h_1(\mathscr{Q}_{1,n}))\}$ are proportional by a fixed constant. Therefore, assumption **A7** allows us to apply lemma 1.7 to reach the desired conclusion on this upper bound. □

Proof of corollary 1.1.

Results 1 and 2 follow from uniform boundedness of $\mathscr{G}_1$. It remains to show result 3 (result 4 is proven in analogous manner). We show result 3 by showing the L^1 convergence, which in turn implies the convergence in probability. Indeed,

$$E\left(\left\|\frac{1}{n}\sum_{i=1}^{n}g_i-g_0^*\right\|_{2,Q_{W,0}}\right)\leq E\left(\frac{1}{n}\sum_{i=1}^{n}\left\|g_i-g_0^*\right\|_{2,Q_{W,0}}\right)=\frac{1}{n}\sum_{i=1}^{n}E\left\|g_i-g_0^*\right\|_{2,Q_{W,0}}.$$

Result 1 implies that the sequence $\left\{E\left\|g_i-g_0^*\right\|_{2,Q_{W,0}}\right\}_i$ converges to 0, then, applying Cesaro's lemma, one may conclude that their partial sums also converge to 0, thus proving the desired L^1 convergence. □

Proof of proposition 1.3.

Define a deterministic function $\mathbf{M}(\varepsilon) = \left| P_{Q_0,g_0^*} D_Y\left(Q_{Y,\beta_0}(\varepsilon), g_0^*\right)\right|$, and a stochastic process $\mathbf{M}_n(\varepsilon) = \left|\frac{1}{n}\sum_{i=1}^n \frac{g_n(A_i|W_i)}{g_i(A_i|W_i)} D_Y\left(Q_{Y,\beta_n}(\varepsilon), g_n\right)\right|$. From definition of ε_0 and differentiability of $L^*(Q_{Y,\beta}(\varepsilon))$ with respect to ε, we know that $\mathbf{M}(\varepsilon_0) = 0$. By construction of TMLE, we also know that $\mathbf{M}_n(\varepsilon_n) = 0$. To apply lemma 1.3, it suffices to show $\sup_{\varepsilon\in\mathscr{E}} |\mathbf{M}_n - \mathbf{M}| = o_P(1)$. Indeed,

$$
\begin{aligned}
&\sup_{\varepsilon\in\mathscr{E}} |\mathbf{M}_n(\varepsilon) - \mathbf{M}(\varepsilon)| \\
&= \sup_{\varepsilon\in\mathscr{E}} \left| \frac{1}{n}\sum_{i=1}^n \frac{2A_i - 1}{g_i(A_i \mid W_i)} \left(Y_i - Q_{Y,\beta_n}(\varepsilon)(A_i,W_i)\right) - P_{Q_0,g_0^*} \frac{2A-1}{g_0^*(A \mid W)} \left(Y - Q_{Y,\beta_0}(\varepsilon)(A,W)\right)\right| \\
&\leq \sup_{\varepsilon\in\mathscr{E}} \left| (P_n - \mathbf{P}_{Q_0,\mathbf{g}_n}) \frac{2A-1}{Z} \left(Y - Q_{Y,\beta_0}(\varepsilon)\right)\right| \qquad (1.16) \\
&+ \sup_{\varepsilon\in\mathscr{E}} \left| \frac{1}{n}\sum_{i=1}^n \frac{2A_i - 1}{g_i(A_i \mid W_i)} \left(Q_{Y,\beta_0}(\varepsilon)(A_i,W_i) - Q_{Y,\beta_n}(\varepsilon)(A_i,W_i)\right)\right| \qquad (1.17)
\end{aligned}
$$

We first study (1.16). Defining $f_\varepsilon(O,Z) = \frac{2A-1}{Z}\left(Y - Q_{Y,\beta_0}(\varepsilon)(A,W)\right)$, the expression (1.16) is equivalent to $\sup_{\varepsilon\in\mathscr{E}} \left|\left(P_n - \mathbf{P}_{Q_0,\mathbf{g}_n}\right) f_\varepsilon\right|$. In order to apply lemma 1.7 (with a fixed class), we shall first control the bracketing entropy of the class $\{f_\varepsilon : \varepsilon \in \mathscr{E}\}$. Using the shorthand notation $q(A,W) \equiv \text{logit}(Q_{Y,\beta_0})(A,W)$, we note that

$$
\begin{aligned}
&|f_{\varepsilon_1}(O,Z) - f_{\varepsilon_2}(O,Z)| \\
&= \left|\frac{2A-1}{Z}\right| \left|\text{expit}\left\{q(A,W) + \varepsilon_1 H(g_0^*)(A,W))\right\} - \text{expit}\left\{q(A,W) - \varepsilon_2 H(g_0^*)(A,W)\right\}\right| \\
&\leq \left|\frac{2A-1}{g_Z(A \mid W)}\right| \left|\exp\left\{-q(A,W) - \varepsilon_1 H(g_0^*)(A,W))\right\} - \exp\left\{-q(A,W) - \varepsilon_2 H(g_0^*)(A,W))\right\}\right| \\
&\lesssim |\varepsilon_1 - \varepsilon_2|.
\end{aligned}
$$

For the last inequality, we first bound the difference of the exponential functions using the mean value theorem argument as before. Then, we apply uniform boundedness of $\mathscr{G}_1$ to conclude that expression on the left-hand-side of the inequality is less or equal, by up to a constant, to the distance between the ε's. This Lipschitz condition, together with the boundedness of $\mathscr{E}$, imply that the parametric class $\{f_\varepsilon : \varepsilon \in \mathscr{E}\}$ indeed satisfies the entropy condition of lemma 1.7 (e.g. example 19.7 in van der Vaart (1998b)). The boundedness conditions in the same lemma is satisfied by uniform boundedness of $\mathscr{G}_1$, Y and $Q_{Y,\beta}(\varepsilon)$. Therefore, we conclude from lemma 1.7 that (1.16) converges to 0 a.s.

Next, we study (1.17). Let us adopt the notations $s_{n,\varepsilon}(O) = \text{logit}(Q_{Y,\beta_n}) + \varepsilon H(g_n)$ and $s_{0,\varepsilon}(O) = \text{logit}(Q_{Y,\beta_0}) + \varepsilon H(g_0^*)$, so that $Q_{Y,\beta_n}(\varepsilon)(O) = \text{expit}\left(s_{n,\varepsilon}(O)\right)$ and $Q_{Y,\beta_0}(\varepsilon)(O) =$

$\text{expit}(s_{0,\varepsilon}(O))$. Then, we have

$$
\begin{aligned}
(1.17) &\leq \sup_{\varepsilon\in\mathscr{E}} \frac{1}{n}\sum_{i=1}^{n}\frac{1}{Z_i}\left|\text{expit}(s_{n,\varepsilon}(O_i)) - \text{expit}(s_{0,\varepsilon}(O_i))\right| \\
&\lesssim \sup_{\varepsilon\in\mathscr{E}} \frac{1}{n}\sum_{i=1}^{n}\frac{1}{Z_i}\left|s_{n,\varepsilon}(O_i) - s_{0,\varepsilon}(O_i)\right| \\
&\leq \sup_{\varepsilon\in\mathscr{E}} \frac{1}{n}\sum_{i=1}^{n}\frac{1}{Z_i}\left\{\left|\text{logit}(Q_{Y,\beta_n}(O_i)) - \text{logit}(Q_{Y,\beta_0}(O_i))\right| + |\varepsilon|\,|(H(g_n)(O_i) - H(g_0^*)(O_i))|\right\} \\
&\lesssim \frac{1}{n}\sum_{i=1}^{n}\frac{\left|\text{logit}(Q_{Y,\beta_n}(O_i)) - \text{logit}(Q_{Y,\beta_0}(O_i))\right|}{Z_i} + \frac{1}{n}\sum_{i=1}^{n}\frac{|1/g_n(O_i) - 1/g_0^*(O_i)|}{Z_i} \\
&\lesssim \frac{1}{n}\sum_{i=1}^{n}\frac{\left|Q_{Y,\beta_n}(O_i) - Q_{Y,\beta_0}(O_i)\right|}{Z_i} + \frac{1}{n}\sum_{i=1}^{n}\frac{|1/g_n(O_i) - 1/g_0^*(O_i)|}{Z_i} \\
&\equiv (P_n - \mathbf{P}_{Q_0,\mathbf{g}_n})\frac{\left|Q_{Y,\beta_n} - Q_{Y,\beta_0}\right|}{Z} + \mathbf{P}_{Q_0,\mathbf{g}_n}\frac{\left|Q_{Y,\beta_n} - Q_{Y,\beta_0}\right|}{Z} \\
&+ (P_n - \mathbf{P}_{Q_0,\mathbf{g}_n})h_2(g_n) + \mathbf{P}_{Q_0,\mathbf{g}_n}h_2(g_n),
\end{aligned}
$$

where $h_2(g)(O,Z) \equiv \frac{\left|1/g(O) - 1/g_0^*(O)\right|}{Z}$. The second inequality follows from applying the mean value theorem argument on the exponential function, as we did in the previous paragraph. The fourth inequality follows from boundedness of $\mathscr{E}$. The fifth inequality follows from yet another mean value theorem argument. We have shown in proof of proposition 1.2 that under the stated assumptions, $\left(P_n - \mathbf{P}_{Q_0,\mathbf{g}_n}\right)\frac{\left|Q_{Y,\beta_n} - Q_{Y,\beta_0}\right|}{Z} + \mathbf{P}_{Q_0,\mathbf{g}_n}\frac{\left|Q_{Y,\beta_n} - Q_{Y,\beta_0}\right|}{Z} = o_P(1)$. To apply lemma 1.7 to conclude that $\left(P_n - \mathbf{P}_{Q_0,\mathbf{g}_n}\right)h_2(g_n) = o_P(1)$, it suffices to note that **A6** and the uniform boundedness of $\mathscr{G}_{1,n}$ and g_0^* implies that $h_2(\mathscr{G}_{1,n})$ satisfies the entropy condition in said lemma. Finally,

$$
\begin{aligned}
&\mathbf{P}_{Q_0,\mathbf{g}_n}h_2(g_n) \lesssim P_{Q_0,g^r}\left|g_n - g_0^*\right| \\
&= E_{Q_{W,0}}\left|g_n(1 \mid W) - g_0^*(1 \mid W)\right| \leq \left\{E_{Q_{W,0}}\left(g_n(1 \mid W) - g_0^*(1 \mid W)\right)^2\right\}^{1/2}.
\end{aligned}
$$

The upper bound converges to 0 in probability by proposition 1.2. We have thus shown that $\sup_{\varepsilon\in\mathscr{E}}|\mathbf{M}_n(\varepsilon) - \mathbf{M}(\varepsilon)| = o_P(1)$. Applying lemma 1.3, we may conclude that ε_n converges to ε_0 in probability.

Define $\mathscr{Q}_1' = \{Q_{Y,\beta_0}^*\}\bigcup_n \mathscr{Q}_{1,n}$, and $\mathscr{G}_1' = \{g_0^*\}\bigcup_n \mathscr{G}_{1,n}$. For the space $\mathscr{Q}_1' \times \mathscr{G}_1' \times \mathscr{E}$, we define a norm that is the sum of the componentwise norms. The previous results imply that $(Q_{Y,\beta_n}, g_n, \varepsilon_n)$ converge to $(Q_{Y,\beta_0}, g_0^*, \varepsilon_0)$ in probability. Define the function $f(Q_Y, g, \varepsilon) \equiv \text{expit}\left(\text{logit}(Q_Y) + \varepsilon\frac{2A-1}{g}\right)$. To see that f is continuous over $\mathscr{Q}_1' \times \mathscr{G}_1' \times \mathscr{E}$,

we note that

$$\|f(Q_{Y,1},g_1,\varepsilon_1)-f(Q_{Y,2},g_2,\varepsilon_2)\|_{2,P_{Q_0,g^r}} \lesssim \left\|\text{logit}(Q_{Y,1})+\varepsilon_1\frac{2A-1}{g_1}-\text{logit}(Q_{Y,2})-\varepsilon_2\frac{2A-1}{g_2}\right\|_{2,P_{Q_0,g^r}}$$
$$\leq \|\text{logit}(Q_{Y,1})-\text{logit}(Q_{Y,2})\|_{2,P_{Q_0,g^r}}+\left\|\varepsilon_2\left(\frac{1}{g_1}-\frac{1}{g_2}\right)\right\|_{2,P_{Q_0,g^r}}+\left\|\frac{1}{g_1}(\varepsilon_1-\varepsilon_2)\right\|_{2,P_{Q_0,g^r}}$$
$$\lesssim \|Q_{Y,1}-Q_{Y,2}\|_{2,P_{Q_0,g^r}}+\|g_1-g_2\|_{2,Q_{W,0}}+|\varepsilon_1-\varepsilon_2|,$$

where the first and the last inequalities follow from the uniform boundedness of the classes $\mathscr{Q}_1'$, $\mathscr{G}_1'$ and $\mathscr{E}$, and from a mean value theorem argument on the exponential function and the log function, respectively. Therefore, continuous mapping theorem applies to f. In particular, we may conclude that $Q^*_{Y,\beta_n}=f(Q_{Y,\beta_n},g_n,\varepsilon_n)$ converges in probability to $Q^*_{Y,\beta_0}=f(Q_{Y,\beta_0},g_0^*,\varepsilon_0)$. This proves our first claim.

The second claim follows directly from property of D_Y:

$$\begin{aligned}0&=P_{Q_0,g_0^*}D_Y(Q^*_{Y,\beta_0},g_0^*)=P_{Q_0,g_0^*}\left(\frac{2A-1}{g_0^*(A\mid W)}\left(Y-Q^*_{Y,\beta_0}\right)\right)\\&=E_{Q_{W,0}}\left(Q_{Y,0}(1,W)-Q_{Y,0}(0,W)\right)-E_{Q_{W,0}}\left(Q^*_{Y,\beta_0}(1,W)-Q^*_{Y,\beta_0}(0,W)\right)\\&=\Psi(Q_0)-\Psi(Q^*_{\beta_0}),\end{aligned}$$

where $Q^*_{\beta_0}\equiv(Q_{W,0},Q^*_{Y,\beta_0})$. □

Proof of lemma 1.2.

The first equality (1.9) follows directly from definition of D^*:

$$\begin{aligned}P_{Q_0,g_0^*}D^*(P_{Q^*_{\beta_n},g_0^*})&=P_{Q_0,g_0^*}\frac{2A-1}{g_0^*}\left(Y-Q^*_{Y,\beta_n}\right)+P_{Q_0,g_0^*}q^*_{Y,\beta_n}-P_nq^*_{Y,\beta_n}\\&=P_{Q_0,g_0^*}q_{Y,0}-P_{Q_0,g_0^*}q^*_{Y,\beta_n}+P_{Q_0,g_0^*}q^*_{Y,\beta_n}-P_nq^*_{Y,\beta_n}\\&=\Psi(Q_0)-\Psi(Q^*_{\beta_n}),\end{aligned}$$

where $q_{Y,0}\equiv Q_{Y,0}(1,W)-Q_{Y,0}(0,W)$.

Let $P_{n,\mathbf{g}_n}$ denote the empirical distribution of $\mathbf{O}_n$ weighted by $\frac{g_n(A_i|W_i)}{g_i(A_i|W_i)}$. To see (1.10) holds, first note that

$$\begin{aligned}\psi_n^*-\psi_0&=-P_{Q_0,g_0^*}D^*(P_{Q^*_{\beta_n},g_0^*})\\&=-\left\{P_{Q_0,g_0^*}D_Y(Q^*_{\beta_n},g_0^*)+P_{Q_0,g_0^*}D_W(Q^*_{\beta_n})\right\}+\left\{P_{n,\mathbf{g}_n}D^*_Y(Q^*_{\beta_n},g_n)+P_nD_W(Q^*_{\beta_n})\right\}\\&=\underbrace{P_nD_W(Q^*_{\beta_n})-P_{Q_0,g_0^*}D_W(Q^*_{\beta_n})}_{(a)}+\underbrace{P_{n,\mathbf{g}_n}D^*_Y(Q^*_{\beta_n},g_n)-P_{Q_0,g_0^*}D_Y(Q^*_{\beta_n},g_0^*)}_{(b)},\end{aligned}$$

where the second equality follows from the fact that both summands in $P_{n,\mathbf{g}_n} D_Y^*(Q^*_{\beta_n}, g_n) + P_n D_W(Q^*_{\beta_n})$ are zero, as implied by construction of Q^*_{Y,β_n}. Our proof of (1.10) relies on the following expansions. Firstly,

$$\begin{aligned}
(a) &= \left(P_n - P_{Q_0,g_0^*}\right) D_W(Q^*_{\beta_n}) + \left(P_n - P_{Q_0,g_0^*}\right) D_W(Q^*_{\beta_0}) - \left(P_n - P_{Q_0,g_0^*}\right) D_W(Q^*_{\beta_0}) \\
&= (P_n - Q_{W,0}) D_W(Q^*_{\beta_0}) + (P_n - Q_{W,0}) \left(D_W(Q^*_{\beta_n}) - D_W(Q^*_{\beta_0})\right) \\
&= (P_n - Q_{W,0}) D_W(Q^*_{\beta_0}) + (P_n - Q_{W,0}) \left(q^*_{Y,\beta_n} - q^*_{Y,\beta_0}\right) + (P_n - Q_{W,0}) \left(\Psi(Q^*_{\beta_0}) - \Psi(Q^*_{\beta_n})\right) \\
&= (P_n - P_{Q_0,\mathbf{g}_n}) D_W(Q^*_{\beta_0}) + (P_n - P_{Q_0,\mathbf{g}_n}) \left(q^*_{Y,\beta_n} - q^*_{Y,\beta_0}\right).
\end{aligned}$$

Secondly,

$$\begin{aligned}
(b) &= \frac{1}{n} \sum_{i=1}^{n} \left(\frac{g_n(A_i \mid W_i)}{g_i(A_i \mid W_i)} \frac{2A_i - 1}{g_n(A_i \mid W_i)} (Y_i - Q^*_{Y,\beta_n}(A_i, W_i)) - P_{Q_0,g_0^*} \frac{2A-1}{g_0^*} (Y - Q^*_{Y,\beta_n}) \right) \\
&= \frac{1}{n} \sum_{i=1}^{n} \left(\frac{2A_i - 1}{g_i(A_i \mid W_i)} (Y_i - Q^*_{Y,\beta_n}(A_i, W_i)) - P_{Q_0,g_i} \frac{2A-1}{g_i(A \mid W)} (Y - Q^*_{Y,\beta_n}) \right) \\
&= (P_n - P_{Q_0,\mathbf{g}_n}) d^*_{Y,\beta_n} \\
&= (P_n - P_{Q_0,\mathbf{g}_n}) d^*_{Y,\beta_0} + (P_n - P_{Q_0,\mathbf{g}_n}) \left(d^*_{Y,\beta_n} - d^*_{Y,\beta_0}\right)
\end{aligned}$$

Adding (a) and (b) yields the desired expression in (1.10). □

Proof of proposition 1.4.

In light of lemma 1.2, in order to show (1.13), it suffices to prove

$$(P_n - \mathbf{P}_{Q_0,\mathbf{g}_n}) \left(q^*_{Y,\beta_n} - q^*_{Y,\beta_0}\right) = o_P(1/\sqrt{n}) \text{ and } (P_n - P_{Q_0,\mathbf{g}_n}) \left(d^*_{Y,\beta_n} - d^*_{Y,\beta_0}\right) = o_P(1/\sqrt{n}).$$

We shall do so via lemma 1.8.

Since all functions in $\mathscr{Q}^*_{1,n}$ are uniformly bounded inside the unit interval, the envelope functions F_n of $\mathscr{Q}^*_{1,n}$ will satisfy the Lindeberg condition. We first wish to show that $J(\delta_n, \mathscr{Q}^*_{1,n}, \|\cdot\|_{2,P_{Q_0,g^r}}) \to 0$ for every $\delta_n \downarrow 0$. Define

$$f(Q_{Y,\beta}, g, \varepsilon) \equiv \operatorname{expit}\left(\operatorname{logit}(Q_{Y,\beta}) + \varepsilon H(g)\right);$$

so that $\mathscr{Q}^*_{1,n} = f\left(\mathscr{Q}_{1,n}, \mathscr{G}_{1,n}, \mathscr{E}\right)$. Given $\alpha > 0$, **A7** and **A6** imply that $N(\alpha, \operatorname{logit}(\mathscr{Q}_{1,n}), \|\cdot\|_{2,P_{Q_0,g^r}})$ and $N(\alpha, 1/\mathscr{G}_{1,n}, \|\cdot\|_{2,P_{Q_0,g^r}})$ are finite. Since $\mathscr{E} \subset \mathbb{R}$ is bounded, $N(\alpha, \mathscr{E}, \|\cdot\|_{2,P_{Q_0,g^r}})$ is also finite. Given $f(Q_{Y,\beta}, g, \varepsilon) \in \mathscr{Q}^*_{1,n}$, let $[l_Q, u_Q]$, $[l_g, u_g]$ and $[l_\varepsilon, u_\varepsilon]$ be the α-brackets for $\operatorname{logit}(Q_{Y,\beta})$, $1/g$ and ε, respectively. Then the bracket

$$\left[\operatorname{expit}\left(l_Q + l_\varepsilon H(l_g)\right), \operatorname{expit}\left(u_Q + u_\varepsilon H(u_g)\right)\right]$$

is a bracket containing $f(Q_{Y,\beta},g,\varepsilon)$, with length upper-bounded, up to a universal constant, by 3α. Therefore,

$$N(\alpha,\mathscr{Q}^*_{1,n},\|\cdot\|_{2,P_{Q_0,g^r}})\lesssim N(k_1\alpha,\mathscr{Q}_{1,n},\|\cdot\|_{2,P_{Q_0,g^r}})N(k_2\alpha,1/\mathscr{G}_{1,n},\|\cdot\|_{2,P_{Q_0,g^r}})N(k_3\alpha,\mathscr{E},\|\cdot\|_{2,P_{Q_0,g^r}}).$$

From the inequality $\sqrt{a+b}\leq\sqrt{a}+\sqrt{b}$, we obtain

$$J(\delta_n,\mathscr{Q}^*_{1,n},\|\cdot\|_{2,P_{Q_0,g^r}})\lesssim J(\delta_n,\mathscr{Q}_{1,n},\|\cdot\|_{2,P_{Q_0,g^r}})+J(\delta_n,1/\mathscr{G}_{1,n},\|\cdot\|_{2,P_{Q_0,g^r}})+J(\delta_n,\mathscr{E},\|\cdot\|_{2,P_{Q_0,g^r}}).$$

Assumptions **A9**, **A10** and boundedness of $\mathscr{E}$ respectively imply that each of the terms converge to 0 for every $\delta_n\downarrow 0$. We conclude that $J(\delta_n,\mathscr{Q}^*_{1,n},\|\cdot\|_{2,P_{Q_0,g^r}})\to 0$ for every $\delta_n\downarrow 0$.

Now, define $q(Q_Y)(W)\equiv Q_Y(1,W)-Q_Y(0,W)$. We wish to show

$$(P_n-\mathbf{P}_{Q_0,\mathbf{g}_n})\left(q^*_{Y,\beta_n}-q^*_{Y,\beta_0}\right)\equiv(P_n-\mathbf{P}_{Q_0,\mathbf{g}_n})\left(q(Q^*_{Y,\beta_n})-q(Q^*_{Y,\beta_0})\right)=o_P(1/\sqrt{n}).$$

Indeed, the functions of $q\left(\mathscr{Q}^*_{1,n}\right)$ are bounded within $(-1,1)$, therefore the corresponding sequence of envelope functions satisfy the Lindeberg condition. For a given α-bracket $[l,u]$ of $\mathscr{Q}^*_{1,n}$, we obtain a $m\alpha$-bracket $[l(1,W)-u(0,W),u(1,W)-l(0,W)]$ of $q(\mathscr{Q}^*_{1,n})$, for some $m>3/\inf_{O\in\mathscr{O}}g^r(A\mid W)$. Therefore, the previous conclusion implies that for every $\delta_n\downarrow 0$, we have $J(\delta_n,q(\mathscr{Q}^*_{1,n}),\|\cdot\|_{2,P_{Q_0,g^r}}))\to 0$. A similar argument shows that $P_{Q_0,g^r}\left(q(Q^*_{Y,\beta_n})-q(Q^*_{Y,\beta_0})\right)^2=o_P(1)$ because $\|Q^*_{Y,\beta_n}-Q^*_{Y,\beta_0}\|_{2,P_{Q_0,g^r}}=o_P(1)$. These two observations allow us to apply lemma 1.8 with $\eta=(Q_{Y,\beta},\varepsilon,g)$, $\eta_n=(Q_{Y,\beta_n},\varepsilon_n,g_n)$, $\eta_0=(Q_{Y,\beta_0},\varepsilon_0,g^*_0)$ and

$$f_{\theta,\eta}(O)\equiv\theta(\eta)(O)\equiv q\left(\operatorname{expit}\left(\operatorname{logit}(Q_{Y,\beta})+\varepsilon H(g)\right)\right).$$

We thus conclude that $\left|\sqrt{n}\left(P_n-\mathbf{P}_{Q_0,\mathbf{g}_n}\right)\left(q^*_{Y,\beta_n}-q^*_{Y,\beta_0}\right)\right|=o_P(1)$.

Next, define $d_Y(Q_Y)(O,Z)=\frac{2A-1}{Z}\left(Y-Q_Y(A,W)\right)$. Then

$$\sqrt{n}(P_n-P_{Q_0,\mathbf{g}_n})\left(d^*_{Y,\beta_n}-d^*_{Y,\beta_0}\right)=\sqrt{n}(P_n-P_{Q_0,\mathbf{g}_n})\left(d_Y(Q^*_{Y,\beta_n})-d_Y(Q^*_{Y,\beta_0})\right).$$

From uniform boundedness of Z, Y and of $\mathscr{Q}^*_{1,n}$ for all n, the envelope functions of $d_Y(\mathscr{Q}^*_{1,n})$ satisfy the Lindeberg condition. The same uniform boundedness of Z also implies that the bracketing number of $\mathscr{Q}^*_{1,n}$ and $d_Y(\mathscr{Q}^*_{1,n})$ differ by a constant. Moreover, the convergence $\|Q^*_{Y,\beta_n}-Q^*_{Y,\beta_0}\|_{2,P_{Q_0,g^r}}=o_P(1)$, established in proposition 1.3, implies that $P_{Q_0,g^r}\left(d_Y(Q^*_{Y,\beta_n})-d_Y(Q^*_{Y,\beta_0})\right)^2\to 0$ in probability. Applying lemma 1.8 with $\theta(\eta)(O,Z)\equiv d_Y(\operatorname{expit}(\operatorname{logit}(Q_Y)+\varepsilon H(g))$ we obtain $\sqrt{n}\left(P_n-P_{Q_0,\mathbf{g}_n}\right)\left(d^*_{Y,\beta_n}-d^*_{Y,\beta_0}\right)=o_P(1)$. This proves (1.13).

To prove the CLT, we wish to apply theorem 9 of Chambaz and van der Laan (2011b). For convenience, we simplify our notation and use c_0 to denote $d^*_{Y,\beta_0} + D_W(Q^*_{Y,\beta_0})$. Firstly,

$$\mathbf{P}_{Q_0,\mathbf{g}_n} c_0^2 \equiv \frac{1}{n}\sum_{i=1}^{n} P_{Q_0,g_i} c_0^2 \tag{1.18}$$

$$= \frac{1}{n}\sum_{i=1}^{n}\left\{ P_{Q_0,g_0^*} \frac{(Y - Q^*_{Y,\beta_0})^2}{g_0^* g_i} \right\} + P_{Q_0,g_0^*}\left(2D_Y(Q^*_{Y,\beta_0}, g_0^*) D_W(Q^*_{Y,\beta_0}) + D_W(Q^*_{Y,\beta_0})^2 \right) \tag{1.19}$$

$$= P_{Q_0,g_0^*}\left\{ \frac{\left(Y - Q^*_{Y,\beta_0}\right)^2}{g_0^*} \frac{1}{n}\sum_{i=1}^{n}\frac{1}{g_i} \right\} + P_{Q_0,g_0^*}\left(2D_Y(Q^*_{Y,\beta_0}, g_0^*) D_W(Q^*_{Y,\beta_0}) + D_W(Q^*_{Y,\beta_0})^2 \right) \tag{1.20}$$

Let us now show that the first term in the last expression converges in probability to $P_{Q_0,g_0^*} D_Y\left(Q^*_{Y,\beta_0}, g_0^*\right)^2$; consequently, $P_{Q_0,\mathbf{g}_n} c_0^2$ converges to $P_{Q_0,g_0^*} D^*\left(P_{Q^*_{\beta_0},g_0^*}\right)^2 \equiv \Sigma_0$ in probability. Indeed,

$$\begin{aligned}
& E\left| P_{Q_0,g_0^*}\left\{ \frac{\left(Y - Q^*_{Y,\beta_0}\right)^2}{g_0^*}\left(\frac{1}{n}\sum_{i=1}^{n}\frac{1}{g_i} - \frac{1}{g_0^*} \right) \right\} \right| \\
& \equiv E\left| E_{Q_{W,0}}\left\{ \sum_{a=0,1} f(a,W)\left(\frac{1}{n}\sum_{i=1}^{n}\frac{1}{g_i(a \mid W)} - \frac{1}{g_0^*(a \mid W)} \right) \right\} \right| \\
& \lesssim E\left\{ E_{Q_{W,0}}\left| \frac{1}{n}\sum_{i=1}^{n}\frac{1}{g_i(1 \mid W)} - \frac{1}{g_0^*(1 \mid W)} \right| \right\} + E\left\{ E_{Q_{W,0}}\left| \frac{1}{n}\sum_{i=1}^{n} g_i(1 \mid W) - g_0^*(1 \mid W) \right| \right\} \\
& \leq E\left\| \frac{1}{n}\sum_{i=1}^{n}\frac{1}{g_i(1 \mid W)} - \frac{1}{g_0^*(1 \mid W)} \right\|_{2,Q_{W,0}} + E\left\| \frac{1}{n}\sum_{i=1}^{n} g_i(1 \mid W) - g_0^*(1 \mid W) \right\|_{2,Q_{W,0}},
\end{aligned}$$

where $f(a,W) \equiv E_{P_{Q_0,g_0^*}}\left((Y - Q^*_{Y,\beta_0})^2 \mid a, W \right)$. The first inequality (up to a universal constant), is due to boundedness of $\mathscr{G}_1$, Y and Q^*_{Y,β_0}, and the second inequality is result of Cauchy-Schwartz inequality on the integrand. The right-hand-side of the final inequality converges to 0 by corollary 1.1. Consequently, $E\left\{ \mathbf{P}_{Q_0,\mathbf{g}_n} c_0^2 \right\}$ converges in probability to Σ_0. A similar argument also shows that $\mathbf{P}_{Q_0,\mathbf{g}_n} c_0^2 - E\mathbf{P}_{Q_0,\mathbf{g}_n} c_0^2$ converges in probability to 0. From assumption **A11**, we know that Σ_0 is strictly greater than 0. Therefore, applying theorem 9 of Chambaz and van der Laan (2011b), $\sqrt{n}\left(P_n - \mathbf{P}_{Q_0,\mathbf{g}_n}\right) c_0$ converges to a normal distribution with variance Σ_0. Moreover, $(P_n - \mathbf{P}_{Q_0,\mathbf{g}_n}) c_0^2$ converge to 0 in probability by strong law of large number for martingales. □

Proof of proposition 1.5.

By definition, the sets $B_{M,n}$ are nested, i.e. $B_{M,n} \subset B_{M,n+1}$. Therefore, the sequence $\left\{P_{Q_0,g^r}L(Q_{Y,\beta_{n,0}})\right\}_n$ is non-increasing and lower bounded by $P_{Q_0,g^r}L(Q_{Y,\beta_0})$. A straightforward set-theoretic argument shows that $P_{Q_0,g^r}L(Q_{Y,\beta_0}) = lim_{n\to\infty}P_{Q_0,g^r}L(Q_{Y,\beta_{n,0}})$. To apply proposition 1.1, it remains to show that $L(\mathscr{Q}_{1,n})$ has bracketing integral of order $o(\sqrt{n})$. Indeed, at each n, $B_{M,n} \subset \mathbb{R}^{d_n}$ and the function $L(Q_{Y,\beta})$ is smooth in β. Therefore, the uniform boundedness of the basis functions and of $\|\beta\|_1$ imply that $\left|L(Q_{Y,\beta}) - L(Q_{Y,\beta'})\right| \lesssim f(O)\|\beta - \beta'\|$ for some function f. We may apply the geometric argument in example 19.7 of van der Vaart (1998b) to conclude that, for n sufficiently large, the number of δ-brackets needed equals, up to a universal constant, to the number of δ-balls to cover the hypercube $\|\beta\|_1 \leq M$ in $\mathbb{R}^{d_n}$. The length between two neighboring vertex on this cube is $L = \sqrt{d_n}M$. This big hypercube can be covered by $\left\lceil L\sqrt{d_n}/(2\delta)\right\rceil^{d_n}$ many small hypercubes whose sides have lengths $2\delta/\sqrt{d_n}$. Each of these cubes can be circumscribed in a ball of diameter $\sqrt{d_n(2\delta/\sqrt{d_n})^2} = 2\delta$. Therefore, we conclude that $N(\delta, L(\mathscr{Q}_{1,n}), \|\cdot\|_{2,P_{Q_0,g^r}}) \lesssim \lceil d_nM/(2\delta)\rceil^{d_n}$. Therefore, $J(1, L(\mathscr{Q}_{1,n}), \|\cdot\|_{2,P_{Q_0,g^r}}) \lesssim O\left(\sqrt{d_n\log(d_n)}\right) = o(\sqrt{n})$ if $d_n = O(n^r)$ for $0 < r < 1$. This completes the proof. □

Proof of proposition 1.6.

To apply proposition 1.2, it suffices to note that by the reverse triangular inequality and Lipschitz continuity of the absolute value function, the same argument as in the previous proof can be applied to conclude that **B4** implies $J(1, h_1(\mathscr{Q}_{1,n}), \|\cdot\|_{2,P_{Q_0,g^r}}) = o(\sqrt{n})$. □

Proof of proposition 1.8.

To apply proposition 1.4, it suffice to show that assumptions **B4** imply $J(\delta_n, \mathscr{Q}_{1,n}, \|\cdot\|_{2,P_{Q_0,g^r}}) \to 0$ for every $\delta_n \downarrow 0$. Indeed,

$$\begin{aligned} J(\delta_n, \mathscr{Q}_{1,n}, \|\cdot\|_{2,P_{Q_0,g^r}}) &\lesssim \int_0^{\delta_n} \sqrt{\log\lceil d_nM/(2\alpha)\rceil^{d_n}}\, d\alpha \leq \int_0^{\delta_n} d_n \log\left(d_nM/\alpha\right) d\alpha \\ &= d_n\delta_n\log(d_nM) - d_n\delta_n\log(\delta_n) + d_n\delta_n = O\left(n^r\delta_n\left(\log(\delta_n) + \log(n)\right)\right), \end{aligned}$$

where the first inequality is explained in the proof of proposition 1.5 and the last equality follow from assumption **B4** that $d_n = O(n^r)$. Therefore, we conclude that $J(\delta_n, \mathscr{Q}^*_{1,n}, \|\cdot\|_{2,P_{Q_0,g^r}}) \to 0$ for every $\delta_n \to 0$ satisfying $n^r\log(n)\delta_n \to 0$. This is complete the proof since such δ_n can be used in obtaining the result in lemma 1.8. □

A.2 Useful lemmas

From here onward, the uncountable supremum will be interpreted as the essential supremum. We will use $\mathbf{1}\{A\}$ to denote the indicator function of the set A.

The following lemma from Van Der Vaart and Wellner (1996) ensures convergence of M-estimators, and consequently, also of Z-estimators.

Lemma 1.3 (**Convergence of M-estimators**, Van Der Vaart and Wellner (1996))**.**
Let $\mathbf{M}_n$ stochastic processes indexed by a metric space $\mathscr{V}$, and let $\mathbf{M} : \mathscr{V} \to \mathbb{R}$ be a deterministic function. Consider a sequence of subsets $\mathscr{V}_n \subset \mathscr{V}$. Suppose the following assumptions hold:

1. *There exists a point $v_0 \in \mathscr{V}$ such that $\mathbf{M}(v_0) < \inf_{v \notin T} \mathbf{M}(v)$ for every open set $T \subset \mathscr{V}$ containing v_0.*
2. *For every n, there exists $v_n^* \in \mathscr{V}_n$ such that $\mathbf{M}(v_n^*) = \inf_{\mathscr{V}_n} \mathbf{M}(v)$. Moreover, $\mathbf{M}(v_n^*) - \mathbf{M}(v_0) = o(1)$.*
3. *$\sup_{v \in \mathscr{V}_n} |\mathbf{M}_n(v) - \mathbf{M}(v)| = o_P(1)$*

If the sequence $v_n \in \mathscr{V}_n$ satisfies $\mathbf{M}_n(v_n) - \mathbf{M}_n(v_n^*) \leq 0$, then v_n converges in probability to v_0.

Proof. Firstly, assumptions 1, 2 and 3 imply that

$$\begin{aligned} 0 &\leq \mathbf{M}(v_n) - \mathbf{M}(v_0) \\ &= (\mathbf{M}(v_n) - \mathbf{M}_n(v_n)) + (\mathbf{M}_n(v_n) - \mathbf{M}_n(v_n^*)) + (\mathbf{M}_n(v_n^*) - \mathbf{M}(v_n^*)) + (\mathbf{M}(v_n^*) - \mathbf{M}(v_0)) \\ &\leq \sup_{\mathscr{V}_n} |\mathbf{M}_n(v) - \mathbf{M}(v)| + (\mathbf{M}(v_n^*) - \mathbf{M}(v_0)) \\ &= o_P(1). \end{aligned}$$

Now, let $d(\cdot,\cdot)$ denote the metric on $\mathscr{V}$. From assumption 1, we have: for each $\eta > 0$, there is a $\delta > 0$ such that $P(d(v_n, v_0) \geq \eta) \leq P(\mathbf{M}(v_n) - \mathbf{M}(v_0) \geq \delta)$. From the observation above, we may thus conclude that $P(d(v_n, v_0) \geq \eta) \to 0$ as $n \to \infty$. □

Much of the results in this chapter concerning uniform laws of large numbers are derived from the maximal inequalities in lemmas 1.4 and 1.5, which are due to van Handel (2011). To draw on those results, we make the following definitions. Let $\phi(x) = e^x - x - 1$. Given a class of functions $\mathscr{F}$, $n \geq 1$, $K > 0$, $\delta > 0$, let $N = N(\delta, \mathscr{F}, \|\cdot\|_{2, P_{Q_0, g^r}})$, and define a $(n, \mathscr{F}, K, \delta)$-bracketing set as a collection $\left\{ \left(\Lambda_i^j, \Gamma_i^j \right) i \leq n \right\}_{j \leq N}$ such that for each $f \in \mathscr{F}$, there exits $j \leq N$ satisfying $\Lambda_i^j \leq f(O_i, Z_i) \leq \Gamma_i^j$, for all $i \leq n$, and such that for all $j \leq N$, $\frac{2K^2}{n} \sum_{i=1}^n E\left\{ \phi\left(\frac{\left|\Lambda_i^j - \Gamma_i^j\right|}{K} \right) | \mathbf{O}_{i-1} \right\} \leq \delta^2$. Let $\mathscr{N}(n, \mathscr{F}, K, \delta)$ denote the cardinality

of the smallest $(n, \mathscr{F}, K, \delta)$-bracketing set. We also define

$$R_{n,K}(f) = \frac{2K^2}{n} \sum_{i=1}^{n} P_{Q_0,g_i} \phi\left(\frac{|f|}{K}\right),$$

for all $f \in \mathscr{F}$ and n.

Lemma 1.4 (**Proposition. A.2 in van Handel (2011)**).

For $i \leq n$, let P_i denote the empirical distribution of the first i observations. Fix $K > 0$. There exists an universal constant $C > 0$ such that for all $n \geq 1$, $R > 0$,

$$\mathbf{P}\left(\sup_{f\in\mathscr{F}} \mathbf{1}\{R_{n,K}(f) \leq R\} \max_{i\leq n} \frac{i}{n} (P_i - \mathbf{P}_{Q_0,\mathbf{g}_i}) f \geq \alpha\right) \leq 2\exp\left(-\frac{n\alpha^2}{C^2(c_1+1)R}\right),$$

for any $\alpha, c_0, c_1 > 0$ such that $c_0^2 \geq C^2(c_1+1)$ and

$$\frac{c_0}{\sqrt{n}} \int_0^{\sqrt{R}} \sqrt{\log \mathscr{N}(n, \mathscr{F}, K, u)} du \leq \alpha \leq \frac{c_1 R}{K}$$

Lemma 1.5 (**Corollary A.8 in van Handel (2011)**).

Consider the same setup as in lemma 1.4. Suppose the class $\mathscr{F}$ is finite.
Fix $K > 0$. For every $R > 0$, and every event C

$$E\left[\max_{f\in\mathscr{F}} \mathbf{1}\{nR_{n,K}(f) \leq R\} \max_{k\leq n} k(P_k - \mathbf{P}_{Q_0,\mathbf{g}_k}) f\right] \leq \sqrt{2R\log\left(1 + \frac{|\mathscr{F}|}{P(C)}\right)} + 8K\log\left(1 + \frac{|\mathscr{F}|}{P(C)}\right).$$

If in addition $\sup_{f\in\mathscr{F}} \|f\|_\infty \leq 3U$, then

$$E\left[\max_{f\in\mathscr{F}} \mathbf{1}\{nR_{n,K}(f) \leq R\} \max_{k\leq n} k(P_k - \mathbf{P}_{Q_0,\mathbf{g}_k}) f\right] \leq \sqrt{2R\log\left(1 + \frac{|\mathscr{F}|}{P(C)}\right)} + 8U\log\left(1 + \frac{|\mathscr{F}|}{P(C)}\right).$$

We summarize a few observations regarding $R_{n,K}$ and $\mathscr{N}(n, \mathscr{F}, K, \delta)$ in lemma 1.6 .

Lemma 1.6 (L^2**-norm version of lemma 7 in Chambaz and van der Laan (2011b)**).

Fix $K > 0$. Suppose $U = \sup_{f\in\mathscr{F}} \|f\|_\infty < \infty$. Then

1. *For each $n \geq 1$, $f \in \mathscr{F}$, $R_{n,4U}(f) \leq \frac{4}{3}\frac{1}{n}\sum_{i=1}^{n} P_{Q_0,g_i} |f|^2$.*
2. *Let $\kappa > 0$, and $g^s \in \mathscr{G}$ be such that $\left\|\frac{g}{g^s}\right\|_\infty \leq \kappa$ for all data-generating treatment assignments g. Then, for each $n \geq 1$, $\delta > 0$, $\log \mathscr{N}(n, \mathscr{F}, 4U, \sqrt{2\kappa}\delta) \leq \log N(\delta, \mathscr{F}, \|\cdot\|_{2,P_{Q_0,g^s}})$.*

Proof. For any $m \geq 2$, $P_{Q_0,g_i}|f|^m \leq U^{m-2}P_{Q_0,g_i}|f|^2 \leq \frac{m!}{2}U^{m-2}P_{Q_0,g_i}|f|^2$. Therefore, for $K = 4U$,

$$2K^2 P_{Q_0,g_i}\phi\left(\frac{|f|^2}{4U}\right) = 2(4U)^2\sum_{m\geq 2}\frac{P_{Q_0,g_i}|f|^m}{m!(4U)^m} \leq 2(4U)^2\sum_{m\geq 2}\frac{\frac{m!}{2}U^{m-2}P_{Q_0,g_i}|f|^2}{m!(4U)^m}$$
$$= 16\sum_{m\geq 2}\frac{P_{Q_0,g_i}|f|^2}{(4)^m} = \frac{4}{3}P_{Q_0,g_i}|f|^2.$$

This proves the first result.

For the second result, fix $\delta > 0$ and let $N = N(\mathscr{F}, \|\cdot\|_{2,P_{Q_0,g^s}}, \delta)$. Suppose we have $(\ell^j, u^j)_{j\leq N}$ a set of δ-brackets, under $\|\cdot\|_{2,P_{Q_0,g^s}}$, covering $\mathscr{F}$. Let $\Lambda_i^j = min(\ell^j(O_i,Z_i), -U)$ and $\Gamma_i^j = max(u^j(O_i,Z_i), U)$. For each $f \in \mathscr{F}$, if it's covered by $[\ell^j, u^j]$, then, for all $i \leq n$, $\Lambda_i^j \leq f(O_i,Z_i) \leq \Gamma_i^j$. We also know that $-U \leq \Lambda_i^j \leq \Gamma_i^j \leq U$ and $\ell^j \leq \Lambda_i^j \leq \Gamma_i^j \leq u^j$. In particular, at a fixed $j \leq N$, for each $i \leq n$ and for all $m \geq 2$,

$$E\left(\left|\Lambda_i^j - \Gamma_i^j\right|^m |\mathbf{O}_{i-1}\right) = P_{Q_0,g_i}\left(\left|\Lambda_i^j - \Gamma_i^j\right|^m\right)$$
$$\leq (2U)^{m-2}P_{Q_0,g_i}\left|\Lambda_i^j - \Gamma_i^j\right|^2 \leq (2U)^{m-2}\kappa P_{Q_0,g^s}\left|\Lambda_i^j - \Gamma_i^j\right|^2 \leq (2U)^{m-2}\kappa P_{Q_0,g^s}\left|\ell^j - u^j\right|^2$$
$$\leq (2U)^{m-2}\kappa\delta^2 \leq \frac{m!}{2}(2U)^{m-2}\kappa\delta^2.$$

Therefore, with the same K as before, we have

$$2K^2 P_{Q_0,g_i}\phi\left(\frac{\left|\Lambda_i^j - \Gamma_i^j\right|}{4U}\right) = 2(4U)^2\sum_{m\geq 2}\frac{P_{Q_0,g_i}\left|\Lambda_i^j - \Gamma_i^j\right|^m}{m!(4U)^m} \leq 32U^2\sum_{m\geq 2}\frac{\frac{m!}{2}(2U)^{m-2}\kappa\delta^2}{m!(4U)^m} = 2\kappa\delta^2.$$

We have thus shown that $\mathscr{N}(n, \mathscr{F}, 4U, \sqrt{2\kappa}\delta) \leq N(\mathscr{F}, \|\cdot\|_{2,P_{Q_0,g^s}}, \delta)$. □

Using lemmas 1.4, one can obtain exponential inequalities needed to establish a uniform law of large numbers. Lemma 1.7 below modifies theorem 8 in Chambaz and van der Laan (2011b) to use an L^2-metric and allow the classes of functions to change with n. In said chapter, one requires that the bracketing integral under sup norm be finite. Here we weaken the condition by controlling the bracketing integral under the L^2-norm. This is possible because the dominated ratio property of $\mathscr{G}_1$ allows one to bound $\mathscr{N}(n, \mathscr{F}, 4U, \sqrt{2\kappa}\delta)$ with $N(\mathscr{F}, \|\cdot\|_{2,P_{Q_0,g^s}}, \delta)$ (See lemma 1.6). Since the classes may grow with n, we control their complexity by controlling the growth of the entropies at speed $o(\sqrt{n})$. If the class is fixed, i.e. $\mathscr{F}_n = \mathscr{F}$ for all n, then the entropy is required to be finite.

Lemma 1.7 (Sieved and L^2-norm version of theorem 8 in Chambaz and van der Laan (2011b)).

Let $U = \sup_n \sup_{f \in \mathscr{F}_n} \|f\|_\infty < \infty$.

If the sequence of entropies satisfies $J\left(U\sqrt{2/(3\kappa)}, \mathscr{F}_n, \|\cdot\|_{2,P_{Q_0,g^r}}\right) = o(\sqrt{n})$ *, then for each* $\alpha > 0$*, there is a constant c and* N_α *such that for all* $n \geq N_\alpha$*,*

$$P\left(\sup_{f \in \mathscr{F}_n} (P_n - \mathbf{P}_{Q_0,\mathbf{g}_n}) f \geq \alpha\right) \leq 2e^{-nc}$$

Consequently, $\sup_{f \in \mathscr{F}_n} \left|\left(P_n - \mathbf{P}_{Q_0,\mathbf{g}_n}\right) f\right|$ *converges to* 0 *almost surely.*

Proof. Given $\alpha > 0$, let $K = 4U$, $R = \frac{4}{3}U^2$, $c_1 = \frac{\alpha K}{R}$, $c_0 = C\sqrt{c_1 + 1}$. Recall that from lemma 1.6, $R_{n,K}(f) \leq R$ for all f. Choose N_α satisfying

$$J\left(\sqrt{R/2\kappa}, \mathscr{F}_n, \|\cdot\|_{2,P_{Q_0,g^r}}\right) \leq \sqrt{n}\frac{\alpha}{c_0\sqrt{2\kappa}} \text{ for all } n > N_\alpha.$$

This is possible due to our assumption on the growth rate of the entropies. Allowing the class $\mathscr{F}$ to change with n in lemma 1.6 does not affect the needed inequality

$$\mathscr{N}(n, \mathscr{F}, 4U, \sqrt{2\kappa}\delta) \leq N(\mathscr{F}, \|\cdot\|_{2,P_{Q_0,g^s}}, \delta).$$

Therefore, after a change of variable, it follows that for all $n > N_\alpha$,

$$\sqrt{n} \geq \sqrt{2\kappa}\frac{c_0}{\alpha}\int_0^{\sqrt{R/2\kappa}} \sqrt{\log N(x, \mathscr{F}, \|\cdot\|_{2,P_{Q_0,g^r}})}dx \geq \frac{c_0}{\alpha}\int_0^{\sqrt{R}} \sqrt{\log \mathscr{N}(n, \mathscr{F}, 4U, x)}dx.$$

Similarly, the proof of lemma 1.4 in van Handel (2011) remains valid if we allow $\mathscr{F}$ to depend on n. Consequently, we obtain

$$P\left(\sup_{f \in \mathscr{F}} (P_n - \mathbf{P}_{Q_0,\mathbf{g}_n}) f \geq \alpha\right) \leq P\left(\sup_{f \in \mathscr{F}} \max_{i \leq n} \frac{i}{n} (P_i - \mathbf{P}_{Q_0,\mathbf{g}_i}) f \geq \alpha\right) \leq 2\exp\left(-n\frac{\alpha^2}{c_0^2 R}\right).$$

□

To obtain the central limit theorems, we require convergences $\sqrt{n}\left(P_n - \mathbf{P}_{Q_0,\mathbf{g}_n}\right) f_n = o_P(1)$, for some random functions $f_n \in \mathscr{F}_n$. Lemma 1.8 specifies sufficient conditions on $\mathscr{F}_n$ for this convergence. The result of this lemma relies on the bracketing entropy bound for the first moment of a supremum of the empirical process. Such a bound was derived in lemma 19.34 of van der Vaart (1998b) for the i.i.d. setting, and we refashion it here in lemma 1.9 to suit the current data generating mechanism. For the next two lemmas, let $\text{Log}(x) \equiv max(1, \log(x))$.

Lemma 1.8 (Convergence empirical process, (van der Vaart and Wellner (2007)).
Let $\mathscr{F}_n = \left\{(O,Z) \mapsto f_{\theta,\eta}(O,Z) : \theta \in \Theta, \eta \in T_n\right\}$ be a class of measurable functions with envelope functions F_n. Suppose the following holds:

1. *The sequence of envelope functions F_n satisfy the Lindeberg condition in (1.11).*
2. *The sequence of entropies satisfy $J\left(\delta_n, \mathscr{F}_n, \|\cdot\|_{2,P_{Q_0,g^r}}\right)$ for every sequence $\delta_n \downarrow 0$.*

Then, for a sequence $\eta_n \in T_n$ satisfying $\sup_{\theta\in\Theta} P_{Q_0,g^r}\left(f_{\theta,\eta_n} - f_{\theta,\eta_0}\right)^2 \to 0$ in probability for some η_0, we have $\sup_{\theta\in\Theta}\left|\sqrt{n}\left(P_n - \mathbf{P}_{Q_0,\mathbf{g}_n}\right)\left(f_{\theta,\eta_n} - f_{\theta,\eta_0}\right)\right| \to 0$ in probability.

Proof. Define $\mathscr{F}'_n \equiv \left\{f_{\theta,\eta_n} - f_{\theta,\eta_0} : \theta \in \Theta\right\}$. For a given $\delta > 0$, the choice of η_n implies that for n sufficiently large, $P_{Q_0,g^r}\left(f_{\theta,\eta_n} - f_{\theta,\eta_0}\right)^2 < \delta^2$ for all $\theta \in \Theta$. We may apply lemma 1.9 to each class $\mathscr{F}'_n$, since n is fixed in the proof of said lemma. Consequently, for $a_n(\delta) \equiv \delta / \sqrt{\mathrm{Log}\left(N(\delta, \mathscr{F}'_n, \|\cdot\|_{2,P_{Q_0,g^r}})\right)}$,

$$\begin{aligned}
&E\left(\sup_{\theta\in\Theta}\left|\sqrt{n}\left(P_n - \mathbf{P}_{Q_0,\mathbf{g}_n}\right)\left(f_{\theta,\eta_n} - f_{\theta,\eta_0}\right)\right|\right) \\
&\lesssim J(\delta, \mathscr{F}'_n, \|\cdot\|_{2,P_{Q_0,g^r}}) + \sqrt{n}\kappa P_{Q_0,g^r} F_n \mathbf{1}\left\{F_n > \sqrt{n}a_n(\delta)\right\} \\
&\lesssim J(\delta, \mathscr{F}_n, \|\cdot\|_{2,P_{Q_0,g^r}}) + \frac{P_{Q_0,g^r} F_n^2 \mathbf{1}\left\{F_n > \sqrt{n}a_n(\delta)\right\}}{a_n(\delta)},
\end{aligned}$$

where the second equality follows the fact that the bracketing number of $\mathscr{F}'_n$ is less than the bracketing number of $\mathscr{F}_n$ for a fixed δ, and the fact that

$$\sqrt{n}a(\delta) P_{Q_0,g^r} F_n \mathbf{1}\left\{F_n > \sqrt{n}a_n(\delta)\right\} < P_{Q_0,g^r} F_n^2 \mathbf{1}\left\{F_n > \sqrt{n}a_n(\delta)\right\}.$$

Assumption 2 implies that $J(\delta, \mathscr{F}'_n, \|\cdot\|_{2,P_{Q_0,g^r}}) = O(1)$ for every $\delta > 0$, therefore, $\delta \mapsto a_n(\delta)$ is bounded away from 0. Consequently, the Lindeberg condition implies that the second term in the right hand side of the last inequality above converges to 0 as $n \to \infty$ for every fixed δ. Assumption 2 also implies that the first term can be arbitrarily small as $n \to \infty$ by choosing a small enough δ.

To complete the proof, suppose one is given η and α. We can pick δ^* such that $J(\delta^*, \mathscr{F}_n, \|\cdot\|_{2,P_{Q_0,g^r}} < \eta\alpha/2$ for all n greater than some fixed N_1. We then apply lemma 1.9 with this given δ^*. Since $P\left(\mathbf{1}\left\{F_n > \sqrt{n}a_n(\delta^*)\right\}\right)$ to 0 as $n \to \infty$, we can pick N_2 such that

$$a_n(\delta^*)^{-1} P_{Q_0,g^r} F_n^2 \mathbf{1}\left\{F_n > \sqrt{n}a(\delta^*)\right\} < \eta\alpha/2$$

for all $n > N_2$. It follows from the Markov inequality and the above conclusion that

$$P\left(\sup_{\theta\in\Theta}\left|\sqrt{n}\left(P_n-\mathbf{P}_{Q_0,\mathbf{g}_n}\right)\left(f_{\theta,\eta_n}-f_{\theta,\eta_0}\right)\right|\geq\eta\right)$$
$$\leq\eta^{-1}E\left(\sup_{\theta\in\Theta}\left|\sqrt{n}\left(P_n-\mathbf{P}_{Q_0,\mathbf{g}_n}\right)\left(f_{\theta,\eta_n}-f_{\theta,\eta_0}\right)\right|\right)\lesssim\alpha$$

for all $n > max(N_1,N_2)$. This completes the proof. □

Lemma 19.34 in van der Vaart (1998b) provides a maximal inequality for the first moment of a supremum of the empirical processes of the form $\sqrt{n}(P_n - P_0)f$ under an i.i.d. setting. The following lemma 1.9 generalizes this maximal inequality for empirical processes of the form $\sqrt{n}(P_n - \mathbf{P}_{Q_0,\mathbf{g}_n})f)$, under the current adaptive RCT sampling.

Lemma 1.9. *(Generalization of lemma 19.34 in van der Vaart (1998b) to martingales)*
Consider a class $\mathscr{F}$ of measurable functions $f : \mathscr{O} \to \mathbb{R}$ with $P_{Q_0,g^r}f^2 < \delta^2$ for every f. Let F be an envelope function of this class. Define $a(\delta) = \delta/\sqrt{\text{Log}N(\delta,\mathscr{F},\|\cdot\|_{2,P_{Q_0,g^r}}}$. The following inequality holds for each $n \geq 1$:

$$E\left\{\sup_{f\in\mathscr{F}}\left|\sqrt{n}\left(P_n-\mathbf{P}_{Q_0,\mathbf{g}_n}\right)f\right|\right\}\lesssim J(\delta,\mathscr{F},\|\cdot\|_{2,P_{Q_0,g^r}})+\sqrt{n}E\left(\frac{1}{n}\sum_{i=1}^{n}P_{Q_0,g_i}F\mathbf{1}\left\{F>\sqrt{n}a(\delta)\right\}\right)$$
$$\leq J(\delta,\mathscr{F},\|\cdot\|_{2,P_{Q_0,g^r}})+\sqrt{n}\kappa P_{Q_0,g^r}F\mathbf{1}\left\{F>\sqrt{n}a(\delta)\right\}$$

Proof. For ease of notation, we will denote the uniform norm of a real-valued operator Π on $\mathscr{F}$ as $\|\Pi\|_{\mathscr{F}} \equiv \sup_{f\in\mathscr{F}}|\Pi(f)|$.

The second inequality in the result follows from the dominated ratio property of $\mathscr{G}_1$ and the fact that $F\mathbf{1}\{F > \sqrt{n}a(\delta)\} \geq 0$. It suffices to show the first inequality. We first divide the class $\mathscr{F}$ by $\sqrt{n}a(\delta)$. On one hand, we have

$$E\left(\left\|\sqrt{n}\left(P_n-\mathbf{P}_{Q_0,\mathbf{g}_n}\right)f\mathbf{1}\left\{F>\sqrt{n}a(\delta)\right\}\right\|_{\mathscr{F}}\right)\leq E\left(\sqrt{n}\left(P_n+\mathbf{P}_{Q_0,\mathbf{g}_n}\right)F\mathbf{1}\left\{F>\sqrt{n}a(\delta)\right\}\right)$$
$$\leq 2\sqrt{n}E\left(\frac{1}{n}\sum_{i=1}^{n}P_{Q_0,g_i}F\mathbf{1}\left\{F>\sqrt{n}a(\delta)\right\}\right),\quad(1.21)$$

where the first inequality follows from the fact that

$$\left|\sqrt{n}\left(P_n-P_{Q_0,\mathbf{g}_n}\right)f\right|\leq\sqrt{n}\left(P_n+P_{Q_0,\mathbf{g}_n}\right)|f|\leq\sqrt{n}\left(P_n+P_{Q_0,\mathbf{g}_n}\right)h,$$

whenever $|f| < h$.

Now, we wish to bound $E\left(\left\|\sqrt{n}\left(P_n-\mathbf{P}_{Q_0,\mathbf{g}_n}\right)\left(f\mathbf{1}\{F\leq\sqrt{n}a(\delta)\}\right)\right\|_{\mathscr{F}}\right)$. The bracketing number for the class $\{f\mathbf{1}\{F \leq \sqrt{n}a(\delta)\} : f \in \mathscr{F}\}$ is smaller than that of $\mathscr{F}$. So,

without loss of generality, we may assume from now on that functions in $\mathscr{F}$ are upper-bounded by $\sqrt{n}a(\delta)$. We proceed using a chaining technique to replace $\mathscr{F}$ with a finite class.

Fix q_0 such that $4\delta \leq 2^{-q_0} \leq 8\delta$. Define a nested sequence of partitions on $\mathscr{F}$, indexed by integers $q \geq q_0$, as follows:

1. For each integer $q \geq q_0$, cover $\mathscr{F}$ with $N_q = N(2^{-q}, \mathscr{F}, \|\cdot\|_{2,P_{Q_0,g^r}})$ many brackets, $[l_{q_i}, u_{q_i}]_{i \leq N_q}$. Define $\mathscr{F}_{q_i} = [l_{q_i}, u_{q_i}] \bigcap \left(\bigcup_{j<i} [l_{q_j}, u_{q_j}]\right)^C$. Then, $\mathscr{F} = \bigcup_{i=1}^{N_q} \mathscr{F}_{q_i}$ is a partition of $\mathscr{F}$ at level q.

2. For each partitioning set $\mathscr{F}_{q_i}$ at level q, define $\Delta_{q_i} = u_{q_i} - l_{q_i}$. Then, the following conditions hold

$$
\begin{aligned}
&(a) \qquad \sum_{q \geq q_0} 2^{-q} \sqrt{\mathrm{Log} N_q} \lesssim \int_0^{\delta} \sqrt{\mathrm{Log} N(\alpha, \mathscr{F}, \|\cdot\|_{2,P_{Q_0,g^r}})} d\alpha; \\
&(b) \qquad \sup_{f,g \in \mathscr{F}_{q_i}} |f-h| \leq \Delta_{q_i} \leq 2F \leq 2\sqrt{n}a(\delta); \\
&(c) \qquad P_{Q_0,g^r} \Delta_{q_i}^2 < 2^{-2q}.
\end{aligned}
$$

3. From this sequence (indexed by q) of partitions, we obtain a nested sequence as follows: At level q that is not a success refinement, we replace each partitioning set in level q by its intersection with all partitioning sets in the previous levels. The new partition has the size at most $\bar{N}_q = N_{q_0} \cdots N_q$. Using the inequality $\sqrt{\log\left(\prod N_p\right)} \leq \sum \sqrt{logN_p}$, the previous condition (a) is still satisfied. From here on, we use N_q to denote $\bar{N}_q$. It should be reminded that it now denotes the number of partitioning sets at level q, not number of brackets.

At each level q, for each $\mathscr{F}_{q_i}$, fix a representative $f_{q_i} \in \mathscr{F}_{q_i}$, and define for any $f \in \mathscr{F}$: if $f \in \mathscr{F}_{q_i}$,

$$
\begin{aligned}
\pi_q f &= f_{q_i} \qquad \text{the representative of its partitioning set} \\
\Delta_q f &= \Delta_{q_i} \quad \text{the envelope for differences in its partitioning set.}
\end{aligned}
$$

We have now given $\mathscr{F}$ a finite representation.

Define for each fixed n, each $q \geq q_0$ and each $f \in \mathscr{F}$:

$$
\begin{aligned}
a_q &= 2^{-q} / \sqrt{\mathrm{Log} N_{q+1}} \\
A_{q-1} f &= \mathbf{1}\left\{\Delta_{q_0} f \leq \sqrt{n} a_{q_0}, \ldots, \Delta_{q-1} f \leq \sqrt{n} a_{q-1}\right\} \\
B_q f &= \mathbf{1}\left\{\Delta_{q_0} f \leq \sqrt{n} a_{q_0}, \ldots, \Delta_{q-1} f \leq \sqrt{n} a_{q-1}, \Delta_q f > \sqrt{n} a_q\right\}
\end{aligned}
$$

By nestedness, $A_q f$ and $B_q f$ are constant in f within each $\mathscr{F}_{q_i}$. Note also that $\delta \leq 2\delta \leq 2^{-(q_0+1)}$ by definition of q_0, $a_{q_0} = 2a(2^{-(q_0+1)})$, and $\delta \mapsto a(\delta)$ is an increasing function. Therefore, $\Delta_{q_0} f \leq 2F \leq 2\sqrt{n}a(\delta) \leq \sqrt{n}a_{q_0}$. Hence $A_{q_0} f = 1$.

Using these ingredients, the function $O \mapsto f(O)$ can be rewritten as:

$$f = \pi_{q_0} f + \sum_{q \geq q_0+1} (f - \pi_q f) B_q f + \sum_{q \geq q_0+1} (\pi_q f - \pi_{q-1} f) A_{q-1} f.$$

Indeed, note that either $B_q f = 0$ for all $q > q_0$ (in which case $A_q f = 1$), or there exists a unique q_1 such that $B_{q_1} f = 1$, and $B_q f = 0$ for all $q \neq q_1$, and $A_q f = 1$ for all $q < q_1$, $A_q f = 0$ for all $q \geq q_1$. In the first case, $f = \pi_{q_0} f + \lim_{q\to\infty} \pi_q f - \pi_{q_0} f$, where $\lim_{q\to\infty} \pi_q f = f$ follows from the fact that both $\pi_q f$ and f are in the bracket $[l_q, u_q]$ and as $q \to \infty$ the size of the bracket $\|u_q - l_q\|_{2,P_{Q_0,g^r}} \to 0$. In the second case, $f = \pi_{q_0} f + (f - \pi_{q_1} f) + \sum_{q_0+1}^{q_1} (\pi_q f - \pi_{q-1} f)$.

Using this representation of f, we obtain the inequality

$$\begin{aligned}
&E \left\| \sqrt{n} (P_n - \mathbf{P}_{Q_0,\mathbf{g}_n}) f \right\|_{\mathscr{F}} \\
&\leq E \left\| \sqrt{n} (P_n - \mathbf{P}_{Q_0,\mathbf{g}_n}) \pi_{q_0} f \right\|_{\mathscr{F}} + E \left\| \sum_{q \geq q_0+1} \sqrt{n} (P_n - \mathbf{P}_{Q_0,\mathbf{g}_n}) (f - \pi_q f) B_q f \right\|_{\mathscr{F}} \\
&+ E \left\| \sum_{q \geq q_0+1} \sqrt{n} (P_n - \mathbf{P}_{Q_0,\mathbf{g}_n}) (\pi_q f - \pi_{q-1} f) A_{q-1} f \right\|_{\mathscr{F}}. \qquad (1.22)
\end{aligned}$$

Our goal is to bound the right-hand-side of the inequality.

To bound the first term on the right-hand side of (1.22), we note that there are N_{q_0} many functions $\pi_{q_0} f$. Applying lemma 1.5 with the functions $\pi_{q_0} f / \sqrt{n}$, $U = 2^{-1} a_{q_0}$ (since $\left|\pi_{q_0} f\right| \leq \sqrt{n}a(\delta) \leq \sqrt{n}a(2^{-(q_0+1)}) = \sqrt{n}2^{-1}a_{q_0}$), $R = \frac{4}{3}\kappa\delta^2$ (by definition of our class $\mathscr{F}$, $P_{Q_0,g^r}(\pi_{q_0} f)^2 \leq \delta^2$), we obtain the inequality

$$\begin{aligned}
&E \left\| \sqrt{n} (P_n - \mathbf{P}_{Q_0,\mathbf{g}_n}) \pi_{q_0} f \right\|_{\mathscr{F}} \lesssim \delta \sqrt{\log(1 + N_{q_0})} + 2^{-1} a_{q_0} \log(1 + N_{q_0}) \\
&\lesssim \delta \sqrt{\mathrm{Log}(N_{q_0})} + a_{q_0} \mathrm{Log}(N_{q_0}) \lesssim \sum_{q \geq q_0+1} 2^{-q} \sqrt{\mathrm{Log}(N_q)} \qquad (1.23)
\end{aligned}$$

The second inequality follows from the fact that $\log(1 + N_q) \leq 2\,\mathrm{Log}(N_q)$: if $N_q \leq e$, then $\log(1 + N_q) \leq \log(1 + e)$; else $\log(1 + N_q) \leq 2\log(N_q)$, for $N_q > e$.

For the second term on the right-hand side of (1.22), note that $|f| < h$ implies that

$$|(P_n - P_{Q_0,\mathbf{g}_n}) f| \leq (P_n + P_{Q_0,\mathbf{g}_n}) h = (P_n - P_{Q_0,\mathbf{g}_n}) h + 2P_{Q_0,\mathbf{g}_n} h.$$

Therefore, $\left|f-\pi_q f\right| \leq \Delta_q f$ implies that

$$E\left\|\sum_{q\geq q_0+1}\sqrt{n}\left(P_n-\mathbf{P}_{Q_0,\mathbf{g}_n}\right)\left(f-\pi_q f\right)B_q f\right\|_{\mathscr{F}} \leq \sum_{q\geq q_0+1} E\left\|\sqrt{n}\left(P_n-\mathbf{P}_{Q_0,\mathbf{g}_n}\right)\Delta_q f B_q f\right\|_{\mathscr{F}}$$
$$+\sum_{q\geq q_0+1} E\left\|2\sqrt{n}\frac{1}{n}\sum_{i=1}^{n} P_{Q_0,g_i}\Delta_q f B_q f\right\|_{\mathscr{F}} \tag{1.24}$$

To bound first term in (1.24), first note that for a fixed level q, there at most N_q many $\Delta_q f B_q f$ functions, each corresponding to the partitioning set $\mathscr{F}_{q_i}$ containing f (recall that the partition is nested, so the partition for f at all previous levels are therefore determined by $\mathscr{F}_{q_i}$). Therefore, $\sup_{f\in\mathscr{F}}\left|\sqrt{n}\left(P_n-\mathbf{P}_{Q_0,\mathbf{g}_n}\right)\Delta_q f B_q f\right|$ can be replaced by taking maximum over a finite set of size N_q. In particular, there are finitely many $\Delta_q f B_q f/\sqrt{n}$. Moreover, $\|\Delta_q f B_q f/\sqrt{n}\|_\infty \leq a_{q-1}$ because $\Delta_q f B_q f \leq \Delta_{q-1} f B_q f \leq \sqrt{n}a_{q-1}$ (first inequality by nestedness and second inequality by definition of $B_q f$), and $P_{Q_0,g_i}\left|\Delta_q f B_q f/\sqrt{n}\right|^2 \leq \kappa P_{Q_0,g^r}\left|\Delta_q f B_q f/\sqrt{n}\right|^2 = \kappa P_{Q_0,g^r}(\Delta_q f B_q f)^2/n \leq \kappa 2^{-2q}/n$. Hence, we can apply lemma 1.5, with $U=a_{q-1}$, $K=4U$, $R=\frac{4}{3}\kappa(2^{-2q})$, to conclude that

$$E\left\|\sqrt{n}\left(P_n-\mathbf{P}_{Q_0,\mathbf{g}_n}\right)\Delta_q f B_q f\right\|_{\mathscr{F}} \leq \sqrt{\frac{32\kappa}{3}2^{-2q}\log(1+N_q)}+8a_{q-1}\log(1+N_q)$$
$$\lesssim 2^{-q}\sqrt{\log(1+N_q)}+\frac{2^{-q}}{\sqrt{\mathrm{Log}(N_q)}}\log(1+N_q) \lesssim 2^{-q}\sqrt{\mathrm{Log}(N_q)}.$$

Now, we bound the second term in (1.24). Since $B_q f=1$ only if $\sqrt{n}a_q < \Delta_q f$, it follows that $\sqrt{n}a_q P_{Q_0,g_i}\Delta_q f B_q f \leq P_{Q_0,g_i}\left(\Delta_q f\right)^2 B_q f \leq 2^{-2q}$. Therefore,

$$\left\|2\sqrt{n}\frac{1}{n}\sum_{i=1}^{n} P_{Q_0,g_i}\Delta_q f B_q f\right\|_{\mathscr{F}} \leq 2\sqrt{n}\frac{2^{-2q}}{\sqrt{n}a_q} = 2\frac{2^{-2q}}{a_q} \simeq 2^{-(q+1)}\sqrt{\mathrm{Log}(N_{q+1})}.$$

Therefore, we obtain a bound for (1.24):

$$E\left\|\sum_{q\geq q_0+1}\sqrt{n}\left(P_n-\mathbf{P}_{Q_0,\mathbf{g}_n}\right)\left(f-\pi_q f\right)B_q f\right\|_{\mathscr{F}} \lesssim \sum_{q\geq q_0+1}\left\{2^{-q}\sqrt{\mathrm{Log}(N_q)}+\frac{2}{a_q}2^{-2q}\right\}$$
$$\lesssim \sum_{q\geq q_0+1} 2^{-q}\sqrt{\mathrm{Log}(N_q)}, \tag{1.25}$$

For the last term on the right-hand side of (1.22), there are at most N_q possibilities for $(\pi_q f-\pi_{q-1}f)A_{q-1}f$. Therefore, we apply lemma 1.5 to the functions $\frac{(\pi_q f-\pi_{q-1}f)A_{q-1}f}{\sqrt{n}}$,

$U = a_{q-1}$ (because $\left|\pi_q f - \pi_{q-1} f\right| A_{q-1} f \leq \Delta_{q-1} f A_{q-1} f \leq \sqrt{n} a_{q-1}$), $K = 4U$, and $R = \frac{4}{3}\kappa(2^{-2q})$, and obtain

$$E \left\| \sum_{q \geq q_0+1} \sqrt{n} \left(P_n - \mathbf{P}_{Q_0,\mathbf{g}_n}\right)\left(\pi_q f - \pi_{q-1} f\right) A_{q-1} f \right\|_{\mathscr{F}} \lesssim \sum_{q \geq q_0+1} \left\{ 2^{-q} \sqrt{\mathrm{Log}(N_q)} \right\} \tag{1.26}$$

Combining (1.22), (1.23),(1.25) and (1.26), we have

$$E \left\| \sqrt{n} \left(P_n - \mathbf{P}_{Q_0,\mathbf{g}_n}\right) f \mathbf{1}\left\{F \leq \sqrt{n} a(\delta)\right\} \right\|_{\mathscr{F}} \lesssim \sum_{q \geq q_0+1} 2^{-q} \sqrt{\mathrm{Log}(N_q)}$$

$$\lesssim \int_0^{\delta} \sqrt{\mathrm{Log} N(\alpha, \mathscr{F}, \|\cdot\|_{2, P_{Q_0,g^r}})} d\alpha$$

This result, together with (1.21), prove the lemma. □

Part II

Semiparametric Inference for Marginal Structural Models

Chapter 2

Targeted Maximum Likelihood Estimation of Dynamic Marginal Structural Models for the Hazard Function

2.1 Introduction

Many research questions in medical and social sciences aim to understand the causal effect of a longitudinal exposure on a time-to-event process. In particular, consider a study where subjects are followed over time; in addition to their baseline covariates, at various time points we also record their time-varying exposure of interest, time-varying covariates, and indicators for the event of interest (say death). Time varying confounding is ubiquitous in these situations: the exposure of interest depends on past covariates, and in turn affects future covariates; right censoring may also be present in a study of this nature, often in response to past covariates and exposure.

To define the causal effect of a longitudinal exposure on mortality, we can use a formal causal framework (Robins (1986)). In it, we first formulate the exposures of interest in terms of interventions to set the values of the exposure and censoring variables, and then compare the distribution of the would-be outcome process (a.k.a. counterfactual outcome process) under different interventions. Depending on the research questions of interest, these interventions can be *static* — assigning the same treatment option to all subjects, or they can be *individualized* (a.k.a dynamic) — assigning treatment options based on past covariates. Since static interventions are a special case of dynamic interventions, we will use the term dynamic for both kinds. One way to assess the effect of the interventions on

mortality is to study how the hazard of the counterfactual outcome changes as a function of time and the interventions; this can be accomplished by summarizing the key features of the hazard function in a marginal structural model (MSM, Robins (1986),Neugebauer and van der Laan (2007)). The MSM not only provides adequate dimension reduction over a potentially complicated hazard function, it also mitigates near positivity violations (due to lack of support for an intervention) by extrapolation.

The parameters of an MSM are often estimated using Inverse Probability Weighted estimation (IPW, van der Laan and Petersen (2007), Robins, Hernan, and Brumback (2000a), Robins (1999), Robins, Orellana, and Rotnitzky (2008)). This estimator is intuitive, easy to implement, and admits influence curve based variance estimates. However, its consistency hinges on correct specification of the conditional treatment (and, if applicable, censoring) probabilities. Moreover, in the presence of strong confounding, the inverse probability weights can become unwieldily large, thus producing very unstable estimates. In the case of static intervention, this instability can be mitigated by introducing marginal kernel weights that down-weight treatment options with little data support; but this solution has limited applicability when the interventions are truly individualized.

Doubly robust and efficient augmented-IPW estimator for a static MSM was proposed in Robins and Rotnitzky (1992), Robins (2000) and Robins, Rotnitzky, and van der Laan (2000b). Under this framework, estimators are defined as solutions to the estimating equation given by the efficient influence curve. While this estimator provides efficiency gain and bias reduction over a misspecified IPW estimator, it still suffers the same general sensitivity to large inverse probability weights. Robins (2000), Robins (2002) and Bang and Robins (2005) proposed an alternative doubly robust estimator based on the innovative insight that the statistical parameters and the corresponding efficient influence curves for the ¡S¡ on the intervention-specific mean can be expressed in terms of iterated conditional expectations. This observation allowed for estimation of only minimal nuisance parameters in addition to the treatment mechanism.

The targeted maximum likelihood estimation (TMLE, van der Laan and Rubin (2006), van der Laan and Rose (2011)) provides a doubly robust and efficient estimator using the substitution principle, therefore it can potentially improve finite sample performance by incorporating global information encoded in the parameter map. A TMLE estimator for longitudinal static MSMs using a stratified approach was proposed by Schnitzer, Moodie, van der Laan, Platt, and Klein (2014). This stratified TMLE uses the longitudinal TMLE for the intervention-specific mean in van der Laan and Gruber (2012) to estimate the mean under each of the static interventions of interest, and then use these to fit MSMs for the hazard and survival functions. While this estimator readily improves upon IPW, it is still vulnerable to insufficient support as it does not take advantage of the extrapolation in the MSM. Most recently, Petersen et al. (2014), building on the results from Robins (2000), Robins (2002) and Bang and Robins (2005), presented a pooled TMLE estimator for lon-

gitudinal dynamic MSM for intervention-specific means. This estimator directly targets the parameters of the MSM, and pools over all interventions of interest in the the updating step, potentially weakening the data support needed to achieve efficiency and double robustness. Petersen et al. (2014) also described how their proposed framework can be generalized to MSM for functions of intervention-specific means.

This chapter builds upon the work of Petersen et al. (2014) to present a targeted maximum likelihood estimator for the marginal structural model for the hazard function under longitudinal dynamic interventions. The proposed estimator is efficient and doubly robust, hence offers an improvement over IPW estimator; it directly targets the MSM parameters by pooling across interventions to update the initial nuisance parameter estimates, hence offers an improvement (in terms of relaxing sensitivity to data support) over the stratified TMLE.

Organization of chapter

This chapter is organized as follows. In section 2.2, we describe the data structure and define the statistical estimand using a nonparametric causal framework. In section 2.3, we first review the non-targeted substitution estimator and the IPW estimator for the estimands of interest, and then present the efficient influence curve, and describe the proposed TMLE estimator. Section 3.5 evaluates the performance of these three estimators in a simulation study mimicking an observational cohort study. This chapter concludes with a summary.

2.2 Defining the Parameter of Interest

Consider a longitudinal data structure

$$O = (L_0, A_0, L_1, A_1, \ldots, L_t, A_t, L_{K+1}),$$

where L_0 denotes the baseline covariates; A_t encodes the exposure variable A_t^1 and the censoring indicator A_t^2, where $A_t^2 = 1$ indicates the subject was right censored at a time $C \leq t$; and L_t denotes all the time-varying covariates measured between A_{t-1} and A_t. In particular, L_t includes the counting outcome process $Y_t \subset L_t$, where $Y_t = 1$ indicates event (say death) has occurred at a time $T \leq t$. We shall use the boldface notation $\mathbf{L}_t \equiv (L_0, \ldots, L_t)$, $\mathbf{L}_{j,t} \equiv (L_j, \ldots, L_t)$ and $\mathbf{L}_{-1} \equiv \emptyset$; similarly for the vector $\mathbf{A}_t$. The observed data consists of n independent and identically distributed (i.i.d.) copies of O drawn from a distribution P_0. Let $\mathscr{M}$ be a statistical model for P_0; the assumptions on this statistical model are limited to true user knowledge, in particular, we avoid strong and restrictive parametric assumptions.

To illustrate these notations (and the subsequent concepts), let us consider the following example from HIV research (extracted from Petersen et al. (2014)). The study population is HIV-infected subjects with immunological failure on first line antiretroviral therapy.

The baseline $t = 0$ is the time of immunological failure, and data is collected on a monthly basis post failure. L_t encompasses time varying covariates at time t, including CD4+ T cell counts and Y_t, an indicator of death by time t. L_0 includes the baseline values of these covariates at $t = 0$, as well as time-independent variables such as patient demography and history prior to first line failure. The treatment variable A_t^1 is the indicator of switching to second line therapy by time t. For simplicity sake, we assume no right censoring in this example.

The time-ordering assumptions implied in the notation can be made explicit by a non-parametric structural equations causal model (NPSEM, Pearl (2009)):

$$L_t = f_{L_t}(\mathbf{L}_{t-1}, \mathbf{A}_{t-1}, U_{L_t});\ A_t = f_{A_t}(\mathbf{L}_t, \mathbf{A}_{t-1}, U_{A_t}), \quad \text{for } k = 0, \ldots K+1.$$

This causal model assumes that each variable X in the observed data structure is an unknown deterministic function f_X of certain observed variables, which we refer to as the *parents of X* and denote by $Pa(X)$, and some unmeasured exogenous random factors U_X. This causal model defines a random variable with distribution $P_{O,U}$ on a unit.

In the HIV example, at each time t, the investigators may specify that CD4 counts, death, and the decision to switch therapy all depend on the patient's entire observed past and unmeasured factors. But if they know that the decision to switch to second line therapy is based only on the most recent CD4 measurement, then $Pa(A_t^1)$ may be restricted to exclude all earlier CD4 measurements.

Parameter of Interest

An *intervention rule d* is a function that deterministically assigns treatment at time t according to $A_t = d(\mathbf{L}_t)$. This rule may be static — assigning the same treatment option a_t regardless of covariate history; or it may be individualized — assigning differential treatment options to different covariate histories. We will use the boldface notation $\mathbf{A}_t = \mathbf{d}(\mathbf{L}_t)$ to denote the vector $(A_k = d(\mathbf{L}_k))_{k=0,\ldots t}$; this means that each of the variables A_k, from $k = 0$ to t, was assigned value according to rule d and its covariate history. In the HIV example, a static rule would be to always switch to second line at some m months after failing first line therapy, and an individualized rule would be to switch at the first time t when the patient's CD4 counts drop below a pre-specified threshold r.

Given a set $\mathscr{D}$ of intervention rules of interest , investigators are often concerned about their comparative causal effect on the outcome process Y_t. To be more precise, consider an ideal experiment where all subjects are assigned treatment under an intervention rule d, and right censoring is also prevented; the covariates, on the other hand, take the value that they may in response to rule d. We call the variables A_t the *intervention variables*, as they are the ones that are subject to manipulation in the ideal experiment. This ideal experiment can be formalized in the NPSEM by setting the equations for A_t to $A_t = d(\mathbf{L}_t)$,

and replacing the input $\mathbf{A}_{t-1}$ into L_t with $\mathbf{d}(\mathbf{L}_{t-1})$, which are treatment assignments from time 0 to $t-1$ under rule d and history $\mathbf{L}_{t-1}$. As a result, the only random endogenous variables of the system are the covariates; we use $L_t(d)$ to denote the time varying covariates that result under the intervention rule d, in particular $Y_t(d)$ is the indicator of death under such a regimen. The comparative causal effect of rule d_1 vs rule d_2 can be assessed by comparing the distributions of the outcome processes $Y_t(d_1)$ vs $Y_t(d_2)$.

Suppose we wish to understand how the outcome process $Y_t(d)$ changes as a function of d, t and some baseline covariate $V \subset L_0$. To this end, we study the intervention-specific hazard function $\lambda(d,t,V) \equiv P(Y_t(d) = 1 \mid Y_{t-1}(d) = 0, V)$ on the space $\mathscr{D} \times \tau \times \mathscr{V}$, where $\mathscr{D}$ is the set of rules we wish to compare, $\tau = \{1, \ldots K+1\}$, and $\mathscr{V}$ is the outcome space of V. Studying the entire function $(d,t,V) \mapsto \lambda(d,t,V)$ may be difficult due to feasibility of computation, amenability to theoretical understanding, and interpretation of results, instead we can study a more tractable, simplified model/summary of this function which captures how λ changes as a function of (d,t,V) in only a few summarizing parameters. More specifically, we consider a marginal structural working model $m_\psi(d,t,V)$ parameterized by $\psi \in \mathbb{R}^J$ for $\lambda(d,t,V)$. In addition, we also consider a user-specified kernel weight function $h(d,t,V)$, and the standard log-likelihood loss. Our causal quantity of interest is the *best* (assessed by the loss function) *weighted* (by h) *approximation* (given by the working model m_ψ) of the hazard function λ . Formally, the causal parameter is defined as

$$\begin{aligned}
\Psi(P_{O,U}) &\equiv \arg\min_\psi E\Big\{ \sum_{t\in\tau, d\in\mathscr{D}} h(d,t,V) I(Y_{t-1}(d)=0) \\
&\qquad \times \log\Big(m_\psi(d,t,V)^{Y_t(d)} \big(1 - m_\psi(d,t,V)\big)^{1-Y_t(d)} \Big) \Big\}. \\
&= \arg\min_\psi E_0 \sum_{t\in\tau, d\in\mathscr{D}} h(d,t,V) P(Y_{t-1}(d)=0 \mid L_0) \\
&\qquad \times \log\Big(m_\psi(d,t,V)^{\lambda(d,t,L_0)} \big(1 - m_\psi(d,t,V)\big)^{1-\lambda(d,t,L_0)} \Big). \qquad (2.1)
\end{aligned}$$

To be concrete, from here on, we will follow common practice and consider a generalized linear working model with logit link, $m_\psi(d,t,V) = \text{expit}(\psi \cdot \phi(d,t,V))$, where $\phi(d,t,V)$ is the vector of linear predictors. However, it is important to note that the methods presented here can be generalized to other working MSM.

Continuing our HIV example, suppose we wish to assess how delay in switching to second line therapy affect mortality. Let d_m denote the rule that dictates switching to second line therapy at m months after first line failure. That is, $A_t = d_m(\mathbf{L}_t) = 0$ for $t < m$ and $Y_t = 0$, and $A_t = d_m(\mathbf{L}_t) = 1$ for $t \geq m$. Let $\mathscr{D}$ be the set of all possible switching times within the study, i.e. $\mathscr{D} = \{d_m : m = 0, \ldots, K, \infty\}$. If we are only interested in the marginal hazard, then we set $V = \emptyset$. Or, if we wish to assess the hazard function stratified by CD4 at the time of first line failure, then we can set $V = \mathbf{1}\{CD4_0 < 50 \text{ cells}/\mu\ell\}$. For simplicity, let us continue this example with the former option, and let the weights be

$h(d,t) = 1$. Same as in Petersen et al. (2014), we choose the logistic working model to be $m_\psi(d_m,t,V) = \text{expit}(\psi_0 + \psi_1 t + \psi_2 max(t-m,0))$, where $max(t-m,0) = 0$ if the patient hasn't switched by time $t-1$, and $max(t-m,0) = t-m$ encodes how long the patient has been on second line therapy if he had readily switched.

Thus far, we have used the NPSEM to formulate the parameter of interest $\Psi(P_{O,U})$ in (2.1). This parameter is a function of the distributions of $L_t(d)$ and $Y_t(d)$, which are generated within an ideal experiment. Unfortunately, such ideal experiments are not always possible in real life; consequently $\Psi(P_{O,U})$ is not always estimable from the observed data. Then, what are the sufficient assumptions on the data-generating process under which the parameter $\Psi(P_{O,U})$ can be identified using the observed data distribution P_0?

To answer this question, we review the expression in (2.1) and note that $P(Y_{t-1}(d) = 0 \mid L_0) = 1 - E(Y_{t-1}(d) \mid L_0)$ and

$$\begin{aligned}\lambda(d,t,L_0) &= \frac{P(Y_t(d)=1,L_0) - P(Y_t(d)=1,Y_{t-1}(d)=1,L_0)}{P(Y_{t-1}(d)=0,L_0)} \\ &= \frac{E(Y_t(d) \mid L_0) - E(Y_{t-1}(d) \mid L_0)}{1 - E(Y_{t-1}(d) \mid L_0)}. \end{aligned} \tag{2.2}$$

Consequently, to identify $\Psi(P_{O,U})$ as a function of P_0, it suffices to establish the identification of the time-dependent causal dose-response curve $\{E(Y_t(d) \mid L_0) : d \in \mathscr{D}, t \in \tau\}$.

To this end, we make the *sequential randomization assumption* (Robins (1986)):

$$A_k \perp \mathbf{L}(d) \mid Pa(A_k), \tag{2.3}$$

and the *positivity assumption*

$$P_0(A_k = d(\mathbf{L}_k) \mid \mathbf{L}_k, \mathbf{A}_{k-1} = \mathbf{d}(\mathbf{L}_{k-1})) > 0 \; a.e., \tag{2.4}$$

for every $d \in \mathscr{D}$ and $k \in \tau$. Assumption (2.3) specified that at each k, the variables A_k are randomized conditional on its parent variables. In our HIV example, this can be satisfied if the parent variables of A_k contain all common determinants of mortality and the decision to switch therapy. Assumption (2.4) requires that under P_0, for each rule d and its compatible covariate history, there is non-zero probability that the subject's exposure will continue to follow this rule at any given time. In our HIV example, given one has not yet switched to second line therapy, there should be non-zero probability that the patient will switch at any given time m in $\mathscr{D}$. Assumption (2.4) would be violated if, say, $\mathscr{D}$ contains the rule to switch at 6 months but no patient in the study population switches therapy at 6 months failing first line. This assumption would also be violated if, say, all patients with $CD4 \geq 50$ cells/$\mu\ell$ at first line failure switch to second line at 1 month after immunological failure. Note that the sequential randomization assumption (2.3) is an untestable assumption on the data-generating process under which we can claim (2.1)

equals the statistical estimand (2.7). Other way around, given the statistical estimand (2.7), assumption (2.3) provides a way to assess causal interpretation of this parameter. On the other hand, the positivity assumption (2.4) is an assumption on the data-generating distribution P_0 (and hence testable using observed data) under which the estimand (2.7) is well-defined.

Under assumptions (2.3) and (2.4), we obtain identification of the intervention-specific mean:

$$E(Y_t(d) \mid L_0) = \sum_{\mathbf{l}_{1,t}} y_t \prod_{j=1}^{t} P_0(l_j \mid L_0, \mathbf{l}_{j-1}, \mathbf{A}_{j-1} = \mathbf{d}(L_0, \mathbf{l}_{1,j-1})) \equiv Q_1^{d,t}(P_0)(L_0). \tag{2.5}$$

This in turn identifies the intervention-specific hazard function:

$$\lambda(d,t,L_0) = \frac{Q_1^{d,t}(P_0)(L_0) - Q_1^{d,t-1}(P_0)(L_0)}{1 - Q_1^{d,t-1}(P_0)(L_0)}.. \tag{2.6}$$

Consequently, the causal parameter of interest $\Psi(P_{O,U})$ in (2.1) is identified as a function of the observed data distribution given by

$$\begin{aligned}\psi_0 \equiv \Psi(P_0) \equiv \arg\min_{\psi} \Bigg\{ & E_{P_0} \sum_{t \in \tau, d \in \mathscr{D}} h(d,t,V) \left(1 - Q_{1,0}^{d,t-1}(L_0)\right) \\ & \times \left(\frac{Q_{1,0}^{d,t}(L_0) - Q_{1,0}^{d,t-1}(L_0)}{1 - Q_{1,0}^{d,t-1}(L_0)} \log m_\psi(d,t,V) + \frac{1 - Q_{1,0}^{d,t}(L_0)}{1 - Q_{1,0}^{d,t-1}(L_0)} \log\left(1 - m_\psi(d,t,V)\right) \right) \Bigg\}. \end{aligned} \tag{2.7}$$

The remainder of this chapter focuses on statistical inference for ψ_0.

2.3 Estimators for ψ_0

Recall that the observed data consist of n i.i.d. copies of $O \sim P_0 \in \mathscr{M}$. Before we introduce the proposed efficient and doubly robust estimator, we will review two available estimators for ψ_0: the non-targeted substitution G-computation estimator and the IPW estimator. The proposed targeted maximum likelihood estimator uses either of these in its updating step. But before we proceed, we shall agree on the following notation.

Notations

We use P_n to denote the empirical distribution of n i.i.d. copies of $O \sim P_0$. Given a function $O \mapsto f(O)$, $P_n f$ denotes the empirical mean $P_n f \equiv \frac{1}{n}\sum_{i=1}^{n} f(O_i)$. More general for any $P \in \mathscr{M}$, $Pf \equiv E_P f(O)$.

For a generic $P \in \mathscr{M}$. We use Q^{L_0} to denote the marginal distribution of L_0. Generalizing (2.5), for a given $t \in \tau$ and $d \in \mathscr{D}$, denote at $P \in \mathscr{M}$:

$$Q_k^{d,t}(P)(\mathbf{L}_{k-1}) \equiv \sum_{\mathbf{l}_{k,t}} y_t \prod_{j=k}^{t} P(l_j \mid \mathbf{L}_{k-1}, \mathbf{l}_{k,j-1}, \mathbf{A}_{j-1} = \mathbf{d}(\mathbf{L}_{k-1}, \mathbf{l}_{k,j-1})), \text{ for } k \leq t, \tag{2.8}$$

and $Q_{t+1}^{d,t}(\cdot) \equiv Y_t$. In the above notation, the superscript t signals the expectant outcome variable Y_t and the subscript k signals the length of the conditioning covariate history. Under assumptions (2.3) and (2.4), $Q_k^{d,t}(P_0)(\mathbf{L}_{k-1}) = E(Y_t(d) \mid \mathbf{L}_{k-1}, \mathbf{A}_{k-1} = \mathbf{d}(\mathbf{L}_{k-1}))$, which is the conditional mean of $Y_t(d)$ given an observed past that has followed the rule d up to time $k-1$. For simplicity, we may sometimes write $Q_k^{d,t}$ instead of $Q_k^{d,t}(P)$ when referring to the functional evaluated at a generic $P \in \mathscr{M}$, and $Q_{k,0}^{d,t}(\mathbf{L}_{k-1})$ instead of $Q_k^{d,t}(P_0)(\mathbf{L}_{k-1})$ when the functional is evaluated at $P_0 \in M$. Bang and Robins (2005) made a key observation that these functionals satisfy the relation

$$Q_k^{d,t}(\mathbf{L}_{k-1}) = E_P\left[Q_{k+1}^{d,t}(\mathbf{L}_k) \middle| \mathbf{L}_{k-1} \right]. \tag{2.9}$$

These functionals are also monotonic in t, i.e. for a given k, $Q_k^{d,t} \leq Q_k^{d,t+1}$ for all $t \geq k$. We will adopt the notation $Q \equiv \left(Q^{L_0}, \left\{ Q_k^{d,t} :, t \in \tau, k \leq t, d \in \mathscr{D} \right\} \right)$. We denote the treatment allocation probabilities $P(A_k \mid \mathbf{L}_k, \mathbf{A}_{k-1})$ as $g(A_k \mid \mathbf{L}_k, \mathbf{A}_{k-1})$, and the product $\prod_{k=1}^{t} g(A_k \mid \mathbf{L}_k, \mathbf{A}_{k-1})$ as $\mathbf{g}(\mathbf{A}_t \mid \mathbf{L}_t)$. For our purposes, the couple (Q, g) readily specifies a distribution $P \in \mathscr{M}$, so sometimes we may abuse notation and write $P = (Q, g)$. At the data generating distribution P_0, we adopt the subscripts Q_0 and g_0.

G-computation Estimator

The identification formula in (2.7) is generally known as the G-computation formula (Robins (1986)). Readily, it delivers a non-targeted substitution estimator (as opposed to the targeted substitution estimator that is TMLE), which is generally known as the parametric G-computation (Gcomp) estimator. More precisely, using the notations in section 5.2, the statistical estimand $\Psi(P_0)$ in (2.7) can be expressed as $\Psi(Q_0)$. Therefore, a non-targeted estimator Q_n of Q_0 will yield a non-targeted substitution estimator $\Psi(Q_n)$ of $\Psi(Q_0)$.

Recall that Q_0 consists of the marginal distribution of L_0, $Q^{L_0}(P_0)$, and the conditional means of $Y_t(d)$, $Q_k^{d,t}(P_0)$ defined in (2.8). To estimate the marginal distribution $Q^{L_0}(P_0)$, we can use the empirical distribution of L_0, denoted $Q_n^{L_0}$. To estimate the conditional means $Q_k^{d,t}(P_0)$, one approach is to estimate the conditional densities for each L_t given its parents and use the definition in (2.8). While this density-based approach ensures that

the monotonicity of in t of $Q_k^{d,t}$ is preserved, the dimension of the nuisance parameter and the computational cost will grow with the the dimension of L_t and the number of time points. One way to implement this approach is to simplifying parametric modeling assumptions on these nuisance parameters. We refer to Taubman, Robins, Mittleman, and Hernan (2009) and Young, Cain, Robins, OReilly, and Hernán (2011) for expositions of this technique.

To minimize estimation of nuisance parameters, one can exploit the recursive relation (2.9) noted by Bang and Robins (2005), and use the many available regression techniques (parametric or data-adaptive) in the literature. We can implement this regression-based approach by running the following two-level algorithm.

1. Starting at $t = K+1$, we estimate the conditional means $\left\{ Q_k^{d,t}(P_0) : k \leq t, d \in \mathscr{D} \right\}$ as follows:

   a) Initiate at $k = t$: Recall that $Q_t^{d,t}(P_0) \equiv E_{P_0}(Y_t \mid \mathbf{L}_{t-1}, \mathbf{A}_{t-1} = \mathbf{d}(\mathbf{L}_{t-1}))$ and $Q_{t+1}^{d,t} \equiv Y_t$. We obtain an estimator $Q_{t,n}^{d,t}$ of $Q_t^{d,t}(P_0)$ by regressing Y_t on the observed values $\mathbf{L}_{t-1}$ and $\mathbf{A}_{t-1}^1$ among observations that remain uncensored by time $t-1$, and evaluate the fitted function at the observed $\mathbf{L}_{t-1}$ and the intervened values $\mathbf{A}_{t-1}^1 = \mathbf{d}(\mathbf{L}_{t-1})$, for these uncensored observations. We can organize the data by having one row for each patient i and each regimen d.

   b) At each subsequent $k = t-1, \ldots, 1$: At the previous step, we have thus far obtained an estimator $Q_{k+1,n}^{d,t}$ of $Q_{k+1}^{d,t}(P_0)$. To obtain an estimator $Q_{k,n}^{d,t}$ of $Q_k^{d,t}(P_0)$, we regress $Q_{k+1,n}^{d,t}(\mathbf{L}_k)$ on the observed values $\mathbf{L}_{k-1}$ and $\mathbf{A}_{k-1}^1$ among those observations that remained uncensored by time $k-1$, and evaluate the fitted function at the observed $\mathbf{L}_{k-1}$ and the intervened values $\mathbf{A}_{k-1}^1 = \mathbf{d}(\mathbf{L}_{k-1})$ on these uncensored observations.

   c) After iterating step (b) in order of decreasing k, we have $\left\{ Q_{k,n}^{d,t} : k \leq t, d \in \mathscr{D} \right\}$.

2. Repeat step 1 in order of decreasing t, from $t = K$ to $t = 1$. At the end, we will have obtained estimators $\left\{ Q_{k,n}^{d,t} : t \in \tau, k \leq t, d \in \mathscr{D} \right\}$

Monotonicity in t, i.e. $Q_{k,n}^{d,t} \leq Q_{k,n}^{t+1}$ for each fixed $k \leq t$, can be enforced in step 1.b with respect to $Q_{k,n}^{t+1}$ obtained at the previous $t+1$ level, or it can be enforced after running the entire algorithm. One simple way to enforce the monotonicity is by sequential truncation; other more sophisticated approaches are available, but they are outside the scope of this chapter.

At the end of this algorithm, we have the conditional means $\left\{ Q_{1,n}^{d,t}(L_0) : t \in \tau, d \in \mathscr{D} \right\}$ for each of the n observations. Pool together these estimates $Q_{1,n}^{d,t}(L_0)$ over all d and t, we will have one row per patient i, rule $d \in \mathscr{D}$ and $t \in \tau$. The G-computation estimator $\psi_n^{Gcomp} \equiv \Psi(Q_n)$ is obtained by fitting a weighted logistic regression of $\frac{Q_{1,n}^{d,t}(L_0) - Q_{1,n}^{d,t-1}(L_0)}{1 - Q_{1,n}^{d,t-1}(L_0)}$ according to the MSM, with weights $h(d,t,V)(1 - Q_{1,n}^{d,t-1}(L_0))$.

Consistency of ψ_n^{Gcomp} relies on consistency of Q_n. In the density-based approach, this means consistent estimation of all the conditional densities of L_t given its parents; in the regression approach, this means consistent estimation of the conditional means $Q_k^{d,t}(P_0)$. In either case, correct specification of Q_0 under a finite dimensional parametric model is possible only in limited applications. Alternatively, we may use machine learning algorithms, such as Super Learner. This option is more enticing, especially when used with the regression-based approach, since there are more data-adaptive techniques available to estimate the conditional mean of a binary variable via regression. However, theoretical results on the asymptotic behavior, such as a central limit theorem, of the resulting estimator $\Psi(Q_n)$ are not available. Moreover, a non-targeted estimator Q_n of Q_0 is obtained by minimizing a global loss function for Q_0, not for $\Psi(Q_0)$. This means, in particular, that the bias-variance tradeoff in Q_n is optimized for the high-dimensional nuisance parameter Q_0, instead of a much lower-dimensional parameter of interest $\Psi(Q_0)$. The proposed targeted estimator in section 5.4 aims to address these two issues by providing a substitution estimator that is asymptotically linear (under appropriate regularity conditions), and optimizes the bias-variance tradeoff of Q_n towards $\Psi(Q_0)$ via an updating step.

Inverse Probability Weighted Estimator

Inverse probability weighted estimation (van der Laan and Petersen (2007),Robins et al. (2008)) is a popular methodology for estimating the parameters of a marginal structural model in the presence of time-varying confounding, due to its ease of implementation and its asymptotic linearity, which allows for construction of Wald confidence intervals.

To begin, we first note that the statistical estimand in (2.7) can be rewritten in an IPW form:

$$\begin{aligned}
&\Psi(P_0) = \Psi^{IPW}(g_0) \\
&\equiv \arg\min_{\psi} \left\{ E_{P_0} \sum_{t \in \tau, d \in \mathscr{D}} h(d,t,V) \frac{I(\mathbf{A}_{t-1} = \mathbf{d}(\mathbf{L}_{t-1}))}{\mathbf{g}_0\left(\mathbf{A}_{t-1} = \mathbf{d}(\mathbf{L}_{t-1}) \mid \mathbf{L}_{t-1}\right)} (1 - Y_{t-1}) \right. \\
&\qquad \left. \times \log\left(m_\psi(d,t,V)^{Y_t} \left(1 - m_\psi(d,t,V)\right)^{1-Y_t} \right) \right\}. \qquad (2.10)
\end{aligned}$$

The IPW estimator $\psi_n^{IPW} \equiv \Psi^{IPW}(g_n)$ is obtained by fitting weighted logistic regression of Y_t according to the MSM, with weights $h(d,t,V)\frac{I(\mathbf{A}_{t-1}=\mathbf{d}(\mathbf{L}_{t-1}))}{\mathbf{g}_n(\mathbf{A}_{t-1}=\mathbf{d}(\mathbf{L}_{t-1})|\mathbf{L}_{t-1})}(1-Y_{t-1})$, where g_n is an estimator for g_0.

The asymptotic theory of the IPW estimator is well understood in the literature. We refer the reader to Robins (1999), van der Laan and Robins (2003b) and van der Laan and Petersen (2007), where the last reference specifically addresses dynamic intervention rules. In summary, ψ_n^{IPW} described above satisfies the estimating equation $P_n D^{IPW}(\psi, g_n) = 0$, where

$$D^{IPW}(\psi,g) = \sum_{t\in\tau, d\in\mathscr{D}} \tilde{h}(d,t,V)\frac{I(\mathbf{A}_{t-1}=\mathbf{d}(\mathbf{L}_{t-1}))}{\mathbf{g}(\mathbf{A}_{t-1}=\mathbf{d}(\mathbf{L}_{t-1})\mid\mathbf{L}_{t-1})}(1-Y_{t-1})\left(Y_t - m_\psi(d,t,V)\right), \quad (2.11)$$

with $\tilde{h}(d,t,V) \equiv h(d,t,V)\phi(d,t,V)$. If the mapping $\Psi^{IPW}: \mathscr{M} \to \mathbb{R}^J$ produces a unique minimizer and hence characterized by the equation $PD^{IPW}(\Psi^{IPW}(g), g) = 0$, then ψ_n^{IPW} is a consistent estimator of ψ_0 provided g_n is a consistent estimator of g_0. Moreover, ψ_n^{IPW} is asymptotically linear with influence curve

$$IC_{IPW}(g_0) = M^{IPW}(\psi_0, g_0, P_0)^{-1} D^{IPW}(\psi_0, g_0), \quad (2.12)$$

where

$$\begin{aligned} M^{IPW}(\psi', g, P) &\equiv -\frac{\partial\ PD^{IPW}(\psi,g)}{\partial\ \psi}\mid_{\psi=\psi'} \\ &= E_P\Big\{\sum_{t\in\tau,d\in\mathscr{D}} h(d,t,V)\phi(d,t,V)\phi(d,t,V)^T \frac{I(\mathbf{A}_{t-1}=\mathbf{d}(\mathbf{L}_{t-1}))}{\mathbf{g}(\mathbf{A}_{t-1}=\mathbf{d}(\mathbf{L}_{t-1})\mid\mathbf{L}_{t-1})}(1-Y_{t-1}) \\ &\quad \times m_{\psi'}(d,t,V)(1-m_{\psi'}(d,t,V))\Big\}. \end{aligned}$$

Let Σ_n^{IPW} denote the sample covariance matrix of $M^{IPW}(\psi_n^{IPW}, g_n, P_n)^{-1} D^{IPW}(\psi_n^{IPW}, g_n)$. The variance of $\sqrt{n}(\psi_n^{IPW} - \psi_0)$ can be estimated using Σ_n^{IPW}. Consequently, we can construct Wald confidence intervals of level $(1-\alpha)$ as $\left[\psi_n^{IPW} \pm \xi_{1-\alpha/2}(\Sigma_n^{IPW}/n)^{1/2}\right]$, where $\xi_{1-\alpha/2}$ is the $(1-\alpha/2)$-quantile of the standard normal distribution.

Though both G-computation estimator and IPW estimator properly account for time-varying confounding, the popularity of IPW over G-computation estimator is apparent from its ease of implementation and its theoretical validity for Wald confidence intervals. Moreover, the treatment probabilities g_0 may arguably be easier to specify correctly than the sequential conditional means $Q_k^{d,t}(P_0)$. However, the G-computation estimator is not entirely without merit compared to the IPW estimator, as the latter is generally more susceptible to near positivity violations due to the inverse probability weighting in (2.10). Often times the kernel functions $h(d,t,V)$ are chosen to stabilize these inverse probability weights; this remedy, however, is less effective when the rules d are individualized, since the marginal function $h(d,t,V)$ cannot depend on the time varying covariates L_t.

Targeted Maximum Likelihood Estimator

As discussed earlier, consistency of the IPW estimator relies on consistency of g_n, while consistency of the G-computation estimator relies on consistency of Q_n. In this section, we propose a semiparametric efficient estimator that is robust against misspecification of either Q_0 or g_0. These theoretical promises hinge on the use of one important ingredient — the efficient influence curve for ψ_0.

The Efficient Influence Curve

Central to our methods is viewing the parameter of interest ψ_0 as the value evaluated at P_0 of the map $\Psi : \mathscr{M} \to \mathbb{R}^J$, where $\Psi(P)$ is given by equation (2.7) after replacing the functionals $Q_{1,0}^{d,t} \equiv Q_1^{d,t}(P_0)$ with $Q_1^{d,t}(P)$ and the marginal expectation under P_0, $E_{P_0}\{\cdot\}$, with the marginal expectation under P, $E_P\{\cdot\}$. It is also straightforward to note from (2.7) that $\Psi(P) = \Psi(Q)$. From its definition, we see that $\psi = \Psi(P)$ satisfies the characterizing equation

$$\begin{aligned} 0 &= U(\psi, Q, P) \\ &= E_P \sum_{t\in\tau, d\in\mathscr{D}} \tilde{h}(d,t,V)\left(1 - Q_1^{d,t-1}(L_0)\right)\left(\frac{Q_1^{d,t}(L_0) - Q_1^{d,t-1}(L_0)}{1 - Q_1^{d,t-1}(L_0)} - m_\psi(d,t,V)\right). \end{aligned} \tag{2.13}$$

Note that $U(\psi, Q, P) = PD^{IPW}(\psi, g)$.

The mapping Ψ is pathwise differentiable on $\mathscr{M}$; its efficient influence curve (EIC) sheds light on the asymptotic properties of all regular and asymptotically linear estimators of $\Psi(P_0)$. The latter statement is formalized in the following lemma. We refer the reader to Bickel, Klaassen, Ritov, and Wellner (1997), van der Laan and Robins (2003b) van der Vaart and Wellner (1996) for definitions and proofs about properties of efficient influence curves in general. The efficient influence curve for the mapping $P \mapsto \Psi(P)$ in (2.7) can be derived using the characterizing equation (2.13) via the functional delta method; this derivation can be found in Petersen et al. (2014) and Schnitzer et al. (2014), we provide it in the appendix for completeness sake.

Lemma 2.1 (Efficient influence curve for Ψ).

Suppose the mapping $\Psi : \mathscr{M} \to \mathbb{R}^J$ is well-defined at P, in the sense that it's a unique minimizer and hence characterized by the equation $U(\Psi(P), Q, P) = 0$. Then, its efficient

influence curve at $P = (Q,g)$, denoted $D^*(Q,g)$, is given by $D^*(\Psi(Q),Q,g)$, where

$$D^*(\psi,Q,g)(O) \equiv M(\psi,Q,P)^{-1} \times \left\{ \sum_{t\in\tau} \sum_{d\in\mathscr{D}} D^{d,t}(\psi,Q,g)(O) + D^{L_0}(\psi,Q)(O) \right\}, \tag{2.14}$$

with

$$M(\psi,Q,P) \equiv -\frac{\partial U(\psi,Q,P)}{\partial \psi}$$

$$= E_P\Big\{ \sum_{t\in\tau,d\in\mathscr{D}} h(d,t,V)\phi(d,t,V)\phi(d,t,V)^T \left(1 - Q_1^{d,t-1}(L_0)\right) m_\psi(d,t,V)(1-m_\psi(d,t,V)) \Big\},$$

$$D^{d,t}(\psi,Q,g)(O) \equiv J_\psi(d,t,V) \sum_{k=1}^{t} \frac{I(\mathbf{A}_{k-1} = \mathbf{d}(\mathbf{L}_{k-1}))}{\mathbf{g}(\mathbf{A}_{k-1} = \mathbf{d}(\mathbf{L}_{k-1}) \mid \mathbf{L}_{k-1})} \left(Q_{k+1}^{d,t}(\mathbf{L}_k) - Q_k^{d,t}(\mathbf{L}_{k-1}) \right),$$

$$D^{L_0}(\psi,Q)(O) \equiv \sum_{t\in\tau,d\in\mathscr{D}} \tilde{h}(d,t,V) \left(1 - Q_1^{d,t-1}(L_0)\right) \left(\frac{Q_1^{d,t}(L_0) - Q_1^{d,t-1}(L_0)}{1 - Q_1^{d,t-1}(L_0)} - m_\psi(d,t,V) \right).$$

with $J_\psi(d,t,V) = \tilde{h}(d,t,V) - I(t \le K)\tilde{h}(d,t+1,V)(1-m_\psi(d,t+1,V)$.

The variance $\mathrm{Var}_P D^*(\Psi(Q),Q,g)(O)$ is a generalized Cramér-Rao lower bound for the asymptotic variance of any regular and asymptotically linear estimator of $\Psi(P)$.

Moreover, if either $Q = Q_0$ or $g = g_0$, then $P_0 D^*(\Psi(Q),Q,g) = 0$ implies that $\Psi(Q) = \Psi(Q_0)$.

Proof. See appendix for derivation of (2.14) and proof of double robustness; see Bickel et al. (1997), van der Laan and Robins (2003b) and van der Vaart and Wellner (1996) for the statement regarding variance bounds. □

Now, we are ready to describe the implementation of a TMLE estimator using D^*.

The Loss Function, the Fluctuation Model, and the Algorithm

In a glimpse, our strategy consists of targetedly updating given initial estimators Q_n of Q_0 by minimizing a pre-specified loss along a least favorable (with respect to ψ_0) submodel through Q_n; iterate this updating procedure until the estimating equation $P_n D^*(Q_n^*, g_n) = 0$ is solved at some final targeted estimate Q_n^* of Q_0, and then evaluate Ψ at this Q_n^*.

More specifically, for each $d \in \mathscr{D}$, $t = K+1,\ldots,1$ and $k \le t$, consider the quasi negative log likelihood loss for $Q_k^{d,t}$, indexed by its expectant $Q_{k+1}^{d,t}$:

$$L\left(Q_k^{d,t}\right)(O) \equiv -I\left(\mathbf{A}_{k-1} = \mathbf{d}\left(\mathbf{L}_{k-1}\right)\right) \log \left(Q_k^{d,t}(\mathbf{L}_{k-1})^{Q_{k+1}^{d,t}(\mathbf{L}_k)} \left(1 - Q_k^{d,t}(\mathbf{L}_{k-1})\right)^{1-Q_{k+1}^{d,t}(\mathbf{L}_k)} \right). \tag{2.15}$$

We suppressed the indexing by $Q_{k+1}^{d,t}$ in the notation. The corresponding least favorable submodel through $Q_k^{d,t}$, parametrized as $\left\{Q_k^{d,t}(\varepsilon) : \varepsilon \in \mathbb{R}^J\right\}$, is chosen to satisfy the score condition $\sum_{k=1}^{t} \frac{\partial L\left(Q_k^{d,t}(\varepsilon)\right)}{\partial \varepsilon} |_{\varepsilon=0} = D^{d,t}(\psi, Q, g)$. In particular, we can choose

$$Q_k^{d,t}(\varepsilon) = \text{expit}\left(\text{logit}\, Q_k^{d,t} + \varepsilon\, H_k^{d,t}(\psi, g)\right), \tag{2.16}$$

where $H_k^{d,t}(\psi, g)(\mathbf{L}_{k-1}) \equiv \frac{J_\psi(d,t,V)}{\mathbf{g}(\mathbf{A}_{k-1}=\mathbf{d}(\mathbf{L}_{k-1})|\mathbf{L}_{k-1})}$. Note that the dependency of $Q_k^{d,t}(\varepsilon)$ on ψ and g are suppressed in the notation.

Before describing the algorithm, we make the following observation. Due to the form of the efficient influence curve, the direction of fluctuation $H_k^{d,t}(\psi, g)$ depends on the estimand itself; consequently, the implementation of the TMLE for the hazard MSM conceals more subtleties than its counterpart for the survival MSM in Schnitzer et al. (2014) and Petersen et al. (2014). For one, in addition to non-targeted initial estimates Q_n and g_n, an initial estimate of ψ_0 is also needed to perform the first update for Q_n. One should choose a consistent initial estimator (albeit not doubly robust), such as the G-computation or the IPW estimator. Since iteration will be performed, choosing either one of these should have the same asymptotic implications, however, it may make a difference in finite sample performance when Q_n is misspecified. The second subtlety is that this TMLE estimator is truly iterative, wherein each iteration uses the previously obtained updated estimate of ψ_0 to steer the direction of fluctuation. The goal of these iterations is to produce targeted estimators Q_n^* of Q_0 that satisfy the efficient score equation $P_n D^*(Q_n^*, g_n) = 0$. However, once the residual term $P_n D^*(Q_n^*, g_n)$ becomes smaller than the standard error of the estimator, computational efforts spent to further minimize it will only yield diminishing returns. Therefore, we use as stopping rule $|P_n D^*(Q_n^*, g_n)| / \hat{SE}_n < 1/\sqrt{n}$, where $\hat{SE}$ is estimated using $\sqrt{\hat{\text{Var}}(D^*(Q_n^*, g_n))/n}$.

With these issues in mind, we are now ready to describe the algorithm.

1. Obtain initial estimators g_n of g_0 and Q_n of Q_0. For the latter, estimate the marginal distribution of L_0 using the empirical distribution, and estimate the conditional means $Q_k^{d,t}(P_0)$ using the regression-based technique described in section 3.4. Obtain initial estimator ψ_n of ψ_0 using either the IPW estimator $\Psi^{IPW}(g_n)$ or the G-computation estimator $\Psi(Q_n)$.

2. Given initial estimators ψ_n, g_n and $Q_n = \left(Q_n^{L_0}, \left\{Q_{k,n}^{d,t}, t \in \tau, k \leq t, d \in \mathscr{D}\right\}\right)$, we sequentially update the conditional means $Q_{k,n}^{d,t}$ in a two-level algorithm as follows:

a) Starting at $t = K+1$, estimate the conditional means $\left\{ Q_k^{d,t}(P_0) : k \leq t, d \in \mathscr{D} \right\}$ as follows:

i. Initiate at $k = t$: Recall that $Q_t^{d,t}(P_0) \equiv E_{P_0}\left(Y_t \mid \mathbf{L}_{t-1}, \mathbf{A}_{t-1} = \mathbf{d}(\mathbf{L}_{t-1})\right)$ and $Q_{t+1}^{d,t} \equiv Y_t$. We update the initial estimator $Q_{t,n}^{d,t}$ of $Q_t^{d,t}(P_0)$ by $Q_{t,n}^{d,t,*} \equiv Q_{t,n}^{d,t}\left(\varepsilon_{t,n}^t\right)$ where

$$\varepsilon_{t,n}^t \equiv \arg\min\ P_n \sum_{d \in \mathscr{D}} L\left(Q_{t,n}^{d,t}(\varepsilon)\right).$$

To implement this, first pool the data by creating one row per patient i and regimen $d \in \mathscr{D}$. Then, regress Y_t on $H_t^{d,t}(\psi_n, g_n)(\mathbf{L}_{t-1})$ with offset $\operatorname{logit} Q_{t,n}^{d,t}(\mathbf{L}_{t-1})$, among observations that remained uncensored by time $t-1$ and assigning weight $I\left(\mathbf{A}_{t-1,i}^1 = \mathbf{d}\left(\mathbf{L}_{t-1,i}\right)\right)$ to row (i,d).

ii. At each subsequent $k = t-1, \ldots, 1$: We have an initial estimator $Q_{k,n}^{d,t}$ of $Q_k^{d,t}(P_0)$, and at the previous step, we have obtained an updated estimator $Q_{k+1,n}^{d,t,*}$ of its expectant $Q_{k+1}^{d,t}(P_0)$. The update estimator $Q_{k,n}^{d,t,*}$ is given by $Q_{k,n}^{d,t,*} \equiv Q_{k,n}^{d,t}(\varepsilon_k^t)$ where

$$\varepsilon_{k,n}^t \equiv \arg\min\ P_n \sum_{d \in \mathscr{D}} L\left(Q_{k,n}^{d,t}(\varepsilon)\right).$$

We can obtain this by regressing $Q_{k+1,n}^{d,t,*}(\mathbf{L}_k)$ on $H_k^{d,t}(\psi_n, g_n)(\mathbf{L}_{k-1})$ with offset $\operatorname{logit} Q_{k,n}^{d,t}(\mathbf{L}_{k-1})$ and weights $I\left(\mathbf{A}_{k-1}^1 = \mathbf{d}\left(\mathbf{L}_{k-1}\right)\right)$, among observations that remained uncensored by time $k-1$.

iii. After iterating step (ii) in order of decreasing k, we have obtained updated estimators $\left(Q_{k,n}^{d,t,*} : k \leq t, d \in \mathscr{D}\right)$.

b) Repeat step (a) in order of decreasing t, from $t = K$ to $t = 1$.

3. At the end of the two-level algorithm in step (2) above, we have obtained targeted estimators $Q_n^* = \left(Q_n^{L_0}, \left\{ Q_{k,n}^{d,t,*}(\cdot) : t \in \tau, k \leq t, d \in \mathscr{D} \right\}\right)$. In particular, we have $\left\{ Q_{1,n}^{d,t,*}(L_0) : t \in \tau, d \in \mathscr{D} \right\}$ for each of the n observations. Pool together these estimates $Q_{1,n}^{d,t,*}(L_0)$ over all d and t, we will have one row per patient i, rule $d \in \mathscr{D}$ and $t \in \tau$. We can then update the parameter estimate using $\psi_n^* \equiv \Psi(Q_n^*)$; this can be implemented by fitting a weighted logistic regression of $\frac{Q_{1,n}^{d,t,*}(L_0) - Q_{1,n}^{d,t-1,*}(L_0)}{1 - Q_{1,n}^{d,t-1,*}(L_0)}$ according to the MSM, with weights $\tilde{h}(d,t,V)(1 - Q_{1,n}^{d,t-1,*}(L_0))$.

4. Repeat step 2 and 3 using ψ_n^* and Q_n^* as initial estimators. Iterate this procedure until the stopping criteria $P_n D^*(Q_n^*, g_n)/\hat{SE}_n < 1/\sqrt{n}$, where $\hat{SE}_n = \sqrt{\hat{\text{Var}}\left(D^*(Q_n^*, g_n)\right)/n}$, is satisfied.

5. The final updates and Q_n^* and ψ_n^* are the targeted maximum likelihood estimators for Q_0 and ψ_0, respectively.

Statistical Inference of TMLE

As we alluded to earlier, given the pre-specified loss function (5.14), the corresponding least favorable submodel (5.15) is chosen so that, by design, the TMLE estimator Q_n^* satisfies $|P_n D^*(Q_n^*, g_n)| \approx 0$. From this property stems the doubly robust and locally efficient properties of ψ_n^*.

Specifically, under regularity and empirical process conditions (see e.g. van der Laan and Rose (2011)), if both Q_n^* and g_n are consistent estimators, then ψ_n^* is asymptotically linear with influence curve $D^*(Q_0, g_0)$; if g_n converges to g_0, but Q_n^* converges to some Q^* (which may be correctly specified or otherwise), then the influence curve of ψ_n^* equals $D^*(Q_0, g_0)$ minus its projection onto the tangent space of $\mathscr{M}$ for g_0. In either case, $\sqrt{n}(\psi_n^* - \psi_0)$ converges weakly to a normal distribution with covariance matrix equal to or greater than the covariance matrix of $D^*(\psi_0, Q^*, g_0)$. Consequently, the asymptotic variance of $\sqrt{n}(\psi_n^* - \psi_0)$ can be conservatively estimated using Σ_n^*, the sample covariance matrix of $D^*(\psi_n^*, Q_n^*, g_n)$. We can construct asymptotically conservative confidence intervals of level $(1-\alpha)$ as $\left[\psi_n^* \pm \xi_{1-\alpha/2}(\Sigma_n^*/n)^{1/2}\right]$, where $\xi_{1-\alpha/2}$ is the $(1-\alpha/2)$-quantile of the standard normal distribution. Note that this influence curve based variance estimate assumes that the weights $h(d,t,V)$ are known functions; when these weights are estimated, this variance estimate should be interpreted as estimating the variance of an estimator of MSM parameters defined by the estimated weights.

2.4 Simulation Study

In this section, we evaluate the relative performances of the G-computation estimator, IPW estimator and TMLE estimator for the parameters of a marginal structural model for the hazard function. For each estimator, we assess the bias, variance, mean squared error (MSE) and coverage estimates for the influence curve based 95% confidence intervals. This data generating process, together with the running example, are extracted from Petersen et al. (2014), only a few parameters are changed to prevent near positivity violations.

Data Generating Process

We use a data generating process that resembles the running example. The goal is to assess the effect of delay in switching to a second line antiretroviral therapy on mortality among HIV infected patients who failed their first line therapy.

A sample consists of n i.i.d. copies of

$$O = (W, CD4_0, A_0, Y_1, CD4_1, A_1, \ldots, Y_K, CD4_K, A_K, Y_{K+1}).$$

The baseline variable W encodes sex, the baseline age and disease stage, the time varying covariate $CD4_t$ encodes the most recent CD4 count measurement at time t, and Y_t is the indicator of death by time t. The treatment variable A_t^1 is the indicator of having switched to second line therapy by time t. The censoring variable A_t^2 consists of (C_t^1, C_t^2), indicating database closer by time t and loss to follow up by time t, respectively.

We briefly summarize the data generating steps here, deferring the details of the distributions to the appendix. First, we draw the baseline time independent covariates W. At each time $t \geq 0$, if uncensored, first draw the CD4 count measurement for time t based on baseline covariate W, prior CD4 counts, and treatment status at previous time point A_{t-1}^1. Then, determine censoring by database closure C_t^1 based on W. If still uncensored, determine censoring by loss to followup C_t^2 based on W, prior CD4 counts, and previous regimen status A_{t-1}^1. If still uncensored and have not switched, determine switching using W and current CD4 counts $CD4_t$. Finally, determine death based on W, current CD4 counts $CD4_t$ and current regimen status A_t^1.

Target Parameter

An intervention rule of interest d_m switches to second line therapy at m months after first line therapy failure and prevents right censoring, given the subject is still alive. The set of regimens of interest are indexed by possible switching times, $\mathscr{D} = \{d_m : m = 0, \ldots, K+1\}$. Note that each d_m is in fact a static rule assigning $A_t = a_t^m$, where $a_t^m = 0$ at $t < m$ and $a_t^m = 1$ at $t \geq m$. We summarize the hazard function $\lambda(d_m, t) = P(Y_t(d_m) = 1 \mid Y_{t-1}(d_m) = 0)$ as

$$m_\psi(d_m, t) = \text{expit}\left(\psi_0 + \psi_1 t + \psi_2 max(t-m, 0)\right),$$

and use kernel weights

$$h(d_m, t) = P_0\left(\mathbf{A}_{t-1} = \mathbf{a}_{t-1}^m\right) I(t < t^*),$$

where t^* is the first time point where all subjects have either died or censored. Both $P_0\left(\mathbf{A}_{t-1} = \mathbf{a}_{t-1}^m\right)$ and t^* are estimated from data, but are treated as given in our parameter definition. This weight mitigates near positivity violation by down-weighting rules with little support in data.

Estimators

We use the regression-based approach in section 3.4 to estimate the initial estimator Q_n of Q_0. A correctly specified $Q_{k,n}^{d,t}$ is obtained by regressing $Q_{k+1,n}^{d,t}$ on W and the entire past covariate history $\mathbf{L}_{k-1} \equiv (CD4_0, \ldots, CD4_{k-1})$ and exposure history $\mathbf{A}_{k-1}^1$. A misspecified $Q_{k,n}^{d,t}$ uses an intercept model. A correctly specified g_n estimates the conditional probability of C_t^1 by adjusting for W and $CD4_0$; estimates conditional probability of C_t^2 by adjusting for W, the last observed $CD4_{t-1}$ and A_{t-1}^1; and estimates treatment allocation probabilities by adjusting for W and $CD4_t$. A misspecified g_n uses an intercept model. Both g_n and Q_n are fitted using Super Learner, under 10-fold cross-validation, with candidate fitting algorithms `glm`, `neural net`, and `stepAIC`.

We will implement the Gcomp and IPW estimators as described in sections 3.4 and 3.4, we will implement two versions of the TMLE estimator, one with Gcomp as initial estimator, denoted as TMLE^{Gcomp}, and one with IPW as initial estimator, denoted as TMLE^{IPW}. In both IPW and TMLE, the product of g in the denominator is truncated below at 0.01. The iterations in TMLE will stop at either $P_n D^*(Q_n^*, g_n)/\hat{SE}_n < 1/\sqrt{n}$, where $\hat{SE}_n = \sqrt{\hat{\text{Var}}\left(D^*(Q_n^*, g_n)\right)/n}$, or at the 17th iteration, whichever comes first.

Results

We assess bias, variance, MSE and coverage probability of the confidence intervals of the estimators over 500 samples of size $n = 500$, with number of time points $K = 2$. The results are displayed in table 2.1. The TMLE implementations in all simulations ended by the stopping rule, and for our purposes are deemed convergent.

Firstly, we compare the two TMLE estimators (with different initial estimators) across the board, and note that, after rounding up to 3 significant digits, the choice of initial estimator has little effect on the performance of the TMLE estimator. Therefore, we will focus on comparative performance of Gcomp, IPW and either TMLE.

When both Q_0 and g_0 are correctly specified, theoretical results predict that TMLE should be more efficient than IPW, and both TMLE and IPW are asymptotically normal with an asymptotic variance that can be estimated by $\hat{\text{Var}}D^*(Q_n^*, g_n)$ and $\hat{\text{Var}}D^{IPW}(g_n)$, respectively; consequently, finite sample variance of these estimators can be approximated using the sample variance of their influence curves divided by $\sqrt{n}$. For the given sample size $n = 500$ and number of time points $K = 2$, however, we see that (i) TMLE only has slight efficiency gain over IPW; (ii) the influence curve (IC) based variance estimates of the TMLE and IPW ("IC-based var.") grossly underestimate the true variance of the estimator; (iii) The coverage of the Wald 95% confidence intervals made using IC based variance estimates ("Coverage: IC var. CI") is very poor. Just how much of the poor

coverage in (iii) is attributed to the poor estimate given by the IC-based variance estimator, as observed in (ii), and how much of it indicates an underlying problem with achieving asymptotic normality? To answer this, we construct a Wald confidence interval using the true variance of the estimators, which can be approximated using the sample variance across the 500 simulations. The coverage of these confidence intervals is reported under "Coverage: true var. CI". We see that this second set of confidence intervals have coverage close to the 95% confidence level, given the moderate sample size. This suggests that the poor coverage problem in (iii) is mostly attributed to poor variance estimate. Indeed, increasing sample size to $n = 5000$ (table 2.2) improves significantly the coverage of the confidence intervals, as the IC based variance estimate more accurately approximates the true variance of the estimators. The efficiency gain of TMLE over IPW is also more apparent in larger sample size. These observations teaches us that even though IPW and TMLE estimators have theoretically valid IC based variance estimates, their accuracy (and the coverage of the corresponding confidence intervals) are generally sensitive to sample size (with respect to the complexity of the data structure) and to potential bias in the truncation of g_n. Consequently, better variance estimate of these estimators, beyond using their first order term in the taylor expansion, is needed, but this is beyond the scope of this chapter. Alternative, instead of employing central limit theorem based inference, one may also obtain variance estimates via bootstrap, when willing to pay the computational expense.

The double robustness properties provide bias reduction as predicted by lemma 2.1. When g_n is misspecified and Q_n is correctly specified, TMLE provides significant bias reduction over the misspecified IPW, with only slight increase in variance, hence leading to overall smaller MSE. The benefit of this bias-variance tradeoff is perhaps more apparent in large sample sizes, since the overall MSE of the TMLE estimators will decrease with sample size, whereas those of IPW will stabilize. When Q_n is misspecified and g_n is correct, TMLE also provides bias reduction over the misspecified Gcomp, but this may come at the price of increased variance.

Table 2.1: $n = 500$, $K = 2$. "IC-based var.": the influence curve based variance estimate; "IC var. CI": the Wald 95% confidence interval using IC-based variance estimate. "true var. CI": Wald confidence interval using true variance of the estimators (assessed through the 500 simulations).

	Q_n **and** g_n **correct**				**correct** Q_n**, misspec.** g_n			**misspec** Q_n**, correct.** g_n		
	Gcomp	IPW	TMLE^{IPW}	TMLE^{Gcomp}	IPW	TMLE^{IPW}	TMLE^{Gcomp}	Gcomp	TMLE^{IPW}	TMLE^{Gcomp}
Bias										
$\hat{\psi}_0$	-8.29e-02	-7.41e-03	1.78e-02	1.78e-02	-1.41e-01	-1.45e-02	-1.45e-02	-7.18e-02	3.78e-02	3.78e-02
$\hat{\psi}_1$	9.95e-04	-2.46e-02	-1.92e-02	-1.92e-02	-2.05e-01	-1.72e-03	-1.72e-03	-3.70e-01	-3.12e-02	-3.12e-02
$\hat{\psi}_2$	4.61e-02	8.54e-03	-8.48e-03	-8.48e-03	5.04e-01	2.90e-03	2.90e-03	7.26e-01	-2.25e-02	-2.25e-02
Variance										
$\hat{\psi}_0$	1.65e-01	1.95e-01	1.92e-01	1.92e-01	7.61e-02	9.18e-02	9.18e-02	7.01e-02	1.84e-01	1.84e-01
$\hat{\psi}_1$	8.53e-02	1.19e-01	1.15e-01	1.15e-01	4.15e-02	4.84e-02	4.84e-02	3.43e-02	1.07e-01	1.07e-01
$\hat{\psi}_2$	3.72e-02	3.27e-02	2.99e-02	2.99e-02	1.72e-02	2.16e-02	2.16e-02	4.14e-28	2.95e-02	2.95e-02
IC-based var.										
$\hat{\psi}_0$	1.43e-01	1.49e-01	1.10e-01	1.10e-01	8.03e-02	7.32e-02	7.32e-02	1.95e-01	1.21e-01	1.21e-01
$\hat{\psi}_1$	8.83e-02	9.16e-02	6.10e-02	6.10e-02	4.37e-02	4.01e-02	4.01e-02	1.35e-01	7.22e-02	7.22e-02
$\hat{\psi}_2$	2.50e-02	3.13e-02	1.94e-02	1.94e-02	1.72e-02	1.74e-02	1.74e-02	4.35e-02	2.19e-02	2.19e-02
MSE										
$\hat{\psi}_0$	1.72e-01	1.95e-01	1.92e-01	1.92e-01	9.58e-02	9.18e-02	9.18e-02	7.52e-02	1.85e-01	1.85e-01
$\hat{\psi}_1$	8.51e-02	1.19e-01	1.16e-01	1.16e-01	8.34e-02	4.83e-02	4.83e-02	1.71e-01	1.08e-01	1.08e-01
$\hat{\psi}_2$	3.93e-02	3.27e-02	2.99e-02	2.99e-02	2.71e-01	2.15e-02	2.15e-02	5.26e-01	3.00e-02	3.00e-02
Coverage: IC var. CI										
$\hat{\psi}_0$	9.16e-01	8.78e-01	8.34e-01	8.34e-01	8.72e-01	9.22e-01	9.22e-01	9.90e-01	8.50e-01	8.50e-01
$\hat{\psi}_1$	8.98e-01	9.08e-01	8.10e-01	8.10e-01	8.50e-01	9.32e-01	9.32e-01	7.92e-01	8.42e-01	8.42e-01
$\hat{\psi}_2$	8.56e-01	9.46e-01	8.66e-01	8.66e-01	3.40e-02	9.22e-01	9.22e-01	4.80e-02	8.86e-01	8.86e-01
Coverage: true var. CI										
$\hat{\psi}_0$	9.44e-01	9.56e-01	9.56e-01	9.56e-01	9.08e-01	9.52e-01	9.52e-01	9.48e-01	9.48e-01	9.48e-01
$\hat{\psi}_1$	9.62e-01	9.42e-01	9.48e-01	9.48e-01	8.42e-01	9.52e-01	9.52e-01	4.76e-01	9.48e-01	9.48e-01
$\hat{\psi}_2$	9.42e-01	9.64e-01	9.58e-01	9.58e-01	3.40e-02	9.50e-01	9.50e-01	0.00e+00	9.56e-01	9.56e-01

Table 2.2: $n = 5000$, $K = 2$. "IC-based var.": the influence curve based variance estimate; "IC var. CI": the Wald 95% confidence interval using IC-based variance estimate. "true var. CI": Wald confidence interval using true variance of the estimators (assessed through the 500 simulations).

	Q_n and g_n correct				**correct Q_n, misspec. g_n**			**misspec Q_n, correct. g_n**		
	Gcomp	IPW	TMLE^{IPW}	TMLE^{Gcomp}	IPW	TMLE^{IPW}	TMLE^{Gcomp}	Gcomp	TMLE^{IPW}	TMLE^{Gcomp}
Bias										
$\hat{\psi}_0$	-7.92e-03	-2.36e-02	-1.25e-02	-1.25e-02	-1.35e-01	-5.78e-03	-5.78e-03	-7.00e-02	3.96e-02	3.96e-02
$\hat{\psi}_1$	-2.06e-02	7.20e-03	9.24e-03	9.24e-03	-2.10e-01	7.10e-04	7.10e-04	-3.71e-01	-4.38e-02	-4.38e-02
$\hat{\psi}_2$	4.66e-03	-7.66e-03	-9.19e-03	-9.19e-03	5.09e-01	-6.64e-03	-6.64e-03	7.26e-01	3.57e-03	3.57e-03
Variance										
$\hat{\psi}_0$	2.67e-02	1.99e-02	1.79e-02	1.79e-02	8.23e-03	1.05e-02	1.05e-02	7.46e-03	1.91e-02	1.91e-02
$\hat{\psi}_1$	2.27e-02	1.30e-02	1.05e-02	1.05e-02	4.56e-03	5.79e-03	5.79e-03	3.51e-03	1.16e-02	1.16e-02
$\hat{\psi}_2$	6.95e-03	3.91e-03	3.35e-03	3.35e-03	1.50e-03	2.01e-03	2.01e-03	4.28e-29	3.52e-03	3.52e-03
IC-based var.										
$\hat{\psi}_0$	1.75e-02	1.89e-02	1.56e-02	1.56e-02	7.95e-03	7.23e-03	7.23e-03	2.39e-02	1.81e-02	1.81e-02
$\hat{\psi}_1$	1.12e-02	1.24e-02	9.62e-03	9.62e-03	4.31e-03	3.95e-03	3.95e-03	1.76e-02	1.21e-02	1.21e-02
$\hat{\psi}_2$	2.96e-03	3.82e-03	2.74e-03	2.74e-03	1.70e-03	1.73e-03	1.73e-03	5.49e-03	3.31e-03	3.31e-03
MSE										
$\hat{\psi}_0$	2.67e-02	2.05e-02	1.80e-02	1.80e-02	2.66e-02	1.05e-02	1.05e-02	1.24e-02	2.06e-02	2.06e-02
$\hat{\psi}_1$	2.31e-02	1.31e-02	1.05e-02	1.05e-02	4.86e-02	5.78e-03	5.78e-03	1.41e-01	1.35e-02	1.35e-02
$\hat{\psi}_2$	6.96e-03	3.96e-03	3.42e-03	3.42e-03	2.60e-01	2.05e-03	2.05e-03	5.26e-01	3.53e-03	3.53e-03
Coverage: IC var. CI										
$\hat{\psi}_0$	9.44e-01	9.16e-01	9.38e-01	9.38e-01	5.80e-01	9.10e-01	9.10e-01	9.96e-01	9.24e-01	9.24e-01
$\hat{\psi}_1$	9.34e-01	9.46e-01	9.30e-01	9.30e-01	1.10e-01	9.10e-01	9.10e-01	7.80e-02	9.42e-01	9.42e-01
$\hat{\psi}_2$	9.22e-01	9.26e-01	8.96e-01	8.96e-01	0.00e+00	9.44e-01	9.44e-01	0.00e+00	9.40e-01	9.40e-01
Coverage: true var. CI										
$\hat{\psi}_0$	9.70e-01	9.52e-01	9.54e-01	9.54e-01	6.78e-01	9.56e-01	9.56e-01	8.62e-01	9.50e-01	9.50e-01
$\hat{\psi}_1$	9.88e-01	9.42e-01	9.52e-01	9.52e-01	1.34e-01	9.52e-01	9.52e-01	0.00e+00	9.32e-01	9.32e-01
$\hat{\psi}_2$	9.80e-01	9.42e-01	9.60e-01	9.60e-01	0.00e+00	9.62e-01	9.62e-01	0.00e+00	9.58e-01	9.58e-01

2.5 Summary

In this chapter, we have presented a doubly robust and efficient substitution estimator for a longitudinal dynamic (or static) marginal structural model for the hazard function. This work builds upon the pooled TMLE methodology in Petersen et al. (2014), as well as earlier work by Bang and Robins (2005).

Unlike TMLE for MSM for survival functions or means, TMLE for MSM for a hazard function is a bona fide iterated estimator and requires an initial estimator of the parameter of interest. There is no theoretical predictions on which initial estimator is more advantageous, the simulations have also offer little evidence in favor of either. This seems to suggests that the choice of initial estimator has little effect, at least in moderate and large sample sizes.

The proposed TMLE estimator offers theoretical advantages over the popular IPW estimator and a non-targeted substitution Gcomp estimator: 1) it offers protection against model misspecifications, 2) it is locally efficient. Moreover, compared to the IPW estimator, TMLE ought to be generally less sensitive (though not immune) to near positivity violations, thanks to the substitution principle. We test these theoretical promises in a simulation study mimicking an observational cohort study. The bias reduction over a misspecified IPW or Gcomp estimator is clear even for a moderate sample size. On the other hand, the efficiency gain of TMLE over IPW is clear only at large sample sizes in this simulation. We also see that at moderate sample sizes, the influence curve based variance estimate of IPW and TMLE is a poor estimate of the true variance. Even though this variance estimate can be improved with increased sample size, its performance is still not satisfactory for real life applications. This suggests that a more sophisticated variance estimate, as well as appropriate diagnostics, are needed for these two estimators. Future research priorities should focus on variance estimation and inference methods that remain resilient in the face of moderate confounding and multiple time points.

2.6 Chapter Appendix

Derivation of Efficient Influence Curve

In order to apply the functional delta method, we first rewrite the characterizing equation (2.13) as

$$\begin{aligned} U(\psi,Q,P) &= E_P \sum_{t\in\tau,d\in\mathscr{D}} \tilde{h}(d,t,V)\left(1-Q_1^{d,t-1}(L_0)\right)\left(\frac{Q_1^{d,t}(L_0)-Q_1^{d,t-1}(L_0)}{1-Q_1^{d,t-1}(L_0)} - m_\psi(d,t,V)\right) \\ &= E_P \sum_{t\in\tau,d\in\mathscr{D}} \tilde{h}(d,t,V)\left(Q_1^{d,t}(L_0)-Q_1^{d,t-1}(L_0) - m_\psi(d,t,V) + m_\psi(d,t,V)Q_1^{d,t-1}(L_0)\right) \\ &= E_P \sum_{t\in\tau,d\in\mathscr{D}} J_\psi(d,t,V)Q_1^{d,t}(L_0) - \tilde{h}(d,t,V)m_\psi(d,t,V). \end{aligned} \tag{2.17}$$

Since $0 = U(\Psi(P),Q,P)$, it follows from implicit differentiation that

$$\frac{d\Psi(P)}{dP} = -\frac{dU(\Psi(P),Q,P)}{d\Psi(P)}^{-1}\frac{dU(\psi,Q,P)}{dP}.$$

Straightforward computations shows that indeed

$$\begin{aligned} &-\frac{dU(\psi,Q,P)}{d\psi} \\ &= E_P\left\{\sum_{d\in\mathscr{D}}\sum_{t=1}^{\tau} h(d,t,V)\left(1-Q_1^{d,t-1}(L_0)\right)m_\psi(d,t,V)(1-m_\psi(d,t,V))\phi(d,t,V)\phi(d,t,V)^T\right\}. \end{aligned}$$

This proves the expression for $M(\psi,Q,P)$ in (2.14).

Next, note that $\frac{dU(\psi,Q,P)}{dP} = \frac{dU(\psi,Q,P)}{dQ^{L_0}}\frac{dQ^{L_0}}{dP} + \sum_{t\in\tau}\sum_{d\in\mathscr{D}}\frac{dU(\psi,Q,P)}{dQ_1^{d,t}}\frac{dQ_1^{d,t}}{dP}$. can be obtained by a simple application of the functional delta method method, using the efficient influence functions for Q^{L_0} and $Q_1^{d,t}$ obtained in Petersen et al. (2014). Therefore, we have

$$\begin{aligned} D^{L_0}(\psi,Q)(O) &\equiv \sum_{t\in\tau}\sum_{d\in\mathscr{D}} \tilde{h}(d,t,V)\left(1-Q_1^{d,t-1}(L_0)\right)\left(\frac{Q_1^{d,t}(L_0)-Q_1^{d,t-1}(L_0)}{1-Q_1^{d,t-1}(L_0)} - m_\psi(d,t,V)\right), \\ D^{d,t}(\psi,Q,g)(O) &= J_\psi(d,t,V)\sum_{k=1}^{t}\frac{I(\mathbf{A}_{k-1}=\mathbf{d}(\mathbf{L}_{k-1}))}{\mathbf{g}_k\left(\mathbf{A}_{k-1}=\mathbf{d}(\mathbf{L}_{k-1})\mid\mathbf{L}_{k-1}\right)}\left(Q_{k+1}^{d,t}(\mathbf{L}_k)-Q_k^{d,t}(\mathbf{L}_{k-1})\right). \end{aligned}$$

Now, we show the robustness property. If $Q = Q_0$, the result is trivial by definition of $\Psi(Q)$ and Q_k^t. We only need to check the second case. When $g = g_0$, at each t, the sum from $k = 1$ to $k = t$ forms a telescopic sum, leaving only the first and the last term.

Using expression in (2.17) to rewrite $D^{L_0}(\Psi(Q),Q)$, we have, up to a constant normalizing matrix,

$$\begin{aligned}
0 = P_0 D^*(\Psi(Q),Q,g_0) &= P_0 \sum_{t=1}^{\tau} \sum_{d\in\mathscr{D}} J_{\Psi(Q)}(d,t,V) Q_{1,0}^{d,t}(L_0) - P_0 \sum_{t=1}^{\tau} \sum_{d\in\mathscr{D}} J_{\Psi(Q)}(d,t,V) Q_1^{d,t}(L_0) \\
&\quad + P_0 \sum_{t=1}^{\tau} \sum_{d\in\mathscr{D}} J_{\Psi(Q)}(d,t,V) Q_1^{d,t}(L_0) - \tilde{h}(d,t,V) m_{\Psi(Q)}(d,t,V) \\
&= P_0 \sum_{t=1}^{\tau} \sum_{d\in\mathscr{D}} J_{\Psi(Q)}(d,t,V) Q_{1,0}^{d,t}(L_0) - \tilde{h}(d,t,V) m_{\Psi(Q)}(d,t,V) \\
&= U(\Psi(Q),Q_0,P_0).
\end{aligned}$$

Therefore, $\Psi(Q) = \Psi(Q_0)$.

Data Generating Distribution for the Simulation Study

The baseline covariate W consists of $W = (W_1,W_2,W_3,W_4)$, where $W_1 = I(30 \leq age \leq 39)$, $W_2 = I(age > 39)$, W_3 indicates sex, and $W_4 =$ disease stage.

To mimic the data generating process in a real life clinical cohort data, we include a monitoring variable M_t, which indicates whether the subject had come into the clinic at time t and thus have their CD4 measured and have the chance to switch regimen. Let $CD4_t^u$ denote the underlying true CD4 counts within a patient; the observed covariate $CD4_t$ equals this underlying $CD4_t^u$ only if a patient was seen at time t (i.e. $M_t = 1$). Subsequent CD4 values $CD4_t^u$ and death Y_t depend on the underlying CD4 counts, not the observed ones. The regimen status A_t^1 has non-zero probability of changing only if the patient was seen at time t, while the decision to change regimen status depends on the observed $CD4$ counts. This monitoring variable is excluded from the adjustment set in the analysis for two reasons: (1) it is an instrumental variable for the effect of switching on mortality; (2) its inclusion would produce positivity violations since subjects who were not seen at the clinic would have zero probability of switching regimen.

Let $\varepsilon_{1,t}$ and $\varepsilon_{2,t}$ be errors drawn from a standard normal distribution. The data gener-

ating distributions are given as follows:

$$W_1 \sim Bern(0.3),\ (W_2 \mid W_1 = 0) \sim Bern(0.5),\ W_3 \sim Bern(0.5),\ W_4 \sim Bern(0.3);$$

$$P(Y_t = 1 \mid W, CD4^u_t) = \text{expit}\left(-1 - 0.1W_1 - 0.1W_2 + 0.1W_3 - 0.2W_4 - 0.7CD4^u_{t-1} - 0.9A^1_{t-1}\right)$$

$$CD4^u_t = \begin{cases} \max\left(\min\left(\varepsilon_{1,t} - W_4, 4\right), -4\right) \text{ if } t = 0, \\ \max\left(\min\left(\varepsilon_{1,t} + 0.1W_1 - 0.1W_2 - 0.1W_3 - 0.5W_4 + 0.9CD4^u_{t-1} + A^1_{t-1}, 4\right), -4\right) \text{ if } t \geq 1; \end{cases}$$

$$P(M_t = 1 \mid W, CD4_{t-1}, A^1_{t-1}) =$$

$$\begin{cases} 1 \text{ if } t = 0, \\ \text{expit}\left(0.4 + 0.1W_1 - 0.2W_2 + 0.3W_3 + 0.1W_4 - 0.1CD4_{t-1} + 0.2A^1_{t-1}\right) \text{ if } t \geq 1; \end{cases}$$

$$CD4_t = \begin{cases} CD4^u_t \text{ if } M_t = 1, \\ CD4_{t-1} \text{ if } M_t = 0; \end{cases}$$

$$P(C^1_t = 1 \mid W, CD4_0) = \begin{cases} 0 \text{ if } t = 0 \\ 1 - \text{expit}\left(2 + 0.1W_1 + 0.2W_2 + 0.1W_3 + 0.1W_4 + 0.1CD4_0\right) \text{ if } t \geq 1; \end{cases}$$

$$C^2_t = I\left(M_{t-2} = 0, M_{t-1} = 0, M_t = 0\right)$$

$$P(A^1_t = 1 \mid M_t, A^1_{t-1}, W, CD4_t) =$$

$$\begin{cases} 1 \text{ if } t \geq 1 \text{ and } A^1_{t-1} = 1; \\ 0 \text{ if } t = 0 \text{ or if } t \geq 1 \text{ and } A^1_{t-1} = 0 \text{ and } M_t = 0; \\ \text{expit}\left(-0.5 + 0.1W_1 + 0.1W_2 + 0.2W_3 + 0.2W_4 - 1.5CD4_t + \varepsilon_{2,t}\right) \text{ otherwise.} \end{cases}$$

Chapter 3

Marginal Structural Models with Counterfactual Effect Modifiers: a twist to a familiar story

3.1 Introduction

In social and medical sciences, research questions often involve systematic comparison of the effectiveness of different exposures on a well-defined outcome of interest. But beyond the overall comparative effectiveness of the exposures on a diverse population, the researchers are also interested in identifying factors that modify the effect of the exposures. Will augmenting a failed citalopram regimen with other medications be beneficial (compared to switching to other medications entirely) for depression patients with certain medical or psychiatric history, but harmful for others? Can an aggressive course of cancer treatment be very effective (compared to a standard treatment) in reducing risk of metastasis on patients with a specific gene mutation, but makes little difference on the general patient population? This type of comparative effectiveness research (CER) often invokes knowledge about the pre-exposure individual characteristics that can potentially change the effect of the exposures. Consequently, a crucial component of CER is evaluating the modification of an exposure's effect by a given set of pre-exposure covariates (*effect modifiers*).

Marginal Structural Models (MSM), introduced by Robins (1997a), model the marginal distributions of an intervention-specific counterfactual mean outcome, possibly conditioning on a subset of pre-treatment covariates — the effect modifiers. The MSMs are useful tools for analyzing the causal effect of a time-varying treatment in the presence of time-varying confounding, as well as for studying the modification of these effects by pre-

treatment covariates (note that though MSM can also be used to study interactions, that application is beyond the scope of this article and we refer to**?** for a discussion on the delineation between effect modification and interaction). Generally, the effect modifiers of interest are cast as variables of the observed past, either as pre-treatment covariates (e.g. see Robins et al. (2000a) for an exposition on pre-treatment effect modifiers) or as variables in an observed history (see van der Laan, Petersen, and Joffe (2005) for a presentation of the so-called history-adjusted MSM). Yet, in some applications the effect modifiers of interest are in fact *counterfactual*. For example, consider an observational study where HIV-infected individuals are followed over time and their CD4 T-cell counts (among other time-varying covariates) are measured at regular intervals. Upon immunological failure of their first-line Antiretroviral Therapy (ART), patients are switched to a second-line ART within a certain period of time. For a specific first-line ART drug, say Zidovudine (AZT), one wishes to evaluate how the effect of delaying switching to second-line ART on mortality is modified by the CD4 counts measured at the time of first-line failure. However, suppose high viral load at diagnosis is highly predictive of receiving AZT as first line treatment. In this case, it is not suffice to perform the effect modification analysis conditioning on the observed first line CD4 and stratify on those who received AZT at first line, lest we assess the effect modification by CD4 counts at failure only among those with high viral load at diagnosis. Instead, one may cast the CD4 counts at failure as counterfactual variables under the intervention to set the first-line treatment to AZT, this way, the statistical estimand accounts for selection bias in the first line treatment assignment by mimicking that an ideal experiment where first line treatment was randomized. In another example, the baseline effect modifiers of interest may be missing at random, wherein the missingness may share common determinants with treatment and the outcome of interest. Consider, for instance, the STAR*D (Sequenced Treatment Alternatives to Relieve Depression) trial, a multi-level, longitudinal pragmatic trial of treatment strategies for major depression (`http://www.edc.gsph.pitt.edu/stard/`). Potential modifiers of a given level's treatment effect include the patient's response to previous level's treatment and the patient's psychiatric history measured at screening. After an initial screening, patients are enrolled into level 1 of the study, where everyone is assigned citalopram. At the start of each subsequent level, if a patient still remains in the study, then he is randomized into one of his chosen treatment strategies. At level 2, the patient can choose between augmenting the citalopram or switching to a new regimen (or to be randomized into either strategy if he chooses so). Suppose we wish to assess the effect modification of 2nd level's treatment by the depression symptoms measured at the exit of level 1. These symptom measures are obtained at clinical visits and level exit surveys. It is reasonable to believe that the more depressed patients and those less satisfied with their level 1 treatments will be less inclined to follow up with the surveys and visits or to report these symptom measures (we acknowledge the high possibly of missing not at random in this example, but that is beyond the

scope of this paper). Consequently, a simple complete-case analysis may fail to adjust for selection bias introduced by this missingness.

In this chapter, we investigate MSMs defined by counterfactual effect modifiers. We aim to make the following contributions to the literature. Firstly, we determine the identification of the causal dose-response curve and MSM parameters in this setting. Secondly, we establish the semiparametric efficiency theory for these statistical parameters, and present a substitution-based, semiparametric efficient and doubly robust estimator using the targeted maximum likelihood estimation methodology (TMLE, e.g. van der Laan and Rubin (2006), van der Laan and Rose (2011)). However, as we shall see, due to the form of the efficient influence curve, the implementation of this estimator may prove arduous in applications where V is high dimensional. To address this problem, our third contribution is a projected influence curve (and the corresponding TMLE estimator), which retains most of the robustness of its efficient peer and can be easily implemented in applications where the use of the efficient influence curve becomes taxing. In addition to these two robust estimators, we also present an inverse-probability-weighted (IPW) estimator (e.g. Robins (1997a), Hernan, Brumback, and Robins (2000)), and a non-targeted G-computation estimator (Robins (1986)).

This chapter is organized as follows. In section 3.2, we use a nonparametric structural equations framework (Pearl (2009)) to formulate the causal inference problem and determine the identifiability of the desired causal parameters from the observed data distribution. In section 3.3, we present the efficient influence curve for the parameter of interest under a saturated semiparametric model, as well as a projected influence curve. The robustness conditions for the efficient and the projected influence curves are established. In section 3.4, we present the construction of two TMLE estimators, one using the projected influence curve (and we call it the Δ-*TMLE*) and one using the efficient influence curve (and we call it the *full TMLE*, an IPW estimator, and a non-targeted G-computation estimator. In section 3.5, a simulation study demonstrates robustness properties of the Δ-TMLE. In section 3.6, we use STAR*D to illustrate the application of Δ-TMLE. A summarizing discussion concludes the chapter .

3.2 Parameters of MSM with Counterfactual Modifier

Consider a longitudinal data structure $O = (W, A_1, L_1, \ldots, A_K, L_K) \sim P_0$, where W encodes baseline covariates, A_t is the variable measured at time t that encodes the exposures of interest and censoring indicators, and L_t encodes covariates (including time-varying confounders) measured between A_t and A_{t+1}, including the outcome process of interest Y_t. Y_K is the final outcome of interest. For the sake of discussion, assume that Y_t is either a binary or a bounded continuous variable (without losing generality, we may assume

it's bounded between (0,1)). Suppose we wish to evaluate effect modification of A_t by a particular $V \subset W$, in an ideal experiment where a variable affecting V, call it Δ , had been set to a given value $\Delta = \delta$. For instance, if V is subject to missingness and Δ indicates whether V is measured, then one may wish to intervene to set $\Delta = 1$. In other applications, Δ may be a first-line treatment and V is a covariate measured between Δ and the second-line treatments A_t, and one is interested in effect modification by V had the first-line treatment been at a particular level $\Delta = \delta$. The other baseline covariates $W \backslash \{V\}$ are divided into those preceding Δ in our time-ordering — we call them W_1 — and those succeding Δ — we call them L_0. Consequently, the observed data structure becomes $O = (W_1, \Delta, V, L_0, A_1, L_1, \ldots, A_K, L_K)$.

In our STAR*D example, regular follow-up visits are conducted throughout each level. At each follow-up visit, covariates are collected, and the patient is subject to dropout, entering remission, or moving onto next level. Suppose we want to compare the effects of switching medication vs augmenting medication at level 2, on the chances of entering remission by the end of that level, among patients who failed the citalopram assigned to everyone at level 1. We use the discrete time scale of weeks. By study protocol, all subjects will have either entered remission (treatment success), moved onto the next level (treatment failure), or dropped out (right censoring), by the end of $K = 23$ weeks. A_1 encodes the treatment strategy received by a patient at level 2: this can be augmenting medication ($A_1 = 0$), switching medication ($A_1 = 1$), or to receive cognitive therapy ($A_1 = 2$); we are only interested in comparing switching medication vs augmenting medication. For $t \geq 2$, A_t is a counting process which drops to 0 if patient was censored by time t. L_t includes time-varying covariates such as visit statistics (time in level thus far, visit frequency, etc), side-effect burden and symptom measures at time t. L_t also contains two counting processes: the outcome process Y_t, which is a binary indicator for entering remission by time t, and a failure process E_t that jumps to 1 if a patient is moved to the next level, in which case the remission status will be zero for this level and the patient is considered non-censored (since the outcome was observed to be unsuccessful). Our final outcome of interest is Y_K — the remission status by end of 23 weeks. Once either the censoring process $A_{t\geq 2}$, the success process Y_t or the failure process E_t jumps, all subsequent variables are encoded by carrying forward the last observation. Baseline covariate W has variables collected at screening, as well as summary of patient's history throughout level 1 (recall that all patients are prescribed citalopram at level 1). To assess effect modification by the patient's response to previous level's treatment, V can be a psychiatric score taken at exit of level 1. Many such effect modifiers of interest are subject to missingness, we use Δ to indicate whether V is measured. In this case, W_1 consists of the covariates (and their missingness status) collected from enrollment up to level 1 exit, as well as level 1 exit summaries, such as frequency of visit and adherence to study protocol, that may affect missingness of V.

The time-ordering assumptions can be captured by a nonparametric structural equations model (NPSEM, Pearl (2009)):

$$W_1 = f_{W_1}(U_{W_1}); \quad \Delta = f_\Delta(W_1, U_\Delta); \quad V = f_V(W_1, \Delta, U_V); \quad L_0 = f_{L_0}(W_1, \Delta, V, U_{L_0});$$
$$A_t = f_{A_t}(W_1, \Delta, V, \mathbf{A}_{t-1}, \mathbf{L}_{t-1}, U_{A_t}); \quad L_t = f_{L_t}(W_1, \Delta, V, \mathbf{A}_t, \mathbf{L}_{t-1}, U_{L_t}), \tag{3.1}$$

where the boldface notation means $\mathbf{X}_s = (X_0, \ldots, X_s)$, for $X = L$ or A (for the latter the indexing start at 1). We will also use the shorthand $\mathbf{X} = \mathbf{X}_K$, and $\mathbf{X}_s^t = (X_t, \ldots, X_s)$. Variables with degenerate indices, such as -1, are empty sets. This framework assumes that each variable X in the observed data structure is an unknown deterministic function of observed variables and some unmeasured exogenous random factors U. From here on, we will refer to the observed variables in the input of f_X as the parents of X. This causal model defines a random variable with distribution $P_{O,U}$ on a unit.

Consider a hypothetical experiment where the research could enforce $\Delta = \delta$ and $\mathbf{A} = \mathbf{a}$ on all units. Δ and $\mathbf{A}$ are the intervention variables. In the Star*D example, where the natural intervention on Δ is 1, this would entail taking precautions to enforce measurement of V, assigning the strategy, say, a_1 equals augmenting medication to all patients, and preventing dropouts throughout the level. Use $V(\delta)$, $L_0(\delta)$, $L_t(\mathbf{a})$ and $Y_K(\mathbf{a})$ to denote the counterfactual modifier, non-intervention covariates and final outcome, respectively, under the intervention of setting $(\Delta, \mathbf{A}) = (\delta, \mathbf{a})$. Note that the intervention of $\Delta = \delta$ on L_t and Y_K are suppressed in this notation. To emphasize that in our question of interest the level δ does not change, we shall give it a constant values, say $\delta = 1$, to simplify notations later on.

For a given exposure of interest $\mathbf{A} = \mathbf{a}$ and an effect modifier value $V = v$, $\rho^F_{\mathbf{a},v}(P_{O,U}) \equiv E(Y_K(\mathbf{a}) \mid V(1) = v)$ is the mean counterfactual outcome under exposure $\mathbf{A} = \mathbf{a}$ for individuals with characteristic $V = v$, if $\Delta = 1$. A causal quantity of interest is the counterfactual conditional dose–response curve $\{\rho^F_{\mathbf{a},v}(P_{O,U}) : (\mathbf{a}, v) \in \mathscr{A} \times \mathscr{V}\}$, where $\mathscr{V}$ is the outcome space of V and $\mathscr{A}$ is a set of interventions on $\mathbf{A}$ we wish to compare. In our STAR*D example, $\mathscr{A} = \{(a_1 = 0, a_{t\geq 2} = 1), (a_1 = 1, a_{t\geq 2} = 1)\}$, even though the outcome space of the treatment node is $a_1 \in \{0, 1, 2\}$. This dose-response curve can be summarized by a working MSM $\{m_\psi(\mathbf{a}, v) : \psi \in S \subseteq \mathbb{R}^d\}$. Since we are considering a final outcome that is either binary or bounded in $(0, 1)$, the range of our working model m falls within the unit interval. For a given kernel weight function $h(\mathbf{a}, v)$, the causal parameter of interest is defined as

$$\Psi^F(P_{O,U}) = \arg\min_{\psi \in S} \Bigg\{ - \sum_{(\mathbf{a},v) \in \mathscr{A} \times \mathscr{V}} p(V(\Delta = 1) = v) h(\mathbf{a}, v) \times \Big\{ \rho^F_{\mathbf{a},v}(P_{O,U}) \log m_\psi(\mathbf{a}, v) + \big(1 - \rho^F_{\mathbf{a},v}(P_{O,U})\big) \log\big(1 - m_\psi(\mathbf{a}, v)\big) \Big\} \Bigg\} \tag{3.2}$$

In words, $\Psi^F(P_{O,U})$ yields the best weighted approximation of the counterfactual conditional dose-response curve, according to the user-specified quasi-loglikelihood loss, kernel weights and MSM. Because we are not assuming that that the MSM is correctly specified (it is only an approximation to the truth), the definition of this parameter hinges upon the choice of model and the kernel weights (Neugebauer and van der Laan (2007)). The main difference between this causal parameter and its analog in the familiar non-counterfactual modifier story (e.g. Robins (1997a), Neugebauer and van der Laan (2007)) resides in the nature of the conditioning variable in the dose-response curve, $E(Y_K(\mathbf{a}) \mid V(1) = v)$ vs $E(Y_K(\mathbf{a}) \mid V = v)$, and the distribution of V used in the definition of optimizing function, $p(V(1) = v)$ vs $p(V = v)$. For this reason, the parameter in (3.2) and its non-counterfactual modifier analog can be staged on a common platform when comparing multiple effect modifiers, some of which are missing.

To identify (3.2) from the data generating distribution P_0, we make a positivity assumption and the Sequential Randomization Assumption (SRA, derived by Robins (1997b)). Specifically, under the positivity assumption, there exists $\{\alpha_\Delta, \alpha_t : t\} \in (0,1)$ such that $\alpha_\Delta \leq p_0(\Delta = 1 \mid W_1)$ and $p_0(A_t = a_t \mid \cdot) < 1 - \alpha_t$, for all t and $\mathbf{a} \in \mathscr{A}$, almost everywhere. The SRA assumes that $\Delta \perp (W_1, V(1), L_0(1), \{L_t(\mathbf{a}) : t\})$, given parents of Δ, and $A_t \perp (W_1, V(1), L_0(1), \{L_t(\mathbf{a}) : t\})$, given parents of A_t. Under these conditions, the joint distribution $(W_1, V(1), L_0(1), \{L_t(\mathbf{a}) : t\})$ is identifiable from the observed data distribution P_0. In our STAR*D example, the plausibility of the SRA can be fortified by measuring enough confounders of the modifier's missingness, the treatment selection, and the censoring mechanism.

By straightforward calculations, the SRA allows us to identify $p(V(1) = v)$ as

$$\gamma_v(P_0) \equiv E_{W_1,0}\{p_0(v \mid \Delta = 1, W_1)\}, \tag{3.3}$$

and the counterfactual mean outcome $\rho^F_{\mathbf{a},v}(P_{O,U})$ as

$$\rho_{\mathbf{a},v}(P_0) \equiv E_{W_1,0}\left\{ \frac{p_0(V = v \mid \Delta = 1, W_1)}{E_{W_1,0}\{p_0(v \mid \Delta = 1, W_1)\}} \times Q^{\mathbf{a},1}_{t=0}(P_0)(V = v, W_1) \right\}, \tag{3.4}$$

where, for $t = 0, \ldots, K$,

$$Q^{\mathbf{a},\delta}_t(P_0)(\mathbf{L}_{t-1}, V, W_1) \equiv \sum_{\mathbf{l}^t_K} y_K \left(\prod_{j=t}^{K} p_0(l_j \mid \mathbf{A}_j = \mathbf{a}_j, \mathbf{L}_{t-1}, \mathbf{l}^t_{j-1}, V, \Delta = \delta, W_1) \right). \tag{3.5}$$

Under the SRA,

$$Q^{\mathbf{a},1}_t(P_0)(\mathbf{L}_{t-1}, V = v, W_1) = E_{P_0}\Big(Y_K(\mathbf{a}) \Big| \mathbf{A}_t = \mathbf{a}_t, \mathbf{L}_{t-1}, V = v, \Delta = 1, W_1 \Big).$$

In other words, $Q^{\mathbf{a},1}_t(P_0)(\cdot)$ equals the counterfactual conditional mean outcome of $Y_K(\mathbf{a})$, given an observed history at time t that has thus far obeyed the intervention of interest.

For sake of interpretation, it is also useful to rewrite (3.4) as

$$\rho_{\mathbf{a},v}(P_0) = E_0\left\{ \frac{1/p_0(\Delta=1 \mid W_1)}{E_{W_1,0}\left(1/p_0(\Delta=1 \mid W_1) \mid V=v,\Delta=1\right)} Q_{t=0}^{\mathbf{a},1}(P_0)(L_0,V=v,W_1) \middle| V=v,\Delta=1 \right\}. \quad (3.6)$$

The weight $\frac{1/p_0(\Delta=1|W_1)}{E_{W_1,0}(1/p_0(\Delta=1|W_1)|V=v,\Delta=1)}$ adjusts for potential selection bias induced by Δ. Indeed, when this weight equals 1, (3.6) and (3.4) are equivalent to the estimands in a complete-case analysis, when Δ represents missingness, or in an unadjusted stratified analysis, when Δ is a first-line treatment.

Combining (3.3) and (3.4), the causal MSM parameter $\Psi^F(P_{O,U})$ in (3.2) identifies to

$$\psi_0 \equiv \Psi(P_0) \equiv \arg\min_{\psi\in S}\left\{ -\sum_{(\mathbf{a},v)\in\mathscr{A}\times\mathscr{V}} \gamma_v(P_0)h(\mathbf{a},v) \times \left\{ \rho_{\mathbf{a},v}(P_0)\log m_\psi(\mathbf{a},v) + (1-\rho_{\mathbf{a},v}(P_0))\log\left(1-m_\psi(\mathbf{a},v)\right)\right\}\right\}. \quad (3.7)$$

At this juncture, for a more concrete discussion we consider the following MSM

$$m_\psi(\mathbf{a},v) = \operatorname{expit}\left(\psi\cdot\phi(\mathbf{a},v)\right), \quad (3.8)$$

where $\phi(\mathbf{a},v)$ is the vector of linear predictors in the generalized linear model. The linear predictors are function of $\mathbf{a}$ and v. We emphasize that the methods in the next sections are easily modified to other MSM.

In the forthcoming sections, we study the statistical inference of $\Psi(P_0)$.

Notations

Before we proceed, let us introduce some useful definitions and notations. Let $\mathscr{M}$ be a saturated semiparametric model containing our data generating distribution P_0. The parameter of interest in (3.7) is the map $P \mapsto \Psi(P)$, from $\mathscr{M}$ to $\mathbb{R}^d$, evaluated at P_0.

Suppose we observe n i.i.d. copies of $O \sim P_0$. Let P_n denote the empirical distribution of this sample. For a function f of O, we will write $P_n f \equiv \frac{1}{n}\sum_{i=1}^n$, and for a distribution P, we will write $Pf = E_P f(O)$.

We generalize the definitions in (3.5) to any $P \mapsto Q_t^{\mathbf{a},1}(P)$ on $\mathscr{M}$, for $t \le K$. At $t = K+1$, we write $Q_{K+1}^{\mathbf{a},1}(P)(O) \equiv Y_K$. Bang and Robins (2005) noted the recursive property

$$Q_t^{\mathbf{a},1}(P)(\mathbf{L}_{t-1},V,W_1) = E_P\left[Q_{t+1}^{\mathbf{a},1}(P)(\mathbf{L}_t,V,W_1) \middle| \mathbf{A}_t=\mathbf{a}_t,\mathbf{L}_{t-1},V,\Delta=1,W_1\right], \quad (3.9)$$

for $t = 0, \ldots, K$. This will prove useful in our upcoming endeavor. We also adopt the notations Q^{W_1} for the marginal distribution of W_1, $Q^V(P)$ for the conditional distribution $P(V \mid W_1, \Delta = 1)$, and $Q \equiv \left(Q^{W_1}(P), Q^V(P), Q_t^{\mathbf{a},1}(P) : t = 0, \ldots K\right)$. We write $g^\Delta(P)$ for the conditional probability $P(\Delta = 1 \mid W_1)$, $g^A(P)$ for the treatment allocation probabilities $P(A_t \mid \mathbf{A}_{t-1}, \mathbf{L}_{t-1}, V, \Delta = 1, W_1)$, and $g \equiv (g^\Delta, g^A)$. When referring to a generic $P \in \mathscr{M}$, we may sometimes write Q and g in place of $Q(P)$ and $g(P)$, similarly for their respective components; when referring to the functions at the data-generating distribution P_0, we may sometimes write Q_0 and g_0, in place of $Q(P_0)$ and $g(P_0)$.

3.3 A Tale of Two Influence Curves

The first leg of our journey is determining the so-called Efficient Influence Curve (EIC) for our parameter of interest. From a fundamental result in Bickel et al. (1997), under standard regularity conditions, the variance of the canonical gradient of Ψ at P_0 provides a generalized Cramer-Rao lower bound for any regular and asymptotically linear estimators of $\Psi(P_0)$. Therefore, this canonical gradient is a vital ingredient in building asymptotically linear and efficient estimators; fittingly, it is also commonly known as the EIC. For parameters in causal inference and missing data applications (such as our case), the EIC also provides insights into the potential robustness against model misspecifications. In section 3.3, we determine the EIC of Ψ under $\mathscr{M}$.

However, as we shall see, in spite their theoretical prowess, estimators which use the EIC will be difficult to implement in practice when the dimension of V is high. To solve this problem, in section 3.3 we present a projection of the EIC onto a model where g^Δ is known; we refer to it as the (projected-IC). This projected-IC retains most of the robustness properties of its efficient peer while altogether avoiding estimation of the components relating to V, hence making a compelling case for trading full efficiency for practically more attainable estimators in the case of high-dimensional V.

Recall that $\Psi(P)$ optimizes a function of $\gamma_v(P)$ and $\rho_{\mathbf{a},v}(P)$. Note also that $\rho_{\mathbf{a},v}(P) = \frac{\eta_{\mathbf{a},v}(P)}{\gamma_v(P)}$, where $\eta_{\mathbf{a},v}(P) = E_P\left\{Q^V(V = v \mid \Delta = 1, W_1) \times Q_{t=0}^{\mathbf{a},1}(V = v, W_1)\right\}$. We will make use of the following useful characterizations for $\Psi(P)$:

Remark 3.1. *For $m_\psi(\mathbf{a},v) = \operatorname{expit}(\psi \cdot \phi(\mathbf{a},v))$ and $\Psi(P)$ defined as in (3.7), we have*

$$\begin{aligned}
0 &= \sum_{(\mathbf{a},v)\in\mathscr{A}\times\mathscr{V}} h(\mathbf{a},v)\phi(\mathbf{a},v)\gamma_v(P)\left\{\frac{\eta_{\mathbf{a},v}(P)}{\gamma_v(P)} - m_{\Psi(P)}(\mathbf{a},v)\right\} \\
&= E_P\left\{\sum_{(\mathbf{a},v)\in\mathscr{A}\times\mathscr{V}} \tilde{h}(\mathbf{a},v)Q^V(V=v\mid\Delta=1,W_1)\left\{Q_{t=0}^{\mathbf{a},1}(V=v,W_1) - m_{\Psi(P)}(\mathbf{a},v)\right\}\right\} \\
&= E_P\left\{\frac{I(\Delta=1)}{g^\Delta(1\mid W_1)}\sum_{\mathbf{a}\in\mathscr{A}} \tilde{h}(\mathbf{a},V)\left\{Q_{t=0}^{\mathbf{a},1}(V,W_1) - m_{\Psi(P)}(\mathbf{a},V)\right\}\right\},
\end{aligned}$$

where $\tilde{h}(\mathbf{a},V) \equiv h(\mathbf{a},V)\phi(\mathbf{a},V)$. The computations are straightforward, and we left them in the Appendix for reference.

Efficient Influence Curve

From first equality in remark 3.1, we can obtain the EIC for $\Psi(P)$ via implicit differentiation. We formally state the result here and leave the proof in the Appendix.

Lemma 3.1 (Efficient Influence Curve).

Consider $\Psi : \mathscr{M} \to \mathbb{R}^d$ as defined in (3.7). Suppose that the following $k \times k$ matrix is invertible at $(\psi,P) = (\Psi(P),P)$:

$$\begin{aligned}
M(\psi,P) &= \frac{\partial}{\partial\psi}\sum_{(\mathbf{a},v)\in\mathscr{A}\times\mathscr{V}} \gamma_v(P)h(\mathbf{a},v)\phi(\mathbf{a},v)\left\{\rho_{\mathbf{a},v}(P) - m_\psi(\mathbf{a},v)\right\} \\
&= P\left\{\frac{I(\Delta=1)}{g^\Delta(1\mid W_1)}\sum_{\mathbf{a}\in\mathscr{A}} h(\mathbf{a},,V)\phi(\mathbf{a},V)\phi(\mathbf{a},V)^\top m_\psi(\mathbf{a},V)\left[1 - m_\psi(\mathbf{a},V)\right]\right\}. \qquad (3.10)
\end{aligned}$$

The efficient influence curve of Ψ at $P \in \mathscr{M}$ is given by

$$M(\Psi(P),P)^{-1} D^*(Q,g,\Psi(P)),$$

where

$$D^*(Q,g,\psi) = \sum_{t=0}^{K} D_t^*(Q,g) + D_V^*(Q,g,\psi) + D_{W_1}^*(Q,\psi) \tag{3.11}$$

with

$$D_t^* \equiv \frac{I(\Delta=1)}{g^{\Delta}(1 \mid W_1)} \sum_{\mathbf{a}\in\mathscr{A}} \tilde{h}(\mathbf{a},V) C_t^{\mathbf{a}} \left\{ Q_{t+1}^{\mathbf{a},1}(\mathbf{L}_t,V,W_1) - Q_t^{\mathbf{a},1}(\mathbf{L}_{t-1},V,W_1) \right\},$$

$$D_V^* \equiv \frac{I(\Delta=1)}{g^{\Delta}(1 \mid W_1)} \sum_{(\mathbf{a},v)\in\mathscr{A}\times\mathscr{V}} \left\{ \tilde{h}(\mathbf{a},v) \left(Q_{t=0}^{\mathbf{a},1}(V=v,W_1) - m_\psi(\mathbf{a},v) \right) \right.$$
$$\left. \times \left(I(V=v) - Q^V(v \mid \Delta, W_1) \right) \right\}$$

$$D_{W_1}^* \equiv \sum_{(\mathbf{a},v)\in\mathscr{A}\times\mathscr{V}} \tilde{h}(\mathbf{a},v) Q^V(v \mid \Delta=1, W_1) \left\{ Q_{t=0}^{\mathbf{a},1}(V=v,W_1) - m_\psi(\mathbf{a},v) \right\},$$

with $C_t^{\mathbf{a}} = \frac{I(\mathbf{A}_t=\mathbf{a}_t)}{\prod_{j=1}^{t} g^A\left(A_j=a_j \mid \mathbf{A}_{j-1}=\mathbf{a}_{j-1}, \mathbf{L}_{j-1}, V, \Delta=1, W_1\right)}$, *for* $t = 1,\ldots,K$, *and* $C_t^{\mathbf{a}} = 1$ *for* $t = 0$.

Proof. The proof is given in the Appendix. □

Three major differences exist between D^* and its analog in the non-counterfactual modifier story. Firstly, per characterization of $\Psi(P)$, the component $D_{W_1}^*$, corresponding to the marginal distribution of the pre-intervention variables, is weighted by Q^V. Secondly, D^* has an additional component D_V^* corresponding to the conditional distribution of V. The weight $\left[Q_{t=0}^{\mathbf{a},1}(v,W_1) - m_\psi(\mathbf{a},v) \right]$ in $D_V^*(P)$ underscores the role of V as an effect modifier, besides being a post-intervention (on Δ) variable. Thirdly, the MSM parameter ψ plays a role in both $D_{W_1}^*$ and D_V^*.

Lemma 3.2 (Double robustness of the efficient influence curve).
If $Q = Q_0$ *or* $g = g_0$, *then* $P_0 D^*(Q,g) = 0$ *implies* $\Psi(Q) = \Psi(Q_0)$.

Proof. The proof is given in the Appendix □

Note that in the condition (a) of lemma 3.2, $Q^V = Q^V(P_0)$ can be relaxed to

$$E_P\left\{ h(\mathbf{a},V)\phi(\mathbf{a},V) \left[Q_{t=0}^{\mathbf{a},1}(V,W_1) - m_\psi(\mathbf{a},V) \right] \middle| \Delta=1, W_1 \right\}\Big\}$$
$$= E_{P_0}\left\{ h(\mathbf{a},V)\phi(\mathbf{a},V) \left[Q_{t=0}^{\mathbf{a},1}(V,W_1) - m_\psi(\mathbf{a},V) \right] \middle| \Delta=1, W_1 \right\}\Big\}.$$

Remark 3.2. ***Dimensionality of*** V ***and Implementation:*** *When V is high-dimensional (or continuous), one may wish to avoid explicit estimation of Q^V in D_V^* and $D_{W_1}^*$ of (3.11). To this end, rewrite $D_V^*(Q,g,\psi)$ and $D_{W_1}^*(Q,\psi)$ as*

$$\begin{aligned}
&D_V^*(Q,g,\psi)\\
&= \frac{I(\Delta=1)}{g^\Delta(1\mid W_1)} \sum_{\mathbf{a}\in\mathscr{A}} \Big\{ \tilde{h}(\mathbf{a},V)\left[Q_{t=0}^{\mathbf{a},1}(V,W_1) - m_\psi(\mathbf{a},V)\right]\\
&\qquad\qquad - E_P\left\{ \tilde{h}(\mathbf{a},V)\left[Q_{t=0}^{\mathbf{a},1}(V,W_1) - m_\psi(\mathbf{a},V)\right] \Big| \Delta, W_1 \right\} \Big\},\\
&D_{W_1}^*(Q,\psi) = \sum_{\mathbf{a}\in\mathscr{A}} E_P\left\{ \tilde{h}(\mathbf{a},V)\left[Q_{t=0}^{\mathbf{a},1}(V,W_1) - m_\psi(\mathbf{a},V)\right] \Big| \Delta=1, W_1 \right\}.
\end{aligned}$$

A regression-based estimator (parametric or data-adaptive) can be used to directly estimate the conditional expectations with respect to Q^V in D_V^* and in $D_{W_1}^*$. For final evaluation of the target parameter ψ, we must solve the estimating equation $D_{W_1}^*$ in the variable ψ. This may be accomplished via powerful numerical tools. However, as this may increase the computational expense as the regression-based estimator becomes more data-adaptive.

This dilemma motivates us to consider trading the fully efficiency D^* for an influence curve that retains most of the robustness properties while altogether avoiding estimating of the components relating to V. We consider this option in section 3.3.

Projected Influence Curve

As motivated by remark 3.2, when V is high-dimensional, we may instead consider a projected influence curve which still retains most of the robustness of D^*.

Lemma 3.3 (Projected Influence Curve).

Consider the setup in lemma 3.1. Up to a normalizing matrix $M(\Psi(P),P)^{-1}$, the following function is a gradient for Ψ at P under the model $\mathscr{M}_{g^\Delta}$, where g^Δ is known:

$$D^\Delta(Q,g,\psi) = \sum_{t=1}^{K} D_t^*(Q,g) + D_W^\Delta(Q,g,\psi),$$

where

$$D_W^\Delta \equiv \frac{I(\Delta=1)}{g^\Delta(1\mid W_1)} \sum_{\mathbf{a}\in\mathscr{A}} \tilde{h}(\mathbf{a},V)\left[Q_{t=1}^{\mathbf{a},1}(L_0,V,W_1) - m_\psi(\mathbf{a},V)\right]. \tag{3.12}$$

In particular, it is a valid estimating function for $\Psi:\mathscr{M}\to\mathbb{R}^d$.

Moreover, if $g^\Delta = g^\Delta(P_0)$ and either $Q_t^{\mathbf{a},1} = Q_t^{\mathbf{a},1}(P_0)$ or $g^A = g^A(P_0)$, then $P_0 D^\Delta(Q,g) = 0$ implies $\Psi(Q) = \Psi(Q_0)$.

Proof. The proof is given in the Appendix. □

At its face value, the proposed D^{Δ} may seem less robust than D^{*}, as it always relies on consistent estimation of $g^{\Delta}(P_0)$. However, as we noted in remark 3.2, when V is high-dimensional, there are more standard machine learning algorithms available for estimation of g^{Δ}.

3.4 Statistical Inference

With the two influence curves D^{*} and D^{Δ} under our belt, in section 5.4 we will build two robust, substitution-based, asymptotically linear estimators via the targeted maximum likelihood estimation (TMLE) methodology. In section 3.4, we describe an inverse-probability-weighted (IPW) estimator that is most commonly used in the literature for estimating coefficients in an MSM. It is easier to implement and may be more intuitive than the robust estimators; however, its consistency relies solely on the correct estimation of g_0, and may suffer stability problems when the weights are extreme. Under standard regularity and empirical process conditions (detailed in e.g. Bickel et al. (1997)), both the TMLE and IPW are asymptotically linear, hence allowing influence curve-based estimate for the standard errors. In section 3.4, we describe a non-targeted substitution estimator which utilizes a non-targeted MLE estimate of Q_0 (or of $Q_t^{\mathbf{a},1}(P_0)$ and $g^{\Delta}(P_0)$). This estimator is biased if these non-targeted MLE are not consistent.

For most of the estimators below, we first need to procure estimators $g_n = \left(g_n^{\Delta}, g_n^{A}\right)$ of g_0. The marginal distribution of W_1 will always be estimated by the empirical marginal distribution. For a given estimator ψ_n of ψ_0, we will use

$$M(\psi_n) \equiv P_n \left\{ \frac{I(\Delta = 1)}{g_n^{\Delta}(1 \mid W_1)} \sum_{\mathbf{a} \in \mathscr{A}} h(\mathbf{a}, V) \phi(\mathbf{a}, V) \phi(\mathbf{a}, V)^{\top} m_{\psi_n}(\mathbf{a}, V) \left[1 - m_{\psi_n}(\mathbf{a}, V)\right] \right\}$$

to estimate the normalizing matrix.

Targeted Maximum Likelihood Estimator

In a traditional non-targeted MLE (like those in section 3.4), relevant parts of the likelihood are estimated either by stratification (nonparametric MLE), by fitting a parametric statistical model, or by using a machine-learning algorithm. These likelihood estimates are then used to evaluate the parameter of interest. As the number of potential confounders increase, these methods may break down due to curse of dimensionality, or yield a bias–variance trade off that is not the most optimal for the parameter of interest (which is a lower-dimensional object than the likelihood components). A targeted MLE adds an updating (targeting) step to the likelihood estimation process that aims to target the fit towards

the parameter of interest, and provide potential robustness and semiparametric efficiency gains. As a result of this targeting step, the final likelihood estimate (coupled with the substitution-based parameter estimate) satisfies a user-chosen score equation, hence also allowing inference based on the Central Limit Theorem. We refer to van der Laan and Rose (2011) for the general methodology. Here, we construct two targeted estimators using D^{Δ} and D^{*}.

Both targeted estimators involve sequentially updating initial estimates of the Q components by finding a best fluctuation along a submodel through a given initial estimate. We gather the following two ingredients before proceeding. Regarding $Q_t^{\mathbf{a},1}$ as a conditional expectation of $Q_{t+1}^{\mathbf{a},1}$, we use the quasi loglikelihood loss function for $Q_t^{\mathbf{a},1}$:

$$L\left(Q_t^{\mathbf{a},1}\right) = -I(\mathbf{A}_t = \mathbf{a}_t)\left\{\log\left(Q_t^{\mathbf{a},1}\right)^{Q_{t+1}^{\mathbf{a},1}} + \log\left(1 - Q_t^{\mathbf{a},1}\right)^{\left(1-Q_{t-1}^{\mathbf{a},1}\right)}\right\}. \tag{3.13}$$

For a given (Q,g), and each $t = 0,\ldots,K$, consider the d-dimensional working submodel $\{Q^t(\varepsilon) : \varepsilon\}$, with

$$Q_t^{\mathbf{a},1}(\varepsilon) = \operatorname{expit}\left(\operatorname{logit} Q_t^{\mathbf{a},1} + \varepsilon \cdot \tilde{h}(\mathbf{a},V)\frac{I(\Delta=1)}{g^{\Delta}(1 \mid W_1)}C_t^{\mathbf{a}}\right). \tag{3.14}$$

This submodel satisfies $\langle \frac{d}{d\varepsilon}\sum_{\mathbf{a}} L\left(Q_t^{\mathbf{a},1}(\varepsilon)\right)|_{\varepsilon=0}\rangle \supset \langle D_t^*(Q,g)\rangle$, where $\langle x\rangle$ represents the linear span of a vector x.

TMLE using D^{Δ}

1. Start at $t = K$: regress Y_K on $(\mathbf{L}_{K-1}, \mathbf{A}_K, L_0, V, W_1)$, among observations with $\Delta = 1$, and then evaluate at $\mathbf{A}_K = \mathbf{a}_K$ to obtain an initial estimator $Q_{t=K,n}^{\mathbf{a},1}$ of $Q_{t=K}^{\mathbf{a},1}(P_0)$. The optimal fluctuation amount around this initial estimate is given by

$$\varepsilon_{K,n} \equiv \arg\min_{\varepsilon} \sum_{\mathbf{a}} P_n L\left(Q_{t=K,n}^{\mathbf{a},1}(\varepsilon)\right).$$

   This can be implemented by creating one row for each individual with $\Delta_i = 1$ and each $\mathbf{a} \in \mathscr{A}$, and fitting a weighted logistic regression of Y_K on the multivariate covariate $\frac{\phi(\mathbf{a},V)}{g_n^{\Delta}(1|W_1)}C_K^{\mathbf{a}}(g_n)$ on these observations with weights $I(\mathbf{A}_K = \mathbf{a}_K)h(\mathbf{a},V)$ and offset $Q_{t=K,n}^{\mathbf{a},1}(\mathbf{L}_{K-1},V,W_1)$. Update the initial estimator using $Q_{t=K,n}^{*,\mathbf{a},1} \equiv Q_{t=K,n}^{\mathbf{a},1}(\varepsilon_{K,n})$.

2. At each subsequent step $t = K-1,\ldots,1$, we have thus far obtained an updated estimator $Q_{t+1,n}^{*,\mathbf{a},1}$ for each individual with $\Delta_i = 1$ and each $\mathbf{a} \in \mathscr{A}$. Regress $Q_{t+1,n}^{*,\mathbf{a},1}$ on $(\mathbf{L}_{t-1}, \mathbf{A}_t, L_0, W_1, V)$ among observations with $\Delta = 1$ and evaluate at $\mathbf{A}_t = \mathbf{a}_t$ to

get an initial estimator $Q_{t,n}^{\mathbf{a},1}$ of $Q_t^{\mathbf{a},1}(P_0)$. The optimal fluctuation amount around this initial estimator is given by $\varepsilon_{t,n} = \arg\min_\varepsilon \sum_{\mathbf{a}} P_n L\left(Q_{t,n}^{\mathbf{a},1}(\varepsilon)\right)$, and can be obtained in an analogous manner to step 1. The updated estimator is $Q_{t,n}^{*,\mathbf{a},1} \equiv Q_{t,n}^{\mathbf{a},1}(\varepsilon_{t,n})$.

3. After sequentially performing step (2) in order of decreasing t, we now have a targeted estimator $Q_{t=1,n}^{*,\mathbf{a},1}$ of $Q_{t=1}^{\mathbf{a},1}(P_0)$. Obtain $\psi_n^{\Delta,TMLE}$ by fitting a weighted logistic regression of $Q_{t=1,n}^{*,\mathbf{a},1}(V,W_1)$ on $\phi(\mathbf{a},V)$, with weights $h(\mathbf{a},V)\frac{I(\Delta=1)}{g_n^\Delta(1|W_1)}$. We call this estimator the Δ-*TMLE*.

By construction, $P_n D^\Delta\left(Q_{t,n}^{*,\mathbf{a},1}, g_n, \psi_n^{\Delta,TMLE}\right) = 0$. From lemma 3.3, $\psi_n^{\Delta,TMLE}$ is an unbiased estimator of ψ_0 if either (1) g_n^Δ and $Q_{t,n}^{*,\mathbf{a},1}$ for $t = 1,\ldots,K$ are consistent, or (2) g_n is consistent. Compared to the full TMLE using D^* in the next section, this estimator is particularly appealing when V is high-dimensional, and still provides more robustness protection than the estimators in sections 3.4 and 3.4. Moreover, under standard regularity and empirical process conditions, $\psi_n^{\Delta,TMLE}$ is asymptotically linear with influence curve $M(\Psi(P_0))^{-1}D^\Delta(P_0)$. The asymptotic covariance of $\sqrt{n}(\psi_n^{\Delta,TMLE} - \psi_0)$ can be estimated by the the sample covariance matrix $\Sigma_n^{\Delta,TMLE}$ of $M(\psi_n^{\Delta,TMLE})^{-1}D^\Delta\left(Q_{t,n}^{*,\mathbf{a},1}, g_n, \psi_n^{\Delta,TMLE}\right)$.

TMLE using D^*

To use D^*, we also consider the loss function $L(Q^V) \equiv -\log Q^V(V \mid \Delta = 1, W_1)$, and a d-dimensional fluctuation model through a given Q^V at $\varepsilon = 0$ given by

$$Q^V(\varepsilon)(V \mid \Delta = 1, W_1) \equiv \frac{Q^V(V \mid \Delta = 1, W_1)\exp\left[\varepsilon \cdot B(Q,\psi)(W_1,V)\right]}{\sum_v Q^V(v \mid \Delta = 1, W_1)\exp\left[\varepsilon \cdot B(Q,g,\psi)(W_1,v)\right]},$$

where $B(Q,\psi)(W_1,V) \equiv \sum_{\mathbf{a}} \tilde{h}(\mathbf{a},V)\left\{Q_{t=0}^{\mathbf{a},1}(V,W_1) - m_\psi(\mathbf{a},V)\right\}$. It is easy to verify that $\langle \frac{d}{d\varepsilon}\frac{I(\Delta=1)}{g_\Delta(1|W_1)} L(Q^V(\varepsilon))|_{\varepsilon=0}\rangle \supset \langle D_V^*(Q,g,\psi)\rangle$. The targeted estimator which uses the D^* will do so via estimators for Q^V, instead of via estimators for a conditional mean with respect to Q^V as discussed in remark 3.2. This way, the estimates for ψ_0 can be easily obtained by fitting a weighted regression.

1. Perform steps (1) and (2) over $t = K,\ldots,0$ in section 3.4 to obtain a targeted estimator $Q_{t=0,n}^{*,\mathbf{a},1}$.

2. Let Q_n^V be an estimator of $Q^V(P_0)$. For each $\mathbf{a} \in \mathscr{A}$, $v \in \mathscr{V}$ and individual i, create a row of data consisting of $Q_{t=0,n}^{*,\mathbf{a},1}(V = v, W_1)$, $h(\mathbf{a},v)$, $\phi(\mathbf{a},v)$ and $Q_n^V(v \mid \Delta = 1, W_1)$.

Obtain a first-iteration estimator ψ_n^1 of ψ_0 by fitting a weighted logistic regression of $Q_{t=0,n}^{*,\mathbf{a},1}(V=v,W_1)$ on $\phi(\mathbf{a},v)$, with weights $h(\mathbf{a},v)\times Q_n^V(v\mid\Delta=1,W_1)$, on this pooled data.

3. Given ψ_n^1 obtained in step (3), we update the estimator for $Q^V(P_0)$ as follows. Using previously obtained $Q_{t=0,n}^{*,\mathbf{a},1}$, g_n, and ψ_n^1, the optimal fluctuation amount around the initial Q_n^V is given by $\varepsilon_n^V=\arg\min_\varepsilon P_n\frac{I(\Delta=1)}{g_n^\Delta(1|W_1)}L\left(Q_n^V(\varepsilon)\right)$. This can be obtained by solving for ε in the equation

$$0=\sum_{i=1}^n\frac{I(\Delta_i=1)}{g_n^\Delta(1\mid W_{1,i})}\times\left\{\hat{B}_n(W_{1,i},V_i)-\frac{\sum_v\hat{B}_n(W_{1,i},v)Q_n^V(v\mid\Delta=1,W_{1,i})\exp\left[\varepsilon\cdot\hat{B}_n(W_{1,i},v)\right]}{\sum_v Q_n^V(v\mid\Delta=1,W_{1,i})\exp\left[\varepsilon\cdot\hat{B}_n(W_{1,i},v)\right]}\right\}$$

where $\hat{B}_n\equiv B\left((Q_{t=0,n}^{*,\mathbf{a},1}:\mathbf{a}),\psi_n^1\right)$. The updated density is given by $Q_{V,n}^1\equiv Q_n^V(\varepsilon_n^V)$.

4. Having obtained an updated density $Q_n^{V,j}$ at the j-th iteration, repeat step (2) and (3) to obtain a targeted estimate of ψ_n^{j+1} and $Q_{V,n}^{j+1}$, until ε_n^V converges to 0. In practice, this convergence can be achieved (close to 0) after a few iterations. We denote the final updates as $\psi_n^{*,TMLE}$, and $Q_n^{*,V}$. We call this estimator the *full TMLE*.

Let $Q_n^*\equiv\left(Q_n^{W_1},Q_n^{*,V},(Q_{t,n}^{*,\mathbf{a},1}:\mathbf{a})\right)$, where $Q_n^{W_1}$ is the empirical distribution of W_1. By design, $P_nD^*(Q_n^*,g_n,\psi_n^{*,TMLE})=0$. From lemma 3.2, we know that $\psi_n^{*,TMLE}$ is unbiased if either $Q_n^*=Q_0$ or $g_n=g_0$. Under standard regularity and empirical process conditions, $\psi_n^{*,TMLE}$ is asymptotically linear with influence curve $M(\Psi(P_0))^{-1}D^*(P_0)$. The asymptotic covariance of $\sqrt{n}(\psi_n^{*,TMLE}-\psi_0)$ can be estimated by the sample covariance matrix $\Sigma_n^{*,TMLE}$ of $\left\{M(\psi_n^{*,TMLE})^{-1}D^*\left(Q_n^*,g_n,\psi_n^{*,TMLE}\right)\right\}$. In particular, since $M(\Psi(P_0))^{-1}D^*(P_0)$ is the canonical gradient of Ψ at P, the estimator $\psi_n^{*,TMLE}$ is asymptotically efficient if all relevant components in D^* are consistently estimated.

Inverse Probability Weighted estimator

From remark 3.1, another valid estimating function for Ψ is given by

$$D^{IPW}(g,\psi)\equiv\frac{I(\Delta=1)}{g^\Delta(1\mid W_1)}\sum_{\mathbf{a}\in\mathscr{A}}\tilde{h}(\mathbf{a},V)C_K^{\mathbf{a}}\left[Y_K-m_\psi(\mathbf{a},V)\right]. \tag{3.15}$$

Up to a normalizing matrix $M(\Psi(P),P)^{-1}$, as defined in (3.10), $D^{IPW}(g,\psi)$ is a gradient for Ψ under a model $\mathscr{M}_g$ where g is known. This is an unbiased estimating function for ψ_0 if $g(P) = g_0$. To implement the IPW estimator, for each $\mathbf{a} \in \mathscr{A}$ and each individual i with $\Delta_i = 1$ and $\mathbf{A}_{i,K} = \mathbf{a}$, we create a row of data consisting of Y_i, $\phi(bfa, V_i)$, $h(\mathbf{a}, V_i)$, $C_K^{\mathbf{a}}$, $g_n^{\Delta}(1 \mid W_{1,i})$. The estimator ψ_n^{IPW} can be obtained by fitting a weighted regression of Y on $\phi(\mathbf{a},V)$, with weights $\frac{1}{g_n^{\Delta}(1|W_1)} h(\mathbf{a},V) C_K^{\mathbf{a}}$. This ψ_n^{IPW} satisfies $P_n D^{IPW}(g_n, \psi_n^{IPW}) = 0$, and it's unbiased if g_n consistently estimates g_0. Under standard regularity and empirical process conditions, ψ_n^{IPW} is asymptotically linear with influence curve $M(\Psi(P_0))^{-1} D^{IPW}(g_0)$. The asymptotic covariance of $\sqrt{n}(\psi_n^{IPW} - \psi_0)$ can be estimated by the sample covariance matrix Σ_n^{IPW} of $\left\{ M(\psi_n^{IPW})^{-1} D^{IPW}\left(g_n, \psi_n^{IPW}\right)\right\}$.

Non-Targeted Substitution Estimator

This is commonly referred to as the G-computation estimator; it utilizes non-targeted MLE estimators for the components of the data generating distribution that are relevant in the definition of Ψ. From (3.7) and remark 3.1, we can express $\Psi(P_0)$ as $\Psi\left(Q_{t=0}^{\mathbf{a},1}(P_0), Q^V(P_0)\right)$ or $\Psi\left(Q_{t=0}^{\mathbf{a},1}(P_0), g_0^{\Delta}\right)$, the latter option opens the door for G-computation estimator in applications with high-dimensional V. Unlike the other estimators discussed so far, there is no theory ensuring a central limit theorem based inference for the G-computation estimator.

To obtain an estimator $Q_{t=0}^{\mathbf{a},1}(P_0)$, we can use a sequential regression approach by performing steps (1) and (2) of section 3.4, starting at $t = K$ and ending at $t = 0$, but without the targeting procedure, i.e. always use $Q_{t+1,n}^{\mathbf{a},1}$ instead of $Q_{t+1,n}^{*,\mathbf{a},1}$ at t. At the end of $K+1$ steps, we have an estimator $Q_{t=0,n}^{\mathbf{a},1}$.

We first consider the representation $\Psi\left(Q_{t=0}^{\mathbf{a},1}(P_0), Q^V(P_0)\right)$. Let Q_n^V be an estimators of $Q^V(P_0)$. For each observation i, each $\mathbf{a} \in \mathscr{A}$ and each $v \in \mathscr{V}$, we create a row of data consisting of $Q_{t=0,n}^{\mathbf{a},1}(V = v, W_1), \phi(\mathbf{a},v), h(\mathbf{a},v), Q_n^V(v \mid \Delta = 1, W_1)$. The estimator $\psi_n^{V,Gcomp}$ can be obtained by a weighted regression $Q_{t=0,n}^{\mathbf{a},1}(V = v, W_1)$ on $\phi(\mathbf{a},v)$, with weights $h(\mathbf{a},v) Q_n^V(v \mid \Delta = 1, W_1)$. This $\psi_n^{V,Gcomp}$ is unbiased if both $Q_{t=0,n}^{\mathbf{a},1}$ and Q_n^V are consistent.

Consider now the alternative representation $\Psi\left(Q_{t=0}^{\mathbf{a},1}(P_0), g_0^{\Delta}\right)$, from the equalities in remark 3.1. For each observation i with $\Delta_i = 1$, and each $\mathbf{a} \in \mathscr{A}$, create a row of data consisting of $Q_{t=0,n}^{\mathbf{a},1}(V, W_1)$, $\phi(\mathbf{a},V)$, $h(\mathbf{a},V)$, $g_n^{\Delta}(1 \mid W_1)$. The estimator $\psi_n^{\Delta,Gcomp}$ can be obtained by a weighted regression $Q_{t=0,n}^{\mathbf{a},1}(V, W_1)$ on $\phi(\mathbf{a},V)$, with weights $\frac{h(\mathbf{a},V)}{g_n^{\Delta}(1|W_1)}$. This $\psi_n^{\Delta,Gcomp}$ is unbiased if both $Q_{t=0,n}^{\mathbf{a},1}$ and g_n^{Δ} are consistent.

3.5 Simulation Study

In this section, we examine the relative performance of the IPW estimator (section 3.4), the Δ-TMLE estimator (section 3.4), and the G-computation estimator (section 3.4) for the parameters of a MSM model.

Data Generating Process and Target Parameter

We consider a survival type example with data structure $O = (W_1, \Delta, V, L_0, (A_t, L_t) : t = 1, \dots K)$ with $K = 3$, where A_1 is the treatment assignment, A_t for $t > 1$ is the indicator of remaining in the study by time t. Time varying covariate L_t consist of L_t^1, L_t^2, and the death indicator Y_t. The data generating process is as follows:

$$
\begin{aligned}
&\left(W_1^1, W_1^2\right) \sim (Bern(0.3), Bern(0.7));\\
&\Delta \sim Bern\left(\text{expit}(1 + 2W_1^1 + 0.1W_1^2)\right);\\
&V \in \{0,1,2\} \sim \Big\{ I(V=1) \sim Bern(\text{expit}(-2 + 1.2W_1^1 + 0.7W_1^2),\\
&\qquad\qquad \{I(V=2) \mid V \neq 1\} \sim Bern(\text{expit}(-0.7 + 1.2W_1^1 + W_1^2)\Big\};\\
&L_0^1 \sim Bern\left(\text{expit}(-0.2 + 2W_1^1 + 0.5W_1^2 + 0.2I(V=1) + 0.4I(V=2))\right),\\
&L_0^2 \sim Bern\left(\text{expit}(-0.8 + W_1^1 + W_1^2 - 0.3I(V=1) - 0.1I(V=2))\right);\\
&A_t \sim \begin{cases}
Bern\Big(\text{expit}(-1 + W_1^1 + 1.3W_1^2 + 0.1I(V=1) + 0.1I(V=2) + 1.2L_0^1 + L_0^2 - 0.7W_1^1 \times L_0^2\\
\qquad -0.5W_1^2 \times L_0^1)\Big), \text{for } t = 1,\\
Bern\Big(\text{expit}\big(2 + W_1^1 + W_1^2 + 0.1I(V=1) + 0.1I(V=2) + 0.6L_0^1 + 1.2L_0^2 - 0.5A - 0.1t\\
\qquad +0.8L_t^1 - 0.3L_t^2 + 0.1L_{t-1}^1 - 0.2L_{t-1}^2 - 0.3A \times L_0^2 + 0.2A \times W_1^1 - 0.3A \times L_t^2\\
\qquad -0.2A \times L_{t-1}^1\big)\Big), \text{for } t > 1.
\end{cases}\\
&L_t^1 \sim Bern\Big(\text{expit}\big(-1 + W_1^1 + 0.1W_1^2 - 0.5I(V=1) - 0.7I(V=2) + L_0^1 + 0.3L_0^2 + 1.5A + 0.4t\\
&\qquad\qquad - L_{t-1}^1 - 0.8A \times I(V=1) - 0.2A \times I(V=2) - 0.3A \times W_1^1\big)\Big);\\
&L_t^2 \sim Bern\Big(\text{expit}\big(-2 + 0.1W_1^1 + 0.1W_1^2 + 0.5I(V=1) + 0.5I(V=2) + 0.7L_0^1 + 0.2L_0^2 - A + 0.2t\\
&\qquad\qquad + L_{t-1}^2 - 0.2A \times I(V=1) - 0.4A \times I(V=2) - 0.3A \times L_0^2\big)\Big);\\
&Y_t \sim Bern\Big(\text{expit}\big(-1.4 + 1.5W_1^1 + W_1^2 - 0.7I(V=1) - 0.8I(V=2) + L_0^1 + L_0^2 + A - L_t^1 - 0.1L_t^2\\
&\qquad\qquad - L_{t-1}^1 - 0.3L_{t-1}^2 + A \times I(V=1) + 0.8A \times I(V=2) - 0.3A \times L_t^1\\
&\qquad\qquad - 0.4A \times L_{t-1}^1 - 0.3A \times L_0^2 - 0.2A \times W_1^1\big)\Big).
\end{aligned}
$$

Once either the censoring jumps to 0 or death process jump to 1, then all subsequent variables are encoded by carrying forward their last observation.

Our interventions of interest are $\Delta = 1$ and $\mathscr{A} = \{(0,1,1),(1,1,1)\}$. Under the above distribution, $0.1 < g_0(A_1 = 1 \mid \cdot) < 0.95$, and $g_0(A_t = 1 \mid \cdot) > 0.5$ for all $t \geq 2$.

We model the dose response $\{\rho_{\mathbf{a},v} : \mathbf{a}, v\}$ by the MSM

$$m_\psi(\mathbf{a},v) = \text{expit}(\psi_1 + \psi_2 a_1 + \psi_3 v_1 + \psi_4 v_2 + \psi_5 a_1 v_1 + \psi_6 a_1 v_2) = \text{expit}(\psi \cdot \phi(\mathbf{a},v)),$$

where $\phi(\mathbf{a},v) = (1, a_1, v_1, v_2, a_1 v_1, a_1 v_2)$, with kernel weights $h(\mathbf{a},v) = p_0(\mathbf{a} \mid v, \Delta = 1)$. Note that in this case, the kernel weights are assumed to be known. The target parameter defined in (3.7) takes value $\psi_0 = (0.825, 0.105, 0.249, -0.046, 1.474, 0.960)$.

Estimators

The Δ mechanism g^Δ is estimated using Super Learner (van der Laan, Polley, and Hubbard (2007)) with candidate fitting algorithms `glm` and `nnet`, adjusting for W_1^1 and W_1^2. Using sample splitting, Super Learner selects a convex combination of the candidate algorithms which yields an estimator with minimal cross validated risk. Theoretical results from van der Vaart, Dudoit, and van der Laan (2006) and van der Laan, Dudoit, and Keleş (2004) showed that this estimator converges to an oracle estimator. We use two estimators for g^A: a correctly specified logistic model (shorthand 'gc'), and a misspecified logistic model ('gm') that omits $W_1^1, W_1^2, L_0^1, L_0^2, L_t^1$. The denominator for each $C_t^{\mathbf{a}}$ is truncated below by 0.025. We use two estimators for $Q_t^{\mathbf{a},1}(P_0)$: both use Super Learner with candidate fitting algorithms `glm` and `nnet`, the correctly specified estimator ('Qc') adjusts for all baseline variables and all time-varying covariates up to one time lag, the misspecified estimator ('Qm') only adjusts for V_1 and V_2 at each time t.

We consider 3 cases of model misspecification on $Q_t^{\mathbf{a},1}$ and g^A: all correct ('Qc, gc'); correct $Q_t^{\mathbf{a},1}$ and misspecified g^A ('Qc,gm'); misspecified $Q_t^{\mathbf{a},1}$ and correct g^A ('Qm, gc'). For all three cases we always use the same correctly specified g^Δ. We implement the second version of the G-comp estimator in 3.4, where the weights are given by g^Δ. The G-computation estimator changes only under specifications 'Qc, gc' and 'Qm, gc'. The IPW estimator changes only under specifications 'Qc, gc' and 'Qc, gm'.

Results

The bias, variance, and coverage probability (for the influence-function-based confidence intervals) are appraised using 500 repetitions.

In table 3.1, we see that when g^A is misspecified, Δ-TMLE using a correct $Q_t^{\mathbf{a},1}$ reduces bias over the misspecified IPW estimator. Similarly, when $Q_t^{\mathbf{a},1}$ is misspecified, Δ-TMLE using the correct g^A reduces bias over the misspecified G-computation estimator. When comparing the correct vs misspecified G-comp, and the correct vs misspecified IPW, coefficients involving the adjusted covariates (V_1, V_2) were still estimated very well by the

misspecified estimator. Under 'Qc, gc', the correct G-computation estimator converges much slower than the IPW and the Δ-TMLE estimators. We posit that this may be due to its sole reliance on the nonparametric likelihood estimates. As expected, G-computation has the smallest sample variance, and IPW has the largest sample variance despite the truncated estimators for g. Under certain regularity conditions, the IPW and Δ-TMLE estimators are asymptotically linear — table 3.2 tabulates the coverage probability of their Influence-Function based confidence intervals. At the correct models (Qc,gc), IPW and Δ-TMLE are asymptotically linear with influence curve D^{IPW} and D^{Δ}, respectively. We used $\sqrt{\hat{var}D_n^{IPW}/n}$ and $\sqrt{\hat{var}D_n^{\Delta}/n}$ to estimate their respective standard errors. As sample size grows, the actual coverage probabilities are quite close to the nominal coverage probability, with IPW having a better coverage. When one of the components is misspecified, the Δ-TMLE still provides very good coverage, even though theoretically D^{Δ} is only part of its influence curve; we postulate that this is because the influence curve based standard error estimates are large relative to the finite sample bias. The misspecified IPW has very good coverage for the covariates that were adjusted for in the misspecified model, but very bad coverage for the confounded coefficients (A and the intercept).

3.6 Data Analysis Example

To illustrate the application of the Δ-TMLE, we revisit our earlier example: the Sequenced Treatment Alternatives to Relieve Depression (STAR*D). After an initial screening process, patients are enrolled into level 1 of the treatment, where everyone was treated with citalopram. At the start of each subsequent level, if a patient still remains in the study, then he is randomized into one of his accepted treatment options. Regular follow-up visits are conducted throughout each level. At each follow-up visit, covariates are collected, and the patient is subject to dropout, entering remission, or moving onto next level. At level 2, we wish to identify potential modifiers of the effect of switching medication vs augmenting medication on the chances of entering remission by the end of level 2. These modifiers are collected prior to the assignment of level 2 treatment and many are subject to missingness, therein lies the need for the tools developed here. Our study population is the set of 1395 patients at level 2 who found medication strategies acceptable. Note that because switching medication and augmenting medication are general treatment strategies that encompasses various treatment options (specific drugs), for most patients these strategies are self-selected.

The data structure was described in detail in section 3.2 as part of our running example. We consider here two types of potential effect modifiers: some are measured at screening and some are measured at exit of level 1. Table 3.3 summarizes percent of missingness, range, and scale of each effect modifier. Table 3.4 tabulate the events in level 2 by strategy

Table 3.1: Results: Bias, Variance, MSE for estimators of ψ_0. Qc = correct $Q_t^{\mathbf{a},1}$, Qm= misspecified $Q_t^{\mathbf{a},1}$, gc=correct g^A, gm=misspecified g^A.

ψ	Intercept		A		V_1		V_2		$A \times V_1$		$A \times V_2$	
n	500	2000	500	2000	500	2000	500	2000	500	2000	500	2000
Bias												
Qc, gc												
Gcomp	0.334	0.364	0.606	0.603	0.82	0.818	0.560	0.552	1.377	1.390	0.906	0.904
IPW	0.016	0.017	0.056	0.063	0.169	0.131	0.070	0.073	0.402	0.117	0.083	0.099
Δ-TMLE	0.036	0.001	0.021	0.002	0.078	0.054	0.014	0.008	0.519	0.018	0.002	0.006
Qc, gm												
IPW	0.492	0.512	0.831	0.836	0.094	0.131	0.075	0.072	0.638	0.160	0.026	0.004
Δ-TMLE	0.022	0.033	0.029	0.028	0.069	0.029	0.02	0.006	0.499	0.015	0.008	0.008
Qm, gc												
Gcomp	0.751	0.773	1.383	1.353	0.66	0.640	0.431	0.414	1.089	1.069	0.752	0.727
Δ-TMLE	0.01	0.024	0.04	0.064	0.123	0.124	0.048	0.062	0.462	0.103	0.041	0.081
Variance												
Qc, gc												
Gcomp	0.082	0.019	0.112	0.023	0.137	0.033	0.084	0.017	0.038	0.008	0.022	0.005
IPW	0.134	0.035	0.226	0.051	0.939	0.118	0.353	0.068	9.605	0.235	0.555	0.112
Δ-TMLE	0.113	0.029	0.187	0.041	0.734	0.091	0.279	0.053	9.149	0.169	0.398	0.092
Qc, gm												
IPW	0.114	0.030	0.196	0.045	0.897	0.093	0.259	0.056	9.298	0.184	0.418	0.094
Δ-TMLE	0.097	0.026	0.16	0.036	0.718	0.078	0.21	0.048	9.062	0.141	0.319	0.078
Qm, gc												
Gcomp	0.095	0.024	0.146	0.035	0.147	0.045	0.091	0.024	0.191	0.057	0.071	0.025
Δ-TMLE	0.123	0.031	0.201	0.044	0.839	0.105	0.316	0.057	9.424	0.207	0.464	0.096
MSE												
Qc, gc												
Gcomp	0.195	0.151	0.479	0.387	0.809	0.701	0.397	0.322	1.934	1.941	0.842	0.822
IPW	0.134	0.035	0.229	0.054	0.967	0.135	0.358	0.073	9.767	0.248	0.562	0.122
Δ-TMLE	0.115	0.029	0.187	0.041	0.740	0.094	0.279	0.053	9.419	0.17	0.398	0.092
Qc, gm												
IPW	0.357	0.292	0.887	0.744	0.906	0.111	0.264	0.062	9.704	0.21	0.419	0.094
Δ-TMLE	0.097	0.027	0.16	0.036	0.723	0.079	0.210	0.048	9.311	0.141	0.319	0.078
Qm, gc												
Gcomp	0.66	0.621	2.059	1.865	0.583	0.455	0.277	0.195	1.377	1.2	0.6378	0.554
Δ-TMLE	0.123	0.031	0.203	0.048	0.854	0.12	0.318	0.061	9.637	0.218	0.466	0.103

received. Note that there are 3 strategies received, but we are only comparing switching medication vs augmenting medication. The multivariate nature of most of the effect modifiers underscores the need for Δ-TMLE. If V is screening covariate, W_1 includes all demographic variables and medical and psychiatric history prior to enrollment, and missing indicator for each of those variables; if V is level 1 exit covariate, then we add to W_1 variables summarizing number of visits, adherence to study protocol, and time spent in level 1.

The MSM is a generalized linear model with logit link. The linear predictor $\phi(\mathbf{a}, v) =$

Table 3.2: Coverage Probability for the Asymptotically Linear Estimators, using Influence-Function based Confidence Interval. Qc = correct $Q_t^{\mathbf{a},1}$, Qm= misspecified $Q_t^{\mathbf{a},1}$, gc=correct g^A, gm=misspecified g^A.

	Intercept		A		V_1		V_2		$A \times V_1$		$A \times V_2$	
n	500	2000	500	2000	500	2000	500	2000	500	2000	500	2000
Qc, gc												
IPW	0.948	0.946	0.944	0.940	0.966	0.930	0.940	0.952	0.912	0.950	0.956	0.960
Δ-TMLE	0.934	0.932	0.934	0.942	0.936	0.924	0.918	0.946	0.906	0.942	0.942	0.938
Qc, gm												
IPW	0.698	0.146	0.556	0.022	0.956	0.928	0.958	0.958	0.928	0.942	0.958	0.964
Δ-TMLE	0.964	0.956	0.954	0.956	0.984	0.984	0.972	0.986	0.932	0.954	0.966	0.964
Qm, gc												
Δ-TMLE	0.942	0.930	0.942	0.948	0.956	0.930	0.936	0.954	0.896	0.942	0.954	0.958

$(1, a_1, v, a_1 \times v)$ and $\psi \in \mathbb{R}^4$ for the binary V, and $\phi(\mathbf{a}, v) = \left(1, a_1, v, v^2, a_1 \times v, a_1 \times v^2\right)$ and $\psi \in \mathbb{R}^6$ for the non-binary V. The kernel weights are $h(\mathbf{a}, V) = p_0(\mathbf{A}_K = \mathbf{a} \mid V)$, to be estimated using Super Learner with fitting algorithms `glm`, `nnet` and `bayesglm`. The initial estimators of g and $Q_t^{\mathbf{a},1}$ adjust for all baseline covariates and time-varying covariates with up to 2 time lag (each covariate is coupled with its missingness indicator). We used Super Learner with the fitting algorithms `glm`, `knnreg`, `nnet` and `bayesglm`; each fitting algorithm is coupled with each of the following screening algorithms: Spearman correlation tests at significance levels 0.01, 0.05, 0.1, 0.2; ranking p-values from the correlation tests and take the top m variables, where m ranges, in increments of 10, from 10% to 90% of the total number of variables being considered.

The measure of treatment heterogeneity is given by

$$\beta(\psi) \equiv \max_v \ell_{OR}(v; \psi) - \min_v \ell_{OR}(v; \psi),$$

where

$$\ell_{OR}(v; \psi) = \log \frac{m(1, v; \psi)}{1 - m(1, v; \psi)} - \log \frac{m((0, v; \psi)}{1 - m(0, v; \psi)} = \psi_2 + \psi_4 v \text{ or } \psi_2 + (\psi_5, \psi_6) \cdot (v, v^2),$$

for binary V or discrete V, respectively. This measure quantifies the most change in log odds ratio between any two values of V. Consider the null hypothesis $H_0 : \beta(\psi_0) = 0$. Using Δ-TMLE, we obtain an estimator $\beta_n = \beta(\psi_n^{\Delta, TMLE})$ of β_0 for each V. An application of the functional delta Method, with the covariance matrix Σ_n of the estimated influence curve D_n^{Δ}, yields a standard error estimate SE_n for β_n. We use the test statistics $T_n = \beta_n / SE_n \sim N(0, 1)$. The false discovery rate (FDR) of the simultaneous comparisons are controlled at 0.05 with the Benjamini-Hochberg procedure. The results of the analysis are summarized in table 3.5

Table 3.3: Effect modifiers of interest for level 2. **hdtot**: Hamilton D17 score; **anxious**: indicator of anxious depression; **atypical**: indicator of atypical depression; **melancholic**: indicator of melancholic depression; **grseb**: side-effect burden score; **qctot**: QIDS total score; **qstot**: self-reported QIDS totoal score; **ictot**: IDS C30 score.

	% missing	Range		type
		min	max	
At Screening				
hdtot	< 1	0.00	42.00	Continuous
anxious	6	0.00	1.00	Binary
atypical	5	0.00	1.00	Binary
melancholic	5	0.00	1.00	Binary
At Level 1 exit				
grseb	1	0.00	6.00	Continuous
qctot	1	0.00	27.00	Continuous
qstot	1	0.00	27.00	Continuous
hdtot	12	0.00	42.00	Continuous
ictot	13	0.00	74.00	Continuous

Table 3.4: Level 2 events by strategy received

	Med-Sw	Med-Aug	any CT
Received	727	565	103
Success	257 (35%)	287 (51%)	53 (51%)
No Success	227 (32%)	132 (23%)	22 (21%)
Dropout	243 (33%)	146 (26%)	28 (27%)

Table 3.5: Data Analysis Results: $\beta_0 = max_v \ell OR(v;\psi_0) - min_v \ell OR(v;\psi_0)$, $H_0 : \beta_0 = 0$. $T_n = \beta_n / SE_n \sim N(0,1)$, where ψ_n given by Δ-TMLE. Control FDR at 0.05 with Benjamini-Hochberg procedure

	P-value	rejectNull	β_n	$\hat{var}\beta_n$	$min_v \ell OR(v;\psi_n)$	$max_v \ell OR(v;\psi_n)$
At Screening						
hdtot	3.0e-01	0e+00	2.1e+00	4.2e+00	-6.0e-01	1.5e+00
anxious	1.0e+00	0e+00	9.7e-04	6.0e-02	-3.6e-01	-3.6e-01
atypical	9.2e-04	1e+00	8.9e-01	7.3e-02	-5.2e-01	3.7e-01
melancholic	7.7e-01	0e+00	7.1e-02	5.8e-02	-3.9e-01	-3.2e-01
At Level 1 exit						
grseb	1.2e-04	1e+00	2.3e+00	3.7e-01	-7.8e-01	1.6e+00
qctot	2.5e-01	0e+00	1.6e+00	1.9e+00	-1.6e+00	1.9e-02
qstot	3.3e-02	0e+00	1.7e+00	6.5e-01	-1.9e+00	-1.4e-01
hdtot	1.4e-02	1e+00	2.3e+00	8.6e-01	-2.2e+00	9.2e-02
ictot	9.0e-02	0e+00	1.6e+00	8.7e-01	-1.3e+00	2.8e-01

3.7 Summary

In this chapter, we studied causal effect modification by a counterfactual modifier. The tools developed here are applicable in situations where the effect modifier is missing at random, or where the effect modifier of interest is between a first and second line treatment, and one is interested in effect modification of the second-line treatment by this variable, had the first line treatment been at a given set value. In the context of missing effect modifier, the parameter of interest teaches us that one should avoid unnecessary assumptions on the type of missingness and should focus the efforts on extracting this information from the data. We established the efficient influence curve (EIC) for the corresponding marginal structural model parameters which provides the semiparametric efficiency bound for all the asymptotically linear estimators. This efficient influence curve is also doubly robust in that it remains an unbiased estimating function if either 1) the outcome expectations and the modifier density, or 2) the intervention densities, are consistently estimated. However, in applications with high-dimensional V, we saw that it may be difficult to fully utilize the EIC without potentially compromising consistency. To solve this problem, we presented a inverse-probability-missingness-weighted influence curve (projected IC), which equals the EIC in a model where the missingness mechanism (or more generally, the assignment of the modifier's intervention) is known. Though not fully efficient under the larger model, this projected IC is robust against misspecification of the outcome models or the exposure mechanisms, whenever the missingness mechanism is consistently estimated. We presented two TMLE estimators using the EIC and the projected IC which also inherit the corresponding robustness properties. We also described an IPW estimator that is unbiased if the intervention probabilities are consistently estimated, and a non-targeted G-computation estimator that is unbiased if the outcome expectations and either the modifier density or the missingness mechanism are consistently estimated. Under standard regularity and empirical process conditions, the TMLE and the IPW estimators are asymptotically linear, thereby allowing CLT-based standard error estimates. Moreover, the TMLE estimator using the EIC will be semiparametric efficient if all the components of the likelihood are consistently estimated.

3.8 Chapter Appendix

Proof of remark 3.1

The first equality in the remark follows from definition of $\Psi(P)$ and choice of $m_\psi(\mathbf{a},v)$, the third equality is trivial. We only show the second one.

$$\begin{aligned}
&\sum_{(\mathbf{a},v)\in\mathscr{A}\times\mathscr{V}} h(\mathbf{a},v)\phi(\mathbf{a},v)\gamma_v(P)\left\{\frac{\eta_{\mathbf{a},v}(P)}{\gamma_v(P)} - m_{\Psi(P)}(\mathbf{a},v)\right\}\\
&= -E_P \sum_{(\mathbf{a},v)\in\mathscr{A}\times\mathscr{V}} h(\mathbf{a},v)\phi(\mathbf{a},v)\left\{Q^V(v\mid\Delta=1,W_1)Q^{\mathbf{a},1}_{t=0}(v,W_1) - \eta_{\mathbf{a},v}(P)\right\}\\
&+ E_P \sum_{(\mathbf{a},v)\in\mathscr{A}\times\mathscr{V}} h(\mathbf{a},v)\phi(\mathbf{a},v)Q^V(v\mid\Delta=1,W_1)\left\{Q^{\mathbf{a},1}_{t=0}(v,W_1) - m_{\Psi(P)}(\mathbf{a},v)\right\}
\end{aligned}$$

The first line in the right-hand-side of the above equation is zero by definition of $\eta_{\mathbf{a},v}(P)$.

Proof of lemma 3.1: Efficient score for $\Psi(P)$ under $\mathscr{M}$

In this appendix, we derive the efficient influence curve at P of the map $\Psi:\mathscr{M}\to\mathbb{R}^d$. For each $P\in\mathscr{M}$, let $\mathscr{H}_P$ denote the Hilbert space of 1-dimensional mean zero measurable functions of O with finite variance, endowed with the covariance inner product. For an $r\in\mathscr{H}_P$, define a 1-dimensional parametric submodel $\{P_r(\alpha):|\alpha|<1/\|r\|_\infty\}$, through P at $\alpha=0$, given by $\frac{dP_r(\alpha)}{dP}=1+\alpha r(O)$. Since we are working under a saturated model $\mathscr{M}$, this submodel is indeed contained in $\mathscr{M}$. We shall consider the the class of all such 1-dimensional submodels indexed by $\mathscr{H}_P$.

For a given $D\in\mathscr{H}_P^d$, we define the vector inner product $E_P(D\times r)$ as the vector of the component-wise inner products $\left(E_P\left(D_j\times r\right):j=1\ldots d\right)$. We wish to how that D^* satisfies

$$\left.\frac{d\Psi(P_r(\alpha))}{d\alpha}\right|_{\alpha=0} = M\left(\Psi(P),P\right)^{-1}E_P\left(D^*\times r\right).$$

From definition of the maps Ψ, γ_v and $\eta_{\mathbf{a},v}$, and our choice of working model $m_\psi(\mathbf{a},v)$, we know that at each $P_r(\alpha)$

$$0=\sum_{(\mathbf{a},v)\in\mathscr{A}\times\mathscr{V}}\gamma_v(P_r(\alpha))h(\mathbf{a},v)\phi(\mathbf{a},v)\left[\frac{\eta_{\mathbf{a},v}(P_r(\alpha))}{\gamma_v(P_r(\alpha))} - m_{\Psi(P_r(\alpha))}(\mathbf{a},v)\right]. \tag{3.16}$$

We perform an implicit differentiation with respect α on the above equation, at $\alpha=0$, to

obtain the equality

$$\frac{d\Psi(P_r(\alpha))}{d\alpha}\bigg|_{\alpha=0} \times \left\{ \sum_{(\mathbf{a},v)\in\mathscr{A}\times\mathscr{V}} \gamma_v(P)h(\mathbf{a},v)\phi(\mathbf{a},v)m_{\Psi(P)}(\mathbf{a},v)\left[1-m_{\Psi(P)}(\mathbf{a},v)\right]\phi(\mathbf{a},v)^\top \right\}$$
$$= \sum_{(\mathbf{a},v)\in\mathscr{A}\times\mathscr{V}} h(\mathbf{a},v)\phi(\mathbf{a},v)\left\{ \frac{d\eta_{\mathbf{a},v}(P_r(\alpha))}{d\alpha}\bigg|_{\alpha=0} - \left(\frac{d\gamma_v(P_r(\alpha))}{d\alpha}\bigg|_{\alpha=0} \right) m_{\Psi(P)}(\mathbf{a},v) \right\}. \quad (3.17)$$

Next, we proceed to express $\frac{d\gamma_v(P_r(\alpha))}{d\alpha}\Big|_{\alpha=0}$ and $\frac{d\eta_{\mathbf{a},v}(P_r(\alpha))}{d\alpha}\Big|_{\alpha=0}$ as $E_P(D_{\gamma_v}\times r)$ and $E_P(D_{\rho_{\mathbf{a},v}}\times r)$, respectively, for some functions D_{γ_v} and $D_{\eta_{\mathbf{a},v}}$ belonging to the Hilbert space $\mathscr{H}_P$.

For convenience of indexing, for a given vector $\mathbf{l}_K$, we shall use the short hand $h_t^{\mathbf{a},v} = (\Delta=1, V=v, \mathbf{L}_{t-1}=\mathbf{l}_{t-1}, \mathbf{A}_t=\mathbf{a}_t)$, for $t=1,\ldots,K$, and $h_t^{\mathbf{a},v}=(\Delta=1,V=v)$ for $t=0$. From definition of $P_r(\alpha)$, it follows that

$$P_r(\alpha)(W_1) = P(W_1)\left(1+\alpha E_P(r\mid W_1)\right),$$
$$P_r(\alpha)(v\mid\Delta=1,W_1) = P(v\mid\Delta=1,W_1)\frac{1+\alpha E_P(r\mid V=v,\Delta=1,W_1)}{1+\alpha E_P(r\mid\Delta=1,W_1)},$$

and $P_r(\alpha)(l_t\mid W_1,H_t^{\mathbf{a},v}) = P(l_t\mid W_1,H_t^{\mathbf{a},v})\frac{1+\alpha E_P\left(r\mid l_t,H_t^{\mathbf{a},v}\right)}{1+\alpha E_P\left(r\mid H_t^{\mathbf{a},v}\right)}$.

Proposition 3.1. *For a given $v\in\mathscr{V}$, $\frac{d\gamma_v(P_r(\alpha))}{d\alpha}\Big|_{\alpha=0} = E_P(D_{\gamma_v}\times r)$, where*

$$D_{\gamma_v}(P) = \frac{I(\Delta=1)}{g^\Delta(1\mid W_1)}\left(I(V=v)-Q^V(v\mid\Delta=1,W_1)\right)+Q^V(v\mid\Delta=1,W_1)-\gamma_v(P). \quad (3.18)$$

Proof.

$$\frac{d\gamma_v(P_r(\alpha))}{d\alpha}\bigg|_{\alpha=0} = \lim_{\alpha\to0}\frac{\gamma_v(P_r(\alpha))-\gamma_v(P)}{\alpha} = \lim_{\alpha\to0}\frac{1}{\alpha}\int_{W_1}\Bigg\{P(W_1)P(v\mid\Delta=1,W_1)\alpha$$
$$\frac{E_P(r\mid v,\Delta=1,W_1)-E_P(r\mid\Delta=1,W_1)+E_P(r\mid W_1)+\alpha E_P(r\mid W_1)E_P(r\mid v,\Delta=1,W_1)}{1+\alpha E_P(r\mid\Delta=1,W_1)}\Bigg\}$$
$$= \int_{W_1} P(W_1)P(v\mid\Delta=1,W_1)\left\{E_P(rI(V=v)\mid\Delta=1,W_1)-E_P(r\mid\Delta=1,W_1)+E_P(r\mid W_1)\right\}$$
$$= E_P\left\{\left(\frac{I(\Delta=1}{g^\Delta(1\mid W_1)}\left(I(V=v)-p(V\mid\Delta=1,W_1)\right)+Q^V(v\mid\Delta=1,W_1)\right)\times r(O)\right\}$$
$$= E_P\left(D_{\gamma_v}(P)\times r(O)\right).$$

In obtaining the last equality, we note that centering the left factor of the integrang by $\gamma_v(P)$ does not change the expression because $E_P(\gamma_v(P)r(O)) = \gamma_v(P)E_P(r(O)) = 0$ by definition of r. It is straightforward to check that indeed $E_P D^*_{\gamma_v}(P) = 0$. Moreover, under our saturated model $D^*_{\gamma_v}(P)$ is in fact in the tangent space. Therefore, it is indeed the efficient influence curve. This concludes the proof of proposition 3.1. □

Proposition 3.2. *For a given* $\mathbf{a} \in \mathscr{A}$, $v \in \mathscr{V}$, $\left.\frac{d\eta_{\mathbf{a},v}(P_r(\alpha))}{d\alpha}\right|_{\alpha=0} = E_P(D_{\eta_{\mathbf{a},v}} \times r)$, *where*

$$
\begin{aligned}
& D_{\eta_{\mathbf{a},v}}(P) \\
& \equiv \frac{I(\Delta=1)}{g^{\Delta}(1 \mid W_1)} I(V=v) \sum_{t=1}^{K} C_t^{\mathbf{a}} \left\{ Q_{t+1}^{\mathbf{a},1}(\mathbf{L}_t, v, W_1) - Q_t^{\mathbf{a},1}(\mathbf{L}_{t-1}, v, W_1) \right\} \\
& + \frac{I(\Delta=1)}{g^{\Delta}(1 \mid W_1)} I(V=v) \left\{ Q_{t=1}^{\mathbf{a},1}(L_0, v, W_1) - Q_{t=0}^{\mathbf{a},1}(v, W_1) \right\} \\
& + \frac{I(\Delta=1)}{g^{\Delta}(1 \mid W_1)} Q_{t=0}^{\mathbf{a},1}(v, W_1) \left\{ I(V=v) - Q^V(v \mid \Delta=1, W_1) \right\} \\
& + Q^V(v \mid \Delta=1, W_1) Q_{t=0}^{\mathbf{a},1}(v, W_1) - \eta_{\mathbf{a},v}(P). \qquad (3.19)
\end{aligned}
$$

where with $C_t^{\mathbf{a}} = \frac{I(\mathbf{A}_t=\mathbf{a}_t)}{\prod_{j=1}^{t} g^A(a_j \mid \mathbf{A}_{j-1}=\mathbf{a}_{j-1}, \mathbf{L}_{j-1}, V, \Delta=1, W_1}$

Proof.

$$
\begin{aligned}
& \left.\frac{d\eta_{\mathbf{a},v}(P_r(\alpha))}{d\alpha}\right|_{\alpha=0} = \lim_{\alpha \to 0} \frac{\eta_{\mathbf{a},v}(P_r(\alpha)) - \eta_{\mathbf{a},v}(P)}{\alpha} \\
& = \lim_{\alpha \to 0} \frac{1}{\alpha} \int_{w_1, l_0, \mathbf{l}_K} y_K P(W_1) P(v \mid \Delta=1, W_1) \prod_{j=0}^{K} P(l_j \mid w_1, h_j^{\mathbf{a},v}) \alpha \\
& \times \left\{ \frac{\sum_{t=0}^{K} E_P(r \mid l_t, w_1, h_t^{\mathbf{a},v}) + E_P(r \mid v, \Delta=1, W_1) + E_P(r \mid W_1)}{1 + \alpha M_P'(w_1, l_0, \mathbf{l}_K)} \right. \\
& \left. \frac{-\sum_{t=0}^{K} E_P(r \mid w_1, h_t^{\mathbf{a},v}) - E_P(r \mid \Delta=1, W_1) + \alpha M_P(w_1, l_0, \mathbf{l}_K)}{1 + \alpha M_P'(w_1, l_0, \mathbf{l}_K)} \right\} \\
& = \int_{w_1, l_0, \mathbf{l}_K} y_K P(w_1) P(v \mid \Delta=1, w_1) \prod_{j=0}^{K} P(l_j \mid w_1, h_j^{\mathbf{a},v}) \\
& \times \left\{ \sum_{t=0}^{K} E_P(r \mid l_t, w_1, h_t^{\mathbf{a},v}) - \sum_{t=0}^{K} E_P(r \mid w_1, h_t^{\mathbf{a},v}) + E_P(r \mid v, \Delta=1, w_1) \right. \\
& \left. - E_P(r \mid \Delta=1, w_1) + E_P(r \mid w_1) \right\}
\end{aligned}
$$

$$
\begin{aligned}
&= E_P\left\{ \frac{I(\Delta=1)}{g^{\Delta}(1 \mid W_1)} I(V=v) \sum_{t=0}^{K} \frac{I(\mathbf{A}_t = \mathbf{a}_t)}{\prod_{j=1}^{t} g^{A}(a_t \mid parents(a_t))} \right. \\
&\left. \times \left[Q_{t+1}^{\mathbf{a},1}(\mathbf{l}_t, v, w_1) - Q_t^{\mathbf{a},1}(\mathbf{l}_{t-1}, v, w_1) \right] \times r(O) \right\} \\
&+ E_P\left\{ \frac{I(\Delta=1)}{g^{\Delta}(1 \mid W_1)} Q_{t=0}^{\mathbf{a},1}(V=v, W_1) \left[I(V=v) - Q^V(v \mid \Delta=1, W_1) \right] \times r(O) \right\} \\
&+ E_P\left\{ Q^V(v \mid \Delta=1, W_1) Q_{t=0}^{\mathbf{a},1}(V=v, W_1) \times r(O) \right\} \\
&= E_P(D_{\eta_{\mathbf{a},v}} \times r)
\end{aligned}
$$

In the first equality, M_P and M_P' are shorthand for the remaining terms in the expansion of the products. This concludes the proof of proposition 3.2 □

Now, we derive the efficient influence curve for Ψ at $P \in \mathscr{M}$. From proposition 3.1 and 3.2, after some simplifications, we conclude that the right-hand-side of (3.17) can be written as

$$
\begin{aligned}
&\sum_{(\mathbf{a},v) \in \mathscr{A} \times \mathscr{V}} h(\mathbf{a}, v) \phi(\mathbf{a}, v) \left\{ \left. \frac{d\eta_{\mathbf{a},v}(P_r(\alpha))}{d\alpha} \right|_{\alpha=0} - \left(\left. \frac{d\gamma_v(P_r(\alpha))}{d\alpha} \right|_{\alpha=0} \right) m_{\Psi(P)}(\mathbf{a}, v) \right\} \\
&= \sum_{(\mathbf{a},v) \in \mathscr{A} \times \mathscr{V}} h(\mathbf{a}, v) \phi(\mathbf{a}, v) \left\{ E_P(D_{\eta_{\mathbf{a},v}}(P) \times r) - E_P(D_{\gamma_v}(P) \times r) m_{\Psi(P)}(\mathbf{a}, v) \right\} \\
&= E_P\left\{ \sum_{(\mathbf{a},v) \in \mathscr{A} \times \mathscr{V}} \tilde{h}(\mathbf{a}, v) \left[D_{\eta_{\mathbf{a},v}}(P) - D_{\gamma_v}(P) m_{\Psi(P)}(\mathbf{a}, v) \right] \times r(O) \right\} \\
&= E_P\left\{ D^*(Q, g, \Psi(P)) \times r \right\},
\end{aligned}
$$

where

$$
\begin{aligned}
D^*(Q, g, \psi) = &\frac{I(\Delta=1)}{g^{\Delta}(1 \mid W_1)} \sum_{\mathbf{a}} \tilde{h}(\mathbf{a}, V) \sum_{t=1}^{K} C_t^{\mathbf{a}} \left\{ Q_{t+1}^{\mathbf{a},1}(\mathbf{L}_t, V, W_1) - Q_t^{\mathbf{a},1}(\mathbf{L}_{t-1}, V, W_1) \right\} \\
&+ \frac{I(\Delta=1)}{g^{\Delta}(1 \mid W_1)} \sum_{\mathbf{a}} \tilde{h}(\mathbf{a}, V) \left\{ Q_{t=1}^{\mathbf{a},1}(L_0, V, W_1) - Q_{t=0}^{\mathbf{a},1}(V, W_1) \right\} \\
&+ \frac{I(\Delta=1)}{g^{\Delta}(1 \mid W_1)} \sum_{(\mathbf{a},v) \in \mathscr{A} \times \mathscr{V}} \left\{ \tilde{h}(\mathbf{a}, V) \left[Q_{t=0}^{\mathbf{a},1}(v, W_1) - m_{\psi}(\mathbf{a}, v) \right] \right. \\
&\qquad \left. \times \left(I(V=v) - Q^V(v \mid \Delta=1, W_1) \right) \right\} \\
&+ \sum_{(\mathbf{a},v) \in \mathscr{A} \times \mathscr{V}} \tilde{h}(\mathbf{a}, V) Q^V(v \mid \Delta=1, W_1) \left\{ Q_{t=0}^{\mathbf{a},1}(v, W_1) - m_{\psi}(\mathbf{a}, v) \right\}.
\end{aligned}
$$

To emphasize the role of P, we shall write $D^*(Q, g, \Psi(P))$ as $D^*(P)$. To see that $D^*(P)$ has zero expectation, we first note that all but the last line are expressed as an expression times a centered

function with respect to P, therefore they will have zero expectation under P; secondly, from remark 3.1, the last line also has zero expectation under P. Since we are operating under a saturated model, each component of $D^*(P)$ is in the tangent space, hence it is in fact the efficient influence curve.

To finish the proof of lemma 3.1, it suffices to see from (3.17) that the normalizing quantity is given by the inverse of

$$\begin{aligned}
&\sum_{(\mathbf{a},v)\in\mathscr{A}\times\mathscr{V}} \gamma_v(P)h(\mathbf{a},v)\phi(\mathbf{a},v)m_{\Psi(P)}(\mathbf{a},v)\left[1-m_{\Psi(P)}(\mathbf{a},v)\right]\phi(\mathbf{a},v)^\top \\
&= \int_{W_1} p(W_1)\sum_{v,\mathbf{a}} Q^V(v \mid \Delta=1,W_1)h(\mathbf{a},v)\phi(\mathbf{a},v)m_{\Psi(P)}(\mathbf{a},v)\left[1-m_{\Psi(P)}(\mathbf{a},v)\right]\phi(\mathbf{a},v)^\top \\
&= P\left\{\frac{I(\Delta=1)}{g^\Delta(1\mid W_1)}\sum_{\mathbf{a}} h(\mathbf{a},v)\phi(\mathbf{a},V)m_{\Psi(P)}(\mathbf{a},v)\left[1-m_{\Psi(P)}(\mathbf{a},v)\right]\phi(\mathbf{a},V)^\top\right\} \\
&= M(\Psi(P),P)
\end{aligned}$$

Proof of lemma 3.2: Robustness of $D^*(P)$

We first consider the case $Q(P)=Q_0$. All but the last line in $D^*(Q_0,g(P),\psi_0)$ are centered about a component of Q_0, therefore they will have expectation zero under P_0. On the other hand, the last line in $P_0D^*(Q_0,g(P),\psi_0)$ is $\int_{w_1} P_0(w_1)\sum_{(\mathbf{a},v)\in\mathscr{A}\times\mathscr{V}} p_0(v\mid\Delta=1,w_1)\tilde{h}(\mathbf{a},v)\left\{\frac{\eta_{\mathbf{a},v}(Q_0)}{\gamma_v(Q_0)}-m_{\psi_0}(\mathbf{a},v)\right\}$, which equals zero by definition of ψ_0.

Next, we consider the case $g(P)=g_0$. By telescoping the sums in (3.11) and applying definition of $Q^{\mathbf{a},1}_{t=0}(P_0)$, we obtain

$$\begin{aligned}
&P_0D^*(Q(P),g_0,\psi_0) \\
&= E_{P_0}\left\{\sum_{(\mathbf{a},v)\in\mathscr{A}\times\mathscr{V}} p_0(V=v\mid(\Delta=1,W_1)\tilde{h}(\mathbf{a},v)Q^{\mathbf{a},1}_{t=0}(P_0)(V=v,W_1)\right\} \\
&- E_{P_0}\left\{\sum_{(\mathbf{a},v)\in\mathscr{A}\times\mathscr{V}} p_0(V=v\mid(\Delta=1,W_1)\tilde{h}(\mathbf{a},v)m_{\psi_0}(\mathbf{a},v)\right\},
\end{aligned}$$

which equals 0 by remark 3.1.

Proof of lemma 3.3

To see that D^Δ is a score, up to a normalizing matrix, on the model where g^Δ is known, we repeat the steps in section 3.8, but this time we only consider the the class of 1-dimensional submodels indexed by

$$R_P = \left\{r = r' - E_P(r'\mid\Delta,W_1)+E_P(r'\mid W_1)-E_P(r') : r'\in\mathscr{H}_P\right\}\subset\mathscr{H}_P.$$

It is straight forward to verify that $E_P(r \mid \Delta = 1, W_1) = 0 = E_P(r \mid W_1)$ for such $r \in R_P$. Therefore, $D_{\gamma_v(P)}$ in this case becomes $\frac{I(\Delta=1)}{g^\Delta(1|W_1)} I(V = v) - \gamma_v(P)$ and

$$\begin{aligned} D_{\eta_{\mathbf{a},v}}(P) = & \frac{I(\Delta = 1)}{g^\Delta(1 \mid W_1)} I(V = v) \sum_{t=1}^{K} C_t^{\mathbf{a}} \left\{ Q_{t+1}^{\mathbf{a},1}(\mathbf{L}_t, v, W_1) - Q_t^{\mathbf{a},1}(\mathbf{L}_{t-1}, v, W_1) \right\} \\ & + \frac{I(\Delta = 1)}{g^\Delta(1 \mid W_1)} I(V = v) \left\{ Q_{t=1}^{\mathbf{a},1}(L_0, v, W_1) - Q_{t=0}^{\mathbf{a},1}(v, W_1) \right\} \\ & + \frac{I(\Delta = 1)}{g^\Delta(1 \mid W_1)} I(V = v) Q_{t=0}^{\mathbf{a},1}(v, W_1) - \eta_{\mathbf{a},v}. \end{aligned}$$

The rest is straightforward from equation (3.16).

To see that $PD^\Delta(Q, g, \Psi(P)) = 0$, it suffices to show that $PD_W^\Delta(Q, g, \Psi(P)) = 0$:

$$\begin{aligned} & P \left\{ \frac{I(\Delta = 1)}{g^\Delta(1 \mid W_1)} \sum_{\mathbf{a}} h(\mathbf{a}, V) \phi(\mathbf{a}, V) \left[Q_{t=1}^{\mathbf{a},1}(L_0, V, W_1) - m_{\Psi(P)}(\mathbf{a}, V) \right] \right\} \\ & = P \left\{ \frac{I(\Delta = 1)}{g^\Delta(1 \mid W_1)} \sum_{\mathbf{a}} h(\mathbf{a}, V) \phi(\mathbf{a}, V) \left\{ Q_{t=1}^{\mathbf{a},1}(L_0, V, W_1) - Q_{t=0}^{\mathbf{a},1}(V, W_1) \right\} \right\} \\ & + P \left\{ \frac{I(\Delta = 1)}{g^\Delta(1 \mid W_1)} \sum_{\mathbf{a}} h(\mathbf{a}, V) \phi(\mathbf{a}, V) \left[Q_{t=0}^{\mathbf{a},1}(V, W_1) - m_{\Psi(P)}(\mathbf{a}, V) \right] \right\}. \end{aligned}$$

The first term in the right hand side of the equality is zero by definition of $Q_1^{\mathbf{a},1}$; the second term is zero by remark 3.1. There D^Δ is a valid estimating function.

To see its robustness properties under $g^\Delta = g^\Delta(P_0)$, first consider the case of $Q_t^{\mathbf{a},1}(P) = Q_t^{\mathbf{a},1}(P_0)$ for $t = 1, \ldots, K$. Trivially, $P_0 D_t^*(Q, g) = 0$ for each $t \geq 1$ by definition of $Q_t^{\mathbf{a},1}$. On the term D_W^Δ, we have

$$\begin{aligned} & P_0 \left\{ \frac{I(\Delta = 1)}{g^\Delta(P_0)(1 \mid W_1)} \sum_{\mathbf{a}} \tilde{h}(\mathbf{a}, V) \left[Q_{t=1}^{\mathbf{a},1}(P_0)(L_0, V, W_1) - m_{\psi_0}(\mathbf{a}, V) \right] \right\} \\ & = P_0 \left\{ \sum_v Q^V(P_0)(v \mid \Delta = 1, W_1) \sum_{\mathbf{a}} \tilde{h}(\mathbf{a}, V) \left[Q_{t=0}^{\mathbf{a},1}(P_0)(V = v, W_1) - m_{\psi_0}(\mathbf{a}, V) \right] \right\} = 0, \end{aligned}$$

per definition of $Q_{t=0}^{\mathbf{a},1}$ and ψ_0.

Next, we consider the case of $g(P) = g_0$. By telescoping the sums in (3.12) and applying definition of $Q_{t=0}^{\mathbf{a},1}$, we again obtain

$$\begin{aligned} P_0 D^\Delta(Q, g_0, \psi_0) = & P_0 \left\{ \sum_{(\mathbf{a},v) \in \mathscr{A} \times \mathscr{V}} p_0(V = v \mid (\Delta = 1, W_1) \tilde{h}(\mathbf{a}, v) Q_{t=0}^{\mathbf{a},1}(P_0)(V = v, W_1) \right\} \\ & - P_0 \left\{ \sum_{(\mathbf{a},v) \in \mathscr{A} \times \mathscr{V}} p_0(V = v \mid (\Delta = 1, W_1) \tilde{h}(\mathbf{a}, v) m_{\psi_0}(\mathbf{a}, v) \right\} = 0, \end{aligned}$$

This completes proof of lemma 3.3.

Part III

Semiparametric Inference for Mediation Analysis

Chapter 4

Targeted Maximum Likelihood Estimation of Natural Direct Effects

4.1 Introduction

The causal effect of an exposure (or *treatment*) on an outcome of interest is often times mediated by intermediate variables (*mediator*). In many causal inference problems, one is interested in the *direct* effect of such exposure on the outcome, not mediated by the effect of the intermediate variables. Robins and Greenland (1992) and Pearl (2001) defined two types of direct effects under the counterfactual framework. The *controlled* direct effect refers to the effect of the exposure on the outcome under an idealized experiment where the mediator is set to a given constant value, whereas the *natural* (or *pure*) direct effect pertains to an experiment where the mediator is set to its would-be value under a reference (null) exposure level. The definition of these causal effects are based on counterfactual outcomes that are not fully observed, therefore they are not always identifiable from the observed data. Identifiability conditions are studied extensively in Robins and Greenland (1992), Pearl (2001), Robins (2003), van der Laan and Petersen (2004), Petersen, Sinisi, and van der Laan (2006), Hafeman and VanderWeele (2010), Imai, Keele, and Yamamoto (2010), Robins and Richardson (2010), and Pearl (2011).

Prior to the formal frameworks developed by Robins and Greenland (1992) and Pearl (2001), the social science literature had proposed the use of parametric linear structural equations in mediation analysis (e.g. Baron and Kenny (1986)), where the outcome response and mediator response are each modeled using linear main term regression on their parent nodes, and the direct and indirect effects are defined and estimated in terms of the coefficients in these regression equations. The limited causal validity of this parameter due to its dependence on model specification (e.g. no-interactions and linearity assumptions) is discussed in Kaufman, Maclehose, and Kaufman (2004). The developments of Robins and Greenland (1992) and Pearl (2001), and the identifiability studies that followed suit, address definition and identification of direct and indirect effects in causal models that do not put restrictions on the distribution of the observed data, allowing one to separate the

identification problem from the estimation problem.

Several approaches to the estimation problem are available in the current literature. A likelihood-based estimator approach (the g-computation formula) builds upon the identifiability results using a substitution estimator plugging in maximum likelihood based estimates of the relevant components of the data generating distribution. The natural direct effect can be identified as a function of the marginal covariate distribution, the conditional mediator density, and the conditional mean outcome (e.g. Robins and Greenland (1992), Pearl (2001), Robins (2003) and van der Laan and Petersen (2004), Petersen et al. (2006)). When all of these components of the data generating distribution are estimated consistently, the resulting g-computation estimate is unbiased and efficient. However, if either of these components is inconsistent, the effect estimate will be biased. VanderWeele and Vansteelandt (2010) illustrated how this approach can be applied to the estimation of natural direct effect odds ratio of rare outcomes. The use of (sequential) g-computation in structural nested models for estimation of controlled direct effects is proposed in Vansteelandt (2009). A second approach to causal effect estimation is based on the estimating equation methodology developed by Robins (1999), Robins and Rotnitzky (2001) and van der Laan and Robins (2003b). Under this approach, a score is expressed as a function of the parameter of interest ψ and a nuisance parameter η (whenever such representation is possible); if the resulting estimating equation, as an equation in the variable ψ, has a unique solution, the parameter estimate is given as the root to this equation. For most parameters arising from causal inference, the efficient influence curve under a nonparametric model is a robust estimating function (i.e. unbiased against mis-specification of specific components of the likelihood), therefore the resulting effect estimate shares the same robustness properties. In van der Laan and Petersen (2008), an application of this approach to a generalized class of direct effects using marginal structural models was discussed. The parameter studied in that work is a population mean of a subject-specific average controlled direct effect, averaged with respect to a user-supplied conditional mediator density given null exposure and individual covariates. If the supplied conditional mediator density is the true conditional mediator density of the data generating process, then the parameter of van der Laan and Petersen (2008) evaluates to the same value as the natural direct effect parameter. However, even in such case, these two parameters are not the same maps on the model since the former is a map indexed by the supplied mediator density and therefore is a function of the outcome expectation and marginal covariate distribution alone. As a consequence, the efficient influence curve of the parameter of van der Laan and Petersen (2008) is not the same as the efficient influence curve of the natural direct effect parameter. VanderWeele (2009) discussed more fully the use of marginal structural models with inverse probability weighting for estimation of the natural direct effect parameter. A third approach to causal effect estimation is the targeted maximum likelihood framework of van der Laan and Rubin (2006) and van der Laan and Rose (2011). For given estimators of relevant components of the likelihood P, one iteratively maximizes the likelihood (or minimize a loss) along a least favorable submodel through the initial estimators. The parameter estimate is given by evaluating the parameter map at the final estimator of the likelihood, thus providing a substitution estimator of the parameter of interest. By construction, the final estimate of the likelihood satisfies the efficient influence curve equation in the variable P. Therefore, the effect estimate also shares the robustness properties of

the efficient influence curve. In addition, the substitution principle incorporates global constraints of the statistical model that do not affect the form of the efficient influence curve; this allows for potential improvement in finite sample performance. van der Laan and Petersen (2008) also applied the targeted MLE procedure to their generalized class of direct effect parameters. Both the estimating equation approach and the targeted MLE approach in van der Laan and Petersen (2008) are robust (with respect to its parameter of interest) against mis-specification of the conditional mean outcome or mis-specification of the treatment mechanism. However, since its parameter of interest is indexed by the user-supplied conditional mediator density, if one is interested in the natural direct effect, then the user-supplied conditional mediator density in the method of van der Laan and Petersen (2008) must be correct. The use of propensity score matching in causal effect estimation was introduced in Rosenbaum and Rubin (1983). Application of propensity score in mediation analysis has also been proposed (e.g. Jo, Stuart, MacKinnon, and Vinokur (2011)).

Most recently, Tchetgen Tchetgen and Shpitser (2011a) derived the efficient influence curves (under a nonparametric model) for the various natural effect parameters, and established their general robustness properties and their implications on efficiency bounds. They also proposed semiparametric efficient, multiply robust estimators based on the estimating equation methodology using the efficient influence curve equation. We also refer the reader to that work for presentation of a sensitivity analysis framework to assess the impact of the ignorability assumption of the mediator variable on inference. In Tchetgen Tchetgen and Shpitser (2011b), the authors extended the theory to the case where one specifies a parametric model for the natural direct (indirect) effect conditional on a subset of baseline covariates.

In this chapter, we apply the targeted MLE framework of van der Laan and Rubin (2006) and van der Laan and Rose (2011) to the estimation of the natural direct effect of a binary exposure. The proposed estimator satisfies the efficient influence curve equation derived in Tchetgen Tchetgen and Shpitser (2011a). However, we note that the robustness conditions in Tchetgen Tchetgen and Shpitser (2011a) may be weakened (lemma 4.1), thereby placing less reliance on the estimation of the mediator density. This weaker version of robustness conditions is of particular interest when the mediator is high-dimensional, since it allows one to replace estimation of the conditional mediator density with objects that are easier (or at least with more available tools) to estimate. More precisely, the proposed estimator is asymptotically unbiased if either one of the following holds: i) the conditional mean outcome given exposure, mediator, and confounders, and the mediated mean outcome difference are consistently estimated; (ii) the exposure mechanism given confounders, and the conditional mean outcome are consistently estimated; or (iii) the exposure mechanism and the mediator density, or the exposure mechanism and the conditional distribution of the exposure given confounders and mediator, are consistently estimated. If all three conditions hold, then the effect estimate is asymptotically efficient. We also extend the results to the estimation of natural indirect effects. In addition, we discuss in detail conditions needed to ensure asymptotic linearity of the resulting estimator. These conditions should provide a guideline for situations where an influence curve based variance estimate is realistic.

This chapter is organized as follows: In section 2 we define formally the natural direct causal effect of a binary treatment on an outcome using the Non-Parametric Structural Equations Model

framework of Pearl (2009), and summarize its identifiability conditions. Based on the identifiability result, one may consider the natural direct effect parameter as a map from the model to the parameter space. We study this map and its efficient influence curve in greater detail in section 2.3. Section 3 describes how to construct a targeted MLE estimator for the natural direct effect of a binary treatment. Asymptotic properties of this estimator are summarized in section 4.3 and proved in the Appendix A. The estimation procedure in section 3 focuses on the targeted estimation of the conditional outcome expectation and the mediated mean outcome difference. An alternative procedure focusing on the conditional outcome expectation and the conditional mediator density is described in Appendix B. This alternative estimator shares the same asymptotic properties as the one proposed in section 3. Section 4 describes in greater detail two alternative estimation methodologies: the estimation equation framework of Robins (1999), and the maximum likelihood based g-computation framework. In section 5, we illustrate with simulations the robustness of the targeted MLE estimator against model mis-specifications. Section 4.6 extends analogously the discussions on identifiability, robustness, and estimation to the case of natural indirect effect. This chapter concludes with a summary and a few remarks.

4.2 Natural Direct Effect of a Binary Treatment

Causal Parameter

Consider n i.i.d observations of $O = (W, A, Z, Y)$, where W represents baseline covariates, A a binary treatment, Z represents a mediator of interest between the treatment and the outcome of interest Y. Let P_0 denote the distribution of O. We apply here the Non-Parametric Structural Equations Model (NPSEM) of Pearl (2009) to encode the causal relations under consideration. The NPSEM on a unit consists of a set of exogenous random variables U which are determined by factors outside the model, a set of endogenous variables X which are determined by variables inside the system $(U \cup X)$, and a set of unspecified deterministic functions $\{f_x : x \in X\}$ which encode for each $x \in X$ the variables that have direct influence on x. More specifically, in the present situation the causal relations are described by the NPSEM

$$\begin{aligned} W &= f_W(U_W) \\ A &= f_A(W, U_A) \\ Z &= f_Z(W, A, U_Z) \\ Y &= f_Y(W, A, Z, U_Y), \end{aligned}$$

where $X = (W, A, Z, Y)$ is the endogenous variable, and $U = (U_W, U_A, U_Z, U_Y)$ is the unobserved exogenous variable. This model defines a random variable (U, X) on the unit of observation, we denote its distribution by $P_{U,X}$.

The counterfactual variables or potential outcomes in the Rubin Causal Model (Rubin (1978), Rosenbaum and Rubin (1983) and Holland (1986)) can be represented as restrictions on the input of the functions f_x. For instance, the counterfactual $Z(a)$ is defined as the random variable

$Z(a) \equiv f_Z(W, A = a, U_Z)$, and can be interpreted as the mediator variable that the unit would have had if the exposure had been a. In particular, $Z(a)$ is a random variable through U_W and U_Z. Similarly, $Y(a', Z(a))$ is the counterfactual outcome that results from setting $Y(a', Z(a)) \equiv f_Y(W, A = a', Z(a), U_Y)$, and can be interpreted as the individual's response if the exposure had been a' while the mediator variable had been identical to the one under exposure a. $Y(a', Z(a))$ is a random variable through U_W, U_Z and U_Y.

Under the NPSEM, a causal parameter of interest is defined as a function of the distribution $P_{U,X}$. More specifically, the *natural direct causal effect* is defined as

$$\Psi(P_{U,X}) = E\left[Y(1, Z(0)) - Y(0, Z(0))\right].$$

This causal parameter can be interpreted from the following hypothetical experiment: one randomly assigns each subject to treatment or control, while always setting the subject's mediator variable to its value under no treatment, and then takes the difference in mean outcome between the treated and control cohort.

Identifiability

We will use the notation $Z(A)$ to denote the unintervened $Z = f_Z(W, A, U_Z)$, which is random through U_W, U_A, U_Z. Similarly, the unintervened $Y(A, Z(A))$ is random through U_W, U_A, U_Z, U_Y. Under experimental or observational studies, for each unit, the investigator only observes the outcome and mediator response under the unit's actual exposure. In other words, the observation is in fact $O = (W, A, Z(A), Y(A, Z(A))$. Hence, the causal parameter $\Psi(P_{U,X})$ is not always identifiable from the observed data.

Conditions under which the natural direct effect (or natural effects in general) will be identifiable were addressed extensively in Robins and Greenland (1992), Pearl (2001), Robins (2003), Petersen et al. (2006), Imai et al. (2010), Robins and Richardson (2010) and Pearl (2011). In particular, Pearl (2001) gave the following identifiability conditions: If randomization assumptions

A1. For all values (a, z), $Y(a, z)$ given W is identifiable,

A2. For all values of a, $Z(a)$ given W is identifiable,

and the conditional independence assumption

A3. For all $a \neq a', z$, $Y(a', z)$ is independent of $Z(a)$ given W

are satisfied, then the causal effect $\Psi(P_{U,X})$ can be expressed as a function of the observed data generating distribution P_0:

$$\begin{aligned}
&\Psi(P_{U,X}) \overset{A1,A2,A3}{=} \Psi(P_0) \\
&\equiv E\left\{ \sum_z \left[E(Y|W, A = 1, Z = z) - E(Y|W, A = 0, Z = z)\right] p(z|W, A = 0) \right\}. \qquad (4.1)
\end{aligned}$$

In the following sections, we will focus on the estimation of this statistical parameter.

Many of these previous authors have established that the randomization assumptions A1 and A2 can be satisfied by requiring that (A,Z) is independent of $Y(a,z)$, given W, and A is independent of $Z(a)$, given W. These can be ensured by measuring sufficient covariates to control for confounding of the effects of treatment on outcome, treatment on mediator, and mediator on outcome. As a result, the distributions of $Y(a,z)$ and $Z(a)$ will be identifiable within covariate stratum.

Petersen et al. (2006) showed that A3 can be weakened to a conditional mean independence

$$E(Y(1,z)-Y(0,z)|W) = E(Y(1,z)-Y(0,z)|W,Z(0)=z).$$

Still, it was recognized in Pearl (2001) that the conditional counterfactual independence is in general difficult to interpret. Imai et al. (2010) offered a stronger version of assumption A3 which is more interpretable: $Y(a',z)$ is independent of Z given W and $A=a$. This new version implies assumption A3, but the converse is not necessarily true. Robins and Richardson (2010) established that in general condition A3 cannot be enforced by randomized experiments. In such cases, what kind of causal interpretations can the statistical parameter in (4.1) still offer? Note that under the randomization assumptions A1 and A2 alone, the statistical parameter (4.1) equals (e.g. Pearl (2001), van der Laan and Petersen (2008)):

$$\Psi(P_0) \stackrel{A1,A2}{=} E_{P_0}\left(\sum_z E(Y(1,z)-Y(0,z)|W)P(Z(0)=z|W)\right).$$

The quantity in the right hand side is the population mean of an average of subject-specific controlled direct effect $E(Y(1,z)-Y(0,z)|W)$, weighted by $P(Z(0)=z|W)$. However, while this quantity serves to provide a causal interpretation for the statistical parameter (4.1) in the absence of condition A3, it is certainly not the natural direct causal effect; therefore one should be cautious about putting it into the context of the traditional total effect decomposition.

The Natural Direct Effect Parameter

Let $\mathscr{M}$ denote a model containing the true data generating distribution P_0. For any $P \in \mathscr{M}$, the likelihood decomposes into

$$P(O) = P_W(W)P_A(A|W)P_Z(Z|W,A)P_Y(Y|W,A,Z).$$

For later convenience, we adopt the notations $g(A|W) \equiv P_A(A|W)$, $Q^W(W) \equiv P_W(W)$, $Q^Z(Z|W,A) \equiv P_Z(Z|W,A)$, and $\bar{Q}^Y(W,A,Z) \equiv E(Y|W,A,Z)$. Moreover, let $Q \equiv (Q^W, Q^Z, \bar{Q}^Y)$. The notations $Q(P_0)$ and $g(P_0)$ are reserved for the corresponding components of the true data generating distribution P_0. For a function $f(O)$, we will use Pf to denote the expectation of $f(O)$ under the probability distribution $P \in \mathscr{M}$. For instance, $P_0 f \equiv \sum_{o\in\mathscr{O}} f(o)dP_0(o)$ denotes the expectation of f under the true data generating distribution, while $P_n f \equiv \frac{1}{n}\sum_{i=1}^n f(o_i)$ denotes the empirical mean of f.

One may consider the natural direct effect parameter Ψ in (4.1) as a map

$$\begin{aligned}\Psi :& \mathscr{M} \to \mathbb{R} \\ & P \mapsto \Psi(P) = \Psi(Q) \equiv E_P\left[E_P\left(\bar{Q}^Y(W,1,Z) - \bar{Q}^Y(W,0,Z)|W,A=0\right)\right].\end{aligned}$$

We refer to the inner expectation above as the *(null level) mediated mean outcome difference*, and denote it by the map $P \mapsto \bar{Q}^Z(P)$, where

$$\bar{Q}^Z(P)(W) \equiv E_P\left(\bar{Q}^Y(W,1,Z) - \bar{Q}^Y(W,0,Z)\big|\, W,A=0\right). \tag{4.2}$$

This way, $\Psi(P) = E_P\left(\bar{Q}^Z(P)(W)\right)$. The parameter of interest (4.1) is this map evaluated at the true data generating distribution: $\psi_0 \equiv \Psi(P_0)$.

Efficient Influence Curve

Under a nonparametric model $\mathscr{M}$, for any $P \in \mathscr{M}$, the *efficient influence curve* (EIC) of Ψ at P, as derived in Tchetgen Tchetgen and Shpitser (2011a), is given by:

$$\begin{aligned}D^*(Q,g,\Psi(Q)) &= \left\{\frac{I(A=1)}{g(1|W)}\frac{Q^Z(Z|W,0)}{Q^Z(Z|W,1)} - \frac{I(A=0)}{g(0|W)}\right\}\left(Y - \bar{Q}^Y(W,A,Z)\right) \\ &+ \frac{I(A=0)}{g(0|W)}\left\{\bar{Q}^Y(W,1,Z) - \bar{Q}^Y(W,0,Z) - E_P\left(\bar{Q}^Y(W,1,Z) - \bar{Q}^Y(W,0,Z)|W,0\right)\right\} \\ &+ E_P\left(\bar{Q}^Y(W,1,Z) - \bar{Q}^Y(W,0,Z)|W,0\right) - \Psi(Q) \\ &= D_Y^* + D_Z^* + D_W^*.\end{aligned}$$

Note that the components D_Y^*, D_Z^*, D_W^* are respectively the projection of D^* onto the tangent subspaces corresponding to the components $P(Y|W,A,Z)$, $P(Z|W,A)$, $P(W)$ of the likelihood.

This efficient influence curve for a nonparametric model can also be derived by first considering $\Psi(P)$ as a function of $P = (Pf : f \in \mathscr{F})$, where $\mathscr{F}$ is a class of indicator functions $\mathscr{F} = \{I(w,a,z,y), I(w,a,z), I(w,a), I(w) : w \in \mathscr{W}, a \in \mathscr{A}, z \in \mathscr{Z}, y \in \mathscr{Y}\}$. For any given "vector" $h = (h(f) : f \in \mathscr{F})$, one can consider a directional derivative $\frac{d}{d\varepsilon}\Psi(P + \varepsilon h)|_{\varepsilon=0}$. The efficient influence curve is given by the directional derivative applied to the direction of $h = (f(O) - Pf : f \in \mathscr{F})$. In other words, it is given by $\sum_{f\in\mathscr{F}} \frac{\partial\Psi(P)}{\partial Pf}(f(O) - Pf)$. A more detailed exposition can be found in van der Laan and Rose (2011).

Robustness of the efficient influence curve

The general robustness conditions of the EIC were given in Tchetgen Tchetgen and Shpitser (2011a): (i) the mediator density $Q^Z(Z|W,A)$ and the conditional mean outcome $\bar{Q}^Y(W,A,Z)$ are both correct; (ii) the conditional mean outcome and the exposure mechanism $g(A|W)$ are both correct; or (iii) the exposure mechanism and the mediator density are both correct. We note below that conditions (i) and (iii) may be weakened to accommodate difficulties in estimation of the mediator

density. In fact, the estimation of Q^Z may be replaced by estimation of a conditional probability of treatment. This is particularly appealing when Z is high dimensional. We summarize these in the following lemma and its subsequent remarks. The proof of this lemma is straightforward from the form of the efficient influence curve, and we refer the interested reader to the Appendix.

Lemma 4.1 (Robustness of the efficient influence curve).

Suppose there exists constants $1 > \delta, \delta' > 0$ such that $g(A = 1|W) < 1 - \delta$ and $Q^Z(Z|W,1) < 1 - \delta'$ a.e. over the support of W and Z. The efficient influence curve is a robust estimating function for the parameter at P_0, in the sense that

$$P_0 D^*(Q, g, \psi_0) = 0,$$

if either of the following holds:

(i) The conditional mean outcome $\bar{Q}^Y = E(Y|W,A,Z)$, and the mediated mean outcome difference $\bar{Q}^Z(P) = E_P(\bar{Q}^Y(W,1,Z) - \bar{Q}^Y(W,0,Z)|W,0)$ are correct.

(ii) The exposure mechanism $g(A|W)$, and the conditional mean outcome are correct.

(iii) The exposure mechanism and conditional mediator density $Q^Z(Z|W,A)$, or the exposure mechanism and the conditional distribution of treatment given mediator and covariates $\gamma(A,W,Z) \equiv p(A|W,Z)$, are correct.

Condition (i) follows from the fact that, given $\bar{Q}^Y$, we only need a conditional expectation of $\bar{Q}^Y(W,1,Z) - \bar{Q}^Y(W,0,Z)$ under $Q^Z(Z|W,0)$. Therefore, consistent estimation of $Q^Z(P_0)$ per se is not necessary to obtain consistent estimator of $\bar{Q}^Z(P_0)$, as long as one has a consistent estimator $\bar{Q}^Y_n$ of $\bar{Q}^Y(P_0)$ and an optimal procedure to regress the difference $\bar{Q}^Y_n(W,1,Z) - \bar{Q}^Y_n(W,0,Z)$ on W among the control observations. Condition (iii) is a consequence of the fact when g is correct, dependence on consistent estimation of Q^Z is only through $\frac{Q^Z(Z|W,0)}{Q^Z(Z|W,1)}$, which can be consistently estimated using either Q^Z or combining ratios of $g(A|W)$ and $p(A|W,Z)$.

When Z is high-dimensional, few tools are available to estimate the conditional mediator density $Q^Z(Z|W,A)$. On the other hand, there is abundant literature addressing estimation of conditional means. This can be used to estimate $\bar{Q}^Z(P)$, and conditional probabilities $\gamma(A|W,Z)$ of a categorical A. Lemma 4.1 implies in particular that estimation of $Q^Z(P_0)$ may be replaced by estimations of $g(P_0)$, $\gamma(P_0)$, and the conditional expectation $\bar{Q}^Z(P_0)$,

4.3 Targeted Maximum Likelihood Estimation for the Natural Direct Effect of a Binary Treatment

In general, under the framework of van der Laan and Rubin (2006) the construction of a targeted MLE (TMLE) estimator of a parameter of interest $\Psi(P_0) = \Psi(Q(P_0))$ calls for two sets of ingredients. For each component $Q_j(P)$ of $Q(P)$, one defines a uniformly bounded (w.r.t. the supremum

norm) *loss function* $L_j : \mathscr{Q}_j \to \mathscr{L}^\infty(K)$ satisfying

$$Q_{j,0} = \arg\min_{\mathscr{Q}_j} P_0 L_j(Q_j),$$

where $\mathscr{L}^\infty(K)$ is the class of functions of O with bounded supremum norm over a set of K containing the support of O under P_0. Given the loss function L_j, one defines a one-dimensional *parametric working submodel* $\{Q_j(P)(\varepsilon_j) : \varepsilon_j\} \subset \mathscr{M}$ passing through $Q_j(P)$ at $\varepsilon_j = 0$ with score $D_j^*(P)$ at $\varepsilon_j = 0$ that satisfies

$$\langle \frac{d}{d\varepsilon_j} L_j(Q_j(P)(\varepsilon_j)) \mid_{\varepsilon_j=0} \rangle \supset \langle D_j^*(P) \rangle,$$

where $\langle h \rangle$ denotes the linear span of a vector h. These result in a least favorable parametric submodel $Q(\varepsilon)$ through Q. For given initial estimator $(Q_n, \hat{g})$ of $(Q(P_0), g(P_0))$, the fluctuation parameter ε is fitted to minimize the empirical risk of $Q_n(\varepsilon)$, providing an updated estimator $Q_n(\hat{\varepsilon})$. This updating process is repeated until $\hat{\varepsilon} \approx 0$. The final estimator Q_n^* of $Q(P_0)$ is then used to obtain a substitution estimator $\Psi(Q_n^*)$ of $\Psi(Q(P_0))$. By its construction, the estimator Q_n^* satisfies the efficient influence curve equation $P_n D^*(Q_n^*, \hat{g}, \Psi(Q_n^*)) = 0$.

To specialize to the natural direct effect, we first note that the parameter of interest and the components D_Z^* and D_W^* of the efficient influence curve depend on Q^Z only through the mediated mean outcome difference $\bar{Q}^Z(P)$ as defined in (4.2). Secondly, the empirical marginal distribution Q_n^W of W is a consistent estimator of $Q^W(P_0)$ that readily solves the equation $P_n D_W^*(\bar{Q}^Z(P), Q_n^W) = 0$ for any $\bar{Q}^Z(P)$. Hence, the proposed estimator will focus on targeted estimation of $\bar{Q}^Y(P_0)(W,A,Z)$, and $\bar{Q}^Z(P_0)(W)$.

An alternative targeted estimation to the one proposed above is to targetedly estimate the conditional mediator density $Q^Z(P_0)$ instead of the mediated mean outcome difference $\bar{Q}^Z(P_0)$. We refer the interested reader to Appendix B for this alternative approach. The key difference between the proposed and the alternative targeting procedures lies in that the former defines a loss function and parametric working submodel for the mediated mean outcome difference $\bar{Q}^Z(P)$, whereas the latter defines a loss function and parametric working submodel for the conditional mediator density Q^Z and then estimates the mediated mean outcome difference $\bar{Q}^Z(P_0)$ by plugging in the targeted mediator density and the targeted $\bar{Q}^Y$. We note that the bias variance trade-off in the proposed targeting procedure is more optimal over the alternative procedure for estimating the ultimate component of interest, which is the mediated mean outcome difference.

Construction of the Targeted MLE

Loss functions and parametric working submodels

Suppose for now that Y is binary or continuous and bounded. In the latter case, without loss of generality we may assume that Y is bounded in $(0,1)$. We consider the minus-loglikelihood loss

function for $\bar{Q}^Y$:

$$L_Y(\bar{Q}^Y)(O) = -\log\left(\bar{Q}^Y(W,A,Z)^Y(1-\bar{Q}^Y(W,A,Z))^{(1-Y)}\right). \tag{4.3}$$

Under this loss function, consider the logistic working submodel

$$\bar{Q}^Y(\varepsilon) \equiv expit\left(logit(\bar{Q}^Y) + \varepsilon C_Y\right),$$

where $C_Y(Q^Z,g)(O) = \left\{\frac{I(A=1)}{g(1|W)}\frac{Q^Z(Z|W,0)}{Q^Z(Z|W,1)} - \frac{I(A=0)}{g(0|W)}\right\}$. Note that this submodel $\bar{Q}^Y(\varepsilon)$ depends on the components Q^Z and g, but we suppress that in the notation. This submodel satisfies

$$\frac{d}{d\varepsilon}L_Y\left(\bar{Q}^Y(\varepsilon)\right)|_{\varepsilon=0} = D_Y^*(\bar{Q}^Y,Q^Z,g). \tag{4.4}$$

For a given $\bar{Q}^Y$, the difference $\bar{Q}^Y(W,Z) \equiv \bar{Q}^Y(W,1,Z) - \bar{Q}^Y(W,0,Z)$ is also bounded. Without loss of generality, we may also assume it is bounded between $(0,1)$. Let the loss function for $\bar{Q}^Z$ be

$$L_Z(\bar{Q}^Z)(O) = \\ -I(A=0)\log\left((\bar{Q}^Z(W))^{\bar{Q}^Y(W,Z)}(1-\bar{Q}^Z(W))^{1-\bar{Q}^Y(W,Z)}\right).$$

Under this loss function, the logistic working submodel

$$\bar{Q}^Z(\varepsilon) \equiv expit\left(logit\left(\bar{Q}^Z\right) + \varepsilon C_Z\right),$$

with $C_Z(g)(O) = \frac{1}{g(0|W)}$, satisfies

$$\frac{d}{d\varepsilon}L_Z\left(\bar{Q}^Z(P)(\varepsilon)\right)|_{\varepsilon=0} = D_Z^*(\bar{Q}^Z(P),\bar{Q}^Y,g). \tag{4.5}$$

The dependence of $\bar{Q}^Z(P)(\varepsilon)$ on g is again suppressed in our notation.

Note that linear transformations onto the unit interval may be needed in order to use the loss functions L_Y and L_Z. However, since the parameter of interest and the components of the efficient influence curve are linear in $\bar{Q}^Y$ and $\bar{Q}^Z(P)$, the necessary linear transformations and their inverse maps do not affect the properties of the estimators.

In settings where Y is not bounded, one may instead use the squared error loss functions

$$L_Y(\bar{Q}^Y)(O) = \left(Y - \bar{Q}^Y(W,A,Z)\right)^2,$$

and

$$L_Z(\bar{Q}^Z)(O) = I(A=0)\left(\bar{Q}^Y(W,Z) - \bar{Q}^Z(W)\right)^2;$$

and corresponding parametric working submodels

$$\bar{Q}^Y(\varepsilon) = \bar{Q}^Y + \varepsilon C_Y$$

and

$$\bar{Q}^Z(\varepsilon) = \bar{Q}^Z + \varepsilon C_Z.$$

However, compared to the minus loglikelihood losses, this choice of loss functions and the corresponding parametric working submodels may result in estimators that are more sensitive to near positivity violations (Gruber and van der Laan (2010), Gruber and van der Laan (2011)). Therefore, in such situations it would be more sensible to bound Y by the range of the observed data, and apply the minus loglikelihood losses above.

Implementation

Let P_n denote the empirical distribution of n i.d.d observations of O. Let g_n, $\bar{Q}_n^Y$ and Q_n^Z, be initial estimators of $g(P_0)$, $\bar{Q}^Y(P_0)$ and $Q^Z(P_0)$, respectively. Let $\varepsilon_n^Y \equiv \arg\min_\varepsilon P_n L_Y\left(\bar{Q}_n^Y(\varepsilon)\right)$ be the optimal ε which minimizes the empirical risk. We are reminded that, though not shown in the notation, the estimators (Q_n^Z, g_n) are used in constructing $\bar{Q}_n^Y(\varepsilon)$. The update

$$\bar{Q}_n^{Y,*} \equiv \bar{Q}_n^Y(\varepsilon_n^Y) \tag{4.6}$$

is the *targeted MLE estimator of* $\bar{Q}^Y(P_0)$.

Next, we obtain an initial estimator $\bar{Q}_n^Z$ of $\bar{Q}^Z(P_0)$, with respect to the targeted expectant $\bar{Q}_n^{Y,*}(W,1,Z) - \bar{Q}_n^{Y,*}(W,0,Z)$, either using the density-based approach, which uses the density estimate Q_n^Z to obtain the conditional mean outcome difference, or using a regression-based approach, which regresses the difference $\bar{Q}_n^Y(W,1,Z) - \bar{Q}_n^Y(W,0,Z)$ on W among control observations. The optimal ε is given by $\varepsilon_n^Z = \arg\min_\varepsilon P_n L_Z\left(\bar{Q}_n^Z(\varepsilon)\right)$. We are reminded that, though not shown in the notation, the estimator g_n is used in constructing $\bar{Q}_n^Z(\varepsilon)$. The update

$$\bar{Q}_n^{Z,*} \equiv \bar{Q}_n^Z(\varepsilon_n^Z) \tag{4.7}$$

is the *targeted MLE estimator of* $\bar{Q}^Z(P_0)$. The *targeted MLE estimator of* $\psi_0 = E_{P_0}\left(\bar{Q}^Z(P_0)(W)\right)$ is thus given by

$$\psi_n^* \equiv \frac{1}{n}\sum_{i=1}^n \bar{Q}_n^{Z,*}(W_i). \tag{4.8}$$

Let $Q_n^* \equiv \left(\bar{Q}_n^{Y,*}, \bar{Q}_n^{Z,*}, Q_n^Z, Q_n^W\right)$, where Q_n^W is the empirical distribution of W. It follows from (4.4) and (4.5) that $P_n D^*(Q_n^*, g_n) = 0$. From this stems the asymptotic properties of ψ_n^* we describe in section 4.3.

Remark 4.1 (implementation)**.** *When Z is high-dimensional, and A is categorical, consistent estimation of $\gamma(A,W,Z) = p(A|W,Z)$ may be more attainable than consistent estimation of $Q^Z(Z|W,A)$. In such case, instead of using an estimator of Q^Z to estimate the ratio $Q^Z(Z|W,0)/Q^Z(Z|W,1)$ in the targeting step of $\bar{Q}^Y$, one can use an estimator $\frac{\gamma_n(A=0|W,Z)}{g_n(A=0|W)}\frac{g_n(A=1|W)}{\gamma_n(A=1|W,Z)}$. By the same reason, one should use a regression-based approach to obtain the initial estimate $\bar{Q}_n^Z$. This way, estimation of Q^Z may be avoided if one has available optimal estimators g_n and $\gamma_n(A|W,Z)$, and a regression-based estimator $\bar{Q}_n^Z$. From lemma 4.1, we see that this still allows for robust estimation.*

Statistical Inference for TMLE

Since the proposed targeted MLE estimator satisfies the efficient influence curve equation, lemma 4.1 implies in particular that the estimator is asymptotically unbiased if either of the following is true: (i) The conditional outcome expectation $\bar{Q}_n^{Y,*}$ and the mediated mean outcome difference $\bar{Q}_n^{Z,*}$ are consistent; (ii) the treatment mechanism g_n and the conditional outcome expectation $\bar{Q}_n^{Y,*}$ are consistent; (iii) the treatment mechanism g_n and the conditional mediator density $Q_n^Z(Z|W,A)$, or the treatment mechanism and $\gamma_n(A|W,Z)$, are consistent. These properties are illustrated in the simulations section below.

Under certain regularity conditions, an estimator that satisfies a given estimating equation will be asymptotically linear with influence curve given by the estimating function (e.g. Bickel et al. (1997), van der Vaart (1998a), van der Laan and Robins (2003b), Tsiatis (2006), Kosorok (2008)). In this case, the central limit theorem implies that one can obtain an asymptotic variance estimate of the said estimator using the variance estimate of its influence curve. In particular, the asymptotic variance of $\sqrt{n}(\psi_n^* - \psi_0)$ can be estimated by the sample variance $\Sigma_n^* \equiv \hat{\mathrm{Var}}(D^*(Q_n^*, g_n))$. We can construct asymptotically conservative confidence intervals of level $(1-\alpha)$ as $\left[\psi_n^* \pm \xi_{1-\alpha/2}(\Sigma_n^*/n)^{1/2}\right]$, where $\xi_{1-\alpha/2}$ is the $(1-\alpha/2)$-quantile of the standard normal distribution.

4.4 Some Existing Estimation Methodologies

In this section, we describe how the estimating equation and the g-computation approaches can be applied to the natural direct effect of a binary exposure, and contrast their theoretical properties with those of the proposed targeted estimator.

Estimating Equation Estimator

Under the estimating equation (EE) based approach (Robins (1999), Robins and Rotnitzky (2001), van der Laan and Robins (2003b)), one may use the efficient influence curve $D^*(P)$ under a nonparametric model as an estimating function of ψ, if i) $D^*(P)$ can be expressed as a function of ψ and some nuisance parameter η, i.e. $D^*(P) = D(\psi(P), \eta(P))$, for some function D, and ii) the solution to the resulting equation in the variable ψ is unique. When these requirements hold, an estimate of the parameter is given by the root of the resulting estimating equation, i.e. ψ_n is defined as the solution to the equation $P_n D^*(\eta_n, \psi_n) = 0$.

An estimator of the natural direct effect under this framework is provided in Tchetgen Tchetgen and Shpitser (2011a). For given estimators $\bar{Q}_n^Y$, Q_n^Z, g_n, and a density-based estimator $\bar{Q}_n^Z$ for

$\bar{Q}^Z(P_0)$, the EE estimator for the natural direct effect is given by

$$\psi_n^{EE} \equiv \frac{1}{n}\sum_{i=1}^{n}\Bigg\{\left(\frac{I(A_i=1)}{g_n(1|W_i)}\frac{Q_n^Z(Z_i|W_i,0)}{Q_n^Z(Z_i|W_i,1)} - \frac{I(A_i=0)}{g_n(0|W_i)}\right)\left(Y_i - \bar{Q}_n^Y(W_i,A_i,Z_i)\right)$$
$$+\frac{I(A_i=0)}{g_n(0|W_i)}\left(\bar{Q}_n^Y(W_i,1,Z_i) - \bar{Q}_n^Y(W_i,0,Z_i) - \bar{Q}_n^Z\right)$$
$$+\bar{Q}_n^Z\Bigg\}$$

We remind the reader again that in the present chapter, $\bar{Q}_n^Z$ may not need to use Q_n^Z, but will surely make use of $\bar{Q}_n^Y$.

For $Q_n \equiv \left(\bar{Q}_n^Y, Q_n^Z, \bar{Q}_n^Z\right)$, this EE estimator solves the efficient influence curve equation

$$P_n D^*\left(Q_n, g_n, \psi_n^{EE}\right) = 0.$$

Therefore, the ψ_n^{EE} estimator and the proposed TMLE estimator share the same asymptotic properties that are inherited from the efficient influence curve. By the same token, they are both sensitive to extreme values of the treatment model, such as in the case of near positivity violations. This was demonstrated in Kang and Schafer (2007). Indeed, in the case of natural direct effect, when $g_n(A_i|W_i)$ is small for some observations, the estimated D_Y^* component of the efficient influence curve will be large; this problem is exacerbated if $A_i = 0$, in which case the estimated D_Z^* is also large. Consequently, the EE estimator may yield estimates that are out of the bounds of the parameter, since constraints such as bounds of the parameter are not reflected in the functional form of the efficient influence curve. The proposed targeted estimator using a logistic working submodel (introduced in Gruber and van der Laan (2010)) aims to provide more stable estimates through the combination of a unit linear transformation, which implicitly estimates the boundary of the parameter domain, and the virtue of the substitution principle.

G-computation Estimator

The sensitivity to near positivity violation of the TMLE estimator and the ψ_n^{EE} estimator stems from the use of inverse probability weightings in the efficient influence curve. A g-computation approach based on the identifiability result in (4.1) avoids this inverse weighting. More specifically, for $\bar{Q}_n^Y$ and Q_n^Z likelihood based estimators of the outcome expectation and mediator density, respectively, consider a g-computation estimator given by:

$$\psi_n^{Gcomp} = \frac{1}{n}\sum_{i=1}^{n}\left(\bar{Q}_n^Y(W_i,1,Z_i) - \bar{Q}_n^Y(W_i,0,Z_i)\right)Q_n^Z(Z_i|W_i,0).$$

This estimator can be similarly defined using a regression-based $\bar{Q}_n^Z$ as $\psi_n^{Gcomp} = \frac{1}{n}\sum_{i=1}^{n}\bar{Q}_n^Z(W_i)$. Unlike the robust TMLE and ψ_n^{EE} estimators, the consistency of the g-computation estimator relies

on correct specification of both the outcome expectation, and mediator density (or the regression procedure for the mediated mean outcome difference). In the case of these likelihood-based estimates being correct, the resulting ψ_n^{Gcomp} is more efficient than the two robust estimators. However, even though this g-computation estimator does not use inverse probability weighting explicitly, it can still be affected by data sparsity, since the quality of the mean outcome estimate (even under the correct specification) is sensitive to the overlap between the empirical covariate distribution of the treated cohort and the empirical covariate distribution of the control cohort.

4.5 Simulation Study

In this section we evaluate the performance of the TMLE estimator, the EE estimator, and the Gcomp estimator under model mis-specification and data sparsity. From lemma 4.1, one expects to see that, in the absence of positivity violations, the TMLE and EE estimator are robust against model mis-specifications.

Simulation Schemes

The following three data generating schemes are used. The mediator variable Z is discrete with three categories: $Z \in \{0,1,2\}$. Each scheme has a version with a binary outcome Y and a version with a continuous and bounded outcome Y. Simulations 2 and 3 differ from simulation 1 in their mediator density and treatment mechanism, respectively.

1. **Simulation 1**: no positivity violations.

$$
\begin{aligned}
&W \sim U(0,2)\\
&P(A=1 \mid W) = \text{expit}\left(-1+2W-0.08W^2\right)\\
&p(Z=0) = \text{expit}(-0.2+0.5A+0.3A\times W+0.7W-1.5W^2),\\
&p(Z=1|Z\neq 0) = \text{expit}(-0.2+0.4A+.8A\times W+0.4W-2.5W^2)\\
&Y \sim Bern\Big(expit(-2+A-W+W^2+Z+0.8A\times W-A\times W^2\\
&\qquad -0.5A\times Z+0.7A\times Z^2)\Big)
\end{aligned}
$$

The treatment probability $g_A(A=1|w)$ is bounded in $(0.26, 0.94)$; the mediator density $Q^Z(z|A=1,w)$ is bounded in $(0.0005, 0.9753)$ for any z and w, whereas the ratio $Q^Z(z|A=0,w)/Q^Z(z|A=1,w)$ is bounded in $(0.13, 2.02)$.

The parameters of interest are $\psi_0 = 0.2585079$. The semiparametric efficiency bounds is $var(D^*(P_0)) \approx 1.157$.

2. **Simulation 2**: larger effect of treatment on the distribution of mediator.

$$p(Z=0) = \text{expit}(-2-2A-0.5A\times W+3W-W^2),$$
$$p(Z=1|Z\neq 0) = \text{expit}(1-4A-A\times W+W+W^2).$$

Distributions for W,A,Y are the same as simulation 1. The mediator density $Q^Z(z|w,A=1)$ ranges in $(0.017,0.081)$ for $Z=0$, ranges in $(0.046,0.697)$ for $Z=1$ and ranges in $(0.256,0.936)$ for $Z=2$. The ratio $\frac{Q^Z(z|w,A=0)}{Q^Z(z|w,A=1)}$ ranges in $(6.583,10.543)$ for $Z=0$, ranges in $(0.717,13.826)$ for $Z=1$ and ranges in $(0.0018,0.253)$ for $Z=2$.

The parameters of interest are $\psi_0 = 0.12556476$ for the binary version, and the semiparametric efficiency bounds are $var(D^*(P_0)) \approx 3.721905$.

3. **Simulation 3**: near positivity violation the treatment mechanism.

$$A \sim Bern\left(expit(-2-3W+5W^2)\right).$$

Distributions for W,Z,Y are the same as simulation 1, therefore the values of the parameters of interest also remain the same. The treatment mechanism is bounded in $g_A(A=1|W) \in (0.0794,0.999994)$. Moreover, $g_A(A=1|W) > 0.99$ for $W > 1.5$.

Estimators

For each data generating distribution, initial maximum likelihood based estimators of the outcome expectation $\bar{Q}^Y(P_0)$, treatment mechanism $g(P_0)$ and mediator density $Q^Z(P_0)$ are obtained according to each of the three cases of model mis-specification in lemma 4.1, as well as the case where all models are correct. The model mis-specifications considered are as follows:

- Mis-specified outcome model is $Y \sim A+W+Z+A\times Z$, with binomial family (with logit link).
- Mis-specified mediator density is multinomial with $p(Z=0|A,W) \sim A$ and $p(Z=1|A,W,Z\neq 0) \sim A$, both from a binomial family with logit link.
- Mis-specified treatment mechanism is $A \sim W^2$ for simulations 1 and 2, and $A \sim W$ for simulation 3, both from a binomial family with logit link.

The estimators ψ_n^{Gcomp} and ψ_n^{EE} are implemented using the density based estimators as described in section 4.4.

The TMLE estimator ψ_n^* are constructed using these initial estimators under logistic working submodels. We consider two implementations of TMLE which differ in their initial estimator of the mediated mean outcome difference $\bar{Q}^{Z,*}(P_0)$. In TMLE 1, the initial estimator $\bar{Q}_n^Z$ is given by a plug-in estimator using Q_n^Z and $\bar{Q}_n^{Y,*}$. In TMLE 2, the initial estimator $\bar{Q}_n^Z$ is obtained by performing a main term parametric regression $(\bar{Q}_n^{Y,*}(W,1,Z) - \bar{Q}_n^{Y,*}(W,0,Z)) \sim W$ among the observations with

$A = 0$. With the data generating distributions under consideration, this initial estimator in TMLE 2 is incorrect regardless of the consistency of $\bar{Q}_n^Y$. However, from lemma 4.1, we expect TMLE 2 to be consistent in the cases (ii) and (iii) of lemma 1, in the absence of positivity violation.

Results

For each data generating distribution, 1000 samples of each size $n = 500$ and $n = 5000$ are generated. Bias, variance and means squared errors (MSE) for each sample size are estimated over the 1000 samples.

Simulation 1: No positivity violation

Recall that the parameters of interest are $\psi_0 = 0.2585079$, and the semiparametric efficiency bound is $var(D^*(P_0)) \approx 1.157$. Therefore, $var(D^*(P_0))/n \approx 2.314e-03$ and $2.314e-04$ for $n = 500$ and 5000, respectively.The results are detailed in table 4.1. When the outcome expectation and the mediator density are correctly specified, the robust estimators TMLE and EE provide little advantage over the Gcomp estimator in terms of bias or efficiency. However, when either the outcome expectation or the mediator density are mis-specified, TMLE and EE using a correct treatment mechanism provide substantial bias correction so that MSE is reducing at rate $1/n$. The two robust estimators behave similarly. Moreover, as predicted by lemma 4.1, TMLE 2, which utilizes a mis-specified initial estimator of the mediated mean outcome difference, behaves as well as TMLE 1 when the treatment mechanism is correct.

Simulation 2: Larger effect of treatment on mediator

Under this simulation scheme, the parameters of interest are $\psi_0 = 0.12556476$ and the efficiency bounds is $var(D^*(P_0)) \approx 3.721905$. Therefore, $var(D^*(P_0)/n$ are approximately $7.444e-03$ and $7.444e-04$ for $n = 500$ and 5000, respectively. In this simulation, the treatment has a moderately larger effect on the mediator distribution. Compared to simulation 1, this simulation scheme has a larger ratio of $Q^Z(z|0,w)/Q^Z(z|1,w)$ for categories of $Z = 0,1$ over a region of the sample space of W (details are explained previously). Results are in table 4.2. We see that in this case all estimators behave as expected as in the previous simulation. When implemented using the correct treatment mechanism, they provide bias reduction over g-computation estimator in the cases when either the mediator density or the outcome model are mis-specified. When the outcome model and mediator density are both correct, then Gcomp is consistent. In this case the TMLE and EE are also consistent but less efficient. In all cases, TMLE and EE behave similarly. We observe again that when the treatment mechanism is correct, TMLE 2, which utilizes a mis-specified initial estimator of the mediated mean outcome difference, behaves as well as TMLE 1.

Table 4.1: Simulation 1: no positivity violations

	Bias		Var		MSE	
n	500	5000	500	5000	500	5000
all correct						
Gcomp	6.35e-04	5.84e-04	2.45e-03	2.26e-04	2.45e-03	2.26e-04
TMLE1	2.39e-04	5.22e-04	2.499e-03	2.29e-04	2.5e-03	2.29e-04
TMLE2	3.10e-04	5.65e-04	2.525e-03	2.29e-04	2.53e-03	2.3e-04
EE	2.01e-04	5.23e-04	2.50e-03	2.29e-04	2.50e-03	2.29e-04
g misspec.						
TMLE	4.45e-04	4.69e-04	2.63e-03	2.37e-04	2.63e-03	2.38e-04
EE	7.29e-04	4.58e-04	2.75e-03	2.45e-04	2.75e-03	2.45e-04
Q^Z misspec.						
Gcomp	4.26e-02	4.08e-02	3.02e-03	2.77e-04	4.8e-03	1.94e-03
TMLE1	2.2e-04	5.69e-04	2.48e-03	2.28e-04	2.48e-03	2.28e-04
TMLE2	2.00e-04	6.2e-04	2.49e-03	2.29e-04	2.49e-03	2.29e-04
EE	2.71e-04	5.47e-04	2.49e-03	2.29e-04	2.49e-03	2.29e-04
$\bar{Q}^Y$ misspec.						
Gcomp	2.83e-02	2.83e-02	2.43e-03	2.26e-04	3.24e-03	1.02e-03
TMLE1	2.07e-04	5.45e-04	2.5e-03	2.29e-04	2.5e-03	2.29e-04
TMLE2	4.05e-04	5.66e-04	2.5e-03	2.3e-04	2.5e-03	2.3e-04
EE	3.72e-04	5.49e-04	2.5e-03	2.29e-04	2.5e-03	2.29e-04

Simulation 3: Near positivity violation.

The parameters of interest are the same as in simulation 1: $\psi_0 = 0.2585079$. Treatment probabilities are bounded in $(0.0794, 0.999994)$, with treatment probability > 0.99 for $W > 1.5$. Estimators using a truncated version of the correct treatment mechanism with an a-priori specified bound of (0.025, 0.975) were also considered.

The results are in table 4.3. When the treatment model values are extreme, the robustness results of lemma 4.1 no longer apply. We observe here that the MSE of TMLE and EE in the case of mis-specification of outcome model or mediator density cease to reduce at a rate proportional to sample size. However, when both of the outcome model and mediator density are correct, TMLE and EE with an incorrect treatment mechanism (either through truncation or incorrect modeling) yield MSE that are proportional to sample size. This last result is predicted by the robustness result (i) of lemma 4.1 since the mis-specified treatment models is bounded away from 1. We observe also that in this simulation scheme, TMLE 2 is less favorable than TMLE 1 across all cases. This may suggest that under data sparsity, the use of plug-in estimator for the mediated mean outcome difference is more beneficial than considerations such as the rate at which it is estimated. We observe an increase in MSE (driven by the increase in variance) as one moves away from the use of substitution principle (with TMLE 1 being the one which uses substitution estimators in all its steps, TMLE 2 which does not use substitution estimator in the initial estimate of the mediated

Table 4.2: Simulation 2: larger effect of treatment on mediator

	Bias		Var		MSE	
n	500	5000	500	5000	500	5000
all correct						
Gcomp	1.99e-03	3.46e-04	6.09e-03	5.7e-04	6.09e-03	5.74e-04
TMLE1	5.46e-03	5.82e-04	8.7e-03	7.87e-04	8.7e-03	7.88e-04
TMLE2	5.23e-03	5.03e-04	8.73e-03	7.89e-04	8.76e-03	7.89e-04
EE	6.05e-03	5.69e-04	8.97e-03	7.86e-04	9.01e-03	7.87e-04
g misspec.						
TMLE	5.12e-03	6.55e-04	8.08e-03	7.34e-04	8.10e-03	7.3e-04
EE	5.1e-03	6.74e-04	8.3e-03	7.69e-04	8.36e-03	7.7e-04
Q^Z misspec.						
Gcomp	1.20e-02	1.31e-02	5.91e-03	5.67e-04	6.05e-03	7.38e-04
TMLE1	3.0e-03	4.96e-04	6.23e-03	5.81e-04	6.2e-03	5.81e-04
TMLE2	2.85e-03	4.20e-04	6.25e-03	5.83e-04	6.25e-03	5.84e-04
EE	2.89e-03	4.71e-04	6.19e-03	5.79e-04	6.20e-03	5.79e-04
$\bar{Q}^Y$ misspec.						
Gcomp	8.81e-03	1.35e-02	5.74e-03	5.82e-04	5.81e-03	7.65e-04
TMLE1	7.60e-03	5.84e-04	8.90e-03	7.96e-04	8.96e-03	7.96e-04
TMLE2	7.8e-03	6.20e-04	8.90e-03	7.95e-04	8.96e-03	7.95e-04
EE	6.8e-03	5.09e-04	8.9e-03	7.92e-04	8.98e-03	7.9e-04

mean outcome difference but uses substitution in the final effect estimate, and EE which does not use substitution at all). This may suggest that in the case of positivity violation, when strict bounds exist on the parameter, the degree at which each step of the estimation procedure respects the bounds affects the stability of the resulting estimate. Nonetheless, rigorous analysis is needed to provide more valid insights.

In this simulation, we observe that TMLE and EE behave differently in some cases. We first consider the version with binary outcome. Since the parameter is an average of probability differences, for a given dataset one would like the effect estimates to be bounded between -1 and 1. However, when using a correctly specified treatment mechanism, the EE estimator exhibits estimates that are out of bound (of magnitude larger than 3 in some cases, and of magnitude 11 and 14 in one dataset). The bias, variance and mse of each estimator are detailed in table 4.3. When outcome model and mediator density are correct, the Gcomp is still consistent despite the positivity violation. Nonetheless, the effect of data-sparsity on g-comp is apparent when comparing this Gcomp estimator with its counterpart in the case of no positivity violation (table 4.1). On the other hand, under correct outcome model and mediator density, TMLE and EE have poor variance when implemented with an untruncated correct treatment mechanism. However, their performances are improved when implemented with a truncated or mis-specified treatment. We also observe that in the case of all models correct, TMLE and EE have a different bias-variance trade-off, with TMLE having smaller variance but larger bias, with respect to EE (which has a larger variance but smaller

bias). This difference in relative bias and variance is also present in the case of mis-specified mediator density but correct outcome and treatment: we observe that using the untruncated correct treatment, TMLE has larger bias and smaller variance than EE; but when the truncated treatment mechanism is used, the two robust estimators behave similarly and provide bias reduction over the Gcomp estimator. When the outcome model is mis-specified, TMLE and EE provide similar bias reduction over g-computation estimator; but TMLE has a smaller variance than EE when the untruncated treatment mechanism is used, while the opposite is true with the truncated treatment mechanism.

Table 4.3: Simulation 3: Binary outcome, positivity violations in $p(A|W)$

	Bias		Var		MSE	
n	500	5000	500	5000	500	5000
all correct						
Gcomp	2.35e-02	2.02e-03	1.09e-02	1.145e-03	1.15e-02	1.15e-03
TMLE1	5.68e-02	3.59e-02	3.45e-02	1.56e-02	3.77e-02	1.69e-02
TMLE2	4.66e-02	7.51e-02	5.92e-02	2.51e-02	6.1e-02	3.08e-02
EE	1.85e-02	3.1e-04	4.69e-02	4.82e-02	4.73e-02	4.82e-02
g truncated						
TMLE	2.59e-02	2.09e-03	1.56e-02	1.59e-03	1.62e-02	1.6e-03
EE	2.39e-02	1.82e-03	1.24e-02	1.25e-03	1.29e-02	1.25e-03
g misspec.						
TMLE	2.32e-02	2.79e-03	1.34e-02	1.38e-03	1.39e-02	1.39e-03
EE	2.64e-02	2.22e-03	1.84e-02	1.57e-03	1.91e-02	1.58e-03
Q^Z misspec.						
Gcomp	5.02e-02	5.85e-02	1.06e-02	1.36e-03	1.32e-02	4.77e-03
TMLE1	1.43e-01	1.13e-01	1.77e-02	6.66e-03	3.83e-02	1.9e-02
TMLE2	4.66e-02	7.7e-02	5.4e-02	2.11e-02	5.66e-02	2.7e-02
EE	5.42e-03	7.11e-03	1.77e-01	5.2e-02	1.77e-01	5.24e-02
Q^Z misspec., g truncated						
TMLE	3.36e-02	1.66e-02	1.53e-02	1.8e-03	1.64e-02	2.07e-03
EE	2.89e-02	3.71e-02	1.39e-02	1.61e-03	1.48e-02	2.98e-03
$\bar{Q}^Y$ misspec.						
Gcomp	8.2e-02	8.26e-02	4.27e-03	4.56e-04	1.1e-02	7.28e-03
TMLE1	4.86e-02	9.41e-03	3.56e-02	1.59e-02	3.79e-02	1.59e-02
TMLE2	1.09e-03	6.62e-02	6.19e-02	2.85e-02	6.19e-02	3.29e-02
EE	3.79e-02	1.16e-02	2.74e-01	1.15e-01	2.75e-01	1.15e-01
$\bar{Q}^Y$ misspec. , g truncated						
TMLE1	6.25e-02	5.5e-02	1.37e-02	1.3e-03	1.76e-02	4.40e-03
EE	7.36e-02	7.08e-02	6.20e-03	6.23e-04	1.16e-02	5.64e-03

4.6 Extension to Natural Indirect Effect

In this section, we extend the above discussions in an analogous fashion to address the natural indirect effect.

In the context of natural effects, the total effect of A on Y can be decomposed into natural indirect and direct effects (Robins and Greenland (1992), Pearl (2001), Robins (2003)):

$$\begin{aligned}&E(Y(1)-Y(0))\\&=[E(Y(1,Z(1))-E(Y(1,Z(0))]+[E(Y(1,Z(0))-E(Y(0,Z(0))],\end{aligned}$$

where $Y(a)\equiv f_Y(W,A=a,Z=Z(a),U_Y)$ on the NPSEM. This decomposition formalizes the concept that the total effect of the exposure on the outcome is a combination of its indirect effect through a mediator Z, and its direct effect not mediated by Z. The quantity $E(Y(1,Z(1))-E(Y(1,Z(0))$ is referred to as the *additive natural indirect effect*. Its identification is studied in the same body of literature (Robins and Greenland (1992), Pearl (2001), Robins (2003), Petersen et al. (2006), Hafeman and VanderWeele (2010), Imai et al. (2010), Robins and Richardson (2010) and Pearl (2011)). More specifically, under the same conditions as those in section 4.2, the natural indirect effect can be identified as

$$\begin{aligned}&E(Y(1,Z(1))-E(Y(1,Z(0))\overset{A1,A2,A3}{=}\Psi_{NIE}(P_0)\\&\equiv E_{P_0}\left(\sum_z \bar{Q}^Y(W,A=1,z)\left[Q^Z(z|W,A=1)-Q^Z(z|W,A=0)\right]\right).\end{aligned}\tag{4.9}$$

The results of Robins and Richardson (2010) thus have the same implications on the difficulty of identifying the natural indirect effect in real experiments, due to the conditional counterfactual independence assumption A3. In such cases, what kind of causal interpretation can the statistical parameter (4.9) still offer? If assumption A3 fails but randomization assumptions A1 and A2 hold, the statistical parameter in (4.9) equals

$$\Psi_{NIE}(P_0)\overset{A1,A2}{=}E_{P_0}\left\{\sum_z E(Y(1,z)|W)\left[p(Z(1)=z|W)-p(Z(0)=z|W)\right]\right\}.$$

The interpretation of the right hand side is not as intuitive as in the natural direct effect case. But since $p(Z(1)=z|W)-p(Z(0)=z|W)$ measures the effect of A on Z, at its face value this alternative effect parameter can be viewed as weighting the different outcomes $E(Y(1,z)|W)$ under z by these effect measures. However, we remind the reader again that this alternative causal parameter only serves to provide a causal interpretation for the statistical parameter (4.9) and one should be cautious about putting it into the context of the traditional total effect decomposition.

The parameter $\Psi_{NIE}(P)$ is also a function of Q alone. To extend the discussions above to the natural indirect effect parameter (4.9), we now consider the *mediated mean outcome* map $P\mapsto\bar{Q}^Z_{NIE}(P)$, where $\bar{Q}^Z_{NIE}(P)(W,A)\equiv E_P\left(\bar{Q}^Y(W,A=1,Z)|W,A\right)$. Let $Q\equiv\left(\bar{Q}^Y,\bar{Q}^Z_{NIE},Q^W,Q^Z\right)$. This way, the parameter can be regarded as $\Psi_{NIE}(P)=\Psi_{NIE}(Q)$.

The efficient influence curve for this parameter (derived inTchetgen Tchetgen and Shpitser (2011a)) is given by

$$\begin{aligned}
&D^*_{NIE}(Q,g,\Psi_{NIE}(Q)) \\
&= \frac{I(A=1)}{g(1|W)}\left\{Y-\bar{Q}^Z_{NIE}(P)(W,1)-\frac{Q^Z(Z|W,0)}{Q^Z(Z|W,1)}\left(Y-\bar{Q}^Y(W,1,Z)\right)\right\} \\
&-\frac{I(A=0)}{g(0|W)}\left(\bar{Q}^Y(W,1,Z)-\bar{Q}^Z_{NIE}(P)(W,0)\right) \\
&+\bar{Q}^Z_{NIE}(P)(W,1)-\bar{Q}^Z_{NIE}(P)(W,0)-\Psi_{NIE}(Q). \qquad (4.10)
\end{aligned}$$

The robustness conditions of Tchetgen Tchetgen and Shpitser (2011a) apply to both natural direct and indirect effects. By the same reasoning (and analogous proof) as that of lemma 1, we note again that conditions (i) and (iii) may be weakened to: (i) the conditional mean outcome $\bar{Q}^Y(W,A,Z)$ and the mediated outcome map $\bar{Q}^Z_{NIE}(P)(W,A)$ are both correct; (iii) the exposure mechanism and mediator density, or the exposure mechanism and the conditional distribution $p(A|W,Z)$, are correct. Therefore, in situations where Z is high dimensional, similar practical implications as those discussed in remarks following lemma 1 apply. However, note that a regression-based estimation procedure for $\bar{Q}^Z_{NIE}(P_0)$ now regresses $\bar{Q}^Y(W,1,Z)$ on W,A.

Since the parameter (4.9) is given by

$$\Psi_{NIE}(P)=E_P\Big(\bar{Q}^Z_{NIE}(W,1)-\bar{Q}^Z_{NIE}(W,0)\Big), \qquad (4.11)$$

the targeted MLE only needs to focus on estimation of the components $Q^W(P_0)$, $\bar{Q}^Y(P_0)$ and $\bar{Q}^Z_{NIE}(P_0)$ of the likelihood. We first rewrite the efficient influence curve in (4.10) as

$$\begin{aligned}
&D^*_{NIE}(Q,g,\Psi_{NIE}(Q)) \\
&= \frac{I(A=1)}{g(1|W)}\left(1-\frac{Q^Z(Z|W,0)}{Q^Z(Z|W,1)}\right)\left(Y-\bar{Q}^Y(W,A,Z)\right) \\
&+\frac{2A-1}{g(A|W)}\left\{\bar{Q}^Y(W,1,Z))-\bar{Q}^Z_{NIE}(P)(W,A)\right\} \\
&+\bar{Q}^Z_{NIE}(P)(W,1)-\bar{Q}^Z_{NIE}(P)(W,0)-\Psi_{NIE}(Q) \\
&\equiv D^*_{NIE,Y}+D^*_{NIE,Z}+D^*_{NIE,W}.
\end{aligned}$$

The reader may have readily noted the parallel between $D^*_{NIE,Z}+D^*_{NIE,W}$ and the efficient influence curve for the familiar additive marginal treatment effect; this is because the indirect effect can viewed as an additive marginal effect of the treatment on $\bar{Q}^Y(W,A=1,Z)$ through its effect on Z, as seen in (4.11). In fact, as we will see shortly, the second part of the implementation of TMLE is very similar to the well-known case of additive marginal effects.

Without loss of generality, we assume that Y is bounded in the unit interval. Under the log-likelihood loss function of (4.3), the least favorable submodel for $\bar{Q}^Y(W,A,Z)$ through a given initial estimator $\bar{Q}^Y_n$ is now given by

$$\bar{Q}^Y_n(\varepsilon)\equiv \text{expit}\left(\text{logit}(\bar{Q}^Y_n)+\varepsilon C_Y\right),$$

where $C_Y \equiv \frac{I(A=1)}{g_n(1|W)}\left(1-\frac{Q_n^Z(Z|W,0)}{Q_n^Z(Z|W,1)}\right)$. Note that the dependence of $\bar{Q}_n^Y(\varepsilon)$ on Q_n^Z and g_n are suppressed in the notation. The targeted MLE of $\bar{Q}^Y(P_0)$ is $\bar{Q}_n^{Y,*} \equiv \bar{Q}_n^Y(\varepsilon_n^Y)$ and is similarly defined as in section 4.3.

Next, consider an estimating $\bar{Q}_{NIE,n}^Z$ for $\bar{Q}_{NIE}^Z(P_0)$, with respect to the updated expectant $\bar{Q}_n^{Y,*}$; it can be a density-based or regressed-based estimator. We use the log-likelihood loss

$$\begin{aligned} &L_Z(\bar{Q}_{NIE}^Z)(O) \\ &= -\log\left\{\bar{Q}_{NIE}^Z(W,A)^{\bar{Q}_n^{Y,*}(W,1,Z)}\left(1-\bar{Q}_{NIE}^Z(W,A)\right)^{1-\bar{Q}_n^{Y,*}(W,1,Z)}\right\}. \end{aligned}$$

The least favorable submodel through the initial estimator $\bar{Q}_{NIE,n}^Z$ is given by

$$\bar{Q}_{NIE,n}^Z(\varepsilon) = \text{expit}\left(\text{logit}\left(\bar{Q}_{NIE,n}^Z\right) + \varepsilon C_Z\right),$$

where $C_Z = \frac{2A-1}{g_n(A|W)}$. The dependence of the submodel on g_n is also suppressed in the notation. In a similar fashion as section 4.3, we obtain the targeted MLE $\bar{Q}_{NIE,n}^{Z,*} \equiv \bar{Q}_{NIE,n}^Z(\varepsilon_n^Z)$. Finally, the targeted MLE of the parameter $\Psi_{NIE}(P_0)$ is given by

$$\psi_{NIE,n}^* \equiv \frac{1}{n}\sum_{i=1}^{n}\left(\bar{Q}_{NIE,n}^{Z,*}(W_i,1) - \bar{Q}_{NIE,n}^{Z,*}(W_i,0)\right).$$

We remind the reader again that the role of the ratio of Q^Z in C_Y may be replaced by ratios of $g(A|W)$ and $p(A|W,Z)$.

The resulting estimator satisfies the efficient influence curve equation, and therefore is asymptotically unbiased if (i) the conditional mean outcome $\bar{Q}^Y$ and the mediated outcome map $\bar{Q}_{NIE}^Z(P)$ are both correct; (ii) the conditional mean outcome and the exposure mechanism $g(A|W)$ are correct; (iii) the exposure mechanism and mediator density $Q^Z(Z|W,A)$, or the exposure mechanism and the conditional distribution $p(A|W,Z)$, are correct. An estimating equation estimator $\psi_{NIE,n}^{EE}$ is also discussed in Tchetgen Tchetgen and Shpitser (2011a). As mentioned in section 4.4, $\psi_{NIE,n}^*$ and $\psi_{NIE,n}^{EE}$ will inherit the same robustness properties from the efficient influence curve, since both satisfy the efficient influence curve equation.

4.7 Summary

In this chapter, we applied the targeted maximum likelihood framework of van der Laan and Rubin (2006) and van der Laan and Rose (2011) to construct a semiparametric efficient, multiply robust, plug-in estimator for the natural direct effect of a binary treatment. This estimator has the property that it satisfies the efficient influence curve equation (derived in Tchetgen Tchetgen and Shpitser (2011a)), and hence also inherits its robustness properties. We noted that the robustness conditions in Tchetgen Tchetgen and Shpitser (2011a) may be weakened (lemma 1), thereby placing less reliance on the estimation of the mediator density. More precisely, the proposed estimator is

asymptotically unbiased if either one of the following holds: i) the conditional mean outcome given exposure, mediator, and confounders, and the mediated mean outcome difference are consistently estimated; (ii) the exposure mechanism given confounders, and the conditional mean outcome are consistently estimated; or (iii) the exposure mechanism and the mediator density, or the exposure mechanism and the conditional distribution of the exposure given confounders and mediator, are consistently estimated. If all three conditions hold, then the effect estimate is asymptotically efficient. We also extended our results analogously to the case of natural indirect effect.

In applications, the components that are difficult to estimate are often times the conditional mean outcome and/or the mediator density. For a high-dimensional Z, few tools are available to estimate the conditional mediator density Q^Z. On the other hand, there is abundant literature addressing estimation of conditional means. This can be used to estimate the mediated mean outcome difference $\bar{Q}^Z(P) \equiv E_P(\bar{Q}^Y(W,1,Z) - \bar{Q}^Y(W,0,Z)|W,A=0)$, and the conditional distributions of a categorical A. Lemma 1 implies that estimation of the mediator density may be replaced by estimations of $g(A|W)$, $p(A|W,Z)$, and the conditional expectation $\bar{Q}^Z(P)$.

We have also described general conditions for the estimator to be asymptotically linear. More specifically, 1) estimators of each component must converge to their respective limits at a reasonable speed, and 2) if there is a component that is not consistently estimated, the consistent estimators of the remaining components must meet stricter asymptotic linearity conditions. These conditions provide a guideline for situations where influence curve based variance estimates are realistic.

Estimators that use the efficient influence curve are robust, but are generally sensitive to practical positivity violations. We refer to Petersen, Porter, S.Gruber, Wang, and van der Laan (2010) for methods of diagnosing and responding to violations of the positivity assumption. The substitution principle and the logistic working submodels in the targeted estimation procedure aim to provide more stable estimates in such situations. However, identification of the parameter depends ultimately on the information available in a given finite sample. A way to improve finite sample robustness is the Collaborative TMLE (C-TMLE) of van der Laan and Gruber (2010), where, instead of estimating the true treatment mechanism, for a given initial estimator of the Q component one estimates a conditional distribution of the treatment, conditioned only on confounders that explain the residual bias of the estimator of Q. We aim to investigate applications of C-TMLE to the effect mediation problem.

4.8 Chapter Appendix

Proof of lemma 4.1

Let $\bar{Q}^Z$ be a map $Q \mapsto \bar{Q}^Z(Q)$, where $\bar{Q}^Z(Q)$ is a function from $\mathscr{W}$ to $\mathbb{R}$. Note that $\bar{Q}^Z(Q)$ may or may not make use of the density Q^Z, but it surely uses $\bar{Q}^Y$. Then

$$
\begin{aligned}
&P_0 D^*(\bar{Q}^Y, \bar{Q}^Z, g, \psi_0) \\
&= P_0\left\{ \frac{g_0(1|W)}{g(1|W)} \sum_z Q^Z(P_0)(z|W,1) \frac{Q^Z(z|W,0)}{Q^Z(z|W,1)} \left(\bar{Q}_0^Y(W,1,z) - \bar{Q}^Y(W,1,z)\right) \right\} && (4.12) \\
&- P_0\left\{ \frac{g_0(0|W)}{g(0|W)} \sum_z Q^Z(P_0)(z|W,0) \left(\bar{Q}_0^Y(W,0,z) - \bar{Q}^Y(W,0,z)\right) \right\} && (4.13) \\
&+ P_0\left\{ \frac{g_0(0|W)}{g(0|W)} \sum_z Q^Z(P_0)(z|W,0) \Big(\bar{Q}^Y(W,1,z) - \bar{Q}^Y(W,0,z)\Big) \right\} && (4.14) \\
&- P_0\left\{ \frac{g_0(0|W)}{g(0|W)} \bar{Q}^Z(W) \right\} && (4.15) \\
&+ P_0\Big\{ \bar{Q}^Z(W) \Big\} - \psi_0 && (4.16)
\end{aligned}
$$

(i) is trivial. Suppose now that (ii) holds. Then (4.12) and (4.13) are each exactly 0. The expression in (4.15) equals $P_0\bar{Q}^Z(W)$, and the expression in (4.14) equals ψ_0. Therefore the mean is zero.

Suppose that (iii) holds. Then, rearranging (4.12) and (4.13) we rewrite the above expectation as

$$
\begin{aligned}
P_0 D^*(Q, g, \psi_0) &= P_0\left\{ \sum_z Q^Z(P_0)(z|W,0) \left(\bar{Q}_0^Y(W,1,z) - \bar{Q}_0^Y(W,0,z)\right) \right\} \\
&- P_0\left\{ \sum_z Q^Z(P_0)(z|W,0) \left(\bar{Q}^Y(W,1,z) - \bar{Q}^Y(W,0,z)\right) \right\} \\
&+ P_0\left\{ \sum_z Q^Z(P_0)(z|W,0) \Big(\bar{Q}^Y(W,1,z) - \bar{Q}^Y(W,0,z)\Big) \right\} \\
&- P_0\bar{Q}^Z(W) + P_0\bar{Q}^Z(W) - \psi_0 \\
&= 0
\end{aligned}
$$

Chapter 5

Targeted Maximum Likelihood Estimation for Longitudinal Mediation Analysis

5.1 Introduction

An exposure often acts on an outcome of interest directly, and/or indirectly through the mediation of some intermediate variables. Identifying and quantifying these two types of effects contribute to further understanding of the underlying causal mechanism. Much of the existing literature on causal mediation is focused on applications in non-longitudinal settings. Causal mediation in a longitudinal setting, by contrast, has received relatively little attention. In this work, we study the effect of a time-varying exposure mediated by a time-varying intermediate variable. More specifically, consider a study where baseline covariates, time-varying treatment, time-varying mediator, time-varying covariates, and an outcome process are observed on subjects that are followed over time. The treatment of interest is influenced by past covariates and mediator, and affects future covariates and mediator. Right censoring, if present, occurs in response to past covariates and treatment. We also allow the outcome to be a time-to-event (say survival) process, in which case, at each time we record whether death has occurred.

The subtlety of longitudinal mediation analysis is best illustrated in the survival setting. Suppose we are interested in the effect of the treatment on the time till death (*failure time*), and the mediator lies on the causal pathway between these two — the risk of one dying at a given time depends on the mediator history, which is also affected by the treatment. Therefore, the treatment can act on the failure time directly, and/or indirectly through its effect on the mediator. The goal of mediation analysis is to quantify these two types of treatment effects on the failure time. The challenge in this setting lies in that the outcome of interest is a process (the *event process*) that happens jointly with the mediator process. In other words, the mediators are also affected by the death process in the sense that they take on default values after death occurs.

One way to assess the direct effect of the treatment on failure time is to compare the distributions of the failure times under different treatments regimens while the mediators are fixed to some common pre-specified values. This is known as the *controlled direct effect* (e.g. Pearl (2001)). Its analysis is very similar to that of a time-dependent deterministic treatment in a non-mediation setting; we refer the reader to existing literature on this topic (e.g. Robins (1997b), Hernan et al. (2000), Stitelman, Gruttola, Wester, and van der Laan (2011)). Controlled direct effects are of interest if the treatment effect under one particular mediator value constitutes a meaningful scientific question. If that is not the case, one may ask a different direct effect question: what would be the effect of treatment on failure time if the treatment had no effect on the mediator (i.e. the mediator takes its value as if treatment were absent)? One way to rigorously formulate this question is using the so-called *natural direct effect* parameter (Robins and Greenland (1992), Pearl (2001)). The natural direct effect has a complementary *natural indirect effect*; together they provide a decomposition of the overall effect of the treatment on the outcome. In this paper, we focus on the natural direct and indirect effects.

In the case of a time-independent mediator that is measured before the onset of the event process (and censoring process), the definition of natural effects and their identifiability (e.g. Lange and Hansen (2011), Tchetgen Tchetgen (2011)) can be extended from the formulations in non-longitudinal setting (e.g. Robins and Greenland (1992), Pearl (2001), Robins (2003), Petersen et al. (2006), Imai et al. (2010)). For inference of these parameters, the use of additive hazard models for the outcome and linear models for the mediator are proposed in Lange and Hansen (2011); the use of accelerated failure time and proportional hazard models are studied in Tein and MacKinnon (2003) and VanderWeele (2011); robust estimators for the natural direct and indirect effects under proportional hazards models and additive hazards models, as well as sensitivity analysis techniques for assessing the impact of the violation of the mediator's ignorability assumption, are developed in Tchetgen Tchetgen (2011). A more complex variation of this setting is when there is a confounder (the *recanting witness* covariate, Avin, Shpitser, and Pearl (2005)) of the mediator–outcome relation that is affected by the treatment of interest. Robins and Richardson (2010) show that with additional conditions regarding independence (or deterministic dependence) of the counterfactual recanting witness under various treatment levels, the resulting mediation parameters can be identified. Tchetgen Tchetgen and VanderWeele (2012) show that these additional conditions can be avoided in some special cases of binary recanting witness, or under additional parametric assumptions on the mediator–recanting witness relation.

In the case of a time-dependent mediator, the probability of the mediator process having a given non-degenerate length would depend on the failure time. This interdependency between the event process and the mediator process poses a challenge when extending the results from the time-independent mediator setting. Firstly, the event history affects both the current mediator (taking a non-degenerate value) and the current event indicator, but it is part of the outcome of interest and thus is not a recanting witness covariate. More specifically, the treatment affects a current mediator both directly and indirectly through its effect on the event history. In asking what is the effect of a given treatment level on the event process not mediated by the mediator process, one must specify how the paths from treatment to mediator should be blocked. If one blocks all paths from

treatment to mediators (both the direct paths of treatment to mediator and the paths through event history), then the parameters defined would be a direct generalization of the definition of mediation formula and natural effects in time-independent mediator settings (by regarding event history as a recanting witness). However, this generalization would yield parameters that are not interpretable for the purpose of effect mediation in this survival setting, since the relation of treatment and event process (outcome of interest) is also altered. In this light, the definition of mediation formula and natural effects should be based on blocking only those paths from the treatment to mediator that are not through survival history (these would be an extension of the path-specific effects discussed in Pearl (2001), Avin et al. (2005), Robins and Richardson (2010)). The direct effect question these parameters would address is: what is the effect of treatment on the survival time, if the treatment had no effect on the mediator process other than through survival history?

Having specified the effects of interest, the second challenge arises in formulating these as parameters in a causal framework. Under the traditional definition of mediation parameters in a non-longitudinal setting (e.g. Robins and Greenland (1992), Pearl (2001), Avin et al. (2005), Robins and Richardson (2010)), the mediator is regarded as intermediate counterfactual outcome. In extending this definition to our effects of interest in the current setting, the time-varying mediators become intermediate counterfactual outcomes under a different treatment level than that of their parent counterfactual survival history. The identifiability conditions for this setting is studied in Shpitser (2013). However, these conditions may prove difficult to satisfy for the purpose of a survival study. We further discuss extending the existing one-time point concepts to the longitudinal setting in appendix A.2. As an alternative, we propose to adopt a stochastic interventions (SI) perspective to causal mediation, introduced by Didelez et al. (2006). Under this formulation, the mediators are regarded as intervention variables, onto which a given counterfactual distribution is enforced. The natural effects can be defined analogously to the ideas in Pearl (2001) and Avin et al. (2005). In particular, they also allow for a total effect decomposition and an interpretation of the natural direct effect as a weighted average of controlled direct effects. Importantly, however, one should note that even though these SI-based parameters and their non-SI-based counterparts in Shpitser (2013) all identify to the same statistical parameters, they are formally different causal parameters defined under different formulations (but aim to answer the same type of mediation questions).

The natural direct and indirect effects in this setting can all be defined in terms of the corresponding version of the mediation functional; we develop a general semiparametric inference framework for this parameter. More specifically, we will derive the efficient influence curves under a locally saturated semiparametric model, and establish their robustness properties. The variances of these functions provide local efficiency bounds, and their robustness properties give information on the types of model mis-specifications that would still allow for unbiased estimation of the parameters. Applying the targeted maximum likelihood methodology (van der Laan and Rose (2011)), we will use these efficient influence curves to construct multiply robust and efficient estimators.

This paper proceeds as follows. Firstly, we establish the definition and determine the identifiability conditions for the parameters of interest. Thereafter, we derive the efficient influence curves of the parameters under a locally saturated semiparametric model, and present the non-targeted

substitution G-computation estimator, the inverse probability weighted (IPW) estimator, and the TMLE estimator for these parameters.

5.2 Data and Parameters of Interest

Consider the data structure $O = (L_0, (A_t, Z_t, L_t) : t = 1, \dots \tau) \sim P_0$, where $L_0 =$ encodes baseline baseline covariates, A_t encodes the time-varying exposure, and possibly a censoring indicator, Z_t denotes the time-varying mediators, and L_t are the time-varying covariates and includes the outcome variable $Y_t \subset L_t$. Without loss of generality, we may assume that it is bounded between 0 and 1. The data consists of n i.i.d. copies of O.

From here on, for any $1 \leq t \leq \tau$ and a time-dependent variable V, we will use the boldface $\mathbf{V}_t$ to denote the vector $(V_1, \dots, V_t)$, use $\mathbf{V}_{s,t}$ to denote the vector $(V_s, \dots, V_t)$. When referring to the entire vector $\mathbf{V}_\tau$, we will also use the shorthand $\mathbf{V}$. Degenerate indices such as $\mathbf{V}_{-1}$ signify the empty set.

The following Non-Parametric Structural Equations Model (NPSEM, Pearl (2009)) encodes the time-ordering assumption on the variables:

$$
\begin{aligned}
L_0 &= f(U_{L_0}), \\
A_t &= f_{A_t}(L_0, \mathbf{A}_{t-1}, \mathbf{Z}_{t-1}, \mathbf{L}_{t-1}, U_{A_t}), \\
Z_t &= f_{Z_t}(L_0, \mathbf{A}_t, \mathbf{Z}_{t-1}, \mathbf{L}_{t-1}, U_{Z_t}), \\
L_t &= f_{L_t}(L_0, \mathbf{A}_t, \mathbf{Z}_t, \mathbf{L}_{t-1}, U_{L_t}).
\end{aligned}
$$

This model assumes that each observed variable X is a deterministic function f_X of other observed variables, referred to as parents of X and denoted as $Pa(X)$, and unobserved exogenous factors U_X. This model defines a distribution $P_{O,U}$ on the unit.

The observed data structure is generated from the above NPSEM without any intervention, and the likelihood of $O \sim P_0$ can be factored according to that time-ordering:

$$
\begin{aligned}
&p_0(O) = p_0(L_0) \\
&\times \prod_{t=1}^{\tau} \Big(p_0(A_t \mid L_0, \mathbf{A}_{t-1}, \mathbf{Z}_{t-1}, \mathbf{L}_{t-1}) p_0(Z_t \mid L_0, \mathbf{A}_t, \mathbf{Z}_{t-1}, \mathbf{L}_{t-1}) p_0(L_t \mid L_0, \mathbf{A}_t, \mathbf{Z}_t, \mathbf{L}_{t-1}) \Big).
\end{aligned}
$$

In the case of a survival outcome or if censoring exists, Z_t and L_t are assigned a default value with probability 1 after censoring or death

Counterfactual outcome

To define the pertinent counterfactual failure time, we propose to use a stochastic interventions (SI) perspective introduced by Didelez et al. (2006). Stochastic interventions (a.k.a stochastic policies, random interventions, randomized dynamic strategies; e.g. Dawid and Didelez (2010), Pearl

(2009), Tian (2008), Robins and Richardson (2010), Diaz and van der Laan (2011)) are generalizations of the traditional static interventions or dynamic regimes where, instead of assigning a deterministic value, one assigns a probability distribution to an intervention variable. Using the non-longitudinal setting as background, Didelez et al. (2006) illustrate how the notions of various direct and indirect effects can be formulated as sequential treatment problems by regarding the mediator as an intervention variable (receiving either a deterministic or stochastic intervention). Though their approach was based on a non-counterfactual causal framework (e.g. Dawid and Didelez (2010)), the essence of their idea remains the same, and is easily generalizable to survival settings.

Let $\mathbf{a}$ and $\mathbf{a}'$ be two possible treatment levels. Let $\mathbf{Z}(\mathbf{a}')$ denote the resulting mediator variable in an ideal experiment that sets $\mathbf{A} = \mathbf{a}'$, and let $p_{\mathbf{Z}(\mathbf{a}')}$ denote the conditional probability of this variable. Consider the following ideal experiment. At baseline, we measure the covariates L_0. At each time $t \geq 1$, censoring (if applicable) is prevented, and we intervene to set $A_t = a_t$; given history $(\ell_0, \mathbf{z}_{t-1}, \ell_{t-1})$, intervene to set Z_t be distributed according to $p_{\mathbf{Z}(\mathbf{a}')}(Z_t \mid \ell_0, \mathbf{z}_{t-1}, \ell_{t-1})$, this way, we draw $Z_t(\mathbf{a}') = z_t$; given $(\ell_0, \mathbf{a}_t, \mathbf{z}_t, \ell_{t-1})$, measure L_t. At the end of the experiment, we denote the final outcome as $Y_\tau(\mathbf{a}, \mathbf{Z}(\mathbf{a}'))$. The difference between this ideal experiment and other ways to extend the mediation formulation from one-time-point setting are discussed in appendix A.2.

Even if one can carry out an intervention on the mediator (separately from the intervention on the treatment), the SI formulation formally requires the external specification of the function $p_{\mathbf{Z}(\mathbf{a}')}(z_t \mid \ell_0, \mathbf{z}_{t-1}, \ell_{t-1})$, which is the conditional distribution of the counterfactual variable $\mathbf{Z}(\mathbf{a}')$. If this conditional distribution is not known, it needs to be ascertained using a separate controlled experiment. Therefore, aside from causal assumptions needed to identify the distribution of the outcome of interest $Y_\tau(\mathbf{a}, \mathbf{Z}(\mathbf{a}'))$ in the main experiment, additional assumptions are needed to identify $p_{\mathbf{Z}(\mathbf{a}')}$ as a function of the data generating distribution. We discuss these in the next section.

Mediation Formula, Natural Direct and Indirect Effects

For concreteness suppose one is interested in the effect of a binary treatment. We refer to the difference $E(Y_\tau(1) - Y_\tau(1, \mathbf{Z}(0)))$ as the *natural indirect effect* (NIE) and $E(Y_\tau(1, \mathbf{Z}(0)) - Y_\tau(0))$ as the *natural direct effect* (NDE), where $Y_\tau(\mathbf{a})$ is the traditional counterfactual outcome under treatment $\mathbf{A} = \mathbf{a}$. Together, these provide a decomposition of the overall effect $E(Y_\tau(1) - Y_\tau(0))$. The identification and estimation of these two effects can be approached through the so-called *mediation formula* (Pearl (2011)):

$$\Psi^{\mathbf{a},\mathbf{a}'}(P_{O,U}) \equiv E(Y_\tau(\mathbf{a}, \mathbf{Z}(\mathbf{a}')). \tag{5.1}$$

It is important to note that while the definition of these parameters are analogous to those in Robins and Greenland (1992), Pearl (2001), Pearl (2011) and Avin et al. (2005), they are ultimately not the same definitions since the mediator variables are conceptualized differently. In this respect, the parameters defined here aim to provide an alternative formulation to questions that arise in mediation analysis in the current survival setting.

The identifiability of these parameters is a consequence of established results regarding stochastic interventions (e.g. Dawid and Didelez (2010), Pearl (2009), Robins and Richardson (2010)). Under the sequential randomization assumption (Robins (1986)) and established results on identification of stochastic interventions (e.g. Dawid and Didelez (2010), Didelez et al. (2006), Pearl (2009), Robins and Richardson (2010)), if the following hold:

- $(\mathbf{L},\mathbf{Z})_{t,\tau}(\mathbf{a}') \perp A_t \mid \mathbf{L}_{t-1},\mathbf{A}_{t-1}=\mathbf{a}'_{t-1},\mathbf{Z}_{t-1}=\mathbf{z}_{t-1}$.
- $\mathbf{L}_{t,\tau}(\mathbf{a},\mathbf{z}) \perp (A_t,Z_t) \mid \mathbf{L}_{t-1},\mathbf{A}_{t-1}=\mathbf{a}_{t-1},\mathbf{Z}_{t-1}=\mathbf{z}_{t-1}$,

then $E(Y_\tau(\mathbf{a},\mathbf{Z}(\mathbf{a}')))$ identifies to

$$\Psi^{\mathbf{a},\mathbf{a}'}(P_0)=\sum_{\mathbf{l},\mathbf{z}} y_\tau p_0(\ell_0)\prod_{t=1}^{\tau} p_0(z_t \mid \mathbf{a}'_t,\mathbf{z}_{t-1},\mathbf{l}_{t-1})p_0(\ell_t \mid \mathbf{a}_t,\mathbf{z}_t,\mathbf{l}_{t-1}). \tag{5.2}$$

Consequently, the natural direct and indirect effects are respectively identified to

$$NDE = E\left(Y_\tau(1,\mathbf{Z}(0))-Y_\tau(0)\right)=\Psi^{1,0}(P_0)-\Psi^{0,0}(P_0) \tag{5.3}$$

$$NIE = E\left(Y_\tau(1)-Y_\tau(1,\mathbf{Z}(0))\right)=\Psi^{1,1}(P_0)-\Psi^{1,0}(P_0). \tag{5.4}$$

Therefore, we will approach estimation of the natural effects through estimation of this mediation functional.

The rest of this chapter is devoted to the statistical inference of $\Psi^{\mathbf{a},\mathbf{a}'}(P_0)$. But before we proceed, we shall agree on the following notation.

Notations

Let $\mathscr{M}$ denote a locally saturated semiparametric model containing the true data generating distribution P_0.

We use P_n to denote the empirical distribution of n i.i.d. copies of $O\sim P_0$. Given a function $O\mapsto f(O)$, $P_n f$ denotes the empirical mean $P_n f\equiv \frac{1}{n}\sum_{i=1}^n f(O_i)$. More general for any $P\in\mathscr{M}$, $Pf\equiv E_P f(O)$.

For a generic $P\in\mathscr{M}$, denote the marginal distribution of L_0 as $Q^{L_0}(P)\equiv P(L_0)$, denote conditional distribution of L_t as $Q^L(P)\left(L_t \mid \mathbf{A}_t,\mathbf{Z}_t,\mathbf{L}_{t-1}\right)\equiv P\left(L_t \mid \mathbf{A}_t,\mathbf{Z}_t,\mathbf{L}_{t-1}\right)$, and denote the mediator density as $Q^Z(Z_t \mid \mathbf{A}_t,\mathbf{Z}_{t-1},\mathbf{L}_{t-1})\equiv P(Z_t \mid \mathbf{A}_t,\mathbf{Z}_{t-1},\mathbf{L}_{t-1})$. Following an important observation by Bang and Robins (2005), we define recursively the following functionals:

$$\begin{aligned}
&\bar{Q}^Z_{\tau+1}\equiv Y_\tau\\
&\bar{Q}^{L,\mathbf{a},\mathbf{a}'}_t(P)\left(\mathbf{L}_{t-1},\mathbf{Z}_t\right)\equiv E_P\left[\bar{Q}^{Z,\mathbf{a},\mathbf{a}'}_{t+1}(P)\left(\mathbf{L}_t,\mathbf{Z}_t\right)\middle|\mathbf{A}_t=\mathbf{a}_t,\mathbf{Z}_t,\mathbf{L}_{t-1}\right]\\
&\bar{Q}^{Z,\mathbf{a},\mathbf{a}'}_t(P)\left(\mathbf{L}_{t-1},\mathbf{Z}_{t-1}\right)\equiv E_P\left[\bar{Q}^{L,\mathbf{a},\mathbf{a}'}_t(P)\left(\mathbf{L}_{t-1},\mathbf{Z}_t\right)\middle|\mathbf{A}_t=\mathbf{a}'_t,\mathbf{Z}_{t-1},\mathbf{L}_{t-1}\right].
\end{aligned} \tag{5.5}$$

It is easy to see that

$$\Psi^{\mathbf{a},\mathbf{a}'}(P_0) = E_{P_0}\left[\bar{Q}^{Z,\mathbf{a},\mathbf{a}'}_{t=1}(L_0)\right]. \tag{5.6}$$

Let $\bar{Q}^{L,\mathbf{a},\mathbf{a}'}(P) \equiv \left(\bar{Q}^{L,\mathbf{a},\mathbf{a}'}_t(P) : 1 \leq t \leq \tau\right)$ and $\bar{Q}^{Z,\mathbf{a},\mathbf{a}'}(P) \equiv \left(\bar{Q}^{Z,\mathbf{a},\mathbf{a}'}_t(P) : 1 \leq t \leq \tau\right)$.

For a generic $P \in \mathscr{M}$, we will sometimes simply write Q^L, Q^Z, $\bar{Q}^{L,\mathbf{a},\mathbf{a}'}$ and $\bar{Q}^{Z,\mathbf{a},\mathbf{a}'}$; for the true P_0, we will sometimes write Q^L_0, Q^Z_0, $\bar{Q}^{L,\mathbf{a},\mathbf{a}'}_0$ and $\bar{Q}^{Z,\mathbf{a},\mathbf{a}'}_0$ in place of $Q^L(P_0)$, $Q^Z(P_0)$, $\bar{Q}^{L,\mathbf{a},\mathbf{a}'}(P_0)$ and $\bar{Q}^{Z,\mathbf{a},\mathbf{a}'}(P_0)$, respectively. Let $Q \equiv \left(Q^{L_0}, Q^L, Q^Z, \bar{Q}^{Z,\mathbf{a},\mathbf{a}'}, \bar{Q}^{L,\mathbf{a},\mathbf{a}'}\right)$. We will also denote $g(A_t \mid \mathbf{A}_{t-1}, \mathbf{Z}_{t-1}, \mathbf{L}_{t-1}) \equiv p(A_t \mid \mathbf{A}_{t-1}, \mathbf{Z}_{t-1}, \mathbf{L}_{t-1})$ and $\mathbf{g}(\mathbf{A}_t \mid \mathbf{Z}_{t-1}, \mathbf{L}_{t-1}) \equiv \prod_{k=1}^t g(A_k \mid \mathbf{A}_{k-1}, \mathbf{Z}_{k-1}, \mathbf{L}_{k-1})$. For our purposes, the couple (Q, g) readily specifies a distribution $P \in \mathscr{M}$, so sometimes we may abuse notation and write $P = (Q, g)$. At the data generating distribution P_0, we adopt the subscripts Q_0 and g_0.

5.3 Efficient Influence Curve

In this section, we develop a general semiparametric inference framework for these parameters. In particular, we derive the Efficient Influence Curves (EIC) of (5.2), (5.3) and (5.4) under a (locally saturated) semiparametric model, and establish their robustness properties. For a given pathwise-differentiable parameter Ψ, under certain regularity conditions, the variance of the EIC of Ψ is a generalized Cramer-Rao lower bound for the variances of the influence curves of asymptotically linear estimators of Ψ. Therefore, the variance of the EIC provides an efficiency bound for the regular and asymptotically linear (RAL) estimators of Ψ. Moreover, under a locally saturated model, the influence curve of any RAL estimator is in fact the EIC. We refer the reader to Bickel et al. (1997) for general theory of efficient semiparametric inference.

The mediation formula in (5.2) can be considered as the value at P_0 of the map $P \mapsto \Psi^{\mathbf{a},\mathbf{a}'}(P) \equiv E_{P_0}\left[\bar{Q}^{Z,\mathbf{a},\mathbf{a}'}_{t=1}(L_0)\right]$ on $\mathscr{M}$. In particular, this map depends on P through Q, i.e. $\Psi^{\mathbf{a},\mathbf{a}'}(P) = \Psi^{\mathbf{a},\mathbf{a}'}(Q)$. Similarly, the natural direct effect in (5.3) and the natural indirect effect in (5.4) are, respectively, the values at P_0 of the maps $P \mapsto \Psi^{NDE}(P) = \Psi^{1,0}(P) - \Psi^{0,0}(P)$ and $P \mapsto \Psi^{NIE}(P) = \Psi^{1,1}(P) - \Psi^{1,0}(P)$.

For $1 \leq t \leq \tau$, $s \leq t$, define the functionals at each $P \in \mathscr{M}$: $\gamma^{1,s,t}(P)(\mathbf{A}_s, \mathbf{Z}_t, \mathbf{L}_{t-1}) \equiv P(A_s \mid \mathbf{A}_{s-1}, \mathbf{Z}_t, \mathbf{L}_{t-1})$ and $\gamma^{2,s,t}(P)(\mathbf{A}_s, \mathbf{Z}_t, \mathbf{L}_t) \equiv P(A_s \mid \mathbf{A}_{s-1}, \mathbf{Z}_t, \mathbf{L}_t)$. We shall write $\gamma^{s,t} = \left(\gamma^{1,s,t}, \gamma^{2,s,t}\right)$ and $\gamma = (\gamma^{s,t} : 1 \leq t \leq \tau, s \leq t)$. Note that these conditional probabilities of A_s differ from the exposure (and possibly censoring) probabilities encoded by g in that the γs are not conditioning on parents of A_s. In particular, they are not orthogonal to Q or g. However, as we shall see in the following lemma, they offer an alternative to obtain robust estimators that are more suitable to real life settings where L_t and/or Z_t may be high dimensional.

Theorem 5.1 (Efficient Influence Curve).

Let $\Psi^{\mathbf{a},\mathbf{a}'} : \mathscr{M} \to \mathbb{R}$ be defined as above. Suppose at $P \in \mathscr{M}$ the conditional probabilities of Q^L, Q^Z and g are all bounded away from 0 and 1. The Efficient influence curve of $\Psi^{\mathbf{a},\mathbf{a}'}$ at P is given by

$D^{*,\mathbf{a},\mathbf{a}'}(P) \equiv D^{*,\mathbf{a},\mathbf{a}'}(Q,g,\Psi^{\mathbf{a},\mathbf{a}'}(Q,g))$, *with*

$$D^{*,\mathbf{a},\mathbf{a}'}(Q,g,\psi^{\mathbf{a},\mathbf{a}'}) \equiv D^{L_0,\mathbf{a},\mathbf{a}'}(Q,g,\psi^{\mathbf{a},\mathbf{a}'}) + \sum_{t\equiv 1}^{\tau} \left(D_t^{L,\mathbf{a},\mathbf{a}'}(Q,g) + D_t^{Z,\mathbf{a},\mathbf{a}'}(Q,g) \right), \tag{5.7}$$

where

$$\begin{aligned}
D_t^{L,\mathbf{a},\mathbf{a}'}(Q,g) &\equiv C_t^{L,\mathbf{a},\mathbf{a}'} \left\{ \bar{Q}_{t+1}^{Z,\mathbf{a},\mathbf{a}'}(\mathbf{L}_t,\mathbf{Z}_t) - \bar{Q}_t^{L,\mathbf{a},\mathbf{a}'}(\mathbf{L}_{t-1},\mathbf{Z}_t) \right\} \\
D_t^{Z,\mathbf{a},\mathbf{a}'}(Q,g) &\equiv C_t^{Z,\mathbf{a},\mathbf{a}'} \left\{ \bar{Q}_t^{L,\mathbf{a},\mathbf{a}'}(\mathbf{L}_{t-1},\mathbf{Z}_t) - \bar{Q}_t^{Z,\mathbf{a},\mathbf{a}'}(\mathbf{L}_{t-1},\mathbf{Z}_{t-1}) \right\} \\
D^{L_0,\mathbf{a},\mathbf{a}'}(Q,g,\psi^{\mathbf{a},\mathbf{a}'}) &\equiv \bar{Q}_1^{Z,\mathbf{a},\mathbf{a}'}(\mathbf{L}_0) - \psi^{\mathbf{a},\mathbf{a}'},
\end{aligned}$$

with

$$\begin{aligned}
C_t^{L,\mathbf{a},\mathbf{a}'} &\equiv \frac{I(\mathbf{A}_t \equiv \mathbf{a}_t)}{\prod_{j\equiv 1}^{t} g(a_j \mid \mathbf{a}_{j-1},\mathbf{Z}_{j-1}),\mathbf{L}_{j-1}} \prod_{j\equiv 1}^{t} \frac{Q^Z(Z_j \mid \mathbf{a}'_j,\mathbf{Z}_{j-1},\mathbf{L}_{j-1})}{Q^Z(Z_j \mid \mathbf{a}_j,\mathbf{Z}_{j-1},\mathbf{L}_{j-1})} \\
&= \frac{I(\mathbf{A}_t \equiv \mathbf{a}_t)}{\prod_{j\equiv 1}^{t} g(a_j \mid \mathbf{a}_{j-1},\mathbf{Z}_{j-1}),\mathbf{L}_{j-1}} \\
&\times \prod_{j\equiv 1}^{t}\prod_{s=1}^{j} \frac{\gamma^{1,s,j}(\mathbf{a}'_s,\mathbf{Z}_j,\mathbf{L}_{j-1})}{\gamma^{1,s,j}(\mathbf{a}_s,\mathbf{Z}_j,\mathbf{L}_{j-1})} \times \frac{g(a_j \mid \mathbf{a}_{j-1},\mathbf{Z}_{j-1},\mathbf{L}_{j-1})}{g(a'_j \mid \mathbf{a}'_{j-1},\mathbf{Z}_{j-1},\mathbf{L}_{j-1})} \prod_{s=1}^{j-1} \frac{\gamma^{2,s,j-1}(\mathbf{a}_s,\mathbf{Z}_{j-1},\mathbf{L}_{j-1})}{\gamma^{2,s,j-1}(\mathbf{a}'_s,\mathbf{Z}_{j-1},\mathbf{L}_{j-1})},
\end{aligned} \tag{5.8}$$

and

$$\begin{aligned}
C_t^{Z,\mathbf{a},\mathbf{a}'} &\equiv \frac{I(\mathbf{A}_t \equiv \mathbf{a}'_t)}{\prod_{j\equiv 1}^{t} g(a'_j \mid \mathbf{a}'_{j-1},\mathbf{Z}_{j-1},\mathbf{L}_{j-1})} \prod_{j\equiv 1}^{t-1} \frac{Q^L(L_j \mid \mathbf{a}_j,\mathbf{Z}_j,\mathbf{L}_{j-1})}{Q^L(L_j \mid \mathbf{a}'_j,\mathbf{Z}_j,\mathbf{L}_{j-1})} \\
&= \frac{I(\mathbf{A}_t \equiv \mathbf{a}'_t)}{\prod_{j\equiv 1}^{t} g(a'_j \mid \mathbf{a}'_{j-1},\mathbf{Z}_{j-1},\mathbf{L}_{j-1})} \prod_{j\equiv 1}^{t-1}\prod_{s=1}^{j} \frac{\gamma^{2,s,j}(\mathbf{a}_s,\mathbf{Z}_j,\mathbf{L}_j)}{\gamma^{2,s,j}(\mathbf{a}'_s,\mathbf{Z}_j,\mathbf{L}_j)} \times \prod_{s=1}^{j} \frac{\gamma^{1,s,j}(\mathbf{a}'_s,\mathbf{Z}_j,\mathbf{L}_{j-1})}{\gamma^{1,s,j}(\mathbf{a}_s,\mathbf{Z}_j,\mathbf{L}_{j-1})}.
\end{aligned} \tag{5.9}$$

Moreover, $D^{*,\mathbf{a},\mathbf{a}'}(P)$ is a multiply robust estimating function of $\Psi^{\mathbf{a},\mathbf{a}'}(P)$ in the sense that if one of the following holds:

R1. Q is correctly specified;

R2. g, $\bar{Q}^{L,\mathbf{a},\mathbf{a}'}$, and either Q^L or γ are correctly specified;

R3. g, $\bar{Q}^{Z,\mathbf{a},\mathbf{a}'}$, and either Q^Z or γ are correctly specified;

then $P_0 D^{*,\mathbf{a},\mathbf{a}'}(Q,g) = 0$ implies $\Psi^{\mathbf{a},\mathbf{a}'}(Q) = \Psi^{\mathbf{a},\mathbf{a}'}(Q_0)$.

Proof. See appendix A1. □

From the robustness statement, we see that even though $\gamma^{s,t}$ is not orthogonal with Q and g, it has very profound implications in the practicality of the robustness properties of the EIC. Specifically, while the functionals g and $\bar{Q}^{L,\mathbf{a},\mathbf{a}'}$ and $\bar{Q}^{Z,\mathbf{a},\mathbf{a}'}$ can be estimated using a variety of regression techniques, estimation of the conditional densities Q^L and Q^Z is very challenging when L and Z are high-dimensional. Statements R2 and R3 allows one to replace estimation of the conditional densities with estimation of $\gamma^{s,t}$, which are conditional probabilities of binary variables. This way, we have traded off estimating Q^Z and Q^L, which are orthogonal components to g, with estimating easier functionals $\gamma^{s,t}$, which are not compatible with the time-ordering decomposition of the likelihood but are estimable from the data nonetheless. In a sense, this is a tradeoff between elegance and practicality.

It is easy to note that if $\mathbf{a} = \mathbf{a}'$, then (5.7) equals the efficient influence curve for the overall treatment effect of a time varying exposure (see e.g. van der Laan and Gruber (2012)). The EICs of both the NDE and NIE can be derived from (5.7) by a simple application of the delta method. We state them in a corollary without proof.

Corollary 5.1. *Suppose the conditions in theorem 5.1 hold for $a, a' \in \{0,1\}$. The efficient influence curve of the natural direct effect is given by*

$$D^{*,NDE}(P)(O) = D^{*,1,0}(P) - D^{*,0,0}(P),$$

and the efficient influence curve of the natural indirect effect is given by

$$D^{*,NIE}(P)(O) = D^{*,1,1}(P) - D^{*,1,0}(P).$$

Moreover, $D^{*,NDE}$ and D^*NIE satisfy the same robustness condition in theorem 5.1 for $\mathbf{a} = 0,1$ and $\mathbf{a}' = 0,1$.

The variances $Var_{P_0}(D^{*,\mathbf{a},\mathbf{a}'}(P_0))$, $Var_{P_0}(D^{*,NDE}(P_0))$, and $Var_{P_0}(D^{*,NIE}(P_0))$ are generalized Cramer-Rao lower bounds for the asymptotic variances of the RAL estimators of $\Psi^{\mathbf{a},\mathbf{a}'}(P_0)$, $\Psi^{NDE}(P_0)$, and $\Psi^{NIE}(P_0)$, respectively.

Estimators which satisfy the EIC equations will also inherit their robustness properties. We will present four estimators in the next section, two of which are robust and locally efficient.

5.4 Estimators

In this section, we develop the G-computation, IPW and TMLE estimators for the mediation functional (5.2); the estimators for the natural direct and indirect effects can be obtained by taking the corresponding differences.

The G-computation and the IPW (inverse probability-of-treatment weighted) estimators are consistent only if the estimates of all the relevant components of P_0 are consistent. On the other

hand, TMLE (targeted maximum likelihood) estimator satisfy the efficient influence curve equation, and hence remain unbiased under the model mis-specifications described in theorem 5.1; under appropriate regularity conditions, it will be asymptotically efficient (e.g. Bickel et al. (1997), van der Laan and Robins (2003b), van der Laan and Rose (2011)).

Let g_n and Q_n denote the estimators of g_0 and Q_0, respectively. In Q_n, $\bar{Q}_n^{L,\mathbf{a},\mathbf{a}'}$ and $\bar{Q}_n^{Z,\mathbf{a},\mathbf{a}'}$ may be density-based estimators that are obtained by plugging in the density estimates Q_n^L and Q_n^Z into the definition of the expectations, or they may be regression-based estimators that are obtained using the relations in (5.5). We will use the latter approach here.

Non-targeted substitution G-computation Estimator

The identification formula in (5.2) which defines that statistical estimand is generally known as the G-computation formula (Robins (1986)). Readily, it delivers a non-targeted substitution estimator (as opposed to the targeted substitution estimator that is TMLE), which is generally known as the G-computation estimator. Rewriting it in terms of $\bar{Q}^{L,\mathbf{a},\mathbf{a}'}$ and $\bar{Q}^{Z,\mathbf{a},\mathbf{a}'}$, as in (5.6), we can obtain a non-targeted substitution estimator $\Psi^{\mathbf{a},\mathbf{a}'}(Q_n)$ of $\Psi^{\mathbf{a},\mathbf{a}'}(Q_0)$, through non-targeted estimates $\bar{Q}_n^{L,\mathbf{a},\mathbf{a}'}$ and $\bar{Q}_n^{Z,\mathbf{a},\mathbf{a}'}$.

To estimate the marginal distribution $Q^{L_0}(P_0)$, we can use the empirical distribution of L_0, denoted $Q_n^{L_0}$. To estimate the conditional means $\bar{Q}^{L,\mathbf{a},\mathbf{a}'}(P_0)$ and $\bar{Q}^{Z,\mathbf{a},\mathbf{a}'}(P_0)$, we can use the following algorithm, which exploits the relations in (5.5), with available regression techniques in the literature.

1. Initiate $\bar{Q}_{\tau+1}^{Z,\mathbf{a},\mathbf{a}'} \equiv Y_\tau$.
2. At each $t = \tau, \dots 1$, in decreasing order, we have obtained estimators $\bar{Q}_{t+1,n}^{Z,\mathbf{a},\mathbf{a}'}$ from a previous step. We obtain $\bar{Q}_{t,n}^{Z,\mathbf{a},\mathbf{a}'}$ and $\bar{Q}_{t,n}^{L,\mathbf{a},\mathbf{a}'}$, in that order, as follows:
   a) Regress $\bar{Q}_{t+1,n}^{Z,\mathbf{a},\mathbf{a}'}(\mathbf{L}_t, \mathbf{Z}_t)$ on observed values $\mathbf{A}_t, \mathbf{Z}_t, \mathbf{L}_{t-1}$ among observations that remained uncensored at time t, and evaluate the fitted function at the observed values $\mathbf{Z}_t, \mathbf{L}_{t-1}$ and the intervened values $\mathbf{A}_t = \mathbf{a}_t$ for these uncensored observations. This results in the estimates $\bar{Q}_{t,n}^{L,\mathbf{a},\mathbf{a}'}(\mathbf{L}_{t-1}, \mathbf{Z}_t)$.
   b) Regress the newly minted $\bar{Q}_{t,n}^{L,\mathbf{a},\mathbf{a}'}(\mathbf{L}_{t-1}, \mathbf{Z}_t)$ on $\mathbf{A}_t, \mathbf{Z}_{t-1}, \mathbf{L}_{t-1}$ among observations that remained uncensored at time t, and evaluate the fitted function at the observed values $\mathbf{Z}_{t-1}, \mathbf{L}_{t-1}$ and the intervened values $\mathbf{A}_t = \mathbf{a}'_t$ for these uncensored observations. This results in the estimates $\bar{Q}_{t,n}^{Z,\mathbf{a},\mathbf{a}'}(\mathbf{L}_{t-1}, \mathbf{Z}_{t-1})$.
3. After running the algorithm in step (2) sequentially from $t = \tau$ down to $t = 1$, we have $\bar{Q}_{t=1,n}^{Z,\mathbf{a},\mathbf{a}'}(L_0)$ for each of the n observations.

The G-computation estimator is given by

$$\psi_n^{Gcomp,\mathbf{a},\mathbf{a}'} \equiv \Psi^{\mathbf{a},\mathbf{a}'}(Q_n) = \frac{1}{n}\sum_{i=1}^{n} \bar{Q}_{t=1,n}^{Z,\mathbf{a},\mathbf{a}'}(L_{0,i}) \tag{5.10}$$

Consistency of $\psi_n^{Gcomp,\mathbf{a},\mathbf{a}'}$ relies on consistency of Q_n. Correct specification of Q_0 under a finite dimensional parametric model is possible only in limited applications. Alternatively, we may use machine learning algorithms, such as Super Learner. This option is more enticing, especially when used with the regression-based approach, since there are more data-adaptive techniques available to estimate the conditional mean of a binary variable via regression. However, theoretical results on the asymptotic behavior, such as a central limit theorem, of the resulting estimator $\Psi^{\mathbf{a},\mathbf{a}'}(Q_n)$ are not available. Moreover, a non-targeted estimator Q_n of Q_0 is obtained by minimizing a global loss function for Q_0, not for $\Psi^{\mathbf{a},\mathbf{a}'}(Q_0)$. This means, in particular, that the bias-variance tradeoff in Q_n is optimized for the high-dimensional nuisance parameter Q_0, instead of a much lower-dimensional parameter of interest $\Psi^{\mathbf{a},\mathbf{a}'}(Q_0)$. The proposed targeted estimator in section 5.4 aims to address these two issues by providing a substitution estimator that is asymptotically linear (under appropriate regularity conditions), and optimizes the bias-variance tradeoff of Q_n towards $\Psi(Q_0)$ via an updating step.

Inverse Probability Weighted Estimator

Instead of estimating the condition expectations $\bar{Q}^{L,\mathbf{a},\mathbf{a}'}$ and $\bar{Q}^{Z,\mathbf{a},\mathbf{a}'}$, one may wish to employ the researcher's knowledge about the treatment assignment and mediation densities. To this end, consider the following function:

$$D^{IPW,\mathbf{a},\mathbf{a}'}(g,Q^Z) \equiv \frac{I(\mathbf{A}_\tau = \mathbf{a}_\tau)}{\prod_{j=1}^{\tau} g(a_j \mid \mathbf{L}_{j-1},\mathbf{Z}_{j-1},\mathbf{a}_{j-1})} \prod_{j=1}^{\tau} \frac{Q^Z(Z_j \mid \mathbf{a}'_j,\mathbf{Z}_{j-1},\mathbf{L}_{j-1})}{Q^Z(Z_j \mid \mathbf{a}_j,\mathbf{Z}_{j-1},\mathbf{L}_{j-1})} Y_\tau. \tag{5.11}$$

It is easy to note that $P_0 D^{IPW,\mathbf{a},\mathbf{a}'}(g_0,Q_0^Z) = \Psi^{\mathbf{a},\mathbf{a}'}(P_0)$. Therefore, given estimators g_n and Q_n^Z, the Inverse Probability Weighted (IPW) Estimator of $\Psi^{\mathbf{a},\mathbf{a}'}(P_0)$ is given by

$$\psi_n^{IPW,\mathbf{a},\mathbf{a}'} \equiv P_n D^{IPW,\mathbf{a},\mathbf{a}'}(g_n,Q_n^Z). \tag{5.12}$$

Consistency of $\psi_n^{IPW,\mathbf{a},\mathbf{a}'}$ relies on consistency of g and Q^Z. As noted in previously, if Z is high dimensional, we may replace estimation of the densities Q^Z with estimation of the conditional probabilities $\gamma^{s,t}$ of A_s given $\mathbf{A}_{s-1}$, $\mathbf{Z}_t$ and $\mathbf{L}_t$ or $\mathbf{L}_{t-1}$. This way, using (5.8), we can rewrite

$$\begin{aligned} &D^{IPW,\mathbf{a},\mathbf{a}'}(g,Q^Z) = D^{IPW,\mathbf{a},\mathbf{a}'}(g,\gamma) \\ &= \frac{I(\mathbf{A}_\tau = \mathbf{a}_\tau)}{\prod_{j=1}^{\tau} g(a_j \mid \mathbf{L}_{j-1},\mathbf{Z}_{j-1},\mathbf{a}_{j-1})} \\ &\times \left\{ \prod_{j=1}^{\tau} \prod_{s=1}^{j} \frac{\gamma^{1,s,j}(\mathbf{a}'_s,\mathbf{Z}_j,\mathbf{L}_{j-1})}{\gamma^{1,s,j}(\mathbf{a}_s,\mathbf{Z}_j,\mathbf{L}_{j-1})} \times \frac{g(a_j \mid \mathbf{a}_{j-1},\mathbf{Z}_{j-1},\mathbf{L}_{j-1})}{g(a'_j \mid \mathbf{a}'_{j-1},\mathbf{Z}_{j-1},\mathbf{L}_{j-1})} \prod_{s=1}^{j-1} \frac{\gamma^{2,s,j-1}(\mathbf{a}_s,\mathbf{Z}_{j-1},\mathbf{L}_{j-1})}{\gamma^{2,s,j-1}(\mathbf{a}'_s,\mathbf{Z}_{j-1},\mathbf{L}_{j-1})} Y_\tau \right\}, \end{aligned} \tag{5.13}$$

and obtain an IPW estimator as $\psi_n^{IPW,\mathbf{a},\mathbf{a}'} = P_n D^{IPW,\mathbf{a},\mathbf{a}'}(g_n, \gamma_n)$.

The asymptotic theory of the IPW estimator is well understood in the literature. We refer the reader to Robins (1999) and van der Laan and Robins (2003b). In particular, it is an asymptotically linear estimator with influence curve $D^{IPW,\mathbf{a}.\mathbf{a}'}$. Consequently, the asymptotic variance of $\sqrt{n}\left(\psi_n^{IPW,\mathbf{a},\mathbf{a}'} - \Psi^{\mathbf{a},\mathbf{a}'}(P_0)\right)$ can be estimated by the sample variance $\hat{\text{Var}} D^{IPW,\mathbf{a}.\mathbf{a}'}(g_n, Q_n^Z)$, or its alternative version in (5.13).

Due to its inverse weighting by treatment and censoring probabilities, this estimator is particularly sensitive to near positivity violations. In particular, if the outcome of interest has a bounded range, the IPW estimator is not guaranteed to stay within this range when the inverse weights become large. Substitution estimators like G-computation and TMLE slightly mitigate this problem by incorporating global information in the parameter map, however, but the effect of near positivity violations still takes form of poor smoothing in these estimators.

Targeted Maximum Likelihood Estimator

To maximize finite sample gain and provide more stable estimates in the presence of near positivity violations, one can make use of the substitution principle. The targeted maximum likelihood estimation (TMLE, van der Laan and Rubin (2006)) provides a substitution-based estimator which also satisfies the EIC equation, thereby remaining unbiased under model mis-specifications.

In a glimpse, our strategy consists of targetedly updating given initial estimators Q_n of Q_0 by minimizing a pre-specified loss along a least favorable (with respect to $\Psi^{\mathbf{a},\mathbf{a}'}(P_0)$) submodel through Q_n, then we obtain a substitution estimator of the parameter by evaluating $\Psi^{\mathbf{a},\mathbf{a}'}(Q_n^*)$. A byproduct of this updating procedure is that the g_n and Q_n^* satisfy $P_n D^{*,\mathbf{a},\mathbf{a}'}(Q_n^*, g_n) = 0$.

More specifically, in light of (5.6), we only need to update $\bar{Q}^{L,\mathbf{a},\mathbf{a}'}$ and $\bar{Q}^{Z,\mathbf{a},\mathbf{a}'}$. From the recursive relations in (5.5), we will use for $\bar{Q}_t^{L,\mathbf{a},\mathbf{a}'}$ and $\bar{Q}_t^{Z,\mathbf{a},\mathbf{a}'}$ the loss functions

$$\begin{aligned} L(\bar{Q}_t^{L,\mathbf{a},\mathbf{a}'}) &= -\left\{\bar{Q}_{t+1}^{Z,\mathbf{a},\mathbf{a}'} \log\left(\bar{Q}_t^{L,\mathbf{a},\mathbf{a}'}\right) + (1-\bar{Q}_{t+1}^{Z,\mathbf{a},\mathbf{a}'}) \log\left(1-\bar{Q}_t^{L,\mathbf{a},\mathbf{a}'}\right)\right\}, \text{ and} \\ L(\bar{Q}_t^{Z,\mathbf{a},\mathbf{a}'}) &= -\left\{\bar{Q}_t^{L,\mathbf{a},\mathbf{a}'} \log\left(\bar{Q}_t^{Z,\mathbf{a},\mathbf{a}'}\right) + (1-\bar{Q}_t^{L,\mathbf{a},\mathbf{a}'}) \log\left(1-\bar{Q}_t^{L,\mathbf{a},\mathbf{a}'}\right)\right\}. \end{aligned} \tag{5.14}$$

We suppressed in the notation the indexing by the functionals at the next time point. The corresponding least favorable submodels through $\bar{Q}_t^{L,\mathbf{a},\mathbf{a}'}$ and $\bar{Q}_t^{Z,\mathbf{a},\mathbf{a}'}$, respectively, are parametrized as $\left\{\bar{Q}_t^{L,\mathbf{a},\mathbf{a}'}(\varepsilon) : \varepsilon \in \mathbb{R}\right\}$ and $\left\{\bar{Q}_t^{Z,\mathbf{a},\mathbf{a}'}(\varepsilon) : \varepsilon \in \mathbb{R}\right\}$, and they ought to satisfy the score conditions $\frac{\partial L\left(\bar{Q}_t^{L,\mathbf{a},\mathbf{a}'}(\varepsilon)\right)}{\partial \varepsilon}|_{\varepsilon=0} = D_t^{L,\mathbf{a},\mathbf{a}'}(Q,g)$ and $\frac{\partial L\left(\bar{Q}_t^{Z,\mathbf{a},\mathbf{a}'}(\varepsilon)\right)}{\partial \varepsilon}|_{\varepsilon=0} = D_t^{Z,\mathbf{a},\mathbf{a}'}(Q,g)$. In particular, we may use the corresponding least favorable submodels

$$\begin{aligned} \bar{Q}_t^{L,\mathbf{a},\mathbf{a}'}(\varepsilon) &= \text{expit}\left(\text{logit}\,\bar{Q}_t^{L,\mathbf{a},\mathbf{a}'} + \varepsilon C_t^{L,\mathbf{a},\mathbf{a}'}\right), \text{ and} \\ \bar{Q}_t^{Z,\mathbf{a},\mathbf{a}'}(\varepsilon) &= \text{expit}\left(\text{logit}\,\bar{Q}_t^{Z,\mathbf{a},\mathbf{a}'} + \varepsilon C_t^{Z,\mathbf{a},\mathbf{a}'}\right). \end{aligned} \tag{5.15}$$

We are now ready to describe the TMLE algorithm, which will targetedly estimate $\bar{Q}_t^{L,\mathbf{a},\mathbf{a}'}$ and $\bar{Q}_t^{Z,\mathbf{a},\mathbf{a}'}$ sequentially in order of decreasing t.

1. Obtain initial estimators g_n and either the densities (Q_n^Z, Q_n^L) or the predictions γ_n defined in section 5.3. These will be used to obtain estimates $C_{t,n}^{Z,\mathbf{a},\mathbf{a}'}$ and $C_{t,n}^{L,\mathbf{a},\mathbf{a}'}$, see (5.8) and (5.9). Let $Q_n^{L_0}$ be the empirical distribution of baseline covariates L_0.

2. Initiate $\hat{Q}_{\tau+1,n}^{*,Z,\mathbf{a},\mathbf{a}'} = Y_\tau$.

3. At each $t = \tau, \ldots 1$, in decreasing order, we have obtained targeted estimator $\bar{Q}_{t+1,n}^{*,Z,\mathbf{a},\mathbf{a}'}$ from a previous step. We obtain targeted estimator $\bar{Q}_{t,n}^{*,Z,\mathbf{a},\mathbf{a}'}$ and $\bar{Q}_{t,n}^{*,L,\mathbf{a},\mathbf{a}'}$, in that order, as follows:

   a) Obtain initial estimator $\bar{Q}_{t,n}^{L,\mathbf{a},\mathbf{a}'}$ by regressing $\bar{Q}_{t+1,n}^{*,Z,\mathbf{a},\mathbf{a}'}(\mathbf{L}_t, \mathbf{Z}_t)$ on observed values $\mathbf{A}_t, \mathbf{Z}_t, \mathbf{L}_{t-1}$ among observations that remained uncensored at time t, and evaluate the fitted function at the observed values $\mathbf{Z}_t, \mathbf{L}_{t-1}$ and the intervened values $\mathbf{A}_t = \mathbf{a}_t$ for these uncensored observations. This results in the estimates $\bar{Q}_{t,n}^{L,\mathbf{a},\mathbf{a}'}(\mathbf{L}_{t-1}, \mathbf{Z}_t)$. Update this estimate using $\bar{Q}_{t,n}^{*,L,\mathbf{a},\mathbf{a}'} \equiv \bar{Q}_{t,n}^{L,\mathbf{a},\mathbf{a}'}(\varepsilon_{t,n}^L)$, where

   $$\varepsilon_{t,n}^L \equiv \arg\min_{\varepsilon} P_n L\left(\bar{Q}_{t,n}^{L,\mathbf{a},\mathbf{a}'}(\varepsilon)\right).$$

   This optimal fluctuation amount can obtained by a weighted logistic regression of the expectant $\bar{Q}_{t+1,n}^{*,Z,\mathbf{a},\mathbf{a}'}(\mathbf{L}_t, \mathbf{Z}_t)$ onto $C_{t,n}^{L,\mathbf{a},\mathbf{a}'}$ with offset logit $\left(\bar{Q}_{t,n}^{L,\mathbf{a},\mathbf{a}'}(\mathbf{L}_{t-1}, \mathbf{Z}_t)\right)$.

   b) Next, obtain an initial estimator $\bar{Q}_{t,n}^{Z,\mathbf{a},\mathbf{a}'}$ by regressing the newly minted targeted estimate $\bar{Q}_{t,n}^{*,L,\mathbf{a},\mathbf{a}'}(\mathbf{L}_{t-1}, \mathbf{Z}_t)$ on $\mathbf{A}_t, \mathbf{Z}_{t-1}, \mathbf{L}_{t-1}$ among observations that remained uncensored at time t, and evaluate the fitted function at the observed values $\mathbf{Z}_{t-1}, \mathbf{L}_{t-1}$ and the intervened values $\mathbf{A}_t = \mathbf{a}_t'$ for these uncensored observations. This results in the estimates $\bar{Q}_{t,n}^{Z,\mathbf{a},\mathbf{a}'}(\mathbf{L}_{t-1}, \mathbf{Z}_{t-1})$. Update this estimate using $\bar{Q}_{t,n}^{*,Z,\mathbf{a},\mathbf{a}'} \equiv \bar{Q}_{t,n}^{Z,\mathbf{a},\mathbf{a}'}(\varepsilon_{t,n}^Z)$, where

   $$\varepsilon_{t,n}^Z \equiv \arg\min_{\varepsilon} P_n L\left(\bar{Q}_{t,n}^{Z,\mathbf{a},\mathbf{a}'}(\varepsilon)\right).$$

   This optimal fluctuation can obtained by a weighted logistic regression of the expectant $\bar{Q}_{t,n}^{*,L,\mathbf{a},\mathbf{a}'}(\mathbf{L}_{t-1}, \mathbf{Z}_t)$ onto $C_{t,n}^{Z,\mathbf{a},\mathbf{a}'}$ with offset logit $\left(\bar{Q}_{t,n}^{Z,\mathbf{a},\mathbf{a}'}(\mathbf{L}_{t-1}, \mathbf{Z}_{t-1}\right)$.

4. After running the algorithm in step (3) sequentially from $t = \tau$ down to $t = 1$, we have targeted estimates $\bar{Q}_{t=1,n}^{*,Z,\mathbf{a},\mathbf{a}'}(L_0)$ for each of the n observations.

Let $Q_n^* \equiv \left(Q_n^{L_0}, Q_n^Z, Q_n^L, \bar{Q}_n^{*,L,\mathbf{a},\mathbf{a}'}, \bar{Q}_n^{*,Z,\mathbf{a},\mathbf{a}'}\right)$. In case that one used γ_n instead of the densities Q_n^Z, Q_n^L, the Q_n^* will be defined accordingly. The TMLE estimator is given by

$$\psi_n^{*,\mathbf{a},\mathbf{a}'} \equiv \Psi^{\mathbf{a},\mathbf{a}'}(Q_n^*) = \frac{1}{n}\sum_{i=1}^{n} \bar{Q}_{t=1,n}^{*,Z,\mathbf{a},\mathbf{a}'}(L_{0,i}) \tag{5.16}$$

By construction, this estimator satisfies $P_n D^{*,\mathbf{a},\mathbf{a}'}(Q_n^*, g_n) = 0$. Consequently, it inherits robustness of EIC described in theorem 5.1. Under certain regularity conditions, it is efficient at P_0, with influence curve $D^*(Q_0, g_0)$. In particular, the asymptotic variance of $\sqrt{n}\left(\psi_n^{*,\mathbf{a},\mathbf{a}'} - \Psi^{\mathbf{a},\mathbf{a}'}(P_0)\right)$ can be estimated using the sample variance $\hat{\text{Var}} D^{*,\mathbf{a},\mathbf{a}'}(Q_n^*, g_n)$.

5.5 Simulation Study

In this section, we conduct a simulation study to evaluate the comparative performance of these three estimators.

Consider the data structure $O = (L_0, A_1, Z_1, L_1, \ldots, A_\tau, Z_\tau, L_\tau)$, with $\tau = 2$. L_0 encodes two baseline covariates L_0^1 and L_0^2, A_t encodes a binary exposure A_t^E and a censoring indicator A_t^C, Z_t is a binary mediator of interest, L_t encodes two time varying covariates L_t^1 and L_t^2, and a death indicator Y_t. These variable are distributed according to the following data generating distribution

$$
\begin{aligned}
&L_0^1 \sim Bern(0.1);\\
&L_0^2 \sim Bern(0.7);\\
&A_t^E \sim Bern\Big(\text{expit}\big(-0.7 + 2L_0^1 - 0.5L_0^2\\
&\quad + I(t>1)\left(-0.5A_{t-1}^E + 0.1L_{t-1}^1 - 0.2L_{t-1}^2 + 0.7A_{t-1}^E \times L_0^1 + 0.2A_{t-1}^E \times L_{t-1}^1\right)\big)\Big)\\
&A_t^C \sim Bern\Big(\text{expit}\big(1 - 0.1t + 1.2L_0^1 - 0.5L_0^2 - 0.5A_t^E + 0.2A_t^E \times L_0^1\\
&\quad + I(t>1)\left(0.1L_{t-1}^1 - 0.2L_{t-1}^2 - 0.2A_t^E \times L_{t-1}^1\right)\big)\Big)\\
&Z_t \sim Bern\Big(\text{expit}\big(-1 + 0.4t + 1L_0^1 - 0.1L_0^2 + 1.5A_t^E - 0.3A_t^E \times L_0^1 + I(t>1)(-Z_{t-1})\big)\Big)\\
&L_t^1 \sim Bern\Big(\text{expit}\big(-1 + 0.4t + 1L_0^1 - 0.1L_0^2 + 1.5A_t^E + Z_t - 0.3Z_t \times A_t^E + I(t>1)\left(-L_{t-1}^1\right)\big)\Big)\\
&L_t^2 \sim Bern\Big(\text{expit}\big(-2 + 0.2t + 1L_0^1 - 0.1L_0^2 - A_t^E + 2Z_t - 0.2A_t^E \times L_t^1 - 0.3Z_t \times A_t^E\\
&\quad + I(t>1)\left(L_{t-1}^2) - 0.4A_{t-1}^E \times L_t^1\right)\big)\Big)\\
&Y_t \sim Bern\Big(\text{expit}\big(-2 + 0.1t + 2L_0^1 - 0.7L_0^2 + A_t^E - 2L_t^1 - 0.1L_T^2 + 0.4Z_t - 0.2A_t^E \times L_t^1 - 0.2A_t^E \times L_0^1\\
&\quad + 0.5A_t^E \times Z_t + I(t>1)\big(-1.5L_{t-1}^1 - 0.3L_{t-1}^2 - 0.4A_{t-1}^E \times L_t^1\\
&\qquad - 0.4A_{t-1}^E \times L_{t-1}^1 + 0.4A_{t-1}^E \times Z_{t-1}\big)\big)\Big).
\end{aligned}
$$

After either censoring or death, all subsequent variables take a default value. The target parameter of interest is $\Psi^{1,0}(P_0) = 0.192$.

Correctly specified g_n, Q_n^Z and Q_n^L are obtained using logistic regression on the parents of the corresponding variables, whereas correctly specified $\bar{Q}^{Z,\mathbf{a},\mathbf{a}'}$ and $\bar{Q}^{L,\mathbf{a},\mathbf{a}'}$ are obtained using Super

Learner with candidate estimators `glm` and `neural net`. the misspecified ones omit important predictors. Their misspecified counterparts omit important predictors.

Results

We consider sample sizes $n = 500$ and $n = 5000$. Bias, variance and mean squared error (MSE) for each sample size are estimated over the 500 datasets. In the table 5.1 below, legend for model specifications are as follows:

Table 5.1: Bias, variance and MSE oer 500 simulations.

		Bias		Var		MSE	
	n	500	5000	500	5000	500	5000
all correct	IPW	0.0065	0.0008	0.0072	0.0013	0.0073	0.0013
	Gcomp	0.0152	0.0136	0.0032	0.0003	0.0035	0.0005
	TMLE	0.0377	0.0074	0.0246	0.0023	0.0260	0.0024
A misspec.	IPW	0.0866	0.0889	0.0052	0.0004	0.0127	0.0083
	TMLE	0.0077	0.0034	0.0041	0.0002	0.0042	0.0002
L misspec.	Gcomp	0.0866	0.0921	0.0036	0.0007	0.0111	0.0092
	TMLE	0.0243	0.0000	0.0216	0.0013	0.0221	0.0013
Z misspec.	IPW	0.0074	0.0000	0.0130	0.0020	0.0131	0.0020
	Gcomp	0.0029	0.0058	0.0059	0.0004	0.0059	0.0004
	TMLE	0.0313	0.0027	0.0292	0.0034	0.0302	0.0034

As predicted by theory, under misspecified A or L components TMLE provides bias reduction over misspecified IPW and Gcomp estimators. The variance of TMLE is also relatively smaller or comparable to IPW and Gcomp, therefore the resulting MSE in these two cases is smaller for TMLE. However, when Z is misspecified, TMLE can still provide bias reduction over Gcomp, but the IPW estimator is relative insensitive to misspecification in Z, Moreover, the variance of TMLE is larger for Gcomp and IPW, relative to the bias, resulting in larger MSE for the TMLE. When all estimators are correct, there is no obvious gain from TMLE.s

5.6 Summary

In this chapter, we proposed to adopt the stochastic interventions (SI) approach of Didelez et al. (2006), where the mediator is considered an intervention variable onto which a given distribution is enforced, to formulate parameters of interest in longitudinal mediation analysis with time varying mediator and exposures. The second contribution of this chapter is a general semiparametric inference framework for the resulting effect parameters. More specifically, efficient influence curves under a locally saturated semiparametric model are derived, and their robustness properties are established. In many applications where the mediator densities are difficult to estimate,

regression-based estimators of these iterated expectations are viable alternatives to substitution-based estimators that rely on consistent estimation of the mediator densities. We also developed the G-computation, IPW and TMLE estimators for the mediational functional.

Under the SI formulation, the treatment of interest as well as the mediator variables are regarded as intervention variables. One can obtain a total effect decomposition and the subsequent definition of natural direct and indirect effects that are analogous to those in Pearl (2001). The natural direct effect (NDE) under this formulation has an intrinsic interpretation as a weighted average of controlled direct effects (CDE), since the CDE can be considered as a deterministic intervention on the treatment and mediator variables. By regarding the mediator variables as intervention variables, the SI formulation requires external specification of a counterfactual mediator distribution. It is important to note that causal mediation, under either SI or non-SI approaches, presupposes that the mediator of interest is amenable to external manipulation. In applications where such manipulations are not conceivable, we should be cautious that causal mediation can only offer answers to purely mechanistic questions defined under hypothetical experiments.

5.7 Chapter Appendix

Appendix A.1: Proof of Theorem 5.1

For any $P \in \mathscr{M}$, we may factor the likelihood according to the time ordering:

$$\begin{aligned} p(O) &= p(L_0) \\ &\times \prod_{t=1}^{\tau} \Big(p(A_t \mid \mathbf{A}_{t-1}, \mathbf{Z}_{t-1}, \mathbf{L}_{t-1}) p(Z_t \mid \mathbf{A}_t, \mathbf{Z}_{t-1}, \mathbf{L}_{t-1}) p(L_t \mid \mathbf{A}_t, \mathbf{Z}_t, \mathbf{L}_{t-1}) \Big). \end{aligned} \tag{5.17}$$

For $O_j \in \{L_0, A_t, Z_t, L_t : t = 1, \ldots, \tau\}$, let P_{O_j} denote the conditional probability of $P_{O_j}(O_j \mid Pa(O_j))$.

Let $L^2(P)$ denote the Hilbert space of mean zero functions of O, endowed with the covariance operator. Consider a rich class of one-dimensional parametric submodels $P(\varepsilon)$ that are generated by only fluctuating P_{O_j}. Under our model, no restrictions are imposed on the conditional probabilities P_{O_j}. As a result, given any function $S_{O_j} \in L^2(P)$ of $(O_j, Pa(O_j))$ with finite variance and $E_P(S_{O_j}(O_j, Pa(O_j)) \mid Pa(O_j)) = 0$, the fluctuation $P_{O_j}(\varepsilon) = (1 + \varepsilon S_{O_j}(O_j, Pa(O_j)))P_{O_j}$ is a valid one-dimensional submodel with score S_{O_j}. Therefore, the tangent subspaces corresponding to fluctuations of each P_{O_j} are given by

$$\begin{aligned} T(P_{L_0}) &= \{S_{L_0}(L_0) : E_P(S_{L_0}) = 0\} \\ T(P_{A_t}) &= \{S_{A_t}(A_t, \mathbf{A}_{t-1}, \mathbf{Z}_{t-1}, \mathbf{L}_{t-1}) : E_P(S_{A_t} \mid \mathbf{A}_{t-1}, \mathbf{Z}_{t-1}, \mathbf{L}_{t-1}) = 0\} \\ T(P_{Z_t}) &= \{S_{Z_t}(Z_t, \mathbf{A}_t, \mathbf{Z}_{t-1}, \mathbf{L}_{t-1}) : E_P(S_{Z_t} \mid \mathbf{A}_t, \mathbf{Z}_{t-1}, \mathbf{L}_{t-1}) = 0\} \\ T(P_{L_t}) &= \{S_{L_t}(L_t, \mathbf{A}_t, \mathbf{Z}_t, \mathbf{L}_{t-1}) : E_P(S_{L_t} \mid \mathbf{A}_t, \mathbf{Z}_t, \mathbf{L}_{t-1}) = 0\}. \end{aligned}$$

Due to the factorization in (5.17), $T(P_i)$ is orthogonal to $T(P_{O_j})$ for $O_i \neq O_j$. Moreover, the tangent space $T(P)$, corresponding to fluctuations of the entire likelihood, is given by the orthogonal sum

of these tangent subspaces, i.e. $T(P) = \bigoplus_j T(P_{O_j})$, and any score $S(O) \in T(P)$ can be decomposed as $\sum_j S_{O_j}(O)$.

Under this generous definition of the tangent subspaces, any function $S(O)$ that has zero mean and finite variance under P is contained in $T(P)$. This implies in particular that any gradient for the pathwise derivative of $\Psi^{\mathbf{a},\mathbf{a}'}(\cdot)$ is contained in $T(P)$, and is thus in fact the canonical gradient. Therefore, it suffices to show that $D^{*,\mathbf{a},\mathbf{a}'}(\cdot)$ in (5.7) is a gradient for the pathwise derivative of $\Psi^{\mathbf{a},\mathbf{a}'}(\cdot)$.

Indeed, for any $S(O) = \sum_j S_{O_j}(O) \in T(P)$, let $P_S(\varepsilon)$ denote the fluctuation of P with score S. Under appropriate regularity conditions, the pathwise derivative at P can be expressed as

$$
\begin{aligned}
&\frac{d}{d\varepsilon}\Psi^{\mathbf{a},\mathbf{a}'}(P_S(\varepsilon))\mid_{\varepsilon=0} \\
&= \frac{d}{d\varepsilon}\Big|_{\varepsilon=0} \left\{ \sum_{\mathbf{l}}\sum_{\mathbf{z}} \Bigg(((1+\varepsilon S_{L_0})P_{L_0})(\ell_0) \right. \\
&\left. \times \prod_{t=1}^{t_0} \left((1+\varepsilon S_{Z_t})P_{Z_t}\right)\left(z_t \mid \ell_0, \mathbf{A}_t = \mathbf{a}'_t, \mathbf{z}_{t-1}, \mathbf{l}_{t-1}\right) \left((1+\varepsilon S_{L_t})P_{L_t}\right)\left(\ell_t \mid \ell_0, \mathbf{A}_t = \mathbf{a}'_t, \mathbf{z}_{t-1}, \mathbf{l}_{t-1}\right) \Bigg) \right\} \\
&= \sum_{\mathbf{l}}\sum_{\mathbf{z}} \left(Q^{L_0}(\ell_0)\mathbf{Q}^Z(\mathbf{z} \mid \ell_0, \mathbf{A} = \mathbf{a}', \mathbf{l})\mathbf{Q}^L(\mathbf{l} \mid \ell_0, \mathbf{A} = \mathbf{a}, \mathbf{z}) \sum_{t=1}^{\tau} S_{L_t} \right) \qquad (5.18) \\
&+ \sum_{\mathbf{l}}\sum_{\mathbf{z}} \left(Q^{L_0}(\ell_0)\mathbf{Q}^Z(\mathbf{z} \mid \ell_0, \mathbf{A} = \mathbf{a}', \mathbf{l})\mathbf{Q}^L(\mathbf{l} \mid \ell_0, \mathbf{A} = \mathbf{a}, \mathbf{z}) \sum_{t=1}^{\tau} S_{Z_t} \right) \qquad (5.19) \\
&+ \sum_{\mathbf{l}}\sum_{\mathbf{z}} \left(Q^{L_0}(\ell_0)\mathbf{Q}^Z(\mathbf{z} \mid \ell_0, \mathbf{A} = \mathbf{a}', \mathbf{l})\mathbf{Q}^L(\mathbf{l} \mid \ell_0, \mathbf{A} = \mathbf{a}, \mathbf{z}) \right) S_{L_0}, \qquad (5.20)
\end{aligned}
$$

where $\mathbf{Q}^Z(\mathbf{z} \mid \ell_0, \mathbf{A} = \mathbf{a}', \mathbf{l}) \equiv \prod_{t=1}^{\tau} Q^Z(z_t \mid \ell_0, \mathbf{A}_t = \mathbf{a}'_t, \mathbf{z}_{t-1}, \mathbf{l}_{t-1})$, similarly for $\mathbf{Q}^L$.

Note firstly that for every $t = 1, \ldots, t_0$,

$$
\begin{aligned}
&E_P\left(D_t^{L,\mathbf{a},\mathbf{a}'}(P)(O) S_{L_t}(L_t, L_0, \mathbf{A}_t, \mathbf{Z}_t, \mathbf{L}_{t-1}) \right) \\
&= \sum_{\mathbf{l}}\sum_{\mathbf{z}} Q^{L_0}(\ell_0)\mathbf{Q}^Z(\mathbf{z} \mid \ell_0, \mathbf{A} = \mathbf{a}', \mathbf{l})\mathbf{Q}^L(\mathbf{l} \mid \ell_0, \mathbf{A} = \mathbf{a}, \mathbf{z}) S_{L_t}.
\end{aligned}
$$

Therefore, (5.18) can be written as

$$
E_P\left(D_t^{L,\mathbf{a},\mathbf{a}'}(P)(O) \sum_{t=1}^{\tau} S_{L_t}(L_t, L_0, \mathbf{A}_t, \mathbf{Z}_t, \mathbf{L}_{t-1}) \right).
$$

Moreover, $D^{L,\mathbf{a},\mathbf{a}'}(P)(O) \in T(P_{L_t|Pa(L_t)})$ by the definition of these tangent subspaces. It thus follows from the orthogonal decomposition of $T(P)$ that

$$
E_P\left\{ D_t^{L,\mathbf{a},\mathbf{a}'}(P)(O) \times S_{L_t}(L_t, L_0, \mathbf{A}_t, \mathbf{Z}_t, \mathbf{L}_{t-1}) \right\} = E_P\left\{ D_t^{L,\mathbf{a},\mathbf{a}'}(P) \left(S_{L_0} + \sum_{t=1}^{\tau} S_{A_t} + S_{Z_t} + S_{L_t} \right) \right\}.
$$

By similar arguments, (5.19) can be written as

$$E_P\left\{D_t^{Z,\mathbf{a},\mathbf{a}'}(P)(O)\times S_{Z_t}(Z_t,L_0,\mathbf{A}_t,\mathbf{Z}_{t-1},\mathbf{L}_{t-1})\right\}=E_P\left\{D_t^{Z,\mathbf{a},\mathbf{a}'}(P)\left(S_{L_0}+\sum_{t=1}^{\tau}S_{A_t}+S_{Z_t}+S_{L_t}\right)\right\}.$$

and (5.20) can be written as

$$E_P\left\{D^{L_0,\mathbf{a},\mathbf{a}'}(P)(O)S_{L_0}(L_0)\right\}=E_P\left\{D^{L_0,\mathbf{a},\mathbf{a}'}(P)\left(S_{L_0}+\sum_{t=1}^{\tau}S_{A_t}+S_{Z_t}+S_{L_t}\right)\right\}.$$

Combining these results, one concludes that

$$\frac{d}{d\varepsilon}\Psi^{\mathbf{a},\mathbf{a}'}(P_S(\varepsilon))\mid_{\varepsilon=0}=E_P\left\{\left(D^{L_0,\mathbf{a},\mathbf{a}'}(P)(O)+\sum_{t=1}^{\tau}D_t^{Z,\mathbf{a},\mathbf{a}'}(P)(O)+D_t^{L,\mathbf{a},\mathbf{a}'}(P)(O)\right)S(O)\right\}$$

Therefore, $D^{*,\mathbf{a},\mathbf{a}'}(P)=D^{L_0,\mathbf{a},\mathbf{a}'}(P)+\sum_{t=1}^{\tau}D_t^{Z,\mathbf{a},\mathbf{a}'}(P)+D_t^{L,\mathbf{a},\mathbf{a}'}(P)$ is a gradient for the pathwise derivative of $\Psi^{\mathbf{a},\mathbf{a}'}$ at P. As discussed above, under the nonparametric model, $D^{*,\mathbf{a},\mathbf{a}'}(P)$ is in fact the canonical gradient.

To rewrite the functions $C_t^{L,\mathbf{a},\mathbf{a}'}$ in (5.8) and $C_t^{Z,\mathbf{a},\mathbf{a}'}$ in (5.9), it suffices to note that

$$\begin{aligned}
\frac{Q^Z(Z_j\mid\mathbf{a}'_j,\mathbf{Z}_{j-1},\mathbf{L}_{j-1})}{Q^Z(Z_j\mid\mathbf{a}_j,\mathbf{Z}_{j-1},\mathbf{L}_{j-1})}&=\frac{p\left(\mathbf{a}'_j,\mathbf{Z}_j,\mathbf{L}_{j-1}\right)}{p\left(\mathbf{a}'_j,\mathbf{Z}_{j-1},\mathbf{L}_{j-1}\right)}\frac{p\left(\mathbf{a}_j,\mathbf{Z}_{j-1},\mathbf{L}_{j-1}\right)}{p\left(\mathbf{a}_j,\mathbf{Z}_j,\mathbf{L}_{j-1}\right)}\\
&=\frac{p\left(\mathbf{a}'_j,\mathbf{Z}_j,\mathbf{L}_{j-1}\right)}{p\left(\mathbf{a}_j,\mathbf{Z}_j,\mathbf{L}_{j-1}\right)}\frac{p\left(\mathbf{a}_j,\mathbf{Z}_{j-1},\mathbf{L}_{j-1}\right)}{p\left(\mathbf{a}'_j,\mathbf{Z}_{j-1},\mathbf{L}_{j-1}\right)}\\
&=\left(\prod_{s=1}^{j}\frac{p(a'_s\mid\mathbf{a}'_{s-1},\mathbf{Z}_j,\mathbf{L}_{j-1})}{p(a_s\mid\mathbf{a}_{s-1},\mathbf{Z}_j,\mathbf{L}_{j-1})}\frac{p(\mathbf{Z}_j,\mathbf{L}_{j-1})}{p(\mathbf{Z}_j,\mathbf{L}_{j-1})}\right)\\
&\times\left(\frac{g(a_j\mid\mathbf{a}_{j-1},\mathbf{Z}_{j-1},\mathbf{L}_{j-1})}{g(a'_j\mid\mathbf{a}'_{j-1},\mathbf{Z}_{j-1},\mathbf{L}_{j-1})}\prod_{s=1}^{j-1}\frac{p(a_s\mid\mathbf{a}_{s-1},\mathbf{Z}_{j-1},\mathbf{L}_{j-1})}{p(a'_s\mid\mathbf{a}'_{s-1},\mathbf{Z}_{j-1},\mathbf{L}_{j-1})}\frac{p(\mathbf{Z}_{j-1},\mathbf{L}_{j-1})}{p(\mathbf{Z}_{j-1},\mathbf{L}_{j-1})}\right)\\
&=\prod_{s=1}^{j}\frac{\gamma^{1,s,j}(\mathbf{a}'_s,\mathbf{Z}_j,\mathbf{L}_{j-1})}{\gamma^{1,s,j}(\mathbf{a}_s,\mathbf{Z}_j,\mathbf{L}_{j-1})}\times\frac{g(a_j\mid\mathbf{a}_{j-1},\mathbf{Z}_{j-1},\mathbf{L}_{j-1})}{g(a'_j\mid\mathbf{a}'_{j-1},\mathbf{Z}_{j-1},\mathbf{L}_{j-1})}\prod_{s=1}^{j-1}\frac{\gamma^{2,s,j-1}(\mathbf{a}_s,\mathbf{Z}_{j-1},\mathbf{L}_{j-1})}{\gamma^{2,s,j-1}(\mathbf{a}'_s,\mathbf{Z}_{j-1},\mathbf{L}_{j-1})},
\end{aligned}$$

and

$$\begin{aligned}
\frac{Q^L(L_j\mid\mathbf{a}_j,\mathbf{Z}_j,\mathbf{L}_{j-1})}{Q^L(L_j\mid\mathbf{a}'_j,\mathbf{Z}_j,\mathbf{L}_{j-1})}&=\frac{p\left(\mathbf{a}_j,\mathbf{Z}_j,\mathbf{L}_j\right)}{p\left(\mathbf{a}_j,\mathbf{Z}_j,\mathbf{L}_{j-1}\right)}\frac{p\left(\mathbf{a}'_j,\mathbf{Z}_j,\mathbf{L}_{j-1}\right)}{p\left(\mathbf{a}'_j,\mathbf{Z}_j,\mathbf{L}_j\right)}=\frac{p\left(\mathbf{a}_j,\mathbf{Z}_j,\mathbf{L}_j\right)}{p\left(\mathbf{a}'_j,\mathbf{Z}_j,\mathbf{L}_j\right)}\frac{p\left(\mathbf{a}'_j,\mathbf{Z}_j,\mathbf{L}_{j-1}\right)}{p\left(\mathbf{a}_j,\mathbf{Z}_j,\mathbf{L}_{j-1}\right)}\\
&=\prod_{s=1}^{j}\frac{p(a_s\mid\mathbf{a}_{s-1},\mathbf{Z}_j,\mathbf{L}_j)}{p(a'_s\mid\mathbf{a}'_{s-1},\mathbf{Z}_j,\mathbf{L}_j)}\frac{p(\mathbf{Z}_j,\mathbf{L}_j)}{p(\mathbf{Z}_j,\mathbf{L}_j)}\times\prod_{s=1}^{j}\frac{p(a'_s\mid\mathbf{a}'_{s-1},\mathbf{Z}_j,\mathbf{L}_{j-1})}{p(a_s\mid\mathbf{a}_{s-1},\mathbf{Z}_j,\mathbf{L}_{j-1})}\frac{p(\mathbf{Z}_j,\mathbf{L}_{j-1})}{p(\mathbf{Z}_j,\mathbf{L}_{j-1})}\\
&=\prod_{s=1}^{j}\frac{\gamma^{2,s,j}(\mathbf{a}_s,\mathbf{Z}_j,\mathbf{L}_j)}{\gamma^{2,s,j}(\mathbf{a}'_s,\mathbf{Z}_j,\mathbf{L}_j)}\times\prod_{s=1}^{j}\frac{\gamma^{1,s,j}(\mathbf{a}'_s,\mathbf{Z}_j,\mathbf{L}_{j-1})}{\gamma^{1,s,j}(\mathbf{a}_s,\mathbf{Z}_j,\mathbf{L}_{j-1})}.
\end{aligned}$$

The first case of the robustness condition is trivial. In the second case, correct $\bar{Q}_t^{L,\mathbf{a},\mathbf{a}'}$ yields $P_0 D_t^{L,\mathbf{a},\mathbf{a}'}(P) = 0$; correct g and the correct ratios $Q^L(\cdot \mid \mathbf{a},\cdot)/Q^L(\cdot \mid \mathbf{a}',\cdot)$, either in terms of Q^L or in terms of γ, produce a telescopic sum over t of $D_t^{Z,\mathbf{a},\mathbf{a}'}$, the ends of which cancel with $D^{L_0,\mathbf{a},\mathbf{a}'}$. Similarly, in the third case, correct $\bar{Q}_t^{Z,\mathbf{a},\mathbf{a}'}$ yields $P_0 D_t^{Z_t,\mathbf{a},\mathbf{a}'}(P) = 0$; the correct g and correct ratios $Q^Z(\cdot \mid \mathbf{a}',\cdot)/Q^Z(\cdot \mid \mathbf{a},\cdot)$ yield a telescopic sum over t of $D_t^{Z,\mathbf{a},\mathbf{a}'}$ whose ends cancel with $D^{L_0,\mathbf{a},\mathbf{a}'}$.

Appendix A.2

In this appendix, we evaluate the various options to formulate the causal mediation problem in the survival setting with time-dependent mediator, without regarding mediators as intervention variables. The first option is a simple extension of the traditional natural effects definition in the existing literature (e.g. van der Laan and Petersen (2004), VanderWeele (2010), Robins and Richardson (2010), Tchetgen Tchetgen and VanderWeele (2012)), where all the paths from the treatment to the mediators are blocked. We shall see that the resulting ideal experiment is not well-defined for the purpose of mediating the effect on the event process. The second option leaves the paths from treatment to mediator through survival history unblocked. However, the sufficient identifiability conditions, while reasonable in other applications, may be too strong for survival study. As a result, we argue that a SI-based perspective of causal mediation offers an attractive alternative to formulate the effect parameters.

We begin by reviewing the one time-point setting. Under the non-SI approach introduced by Robins and Greenland (1992) and Pearl (2001), one defines a counterfactual event indicator $Y(a,Z(a'))$ according to the following experiment

$$\begin{aligned} W &= f_W(U_W) \\ Z(a') &= f_Z(W, A = a', U_Z), \\ Y(a,Z(a')) &= f_Y(W, A = a, Z(a'), U_Y). \end{aligned} \tag{5.21}$$

$Y(a,Z(a'))$ is the event indicator in an ideal experiment where A is set to a, and the intermediate variable Z takes its value under the influence of $A = a'$. The identifiability conditions (Pearl (2001)) for $P(Y(a,Z(a')) = 0)$ are $Y(a,z) \perp (A,Z) \mid W$, $Y(a,z) \perp Z(a') \mid W$, and $Z(a') \perp A \mid W$.

A2.1: Blocking all paths from treatment to mediators

A direct extension of (5.21) is to conceptualize the mediator process as being defined entirely in a world with $A = a'$. The hypothetical experiment generating the outcome is:

$$\begin{aligned} W &= f_W(U_W) \\ A &= a \\ Z_t(a') &= f_{Z_t}(W, A = a', \mathbf{Z}_{t-1}(a'), Y_{t-1}(a'), U_{Z_t}), \\ Y_t(a') &= f_{Y_t}(W, A = a', \mathbf{Z}_t(a'), Y_{t-1}(a'), U_{Y_t}), \\ Y_t(a, \mathbf{Z}(a')) &= f_{Y_t}(W, A = a, \mathbf{Z}_t(a'), Y_{t-1}(a, \mathbf{Z}(a')), U_{Y_t}). \end{aligned} \tag{5.22}$$

The experiment can be run by either drawing variables subsequently according to the order above, or by first drawing $(\mathbf{Z}(a'), \mathbf{Y}(a'))$, and then draw $\mathbf{Y}(a, \mathbf{z})$ with the given realization of $\mathbf{Z}(a') = \mathbf{z}$. Either way, if one draws $Y_t(a') = 1$, i.e. event happens at time t under treatment $A = a'$, the next mediator $Z_{t+1}(a')$ is assigned the degenerate value at time $t+1$. Then, drawing $Y_{t+1}(a, \mathbf{Z}(a'))$, when the latest mediator has the degenerate value but $Y_t(a, \mathbf{Z}(a')) = 0$, is not defined. One could deterministically set $Y_{t+1}(a, \mathbf{Z}(a')) = 1$ in this case and still obtain a well-defined survival time, but this would allow the effect of treatment $A = a'$ on survival to influence the effect of treatment $A = a$ on survival, which is contrary to the purpose of mediating the effect of $A = a$ on the survival process.

In this light, well-defined mediation formulas and natural effects in the current setting should not block the paths of treatment to mediator through the survival history. In other words, the direct effect questions should be rephrased to "what is the effect of treatment on survival, if treatment had no other effect on the mediators other than through the survival history?".

A2.2: Only blocking those paths from treatment to mediators that are not through survival history

Due to the considerations above, we wish to define mediation effects where the paths from treatment to mediator through the outcome process is left unblocked. These effects of interest are extension of the path-specific effects discussed in Pearl (2001), Avin et al. (2005) and Robins and Richardson (2010). Consider the following hypothetical experiment:

$$\begin{aligned} W &= f_W(U_W) \\ A &= a \\ Z_t(a', Y(a)) &\equiv f_{Z_t}(W, A = a', \mathbf{Z}_{t-1}(a', \mathbf{Y}(a)), \mathbf{Y}_{t-1}(a, Z(a')), U_{Z_t}), \\ Y_t(a, Z(a')) &\equiv f_{Y_t}(W, A = a, \mathbf{Z}_{t-1}(a', \mathbf{Y}(a)), \mathbf{Y}_{t-1}(a, Z(a')), U_{Y_t}). \end{aligned} \tag{5.23}$$

Note the simplified notation for $Z_t(a', Y(a))$ and $Y_t(a, Z(a'))$. For $t = 2$, $Z_2(a', Y(a))$ is in fact $Z_2(a', Y_1(a, Z_1(a')))$, and $Y_2(a, Z(a'))$ is in fact $Y_2(a, (Z_2(a', Y_1(a, Z_1(a'))), Z_1(a')))$. This experiment differs from the (5.22) in that the event process affecting each mediator response is the outcome process of interest. More specifically, under (5.23) the experiment first sets $A = a$; at each

visit, given realization $(W = w, A = a, \mathbf{Z}_{t-1}(a', Y(a)) = \mathbf{z}_{t-1}, Y_{t-1}(a, Z(a')) = y_{t-1})$, it measures the response Z_t would have had if the treatment were $A = a'$ while the rest of the history remained the same; then, with given realization $(W = w, A = a, \mathbf{Z}_t(a', a) = \mathbf{z}_t, Y_{t-1}(a, Z(a')) = y_{t-1})$, it measures the event indicator Y_t.

The difference between the experiment in (5.23) and the SI-based experiment in the main text lies in that under the SI formulation, one only needs to identify the distribution $P_{\mathbf{Z}(\mathbf{a}')}$, which remains in the real of experiment with intervention $\mathbf{A}'_{\mathbf{a}}$. Under (5.23), the conditional probability $P(Z_t(a', Y(a)) \mid W, \mathbf{Z}_{t-1}(a', Y(a)), Y_{t-1}(a, Z(a')))$ is part of a cross world process. Therefore, even though the SI-based parameter would identify to the same statistical parameter (5.2), they are differently formulated causal parameters.

In order to identify $E(Y_\tau(a, Z(a')))$ defined in (5.23), one would require that the death at a given time t be independent of future potential mediators — this condition is too strong for the purpose of effect mediation in a survival study.

5.8 Bibliography

A. C. Atkinson and A. Biswas. Adaptive biased-coin designs for skewing the allocation proportion in clinical trials with normal responses. *Stat. Med.*, 24(16):2477–2492, 2005.

C. Avin, I. Shpitser, and J. Pearl. Identifiability of path-specific effects. 2005.

U. Bandyopadhyay and A. Biswas. Adaptive designs for normal responses with prognostic factors. *Biometrika*, 88(2):409–419, 2001.

H. Bang and J. M. Robins. Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61:962–972, 2005.

R.M. Baron and D.A. Kenny. The moderator-mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *Journal of Penalty and Social Psychology*, 51(6):1173–1182, 1986.

P. J. Bickel, C. A. J. Klaassen, Y. Ritov, and J. A. Wellner. *Efficient and adaptive estimation for semiparametric models*. Springer-Verlag, New York, 1998. Reprint of the 1993 original.

P.J. Bickel, C.A.J. Klaassen, Y. Ritov, and J. Wellner. *Efficient and Adaptive Estimation for Semiparametric Models*. Springer-Verlag, 1997.

A. Biswas, R. Bhattacharya, and E. Park. On a class of optimal covariate-adjusted response-adaptive designs for survival outcomes. *Statistical methods in medical research*, 2014.

A. Chambaz and M. J. van der Laan. Targeting the optimal design in randomized clinical trials with binary outcomes and no covariate: Simulation study. *Int. J. Biostat.*, 7(1), 2011a.

A. Chambaz and M. J. van der Laan. Inference in targeted goup sequential covariate-adjusted randomized clinical trials. *Scandinavian Journal of Statistics*, 41(1):104–140, 2013.

Antoine Chambaz and Mark J van der Laan. Estimation and testing in targeted group sequential covariate-adjusted randomized clinical trials. 2011b.

Y. I. Chang and E. Park. Sequential estimation for covariate-adjusted response-adaptive designs. *J. Korean Statistical Society*, 42(1):105–116, 2013.

A.P. Dawid and V. Didelez. Identifying the consequences of dynamic treatment strategies: A decision-theoretic overview. *Statistics Surveys*, 4:184–231, 2010.

I. Diaz and M. van der Laan. Population intervention causal effects based on stochastic interventions. *Biometrics*, 2011.

V. Didelez, A.P. Dawid, and S. Geneletti. Direct and indirect effects of sequential treatments. In *Proceedings of the 22nd Annual Conference on Uncertainty in Artifical Intelligence*, pages 138–146, 2006.

J. Friedman, T. Hastie, and R. Tibshirani. Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 2010.

S. Gruber and M.J. van der Laan. A targeted maximum likelihood estimator of a causal effect on a bounded continuous outcome. *International Journal of Biostatistics*, 6, 2010.

S. Gruber and M.J. van der Laan. Bounded continuous outcomes. In M.J. van der Laan and S. Rose, editors, *Targeted Learning: Causal Inference for Observational and Experimental Data*. Springer, 2011.

D.M. Hafeman and T.J. VanderWeele. Alternative assumptions for the identification of direct and indirect effects. *Epidemiology*, 2010.

M.A. Hernan, B. Brumback, and J.M. Robins. Marginal structural models to estimate the causal effect of zidovudine on the survival of HIV-positive men. *Epidemiology*, 11(5):561–570, 2000.

P. Holland. Statistics and causal inference. *Journal of the American Statistical Association*, 81: 945–960, 1986.

F. Hu and W. F. Rosenberger. *The theory of response-adaptive randomization in clinical trials*, volume 525. John Wiley & Sons, 2006.

T. Huang, Z. Liu, and F. Hu. Longitudinal covariate-adjusted response-adaptive randomized designs. *J. Statistical Planning and Inference*, 143(10):1816–1827, 2013.

K. Imai, L. Keele, and T. Yamamoto. Identification, inference and sensitivity analysis for causal mediation effects. *Statistical Science*, 25(1):51–71, 2010.

C. Jennison and B. W. Turnbull. *Group Sequential Methods with Applications to Clinical Trials*. Chapman & Hall/CRC, Boca Raton, FL, 2000.

B. Jo, E.A. Stuart, D.P. MacKinnon, and A.D. Vinokur. The use of propensity scores in mediation analysis. *Multivariate Behavioral Research*, 46(3):425–452, 2011.

J. Kang and J. Schafer. Demystifying double robustness: A comparison of alternative strategies for estimating a population mean from incomplete data (with discussion). *Statistical Science*, 22: 523–39, 2007.

J.S. Kaufman, R.F. Maclehose, and S. Kaufman. A further critique of the analytic strategy of adjusting for covariates to identify biologic mediation. *Epidemiologic Perspectives & Innovations*, page 1:4, 2004.

M. Kosorok. *Introduction to Empirical Processes and Semiparametric Inference*. Springer-Verlag, 2008.

T. Lange and J.V. Hansen. Direct and indirect effects in a survival context. *Epidemiology*, 22(4): 575, 2011.

R. Neugebauer and M. van der Laan. Nonparametric causal effects based on marginal structural models. *Journal of Statistical Planning and Inference*, 137(2):419–434, 2007.

J. Pearl. Direct and indirect effects. In *Proceedings of the seventeenth conference on uncertainty in artificial intelligence*, pages 411–420. Citeseer, 2001.

J. Pearl. *Causality: Models, Reasoning and Inference*. Cambridge University Press, New York, 2nd edition, 2009.

J. Pearl. The mediation formula: A guide to the assessment of causal pathways in nonlinear models. In C. Berzuini, P. Dawid, and L. Bernardinelli, editors, *Causality: Statistical Perspectives and Applications*. 2011.

Judea Pearl. Causal diagrams for empirical research. *Biometrika*, 82(4):669–688, 1995.

M. Petersen, K. Porter, S.Gruber, Y. Wang, and M.J. van der Laan. Diagnosing and responding to violations in the positivity assumption. Technical report 269, Division of Biostatistics, University of California, Berkeley, 2010. URL `http://www.bepress.com/ucbbiostat/paper269`.

M. Petersen, J. Schwab, S. Gruber, N. Blaser, M. Schomaker, and M.J. van der Laan. Targeted maximum likelihood estimation for dynamic and static longitudinal marginal structural working models. *J Causal Inference*, 2014.

M.L. Petersen, S.E. Sinisi, and M.J. van der Laan. Estimation of direct causal effects. *Epidemiology*, 17(3):276–284, 2006.

David Pollard. *Convergence of stochastic processes*. David Pollard, 1984.

R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2014.

J. M. Robins. Marginal structural models. *Proceedings of the American Statistical Association. Section on Bayesian Statistical Science*, pages 1–10, 1997a.

J. M. Robins. Commentary on "using inverse weighting and predictive inference to estimate the effects of time-varying treatments on the discrete-time hazard. *Stat Med*, 210:1663–80, 2002.

J. M. Robins and T. S. Richardson. Alternative graphical causal models and the identification of direct effects. In P. Shrout, editor, *In Causality and Psychopathology: Finding the Determinants of Disorders and Their Cures*. Oxford University Press., 2010.

J. M. Robins, M. A. Hernan, and B. Brumback. Marginal structural models and causal inference in epidemiology. *Epidemiology*, 11(5):550–560, 2000a.

J. M. Robins, L. Orellana, and A. Rotnitzky. Estimaton and extrapolation of optimal treatment and testing strategies. *Statistics in Medicine*, 27(23):4678–4721, 2008.

J.M. Robins. A new approach to causal inference in mortality studies with sustained exposure periods - application to control of the healthy worker survivor effect. *Mathematical Modelling*, 7:1393–1512, 1986.

J.M. Robins. Causal inference from complex longitudinal data. In Editor M. Berkane, editor, *Latent Variable Modeling and Applications to Causality*, pages 69–117. Springer Verlag, New York, 1997b.

J.M. Robins. Marginal structural models versus structural nested models as tools for causal inference. In *Statistical models in epidemiology: the environment and clinical trials*, pages 95–134. Springer-Verlag, 1999.

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

J.M. Robins. Semantics of causal dag models and the identification of direct and indirect effects. In N. Hjort P. Green and S. Richardson, editors, *Highly Structured Stochastic Systems*, pages 70–81. Oxford University Press, Oxford, 2003.

J.M. Robins and S. Greenland. Identifiability and exchangeability for direct and indirect effects. *Epidemiology*, 3(0):143–155, 1992.

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

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

J.M. Robins, A. Rotnitzky, and M.J. van der Laan. Comment on "on profile Likelihood" by S.A. Murphy and A.W. van der Vaart. *Journal of the American Statistical Association – Theory and Methods*, 450:431–435, 2000b.

P.R. Rosenbaum and Donald B. Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70:41–55, 1983.

W. F. Rosenberger. New directions in adaptive designs. *Statist. Sci.*, 11:137–149, 1996.

W. F. Rosenberger, A. N. Vidyashankar, and D. K. Agarwal. Covariate-adjusted response-adaptive designs for binary response. *Journal of biopharmaceutical statistics*, 11(4):227–236, 2001.

W. F. Rosenberger, O. Sverdlov, and F. Hu. Adaptive randomization for clinical trials. *J Biopharm Stat*, 22(4):719–36, 2012.

D.B. Rubin. Bayesian inference for causal effects: the role of randomization. *Annals of Statistics*, 6:34–58, 1978.

Mireille E Schnitzer, Erica EM Moodie, Mark J van der Laan, Robert W Platt, and Marina B Klein. Modeling the impact of hepatitis c viral clearance on end-stage liver disease in an hiv co-infected cohort with targeted maximum likelihood estimation. *Biometrics*, 70(1):144–152, 2014.

J. Shao and X. Yu. Validity of tests under covariate-adaptive biased coin randomization and generalized linear models. *Biometrics*, 69(4):960–969, 2013.

J. Shao, X. Yu, and B. Zhong. A theory for testing hypotheses under covariate-adaptive randomization. *Biometrika*, 97(2):347–360, 2010.

Ilya Shpitser. Counterfactual graphical models for longitudinal mediation analysis with unobserved confounding. *Cognitive science*, 37(6):1011–1035, 2013.

O.M. Stitelman, V. De Gruttola, C.W. Wester, and M.J. van der Laan. Rcts with time-to-event outcomes and effect modification parameters. In M. J. van der Laan and S. Rose, editors, *Targeted Learning: Causal Inference for Observational and Experimental Data*. Springer, 2011.

O. Sverdlov, W. F. Rosenberger, and Y. Ryeznik. Utility of covariate-adjusted response-adaptive randomization in survival trials. *Statistics in Biopharmaceutical Research*, 5(1):38–53, 2013.

Sarah L Taubman, J.M. Robins, Murray A. M. Mittleman, and M.A. Hernan. Intervening on risk factors for coronary heart disease: an application of the parametric g-formula. *Int. J. Epidemiology*, 38(6):1599–1611, 2009.

E.J. Tchetgen Tchetgen. On causal mediation analysis with a survival outcome. *The International Journal of Biostatistics*, 7(1):1–38, 2011.

E.J. Tchetgen Tchetgen and I. Shpitser. Semiparametric theory for causal mediation analysis: efficiency bounds, multiple robustness, and sensitivity analysis. Technical report 130, Biostatistics, Harvard University, June 2011a. URL `http://www.bepress.com/harvardbiostat/paper130`.

E.J. Tchetgen Tchetgen and I. Shpitser. Semiparametric estimation of models for natural direct and indirect effects. Technical report 129, Biostatistics, Harvard University, June 2011b. URL `http://www.bepress.com/harvardbiostat/paper129`.

Eric J. Tchetgen Tchetgen and Tyler J. VanderWeele. On identification of natural direct effects when a confounder of the mediator is directly affected by exposure. Harvard University Biostatistics Working Paper Series, http://biostats.bepress.com/harvardbiostat/paper148, 2012.

JY Tein and DP MacKinnon. Estimating mediated effects with survival data. *New Developments on Psychometrics. Tokyo, Japan: Springer-Verlag Tokyo Inc*, pages 405–12, 2003.

J. Tian. Identifying dynamic sequential plans. In *Proceedings of the Twenty-Fourth Annual Conference on Uncertainty in Artificial Intelligence (UAI-08)*, pages 554–561, 2008.

R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.

A.A. Tsiatis. *Semiparametric Theory and Missing Data*. Springer, New York, 2006.

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

M. J. van der Laan and M. Petersen. Direct effect models. *The International Journal of Biostatistics*, 4(1), 2008.

M. J. van der Laan and J. M. Robins. *Unified methods for censored longitudinal data and causality*. Springer Series in Statistics. Springer-Verlag, New York, 2003a.

M. J. van der Laan and D. Rubin. Targeted maximum likelihood learning. *Int. J. Biostat.*, 2(1), 2006.

Mark J van der Laan and Susan Gruber. Targeted minimum loss based estimation of causal effects of multiple time point interventions. *The international journal of biostatistics*, 8(1), 2012.

M.J. van der Laan and S. Gruber. Collaborative double robust penalized targeted maximum likelihood estimation. *The International Journal of Biostatistics*, 6(1), 2010.

M.J. van der Laan and M. Petersen. Estimation of direct and indirect causal effects in longitudinal studies. Technical report 155, Division of Biostatistics, University of California, Berkeley, August 2004.

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

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

M.J. van der Laan and S. Rose. *Targeted Learning: Causal Inference for Observational and Experimental Data*. Springer Series in Statistics. Springer, first edition, 2011.

M.J. van der Laan, S. Dudoit, and S. Keleş. Asymptotic optimality of likelihood-based cross-validation. *Statistical Applications in Genetics and Molecular Biology*, 3, 2004.

M.J. van der Laan, M.L. Petersen, and M.M. Joffe. History-adjusted marginal structural models and statically-optimal dynamic treatment regimens. *The International Journal of Biostatistics*, 1(1):10–20, 2005.

M.J. van der Laan, E.C. Polley, and A.E. Hubbard. Super learner. *Statistical Applications in Genetics and Molecular Biology*, 6(25), 2007. ISSN 1.

A. W. van der Vaart. *Asymptotic Statistics*. Cambridge University Press, 1998a.

A. W. van der Vaart. *Asymptotic statistics*, volume 3 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge, 1998b.

Aad W Van Der Vaart and Jon A Wellner. *Weak Convergence*. Springer, 1996.

Aad W van der Vaart and Jon A Wellner. Empirical processes indexed by estimated functions. *Lecture Notes-Monograph Series*, pages 234–252, 2007.

A.W. van der Vaart and J.A. Wellner. *Weak Convergence and Empirical Processes*. Springer-Verlag, New York, 1996.

A.W. van der Vaart, S. Dudoit, and M.J. van der Laan. Oracle inequalities for multi-fold cross-validation. *Statistics and Decisions*, 24(3):351–371, 2006.

R. van Handel. On the minimal penalty for markov order estimation. *Probability Theory and Related Fields*, 150:709–738, 2011.

T.J. VanderWeele. Marginal structural models for the estimation of direct and indirect effects. *Epidemiology*, 20:18–26, 2009.

T.J. VanderWeele. Direct and indirect effects for neighborhood-based clustered and longitudinal data. *Sociological Methods & Research*, 38(4):515–544, 2010.

T.J. VanderWeele. Causal mediation analysis with survival data. *Epidemiology*, 22(4):582, 2011.

T.J. VanderWeele and S. Vansteelandt. Odds ratios for mediation analysis for a dichotomous outcome. *Am. J. of Epidemiology*, 172:1339–1348, 2010.

S. Vansteelandt. Estimating direct effects in cohort and case control studies. *Epidemiology*, 20: 851–860, 2009.

Jessica G Young, Lauren E Cain, James M Robins, Eilis J OReilly, and Miguel A Hernán. Comparative effectiveness of dynamic treatment regimes: an application of the parametric g-formula. *Statistics in biosciences*, 3(1):119–143, 2011.

L-X. Zhang and F-F. Hu. A new family of covariate-adjusted response-adaptive designs and their properties. *Appl. Math. J. Chinese Univ. Ser. B*, 24(1):1–13, 2009.

L-X. Zhang, F. Hu, S. H. Cheung, and W. S. Chan. Asymptotic properties of covariate-adjusted response-adaptive designs. *Ann. Statist.*, 35(3):1166–1182, 2007.