University of California, Berkeley
U.C. Berkeley Division of Biostatistics Working Paper Series

Year 2016 *Paper* 348

Marginal Structural Models with Counterfactual Effect Modifiers

Wenjing Zheng* Zhehui Luo†
Mark J. van der Laan‡

*University of California, Berkeley, Division of Biostatistics, wenjing.zheng@berkeley.edu
†Michigan State University
‡University of California, Berkeley, University of California, Berkeley, Division of Biostatistics, laan@berkeley.edu

Marginal Structural Models with Counterfactual Effect Modifiers

Wenjing Zheng, Zhehui Luo, and Mark J. van der Laan

Abstract

In health and social sciences, research questions often involve systematic assessment of the modification of treatment causal effect by patient characteristics, in longitudinal settings with time-varying or post-intervention effect modifiers of interest. In this work, we investigate the robust and efficient estimation of the so-called Counterfactual-History-Adjusted Marginal Structural Model (van der Laan and Petersen (2007)), which models the conditional intervention-specific mean outcome given modifier history in an ideal experiment where, possible contrary to fact, the subject was assigned the intervention of interest, including the treatment sequence in the conditioning history. We establish the semiparametric efficiency theory for these models, 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)). To facilitate implementation in applications where the effect modifier is high dimensional, 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)). The comparative performance of these estimators are assessed in a simulation study. The use of the TMLE estimator (based on the projected influence curve) is illustrated in a secondary data analysis for the Sequenced Treatment Alternatives to Relieve Depression (STAR*D) trial.

1 Introduction

In social and health sciences, research questions often involve systematic comparison of the effectiveness of different longitudinal exposures or treatment strategies on an outcome of interest. Specifically, consider a study where subjects are followed over time. In addition to their baseline covariates, we record their time-varying treatment of interest, time-varying covariates, and the outcomes of interest. Time-varying confounding is ubiquituous in these situations: the treatment of interest depends on past covariates and in turn affects future covariates; right censoring is often present, in response to past covariates and treatment. It has been widely recognized that in these cases, conventional analytic methods, such as multiple regression, often fail to properly account for the time-varying confounding of the treatment effect (e.g. Robins, Hernan, and Brumback (2000)). Marginal Structural Models (MSMs), introduced by Robins (1997a), are well-established and widely used tools to address this problem of time-varying confounding; these models specify the marginal expectation of an *intervention-specific counterfactual outcome* (i.e. the mean outcome of a subject in an ideal experiment where he/she was assigned to followed a given intervention).

To assess effect modification, MSMs are traditionally used to model the conditional counterfactual mean outcome given an observed history. Yet, in many settings one may wish to model the conditional counterfactual mean outcome given a counterfactual history. Consider the simple case of effect modification by baseline covariates. A traditional *observed baseline MSM* conditions on the observed baseline modifiers (Robins et al. (2000)) and allows one to assess how the treatment effect changes as a function of the observed covariate values. *(Observed) History-Adjusted MSMs* (HA-MSMs), introduced in van der Laan, Petersen, and Joffe (2005) and Petersen, Deeks, Martin, and van der Laan (2007a), generalize the *observed baseline MSMs* by modeling the counterfactual mean outcome given the observed history of treatment and modifiers of interest up till a time point. However, since the modifiers of interest may be affected by their preceding treatment assignments, which may in turn depend on past modifiers and other covariates, the counterfactual mean within each strata of this history will also be affected by the *observed treatment mechanism* (i.e. the way treatment assignments based on past covariates were made in the observed data. e.g. randomized assignment vs assignment based on specific determinants of outcome). In this case, the parameters of the HA-MSM would not be generalizable to an equivalent population with different treatment mechanism (Petersen, Deeks, and van der Laan (2007b)). Instead, the the true outcome that one wishes to model is in fact the the conditional mean outcome given modifier history in an ideal experiment where subject were assigned a given intervention on interest, including the treatment sequence in the conditioning

history (i.e. the conditional counterfactual mean outcome given a *counterfactual history* of the time-varying modifiers of interest up to a given time point).

To model these conditional counterfactual mean outcome given a *counterfactual history*, we can use the so-called Counterfactual-History-Adjusted MSMs (CHA-MSM), introduced by van der Laan and Petersen (2007). Inverse Probability of Treatment Weighted (IPTW) estimators for CHA-MSM were proposed in van der Laan and Petersen (2007). These estimators are very intuitive, can be easily implemented using standard software, and offers influence-function based standard error estimates. However, their consistency rely on consistent estimation of all the treatment weights. Doubly Robust IPTW (DR-IPTW) estimators for CHA-MSM were also described in van der Laan and Petersen (2007); they were based estimating equations derived by orthogonalizing the IPTW estimating function with respect to the treatment mechanism. Contrary to IPTW, these DR-IPTW estimators are consistent if the treatment weights or the conditional covariate and outcome densities are consistently estimated. Moreover, they are solutions to the estimating equation defined by the efficient influence function, and thus are asymptotically semi-parametric efficient.

Despite these advances, there are still many gaps in the efficiency theory and robust estimation of CHA-MSMs. Firstly, even with the efficient influence function being a key actor in semi-parametric estimation, there still lacks an explicit representation of it as an orthogonal decomposition of the nuisance parameters corresponding to the time-varying covariates. Compared to the IPTW-orthogonalized representation in van der Laan and Petersen (2007), such an explicit representation would provide a comprehensive picture of the efficiency theory for CHA-MSM. In particular, by shining a light on the role of the nuisance parameters in the efficient influence function, such an explicit representation can inform the study of semi-parametric estimation for these models, advise on the trade-offs in estimating different nuisances parameters, and provide insights on the challenges and solutions to handling high-dimensional covariates. Secondly, as estimating equation based estimators, both the IPTW and the DR-IPTW may be unstable in the presence of near positivity violations (Petersen, Porter, S.Gruber, Wang, and van der Laan (2010)), resulting in biased point and standard error estimates in these settings. In applications with dynamic treatment regimes, this instability is especially difficult to circumvent due to the limitations of effective weight stabilization. By contrast, a substitution-based estimator for these models can provide a way to maximize finite sampler performance by preserving global information about the model and the parameters.

This paper aims to fill these gaps in the literature by establishing the efficiency theory for CHA-MSM and providing semi-parametric, substitution-based, efficient and robust estimators. Firstly, we describe the identification of the con-

ditional counterfactual intervention-specific mean outcome given a counterfactual history up to a given time point, and the identification of the corresponding MSM parameters of interest. Secondly, we determine the efficient influence function for these statistical parameters as an orthogonal decomposition of the nuisance parameters. This efficient influence function is used to construct a substitution-based, semi-parametric 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 function, the computations in 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 function (and the corresponding TMLE estimator), which retains most of the robustness of its efficient peer and can be used in applications where the use of the efficient influence function becomes taxing. In addition to these two robust estimators, we also present a non-targeted substitution estimator (Robins (1986)) which is also applicable in high-dimensional data.

1.1 Illustrative example

Throughout this paper, we will illustrate our presentation using an example from mental health research. 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/`). 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 to an option within one of his chosen treatment strategies. At level 2, the patient can choose between augmenting the citalopram with multiple options or switching to a new regimen with multiple options, or both. We are interested in the comparative effect of augmenting vs switching medication; because these two strategies are not randomized, this analysis is analogous to an observational study. Suppose we wish to assess the effect modification of 2nd level's treatment strategy (augmenting vs switching) by the depression symptoms measured prior to entering level 2 (either measured at enrollment or level 1 exit, for our purpose, they are both considered baseline modifiers). 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 treatment may be less inclined to follow up with the surveys and visits or to report these symptom measures. In this case, a simple complete-case analysis of effect modification may only provide results applicable to patients with less

severe symptoms or more satisfied with their level 1 treatment. If we assume that this missingness, while associated with outcome, can be predicted using covariate history collected up to level 1 exit, then we can regain some of this generalizability. Indeed, we cast these symptom measures as counterfactual variables under an intervention on their missingness indicator. This way, the target parameter is formulated in terms of an ideal experiment where the report of the symptom measures was always ensured.

1.2 Organization of this article

This paper is organized as follows. In section 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, we present the efficient influence function for the parameter of interest under a saturated semiparametric model, as well as a projected influence function. The robustness conditions for the efficient and the projected influence functions are also established. In section 4, we present the construction of two TMLE estimators, one using the projected influence function (we call it the *projected TMLE*) and one using the efficient influence function (we call it the *full TMLE*), and a non-targeted substitution estimator. In section 5, a simulation study demonstrates robustness properties of the *projected* TMLE. In section 6, we use STAR*D to illustrate the application of *projected* TMLE. A final discussion concludes this article.

2 Data Structure and Parameters of Interest

For the simplicity of exposition, we consider the case where one wish to assess the conditional intervention-specific mean outcome given the counterfactual history *up to the first time point*.

Specifically, consider a longitudinal data structure

$$O = (W, A_0, V, L_0, A_1, L_1, \ldots, A_K, L_K) \sim P_0,$$

where W encodes baseline covariates, A_t is the variable measured at time t that encodes the exposures of interest and censoring indicators, V encodes the time-varying history (between treatment A_0 and A_1) that one wishes to condition on, 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 in (0,1)).

We illustrate this notation with the data example introduced in section 1.1. Our goal in this example is to assess the effect modification of 2nd level's treatment strategy (augmenting vs switching, measured at start of level 2) by the depression symptoms measured prior to entering level 2; yet many of these baseline effect modifiers were missing at random. 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 23 weeks. For our study goal, the variables that we would enforce/randomize in an ideal experiment are the missingness status of the baseline effect modifier, the treatment assignment at start of level 2, and the censoring status at each week of the level. To this end, let V be the baseline effect modifier of interest. Let A_0 be the indicator for measurement of the baseline effect modifier of interest, $A_1 = A_1^{rx}$ be the treatment strategy received by a patient at level 2 ($A_1 = 0$ for augmenting medication and $A_1 = 1$ for switching medication). We use W to encode the baseline covariates that may affect the measurement status A_0, and L_0 to encode other baseline covariates (beside the modifier of interest) measured prior to treatment assignment A_1. Hence, under this indexing, $K = 24$ and the first week of level 2 starts at $t = 2$. For $t \geq 2$, $A_t = A_t^C$ is a counting process which drops to 0 if patient was censored by time t, and L_t are time-varying covariates such as visit statistics (duration in level thus far, visit frequency, etc), side-effect burden and symptom measures at time $t \geq 2$. The variable L_t also contains two counting processes: the outcome process Y_t, which is a binary indicator for entering remission by time t, and an exit 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 — the remission status by end of 23 weeks.

For later convenience, we introduce here some useful vector notations. For an time-varying variable X_s, let $\mathbf{X}_s \equiv (X_1, \ldots, X_s)$ and $\underline{\mathbf{X}}_s \equiv (X_s, \ldots, X_K)$; we will also use the shorthand $\mathbf{X} = \mathbf{X}_K$. The equivalent notations in lower case will apply to the realizations of the corresponding variable. Finally, variables with degenerate indices, such as $t = -1$, are empty sets.

2.1 Causal Model and Parameter of Interest

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

$$W = f_W(U_W); \quad A_0 = f_{A_0}(W, U_{A_0}); \quad V = f_V(W, A_0, U_V); \quad L_0 = f_{L_0}(W, A_0, V, U_{L_0});$$
$$A_t = f_{A_t}(W, A_0, V, \mathbf{A}_{t-1}, \mathbf{L}_{t-1}, U_{A_t}); \quad L_t = f_{L_t}(W, A_0, V, \mathbf{A}_t, \mathbf{L}_{t-1}, U_{L_t}), \tag{1}$$

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 and denote this set as $Pa(X)$. This causal model defines a random variable with distribution $P_{O,U}$ on a unit.

To assess the effect of the interventions, we can study the *intervention-specific counterfactual mean outcomes* as a function of the interventions. These counterfactual mean outcomes can be obtained from an ideal experiment where one intervenes to assign $\mathbf{A} = \mathbf{a}$, and measures the resulting covariates (including effect modifiers are interest) and the outcome of interest. In our example, an ideal experiment would set $A_0 = a_0'$ (always measure baseline modifier), $A_1 = a_1$ (switching or augmenting), and $A_{t\geq 2} = 1$ (always prevent drop outs). This data-generating process can be formalized using (1) by setting $A_t = a_t$ in the equations for A_t and in the parents of the non-intervention variables V, L_t and Y. We denote these resulting, possibly *counterfactual*, covariates as $V(a_0'), L_t(\mathbf{a}_t)$ and $Y(\mathbf{a})$.

Now, we wish to assess effect modification of such treatment effect by the changing values of V. That is, we want ask 'how does the differential effect of $(a_0, \underline{\mathbf{a}}_1')$ vs $(a_0, \underline{\mathbf{a}}_1)$, or of $(a_0', \underline{\mathbf{a}}_1')$ vs $(a_0, \underline{\mathbf{a}}_1')$, differ as a function of V, which is affected by the treatment assignment A_0?'. To answer this question, our parameters of interests are the so-called *Counterfactual-History-Adjusted* mean outcomes (van der Laan and Petersen (2007))

$$\rho_{a_0', \underline{\mathbf{a}}_1, v}(P_{O,U}) \equiv E\left(Y(a_0', \underline{\mathbf{a}}_1) \mid V(a_0') = v\right), \tag{2}$$

and our goal is to study how these mean outcomes change as a function of a_0', $\underline{\mathbf{a}}_1$ and v. As discussed in section 1, (2) differs from the traditional History-Adjusted mean outcomes (van der Laan et al. (2005)) $E\left(Y(A_0, \underline{\mathbf{a}}_1) \mid V = v, A_0 = a_0'\right)$ in that these traditional parameters condition on the observed treatment values, hence within each such strata, the conditional mean outcome may still depend on the observed treatment mechanism (thus affected by potential selection bias). In our example, the conditional mean outcomes we wish to study are $E(Y(1, \underline{\mathbf{a}}_1) \mid V(1) = v)$, with $\{\underline{\mathbf{a}}_1 = (a_1, a_2 = \cdots = a_K = 1) : a_1 = 0, 1\}$.

A challenge to study of (2) is that the curse of dimensionality of arises in applications with more than two time points and/or with categorical or continuous V. To address this issue, it is useful to summarize (2) using a working marginal structural model, $m_\psi(a_0', \underline{\mathbf{a}}_1, v)$, which is parametrized by a finite dimensional $\psi \in S \subset \mathbb{R}^d$. We refer to this model as the *Counterfactual-History-Adjusted Marginal Structural Model* (CHA-MSM). Since our Y is binary (or bounded within the unit interval), for concreteness sake, we consider a logistic MSM

$$m_\psi(a_0', \underline{\mathbf{a}}_1, v) = expit\left(\psi \cdot \phi(a_0', \underline{\mathbf{a}}_1, v)\right), \tag{3}$$

where $\phi(a_0', \underline{\mathbf{a}}_1, v)$ is the vector of linear predictors in the generalized linear model. We emphasize that the methods presented here are easily modified to other forms of MSM.

Given a CHA-MSM (3) for the conditional mean outcomes (2), the true MSM parameter ψ can be interpreted as the best summary measure of the conditional mean outcomes $\rho_{a_0', \underline{\mathbf{a}}_1, v}(P_{O,U})$ as a function of a_0, $\underline{\mathbf{a}}_1'$ and v. Formally, let $\mathscr{A}$ and $\mathscr{A}'$ denote the set of interventions of interest and $\mathscr{V}$ denote the set of modifier values of interest. For a given kernel weight function $h(a_0', \underline{\mathbf{a}}_1, v)$, the true MSM parameter ψ in (3) is defined as

$$\Psi(P_{O,U}) = \arg\min_{\psi \in S} \left\{ - \sum_{(a_0', \underline{\mathbf{a}}_1, v) \in \mathscr{A}' \times \mathscr{A} \times \mathscr{V}} P(V(a_0') = v) h(a_0', \underline{\mathbf{a}}_1, v) \right.$$
$$\left. \times \left\{ \rho_{a_0', \underline{\mathbf{a}}_1, v}(P_{O,U}) \log m_\psi(a_0', \underline{\mathbf{a}}_1, v) + \left(1 - \rho_{a_0', \underline{\mathbf{a}}_1, v}(P_{O,U})\right) \log\left(1 - m_\psi(a_0', \underline{\mathbf{a}}_1, v)\right) \right\} \right\} \quad (4)$$

In words, $\Psi(P_{O,U})$ yields the best weighted approximation of the counterfactual conditional dose-response curve $\rho_{a_0', \underline{\mathbf{a}}_1, v}(P_{O,U})$, according to the quasi-loglikelihood loss, kernel weights and working model $m_\psi(a_0', \underline{\mathbf{a}}_1, v)$.

The rest of this paper is devoted to the identification and inference of (4) from the observed data.

2.2 Statistical Estimand

To identify (4) from the data generating distribution P_0, we make a Positivity Assumption (PA) and the Sequential Randomization Assumption (SRA, derived by Robins (1997b)). Specifically, under the PA, there exists $\alpha_t > 0$ such that $\alpha_t \leq P_0(A_t = a_t \mid Pa(A_t))$, for all t and $\mathbf{a} \in \mathscr{A}$, almost everywhere. The SRA assumes that $A_t \perp (W, V(a_0), \{L_t(\mathbf{a}) : t\})$, given parents of A_t. Under these conditions, the joint distribution $(W, V(a_0), \{L_t(\mathbf{a}) : t\})$ 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 $P(V(a_0') = v)$ as

$$\gamma_{a_0', v}(P_0) \equiv E_{P_0}\left\{P_0(V = v \mid A_0 = a_0', W)\right\}, \quad (5)$$

and the counterfactual mean outcome $\rho_{a_0', \underline{\mathbf{a}}_1, v}(P_{O,U})$ as

$$\rho_{a_0', \underline{\mathbf{a}}_1, v}(P_0) \equiv E_{P_0}\left\{ \frac{P_0(V = v \mid A_0 = a_0', W)}{E_{P_0}\left\{P_0(V = v \mid A_0 = a_0', W)\right\}} \times Q_{t=0}^{a_0', \underline{\mathbf{a}}_1}(P_0)(V = v, W) \right\}, \quad (6)$$

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

$$Q_t^{a_0', \underline{\mathbf{a}}_1}(P_0)(\mathbf{L}_{t-1}, V, W) \equiv \sum_{\ell_{t,K}} y_K \left(\prod_{j=t}^{K} P_0(l_j \mid \mathbf{A}_j = \mathbf{a}_j, \mathbf{L}_{t-1}, \ell_{t,j-1}, V, A_0 = a_0', W) \right). \quad (7)$$

It is also useful to rewrite (6) as

$$\rho_{a_0', \underline{\mathbf{a}}_1, v}(P_0) = E_{P_0} \left\{ \frac{1/P_0(A_0 = a_0' \mid W)}{E_{P_0}\left(1/P_0(A_0 = a_0' \mid W) \mid V = v, A_0 = a_0'\right)} Q_{t=1}^{a_0', \underline{\mathbf{a}}_1}(P_0)(L_0, V = v, W) \,\middle|\, V = v, A_0 = a_0' \right\}. \quad (8)$$

The weight $\frac{1/P_0(A_0 = a_0' \mid W)}{E_{P_0}\left(1/P_0(A_0 = a_0' \mid W) \mid V = v, A_0 = a_0'\right)}$ adjusts for potential selection bias introduced by treatment assignment$A_0 = a_0'$. Indeed, if A_0 does not depend on W (in our STAR*D example, this means missingness is completely at random), then this weight equals 1, in which case (8) and (6) are equivalent to the estimands in an analysis which simply stratifies by A_0.

Combining (5) and (6), the causal MSM parameter $\Psi(P_{O,U})$ in (4) identifies to

$$\psi_0 \equiv \Psi(P_0) \equiv \arg\min_{\psi \in S} \left\{ - \sum_{(a_0', \underline{\mathbf{a}}_1, v) \in \mathscr{A}' \times \mathscr{A} \times \mathscr{V}} \gamma_{a_0', v}(P_0) h(a_0', \underline{\mathbf{a}}_1, v) \times \left\{ \rho_{a_0', \underline{\mathbf{a}}_1, v}(P_0) \log m_\psi(a_0', \underline{\mathbf{a}}_1, v) + \left(1 - \rho_{a_0', \underline{\mathbf{a}}_1, v}(P_0)\right) \log\left(1 - m_\psi(a_0', \underline{\mathbf{a}}_1, v)\right) \right\} \right\}. \quad (9)$$

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

$$m_\psi(a_0', \underline{\mathbf{a}}_1, v) = expit\left(\psi \cdot \phi(a_0', \underline{\mathbf{a}}_1, v)\right), \quad (10)$$

where $\phi(a_0', \underline{\mathbf{a}}_1, v)$ is the vector of linear predictors in the generalized linear model. 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 $\Psi(P_0)$.

2.3 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 (9) 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 $P_n f \equiv \frac{1}{n}\sum_{i=1}^{n} f$, and for a distribution P, we will write $Pf = E_P f(O)$.

We generalize the definitions in (7) to any $P \in \mathscr{M}$, for $t \leq K$. At $t = K+1$, we write $Q_{K+1}^{a_0',\underline{\mathbf{a}}_1}(P)(O) \equiv Y_K$. Bang and Robins (2005) noted the following recursive property when dealing with longitudinal intervention-specific mean outcomes:

$$Q_t^{a_0',\underline{\mathbf{a}}_1}(P)(\mathbf{L}_{t-1},V,W) = E_P\left[Q_{t+1}^{a_0',\underline{\mathbf{a}}_1}(P)(\mathbf{L}_t,V,W) \middle| \mathbf{A}_t = \mathbf{a}_t, \mathbf{L}_{t-1}, V, A_0 = a_0', W \right], \quad (11)$$

for $t = 0,\ldots,K$. This will prove useful in our upcoming endeavor. We also adopt the notations $Q^W(P)$ for the marginal distribution of W, $Q^V(P)$ for the conditional distribution $P(V \mid W, A_0)$, and $Q(P) \equiv \left(Q^W(P), Q^V(P), Q_t^{a_0',\underline{\mathbf{a}}_1}(P) : t = 0,\ldots K \right)$. We write g for the treatment allocation probabilities $P(A_t \mid \mathbf{A}_{t-1}, \mathbf{L}_{t-1}, V, W)$, and g^{A_0} for the one pertaining to A_0. 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 functionals 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 A Tale of Two Influence functions

The first leg of our journey is determining the so-called Efficient influence function (EIC) for our parameter of interest. From a fundamental result in Bickel, Klaassen, Ritov, and Wellner (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 those in our examples), the EIC also provides insights into the potential robustness against model misspecifications. In section 3.1, 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 3.2 we present a projection of the EIC onto a model where g^{A_0} is known; we refer to it as the *projected Influence Function* (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 $\gamma_{a_0',v}(P)$ and $\rho_{a_0',\underline{\mathbf{a}}_1,v}(P)$ (see (5) and (6) for definitions). Note also that $\rho_{a_0',\underline{\mathbf{a}}_1,v}(P) = \frac{\eta_{a_0',\underline{\mathbf{a}}_1,v}(P)}{\gamma_{a_0',v}(P)}$, where

$$\eta_{a_0',\underline{\mathbf{a}}_1,v}(P) = E_P\left\{Q^V(V=v \mid A_0=a_0', W) \times Q_{t=0}^{a_0',\underline{\mathbf{a}}_1}(V=v,W)\right\}.$$

We will make use of the following useful characterizations for $\Psi(P)$:

Remark 1. *For $m_\psi(a_0',\underline{\mathbf{a}}_1,v) = expit\left(\psi\cdot\phi(a_0',\underline{\mathbf{a}}_1,v)\right)$ and $\Psi(P)$ defined as in (9), we have*

$$\begin{aligned}
0 = U(\Psi(P),P) &\equiv \sum_{(a_0',\underline{\mathbf{a}}_1,v)\in\mathscr{A}'\times\mathscr{A}\times\mathscr{V}} h(a_0',\underline{\mathbf{a}}_1,v)\phi(a_0',\underline{\mathbf{a}}_1,v)\gamma_{a_0',v}(P)\left\{\frac{\eta_{a_0',\underline{\mathbf{a}}_1,v}(P)}{\gamma_{a_0',v}(P)} - m_{\Psi(P)}(a_0',\underline{\mathbf{a}}_1,v)\right\} \\
&= E_P\left\{\sum_{(a_0',\underline{\mathbf{a}}_1,v)\in\mathscr{A}'\times\mathscr{A}\times\mathscr{V}} \tilde{h}(a_0',\underline{\mathbf{a}}_1,v)Q^V(V=v\mid A_0=a_0',W)\left\{Q_{t=0}^{a_0',\underline{\mathbf{a}}_1}(V=v,W) - m_{\Psi(P)}(a_0',\underline{\mathbf{a}}_1,v)\right\}\right\} \\
&= E_P\left\{\frac{I(A_0=a_0')}{g^{A_0}(1\mid W)}\sum_{(a_0',\underline{\mathbf{a}}_1,v)\in\mathscr{A}'\times\mathscr{A}\times\mathscr{V}} \tilde{h}(a_0',\underline{\mathbf{a}}_1,V)\left\{Q_{t=0}^{a_0',\underline{\mathbf{a}}_1}(V,W) - m_{\Psi(P)}(a_0',\underline{\mathbf{a}}_1,V)\right\}\right\},
\end{aligned}$$

where $\tilde{h}(a_0',\underline{\mathbf{a}}_1,v) \equiv h(a_0',\underline{\mathbf{a}}_1,v)\phi(a_0',\underline{\mathbf{a}}_1,v)$. The computations are straightforward, and we left them in the Appendix for reference.

3.1 Efficient Influence Function

From the first equality in remark 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 1 (Efficient Influence Function).

Consider $\Psi:\mathscr{M}\to\mathbb{R}^d$ as defined in (9). Suppose that the following $k\times k$ normalizing matrix is invertible at $(\psi,P) = (\Psi(P),P)$:

$$\begin{aligned}
M(\psi,P) &= \frac{\partial}{\partial\psi}\sum_{(a_0',\underline{\mathbf{a}}_1,v)\in\mathscr{A}'\times\mathscr{A}\times\mathscr{V}} \gamma_{a_0',v}(P)h(a_0',\underline{\mathbf{a}}_1,v)\phi(a_0',\underline{\mathbf{a}}_1,v)\left\{\rho_{a_0',\underline{\mathbf{a}}_1,v}(P) - m_\psi(a_0',\underline{\mathbf{a}}_1,v)\right\} \\
&= P\left\{\frac{I(A_0=a_0')}{g^{A_0}(1\mid W)}\sum_{a_0',\underline{\mathbf{a}}_1\in\mathscr{A}'\times\mathscr{A}} h(\mathbf{a},,V)\phi(a_0',\underline{\mathbf{a}}_1,V)\phi(a_0',\underline{\mathbf{a}}_1,V)^\top m_\psi(a_0',\underline{\mathbf{a}}_1,V)\left[1-m_\psi(a_0',\underline{\mathbf{a}}_1,V)\right]\right\}.
\end{aligned} \tag{12}$$

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

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

where

$$D^*(Q,g,\psi) = \sum_{t=0}^{K} D_t^*(Q,g) + D_V^*(Q,g,\psi) + D_W^*(Q,\psi), \tag{13}$$

with

$$D_t^*(Q,g) \equiv \frac{I(A_0 = a_0')}{g^{A_0}(1 \mid W)} \sum_{a_0', \underline{\mathbf{a}}_1 \in \mathscr{A}' \times \mathscr{A}} \tilde{h}(a_0', \underline{\mathbf{a}}_1, V) C_t^{\underline{\mathbf{a}}_1} \left\{ Q_{t+1}^{a_0', \underline{\mathbf{a}}_1}(\mathbf{L}_t, V, W) - Q_t^{a_0', \underline{\mathbf{a}}_1}(\mathbf{L}_{t-1}, V, W) \right\},$$

$$D_V^*(Q,g,\psi) \equiv \frac{I(A_0 = a_0')}{g^{A_0}(1 \mid W)} \sum_{a_0', \underline{\mathbf{a}}_1 \in \mathscr{A}' \times \mathscr{A}} \left\{ \tilde{h}(a_0', \underline{\mathbf{a}}_1, V) \left[Q_{t=0}^{a_0', \underline{\mathbf{a}}_1}(V,W) - m_\psi(a_0', \underline{\mathbf{a}}_1, V) \right] \right.$$
$$\left. - E_P \left(\tilde{h}(a_0', \underline{\mathbf{a}}_1, V) \left[Q_{t=0}^{a_0', \underline{\mathbf{a}}_1}(V,W) - m_\psi(a_0', \underline{\mathbf{a}}_1, V) \right] \Big| A_0, W \right) \right\},$$

$$D_W^*(Q,\psi) \equiv \sum_{\mathbf{a} \in \mathscr{A}} E_P \left\{ \tilde{h}(a_0', \underline{\mathbf{a}}_1, V) \left[Q_{t=0}^{a_0', \underline{\mathbf{a}}_1}(V,W) - m_\psi(a_0', \underline{\mathbf{a}}_1, V) \right] \Big| A_0 = a_0', W \right\},$$

where $C_t^{\underline{\mathbf{a}}_1} = \frac{I(\mathbf{A}_t = \mathbf{a}_t)}{\prod_{j=1}^{t} g^A\left(A_j = a_j | \mathbf{A}_{j-1} = \mathbf{a}_{j-1}, \mathbf{L}_{j-1}, V, A_0 = a_0', W\right)}$, *for* $t = 1, \ldots, K$, *and* $C_t^{\underline{\mathbf{a}}_1} = 1$ *for* $t = 0$.

Moreover, 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 first robustness condition of lemma 1, $Q^V = Q^V(P_0)$ can be relaxed to

$$E_P \left\{ \tilde{h}(a_0', \underline{\mathbf{a}}_1, V) \left[Q_{t=0}^{a_0', \underline{\mathbf{a}}_1}(V,W) - m_\psi(a_0', \underline{\mathbf{a}}_1, V) \right] \Big| A_0 = a_0', W \right\} \Big\}$$
$$= E_{P_0} \left\{ \tilde{h}(a_0', \underline{\mathbf{a}}_1, V) \left[Q_{t=0}^{a_0', \underline{\mathbf{a}}_1}(V,W) - m_\psi(a_0', \underline{\mathbf{a}}_1, V) \right] \Big| A_0 = a_0', W \right\} \Big\}.$$

Remark 2. ***Dimensionality of*** V ***and Implementation:*** *When* V *is high-dimensional (or continuous), a regression-based estimator (parametric or data-adaptive) can be used to directly estimate the conditional expectations with respect to* Q^V *that appear in* D_V^* *and in* D_W^*. *However, for final evaluation of the target parameter* ψ, *we must solve the estimating equation* D_W^* *in the variable* ψ. *One way to implement this is using standard generalized linear modeling software package to regress* $Q_{t=0}^{a_0', \underline{\mathbf{a}}_1}(v,W)$ *onto the model* m_ψ *with weights* $I(A_0 = a_0') h(a_0', \underline{\mathbf{a}}_1, V) Q^V(v \mid A_0 = a_0', W)$, *for a pooled data with every value in the outcome space of* V. *When*

V is high-dimensional, this may be difficult to accomplish. Instead, using the given representation of D^*_W, one may employ numerical tools to solve for ψ in the corresponding estimating equation.

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

3.2 Projected Influence Function

As motivated by remark 2, when V is high-dimensional, we may instead consider a projected influence function which still retains most of the robustness of D^*.

Lemma 2 (Projected Influence Function).
Consider the setup in lemma 1. Up to a normalizing matrix $M(\Psi(P),P)^{-1}$, the following function is a gradient for Ψ at P under the model $\mathscr{M}_{A_0}$, in which g^{A_0} is known:

$$D^{A_0}(Q,g,\psi) = \sum_{t=1}^{K} D^*_t(Q,g) + D^{A_0}_W(Q,g,\psi),$$

where

$$D^{A_0}_W \equiv \frac{I(A_0 = a_0')}{g^{A_0}(1 \mid W)} \sum_{a_0',\underline{\mathbf{a}}_1 \in \mathscr{A}' \times \mathscr{A}} \tilde{h}(a_0',\underline{\mathbf{a}}_1,V) \left[Q_{t=1}^{a_0',\underline{\mathbf{a}}_1}(L_0,V,W) - m_\psi(a_0',\underline{\mathbf{a}}_1,V) \right]. \qquad (14)$$

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

Moreover, if $g^{A_0} = g^{A_0}(P_0)$ and either $Q_t^{a_0',\underline{\mathbf{a}}_1} = Q_t^{a_0',\underline{\mathbf{a}}_1}(P_0)$ or $g^A = g^A(P_0)$, then $P_0 D^{A_0}(Q,g) = 0$ implies $\Psi(Q) = \Psi(Q_0)$.

Proof. The proof is given in the Appendix. □

At its face value, the proposed D^{A_0} may seem less robust than D^*, as it always relies on consistent estimation of $g^{A_0}(P_0)$. However, as we noted in remark 2, when V is high-dimensional, there are more standard machine learning algorithms available for estimation of g^{A_0}. Moreover, the estimators that utilizes D^{A_0} are also easier to implement, since standard software packages can be used to solve for ψ in the corresponding estimating equation of $D^{A_0}_W$.

4 Statistical Inference

With the two influence curves D^* and D^{A_0} under our belt, in section 4.1 we will build two robust, substitution-based, asymptotically linear estimators via the targeted maximum likelihood estimation (TMLE) methodology. In section 4.2, 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 estimators for the standard errors. In section 4.3, we describe a non-targeted substitution estimator which utilizes a non-targeted MLE estimate of Q_0 (or of $Q_t^{a_0',\underline{\mathbf{a}}_1}(P_0)$ and $g^{A_0}(P_0)$). This estimator is biased if these non-targeted MLE are not consistent.

For most of the estimators below, we first need to procure estimators $g_n = \left(g_n^{A_0}, g_n^A\right)$ of g_0. The marginal distribution of W will always be estimated by the empirical distribution. For a given estimator ψ_n of ψ_0, we will use

$$M(\psi_n) \equiv P_n \left\{ \frac{I(A_0 = a_0')}{g_n^{A_0}(1 \mid W)} \sum_{a_0', \underline{\mathbf{a}}_1 \in \mathscr{A}' \times \mathscr{A}} h(a_0', \underline{\mathbf{a}}_1, V) \phi(a_0', \underline{\mathbf{a}}_1, V) \phi(a_0', \underline{\mathbf{a}}_1, V)^\top m_{\psi_n}(a_0', \underline{\mathbf{a}}_1, V) \left[1 - m_{\psi_n}(a_0', \underline{\mathbf{a}}_1, V)\right] \right\}$$

to estimate the normalizing matrix.

4.1 Targeted Maximum Likelihood Estimator

In a traditional non-targeted MLE (like those in section 4.3), 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 increases, these methods may break down due to the 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 Rubin

(2006) and van der Laan and Rose (2011) for the general methodology. Here, we construct two targeted estimators using D^{A_0} and D^*, respectively.

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 $Q_t^{a_0',\underline{\mathbf{a}}_1}$ as a conditional expectation of $Q_{t+1}^{a_0',\underline{\mathbf{a}}_1}$, as given in (11), we use the quasi loglikelihood loss function for $Q_t^{a_0',\underline{\mathbf{a}}_1}$:

$$L\left(Q_t^{a_0',\underline{\mathbf{a}}_1}\right) = -I(\mathbf{A}_t = \mathbf{a}_t)\left\{\log\left(Q_t^{a_0',\underline{\mathbf{a}}_1}\right)^{Q_{t+1}^{a_0',\underline{\mathbf{a}}_1}} + \log\left(1 - Q_t^{a_0',\underline{\mathbf{a}}_1}\right)^{\left(1-Q_{t-1}^{a_0',\underline{\mathbf{a}}_1}\right)}\right\}. \tag{15}$$

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

$$Q_t^{a_0',\underline{\mathbf{a}}_1}(\varepsilon) = expit\left(logit Q_t^{a_0',\underline{\mathbf{a}}_1} + \varepsilon \cdot \tilde{h}(a_0',\underline{\mathbf{a}}_1,V)\frac{I(A_0 = a_0')}{g^{A_0}(1 \mid W)}C_t^{\underline{\mathbf{a}}_1}\right). \tag{16}$$

This submodel satisfies $\langle \frac{d}{d\varepsilon}\sum_{\mathbf{a}} L\left(Q_t^{a_0',\underline{\mathbf{a}}_1}(\varepsilon)\right)|_{\varepsilon=0}\rangle \supset \langle D_t^*(Q,g)\rangle$, where $\langle x\rangle$ represents the linear span of a vector x.

4.1.1 *projected* TMLE using projected-IC D^{A_0}

1. Start at $t = K$: regress Y_K on $(\mathbf{L}_{K-1},\mathbf{A}_K,L_0,V,W)$, and then evaluate at $A_0 = a_0'$ and $\underline{\mathbf{A}}_1 = \underline{\mathbf{a}}_1$ to obtain an initial estimator $Q_{t=K,n}^{a_0',\underline{\mathbf{a}}_1}$ of $Q_{t=K}^{a_0',\underline{\mathbf{a}}_1}(P_0)$. The optimal fluctuation amount around this initial estimate is given by

$$\varepsilon_{K,n} \equiv \arg\min_{\varepsilon}\sum_{\mathbf{a}} P_n L\left(Q_{t=K,n}^{a_0',\underline{\mathbf{a}}_1}(\varepsilon)\right).$$

   This can be implemented by creating one row for each individual and each $(a_0',\underline{\mathbf{a}}_1) \in \mathscr{A}' \times \mathscr{A}$, and fitting a weighted logistic regression of Y_K on the multivariate covariate $\frac{\phi(a_0',\underline{\mathbf{a}}_1,V)}{g_n^{A_0}(1|W)}C_K^{\underline{\mathbf{a}}_1}(g_n)$ on these observations with weights $I(A_0 = a_0')I(\underline{\mathbf{A}}_1 = \underline{\mathbf{a}}_1)h(a_0',\underline{\mathbf{a}}_1,V)$ and offset $Q_{t=K,n}^{a_0',\underline{\mathbf{a}}_1}(\mathbf{L}_{K-1},V,W)$. Update the initial estimator using $Q_{t=K,n}^{*,a_0',\underline{\mathbf{a}}_1} \equiv Q_{t=K,n}^{a_0',\underline{\mathbf{a}}_1}(\varepsilon_{K,n})$.
2. At each subsequent step $t = K-1,\ldots,1$, we have thus far obtained an updated estimator $Q_{t+1,n}^{*,a_0',\underline{\mathbf{a}}_1}$ for each individual with $(A_0)_i = 1$ and each $\mathbf{a} \in \mathscr{A}$. Regress $Q_{t+1,n}^{*,a_0',\underline{\mathbf{a}}_1}$ on $(\mathbf{L}_{t-1},\mathbf{A}_t,L_0,W,V)$ among observations with $A_0 = a_0'$ and evaluate at $\mathbf{A}_t = \mathbf{a}_t$ to get an initial estimator $Q_{t,n}^{a_0',\underline{\mathbf{a}}_1}$ of $Q_t^{a_0',\underline{\mathbf{a}}_1}(P_0)$. The

optimal fluctuation amount around this initial estimator is given by $\varepsilon_{t,n} = \arg\min_\varepsilon \sum_{\mathbf{a}} P_n L\left(Q_{t,n}^{a_0',\underline{\mathbf{a}}_1}(\varepsilon)\right)$, and can be obtained in an analogous manner to step 1. The updated estimator is $Q_{t,n}^{*,a_0',\underline{\mathbf{a}}_1} \equiv Q_{t,n}^{a_0',\underline{\mathbf{a}}_1}(\varepsilon_{t,n})$.

3. After sequentially performing step (2) in order of decreasing t, we now have a targeted estimator $Q_{t=1,n}^{*,a_0',\underline{\mathbf{a}}_1}$ of $Q_{t=1}^{a_0',\underline{\mathbf{a}}_1}(P_0)$. Obtain an estimator ψ_n^{pTMLE} of ψ_0 by fitting a weighted logistic regression of $Q_{t=1,n}^{*,a_0',\underline{\mathbf{a}}_1}(V,W)$ on $\phi(a_0',\underline{\mathbf{a}}_1,V)$, with weights $h(a_0',\underline{\mathbf{a}}_1,V)\frac{I(A_0=a_0')}{g_n^{A_0}(1|W)}$. We call this estimator the *projected TMLE*.

By construction, $P_n D^{A_0}\left(Q_{t,n}^{*,a_0',\underline{\mathbf{a}}_1}, g_n, \psi_n^{pTMLE}\right) = 0$. From lemma 2, ψ_n^{pTMLE} is an unbiased estimator of ψ_0 if either (1) $g_n^{A_0}$ and $Q_{t,n}^{*,a_0',\underline{\mathbf{a}}_1}$ 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 4.2 and 4.3. Moreover, under standard regularity and empirical process conditions, ψ_n^{pTMLE} is asymptotically linear with influence function $M(\Psi(P_0))^{-1}D^{A_0}(P_0)$. The asymptotic covariance of $\sqrt{n}(\psi_n^{pTMLE} - \psi_0)$ can be estimated by the sample covariance matrix Σ_n^{pTMLE} of $M(\psi_n^{pTMLE})^{-1}D^{A_0}\left(Q_{t,n}^{*,a_0',\underline{\mathbf{a}}_1}, g_n, \psi_n^{pTMLE}\right)$.

4.1.2 Full TMLE using EIC D^*

Motivated by remark 2, we rewrite D_V^* and D_W^* in (13) as

$$D_V^*(Q,g,\psi) \equiv \frac{I(A_0 = a_0')}{g^{A_0}(1 \mid W)} \sum_{(a_0',\underline{\mathbf{a}}_1,v)\in\mathscr{A}'\times\mathscr{A}\times\mathscr{V}} \left\{\tilde{h}(a_0',\underline{\mathbf{a}}_1,v)\left(Q_{t=0}^{a_0',\underline{\mathbf{a}}_1}(V=v,W) - m_\psi(a_0',\underline{\mathbf{a}}_1,v)\right)\right.$$
$$\left.\times\left(I(V=v) - Q^V(v \mid A_0,W)\right)\right\}$$
$$D_W^*(Q,\psi) \equiv \sum_{(a_0',\underline{\mathbf{a}}_1,v)\in\mathscr{A}'\times\mathscr{A}\times\mathscr{V}} \tilde{h}(a_0',\underline{\mathbf{a}}_1,v)Q^V(v \mid A_0 = a_0',W)\left\{Q_{t=0}^{a_0',\underline{\mathbf{a}}_1}(V=v,W) - m_\psi(a_0',\underline{\mathbf{a}}_1,v)\right\},$$

To use D^*, we also consider the loss function $L(Q^V) \equiv -\log Q^V(V \mid A_0 = a_0',W)$, and a d-dimensional fluctuation model through a given Q^V at $\varepsilon = 0$ given by

$$Q^V(\varepsilon)(V \mid A_0 = a_0',W) \equiv \frac{Q^V(V \mid A_0 = a_0',W)\exp\left[\varepsilon \cdot B(Q,\psi)(W,V)\right]}{\sum_v Q^V(v \mid A_0 = a_0',W)\exp\left[\varepsilon \cdot B(Q,g,\psi)(W,v)\right]},$$

where $B(Q,\psi)(W,V) \equiv \sum_{\mathbf{a}} \tilde{h}(a_0',\underline{\mathbf{a}}_1,V)\left\{Q_{t=0}^{a_0',\underline{\mathbf{a}}_1}(V,W) - m_\psi(a_0',\underline{\mathbf{a}}_1,V)\right\}$. It is easy to verify that $\langle \frac{d}{d\varepsilon}\frac{I(A_0=a_0')}{g_{A_0}(1|W)}L(Q^V(\varepsilon))|_{\varepsilon=0}\rangle \supset \langle D_V^*(Q,g,\psi)\rangle$. The targeted estimator

which uses 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 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 4.1.1 to obtain a targeted estimator $Q^{*,a_0',\underline{\mathbf{a}}_1}_{t=0,n}$.
2. Let Q^V_n 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 $Q^{*,a_0',\underline{\mathbf{a}}_1}_{t=0,n}(V = v, W)$, $h(a_0', \underline{\mathbf{a}}_1, v)$, $\phi(a_0', \underline{\mathbf{a}}_1, v)$ and $Q^V_n(v \mid A_0 = a_0', W)$. Obtain a first-iteration estimator ψ^1_n of ψ_0 by fitting a weighted logistic regression of $Q^{*,a_0',\underline{\mathbf{a}}_1}_{t=0,n}(V = v, W)$ on $\phi(a_0', \underline{\mathbf{a}}_1, v)$, with weights $h(a_0', \underline{\mathbf{a}}_1, v) \times Q^V_n(v \mid A_0 = a_0', W)$, on this pooled data.
3. Given ψ^1_n obtained in previous step, we update the estimator for $Q^V(P_0)$ as follows. Using previously obtained $Q^{*,a_0',\underline{\mathbf{a}}_1}_{t=0,n}$, g_n, and ψ^1_n, the optimal fluctuation amount around the initial Q^V_n is given by $\varepsilon^V_n = \arg\min_\varepsilon P_n \frac{I(A_0 = a_0')}{g^{A_0}_n(1|W)} L\left(Q^V_n(\varepsilon)\right)$. This can be obtained by solving for ε in the equation

$$0 = \sum_{i=1}^{n} \frac{I((A_0)_i = 1)}{g^{A_0}_n(1 \mid W_{1,i})} \times \left\{ \hat{B}_n(W_{1,i}, V_i) - \frac{\sum_v \hat{B}_n(W_{1,i}, v) Q^V_n(v \mid A_0 = a_0', W_{1,i}) \exp\left[\varepsilon \cdot \hat{B}_n(W_{1,i}, v)\right]}{\sum_v Q^V_n(v \mid A_0 = a_0', W_{1,i}) \exp\left[\varepsilon \cdot \hat{B}_n(W_{1,i}, v)\right]} \right\}$$

   where $\hat{B}_n \equiv B\left((Q^{*,a_0',\underline{\mathbf{a}}_1}_{t=0,n} : \mathbf{a}), \psi^1_n\right)$. The updated density is given by $Q^1_{V,n} \equiv Q^V_n(\varepsilon^V_n)$.
4. Having obtained an updated density $Q^{V,j}_n$ at the j-th iteration, repeat step (2) and (3) to obtain a targeted estimate of ψ^{j+1}_n and $Q^{j+1}_{V,n}$, until ε^V_n converges to 0. In practice, this convergence can be achieved (close to 0) after a few iterations. We denote the final updates as $\psi^{*,TMLE}_n$, and $Q^{*,V}_n$. We call this estimator the *full TMLE*.

Let $Q^*_n \equiv \left(Q^W_n, Q^{*,V}_n, fQ^{*,a_0',\underline{\mathbf{a}}_1}_{t,n}\right)$, where Q^W_n is the empirical distribution of W. By design, $P_n D^*(Q^*_n, g_n, \psi^{*,TMLE}_n) = 0$. From lemma 1, we know that $\psi^{*,TMLE}_n$ is unbiased if either $Q^*_n = Q_0$ or $g_n = g_0$. Under standard regularity and empirical process conditions, $\psi^{*,TMLE}_n$ is asymptotically linear with influence function $M(\Psi(P_0))^{-1} D^*(P_0)$. The asymptotic covariance of $\sqrt{n}(\psi^{*,TMLE}_n - \psi_0)$ can be estimated by the sample covariance matrix $\Sigma^{*,TMLE}_n$ of $\left\{ M(\psi^{*,TMLE}_n)^{-1} D^*\left(Q^*_n, g_n, \psi^{*,TMLE}_n\right) \right\}$.

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.

4.2 Inverse Probability Weighted estimator

From remark 1, another valid estimating function for Ψ is given by

$$D^{IPW}(g,\psi) \equiv \frac{I(A_0 = a_0')}{g^{A_0}(1 \mid W)} \sum_{a_0',\underline{\mathbf{a}}_1 \in \mathscr{A}' \times \mathscr{A}} \tilde{h}(a_0',\underline{\mathbf{a}}_1,V) C_K^{\underline{\mathbf{a}}_1} \left[Y_K - m_\psi(a_0',\underline{\mathbf{a}}_1,V) \right]. \quad (17)$$

Up to a normalizing matrix $M(\Psi(P),P)^{-1}$, as defined in (12), $D^{IPW}(g,\psi)$ 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 $a_0', \underline{\mathbf{a}}_1$ and each individual i with$A_0 = a_0'$ and $\underline{\mathbf{A}}_1 = \underline{\mathbf{a}}_1$, we create a row of data consisting of $Y_{K,i}$, $\phi(a_0',\underline{\mathbf{a}}_1,V_i)$, $h(a_0',\underline{\mathbf{a}}_1,V_i)$, $C_{K,i}^{\underline{\mathbf{a}}_1}$, $g_n^{A_0}(1 \mid W_{1,i})$. The estimator ψ_n^{IPW} can be obtained by fitting a weighted regression of Y_K on $\phi(a_0',\underline{\mathbf{a}}_1,V)$, with weights $\frac{1}{g_n^{A_0}(1|W)} h(a_0',\underline{\mathbf{a}}_1,V) C_K^{\underline{\mathbf{a}}_1}$. 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 function $M(\Psi(P_0))^{-1}D^{IPW}(g_0)$. The asymptotic covariance of $\sqrt{n}(\psi_n^{IPW} - \psi_0)$ can be estimated by the sample covariance matrix Σ_n^{IPW} of $\left\{ M(\psi_n^{IPW})^{-1} D^{IPW}\left(g_n,\psi_n^{IPW}\right) \right\}$.

4.3 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 (9) and remark 1, we can express $\Psi(P_0)$ as $\Psi\left(Q_{t=0}^{a_0',\underline{\mathbf{a}}_1}(P_0), Q^V(P_0)\right)$ or $\Psi\left(Q_{t=0}^{a_0',\underline{\mathbf{a}}_1}(P_0), g_0^{A_0}\right)$, the latter option opening 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 $Q_{t=0}^{a_0',\underline{\mathbf{a}}_1}(P_0)$, we can use a sequential regression approach by performing steps (1) and (2) of section 4.1.1, starting at $t = K$ and ending at $t = 0$, but without the targeting procedure, i.e. always use $Q_{t+1,n}^{a_0',\underline{\mathbf{a}}_1}$ instead of $Q_{t+1,n}^{*,a_0',\underline{\mathbf{a}}_1}$ at t. At the end of $K+1$ steps, we have an estimator $Q_{t=0,n}^{a_0',\underline{\mathbf{a}}_1}$.

We first consider the representation $\Psi\left(Q_{t=0}^{a'_0,\underline{\mathbf{a}}_1}(P_0), Q^V(P_0)\right)$. Let Q_n^V be an estimators of $Q^V(P_0)$. For each observation i, each $a_0, \underline{\mathbf{a}}_1$ and each $v \in \mathscr{V}$, we create a row of data consisting of $Q_{t=0,n}^{a'_0,\underline{\mathbf{a}}_1}(V=v,W), \phi(a'_0,\underline{\mathbf{a}}_1,v), h(a'_0,\underline{\mathbf{a}}_1,v), Q_n^V(v \mid A_0=a'_0,W)$. The estimator $\psi_n^{V,Gcomp}$ can be obtained by a weighted regression $Q_{t=0,n}^{a'_0,\underline{\mathbf{a}}_1}(V=v,W)$ on $\phi(a'_0,\underline{\mathbf{a}}_1,v)$, with weights $h(a'_0,\underline{\mathbf{a}}_1,v)Q_n^V(v \mid A_0=a'_0,W)$. This $\psi_n^{V,Gcomp}$ is unbiased if both $Q_{t=0,n}^{a'_0,\underline{\mathbf{a}}_1}$ and Q_n^V are consistent.

Consider now the alternative representation $\Psi\left(Q_{t=0}^{a'_0,\underline{\mathbf{a}}_1}(P_0), g_0^{A_0}\right)$, from the equalities in remark 1. For each observation i and each $a_0, \underline{\mathbf{a}}_1$, create a row of data consisting of $Q_{t=0,n}^{a'_0,\underline{\mathbf{a}}_1}(V,W)$, $\phi(a'_0,\underline{\mathbf{a}}_1,V)$, $h(a'_0,\underline{\mathbf{a}}_1,V)$, $g_n^{A_0}(1 \mid W)$. The estimator $\psi_n^{A_0,Gcomp}$ can be obtained by a weighted regression $Q_{t=0,n}^{a'_0,\underline{\mathbf{a}}_1}(V,W)$ on $\phi(a'_0,\underline{\mathbf{a}}_1,V)$, with weights $\frac{h(a'_0,\underline{\mathbf{a}}_1,V)}{g_n^{A_0}(1|W)}$. This $\psi_n^{A_0,Gcomp}$ is unbiased if both $Q_{t=0,n}^{a'_0,\underline{\mathbf{a}}_1}$ and $g_n^{A_0}$ are consistent.

5 Simulation Study

In this section, we examine the relative performance of the *projected* TMLE estimator (section 4.1.1), the IPW estimator (section 4.2), and the G-computation estimator (section 4.3) for the parameters of an MSM model.

5.1 Data Generating Process and Target Parameter

We consider a survival type example with data structure $O=(W,A_0,V,L_0,(A_t,L_t): t=1,\ldots K)$ 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:

$$
\begin{aligned}
&\left(W^1, W^2\right) \sim \left(Bern(0.3), Bern(0.7)\right);\\
&A_0 \sim Bern\left(expit(1+2W^1+0.1W^2)\right);\\
&V \in \{0,1,2\} \sim \Big\{ I(V=1) \sim Bern\left(expit(-2+1.2W^1+0.7W^2)\right),\\
&\qquad\qquad \{I(V=2) \mid V \neq 1\} \sim Bern\left(expit(-0.7+1.2W^1+W^2)\right)\Big\};\\
&L_0^1 \sim Bern\left(expit(-0.2+2W^1+0.5W^2+0.2I(V=1)+0.4I(V=2))\right),\\
&L_0^2 \sim Bern\left(expit(-0.8+W^1+W^2-0.3I(V=1)-0.1I(V=2))\right);\\
&A_t \sim \begin{cases}
Bern\Big(expit(-1+W^1+1.3W^2+0.1I(V=1)+0.1I(V=2)+1.2L_0^1+L_0^2-0.7W^1\times L_0^2\\
\qquad -0.5W^2\times L_0^1)\Big), \text{for } t=1,\\
Bern\Big(expit\big(2+W^1+W^2+0.1I(V=1)+0.1I(V=2)+0.6L_0^1+1.2L_0^2-0.5A-0.1t\\
\qquad +0.8L_t^1-0.3L_t^2+0.1L_{t-1}^1-0.2L_{t-1}^2-0.3A\times L_0^2+0.2A\times W^1-0.3A\times L_t^2\\
\qquad -0.2A\times L_{t-1}^1\big)\Big), \text{for } t>1.
\end{cases}\\
&L_t^1 \sim Bern\Big(expit\big(-1+W^1+0.1W^2-0.5I(V=1)-0.7I(V=2)+L_0^1+0.3L_0^2+1.5A+0.4t\\
&\qquad -L_{t-1}^1-0.8A\times I(V=1)-0.2A\times I(V=2)-0.3A\times W^1\big)\Big);\\
&L_t^2 \sim Bern\Big(expit\big(-2+0.1W^1+0.1W^2+0.5I(V=1)+0.5I(V=2)+0.7L_0^1+0.2L_0^2-A+0.2t\\
&\qquad +L_{t-1}^2-0.2A\times I(V=1)-0.4A\times I(V=2)-0.3A\times L_0^2\big)\Big);\\
&Y_t \sim Bern\Big(expit\big(-1.4+1.5W^1+W^2-0.7I(V=1)-0.8I(V=2)+L_0^1+L_0^2+A-L_t^1-0.1L_t^2\\
&\qquad -L_{t-1}^1-0.3L_{t-1}^2+A\times I(V=1)+0.8A\times I(V=2)-0.3A\times L_t^1\\
&\qquad -0.4A\times L_{t-1}^1-0.3A\times L_0^2-0.2A\times W^1\big)\Big).
\end{aligned}
$$

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 $A_0 = a_0'$ 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 $\{\rho_{a_0',\underline{\mathbf{a}}_1,v} : \mathbf{a}, v\}$ by the MSM

$$m_\psi(a_0', \underline{\mathbf{a}}_1, v) = expit(\psi_1 + \psi_2 a_1 + \psi_3 v_1 + \psi_4 v_2 + \psi_5 a_1 v_1 + \psi_6 a_1 v_2) = expit(\psi \cdot \phi(a_0', \underline{\mathbf{a}}_1, v)),$$

where $\phi(a_0', \underline{\mathbf{a}}_1, v) = (1, a_1, v_1, v_2, a_1 v_1, a_1 v_2)$, with kernel weights $h(a_0', \underline{\mathbf{a}}_1, v) = P_0(\mathbf{a} \mid v, A_0 = a_0')$. Note that in this case, the kernel weights are assumed to be known. The target parameter defined in (9) takes value $\psi_0 = (0.825, 0.105, 0.249, -0.046, 1.474, 0.960)$.

5.2 Estimators

The A_0 mechanism g^{A_0} is estimated using Super Learner (van der Laan, Polley, and Hubbard (2007)) with candidate fitting algorithms `glm` and `nnet`, adjusting for W^1 and W^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 $W^1, W^2, L_0^1, L_0^2, L_t^1$. The denominator for each $C_t^{\underline{\mathbf{a}}_1}$ is truncated below by 0.025. We use two estimators for $Q_t^{a_0',\underline{\mathbf{a}}_1}(P_0)$: 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 $Q_t^{a_0',\underline{\mathbf{a}}_1}$ and g^A: all correct ('Qc, gc'); correct $Q_t^{a_0',\underline{\mathbf{a}}_1}$ and misspecified g^A ('Qc,gm'); misspecified $Q_t^{a_0',\underline{\mathbf{a}}_1}$ and correct g^A ('Qm, gc'). For all three cases we always use the same correctly specified g^{A_0}. We implement the second version of the G-comp estimator in 4.3, where the weights are given by g^{A_0}. 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'.

5.3 Results

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

In table 1, we see that when g^A is misspecified, *projected* TMLE using a correct $Q_t^{a_0',\underline{\mathbf{a}}_1}$ reduces bias over the misspecified IPW estimator. Similarly, when $Q_t^{a_0',\underline{\mathbf{a}}_1}$ is misspecified, *projected* TMLE using the correct g^A reduces bias over the misspecified G-computation estimator. When comparing the correct vs misspecified G-computation estimators, 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 *projected* 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 *projected* TMLE estimators are asymptotically linear — table 2 tabulates the coverage probability of their Influence-Function based confidence intervals. At the correct models (Qc,gc), IPW and *projected* TMLE are asymptotically linear with influence function D^{IPW} and D^{A_0}, respectively. We used $\sqrt{\hat{var}D_n^{IPW}/n}$ and $\sqrt{\hat{var}D_n^{A_0}/n}$ to estimate their respective standard errors. As sample size grows, the actual coverage probability is quite close to the nominal coverage probability, with IPW having a better coverage. When one of the components is misspecified, the *projected* TMLE still provides very good coverage, even though theoretically D^{A_0} is only part of its influence curve; we postulate that this is because the influence function 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).

6 Data Analysis Example

To illustrate the application of the *projected* 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 to an option (i.e. a particular mediation) within one of his accepted treatment strategies (augmenting vs switching). 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.

In the field of depression research, there is very few literature concerning individual baseline characteristics that may differentially modify the effect of the treatment strategies. For this data analysis example, suppose 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. Because the strategies themselves are not randomized, this analysis is *de facto* equivalent to that of an observational study, and our estimators will correspondingly account for baseline and time-varying confounding. All candidate modifiers are *a priori* selected through a literature review of previous START*D publications. These covariates are measured prior to the assignment of level 2 treatment, either at study screening, clinical visits or exit surveys at the end of level 1. Many of these candidate modifiers are subject to missingness, often times due to missing visits and surveys, 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 medications), for most patients these strategies are self-selected. Table 3 tabulates the events in level 2 by strategy received. Note that there are three strategies received, but we are only comparing switching medication vs augmenting medication. The data structure was described in detail in section 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 4 summarizes percent of missingness, range, and scale of each effect modifier. Of all the candidate modifiers proposed by literature review, we exclude from our analysis the history of amphetamine use at baseline and history of drug abuse, since these two variables are missing for more than 70% of the patients. The multivariate nature of most of the effect modifiers underscores the need for *projected* TMLE. If V is screening covariate, W includes all demographic variables and medical and psychiatric history prior to enrollment; if V is level 1 exit covariate, then we add to W 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 $\phi(a_0', \underline{\mathbf{a}}_1, v) = (1, a_1, v_1, \ldots, v_k, a_1 \times v_1, \ldots, a_1 \times v_k)$ and $\psi \equiv \left(\psi_1, \psi_2, \psi_{3,1}, \ldots, \psi_{3,k}, \psi_{4,1}, \ldots, \psi_{4,k}\right) \in \mathbb{R}^{2k+2}$ for the categorical V with $k+1$ categories, and $\phi(a_0', \underline{\mathbf{a}}_1, v) = \left(1, a_1, v, v^2, a_1 \times v, a_1 \times v^2\right)$ and $\psi \equiv (\psi_1, \psi_2, \psi_{3,1}, \psi_{3,2}, \psi_{4,1}, \psi_{4,2}) \in \mathbb{R}^6$ for the semi-continuous V. The kernel weights are $h(a_0', \underline{\mathbf{a}}_1, V) = P_0(A_1 = a_1 \mid V)$, to be estimated using Super Learner with fitting algorithms `glm, knnreg, nnet` and `bayesglm`. The Super Learner will use 10-fold cross validation to select the best weighted combination of these fitting algorithms to estimate $h(a_0', \underline{\mathbf{a}}_1, V)$. The initial estimators of g and $Q_t^{a_0', \underline{\mathbf{a}}_1}$ 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 k variables, where k ranges from 10 variables up to 30% of the total number of variables being considered, in increments of 20. The Super Learner selects the best linear combination of all screening-fitting couples.

The goal of the analysis is to identify potential effect modifiers from a pool of candidate pre-treatment covariates. Our strategy is to measure the treatment heterogeneity for each of these covariates using an MSM, and then apply multiple hypothesis testing methods to identify those for which the treatment heterogeneity is significant. A common way to assess treatment heterogeneity across strata of V is to test whether the coefficients of the interaction terms $\psi_4 = (\psi_{4,1}, \ldots \psi_{4,k})$ are

different from 0. We perform the corresponding Wald test in table 5, with the null hypothesis $\psi_{\mathbf{4}} = 0$ and test statistics $T_n = {\psi_{\mathbf{4},\mathbf{n}}}^{\top} \Sigma_n^{-1} \psi_{\mathbf{4},\mathbf{n}} \sim \chi^2_k$, where $\psi_{\mathbf{4},\mathbf{n}}$ is the $A_0 - TMLE$ estimator, $\Sigma_n = \hat{Cov}(D^{A_0}_{\psi_4})/n$, and k is the number of interaction terms. The false discovery rate (FDR) of the simultaneous comparisons are controlled at 0.05 with the Benjamini-Hochberg procedure.

Since these candidate covariates differ in their type (semi-continuous vs binary vs multiple categories), the corresponding MSM parameters are of different dimensions, we also consider the measure of treatment heterogeneity

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

where

$$\ell_{OR}(v;\psi) \quad = \log \tfrac{m(1,v;\psi)}{1-m(1,v;\psi)} - \log \tfrac{m((0,v;\psi)}{1-m(0,v;\psi)} = \begin{cases} \psi_2 + \psi_{4,1} v_1 + \cdots + \psi_{4,k} v_k, \text{ for categorical } V, \\ \psi_2 + \psi_{4,1} v + \psi_{4,2} v^2, \text{ for semi-continuous } V. \end{cases}$$

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 *projected* TMLE, we obtain an estimator $\beta_n = \beta(\psi_n^{pTMLE})$ of β_0 for each V. An application of the functional delta Method, with the covariance matrix Σ_n of the estimated influence function $D_n^{A_0}$, 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 results of the analysis are summarized in table 6. Interestingly, the two approaches yield almost all the same potential effect modifiers.

Using our methods, we identified six factors, two at screening and four at level 1 exit, that modified the comparative treatment effect of augmenting vs switching medication. These effect modifiers are: prior suicide attempts, atypical depression, level 1 side-effect frequencies and burden (impairment), SF-36 physical functionality, and quality of life at level 1 exit. In general, the heterogeneity of the treatment effects are most pronounced for the level 1 exit modifiers than for those from screening, indicating opportunities for close monitoring and medication management.

There has only been one publication that addressed the heterogeneous treatment effects between switching and augmenting in level 2 treatment in STAR*D (Ellis, Dusetzina, Hansen, Gaynes, Farley, and Stumer (2013)). The authors employed propensity score and weighting methods to examine heterogeneity of the treatment effect among the treated (medication augmentation) through stratification by propensity score decile. This approach does not explicitly identify factors contributing to heterogeneity. Missing data is less problematic in this case, and is handled by multiple imputation.

Table 1: Results: Bias, Variance, MSE for estimators of ψ_0. Qc = correct $Q_t^{a_0',\underline{\mathbf{a}}_1}$, Qm= misspecified $Q_t^{a_0',\underline{\mathbf{a}}_1}$, 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
projected 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
projected 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
projected 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
projected 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
projected 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
projected 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
projected 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
projected 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
projected 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

Table 2: Coverage Probability for the Asymptotically Linear Estimators, using Influence-Function based Confidence Interval. Qc = correct $Q_t^{a_0',\underline{\mathbf{a}}_1}$, Qm= misspecified $Q_t^{a_0',\underline{\mathbf{a}}_1}$, 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
projected 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
projected 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												
projected 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

Table 3: Level 2 events by acceptable 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 4: Candidate effect modifiers of interest for level 2. Candidates with more than 50% missing in the study population will be discarded from the analysis.

	% missing	type (# of categories, or range of semi-continuous variable)
At screening		
race	0	Categorical (3)
years of schooling completed	0	Semi-continuous (0,27)
employment status	0	Categorical (3)
education level	0	Categorical (3)
attempted suicide	0	Categorical (2)
duration of current major depression episode (in units of 6 months)	1	Semi-continuous (0,111)
AxI: Amphetamine	73 (will not use)	Categorical (2)
history of drug abuse	76 (will not use)	Categorical (2)
# of Cumulative Illness Rating (CIRS) categories	0	Semi-continuous (0,12)
CIRS severity score	0	Semi-continuous (0,4)
Have AxI Obsessive-Compulsive Disorder	1	Categorical (2)
Have AxI Panic Disorder	1	Categorical (2)
Have AxI Post Traumatic Stress Disorder	2	Categorical (2)
Have AxI Agoraphobia	1	Categorical (2)
Have AxI Social Anxiety Disorder	2	Categorical (2)
Have AxI Generalized Anxiety Disorder	2	Categorical (2)
Anxious Depression at screening	6	Categorical (2)
Atypical Depression at screening	5	Categorical (2)
Melancholic Depression at screening	5	Categorical (2)
At level 1 exit		
Maximum Side Effects	1	Categorical (7)
SF-36 Physical Health (in units of 5)	23	Semi-continuous (0,12)
SF-36 Mental Health (in units of 5)	23	Semi-continuous ((0,12)
Quality of Life	23	Semi-continuous (0,98)
Side effect burden	1	Categorical (7)
Quick Inventory of Depressive Symptom (QIDS) total score (self-reported)	1	Semi-continuous (0,27)
Hamilton rating scale for depression (17-time) score	12	Semi-continuous (0,43)
IDS (30-item) score	13	Semi-continuous (0,74)

Table 5: Data Analysis Results: Multivariate Wald Test. $\psi_{\mathbf{4}} = (\psi_{4,1}, \ldots \psi_{4,k})$ are the coefficients of the interaction terms in the MSM; if V is semi-continuous, interaction terms are $\psi_{\mathbf{4}} \cdot (v, v^2)$, if V is categorical with $k+1$ categories, interaction terms are $\psi_{\mathbf{4}} \cdot (v_1, \ldots, v_k)$. $H_0 : \psi_{\mathbf{4}} = 0$. $T_n = \psi_{\mathbf{4},\mathbf{n}}{}^{\top} \Sigma_n^{-1} \psi_{\mathbf{4},\mathbf{n}} \sim \chi_k^2$, where ψ_n is given by *projected* TMLE, and Σ_n is the corresponding covariance estimator based on the covariance matrix of the projected-IC. Control false discovery rate at 0.05 with Benjamini-Hochberg procedure

	P-value	reject	$\psi_{\mathbf{4}}$	$SE(\psi_{\mathbf{4}})$
At screening				
race	4.97e-01	F	(-1.68e-01, -4.62e-01)	(2.62e-01, 4.27e-01)
years of schooling completed	2.89e-01	F	(2.50e-01, -1.10e-02)	(1.71e-01, 7.07e-03)
employment status	1.50e-01	F	(3.92e-01, -3.94e-01)	(2.30e-01, 5.89e-01)
education level	1.13e-01	F	(4.89e-01, -5.69e-03)	(3.13e-01, 3.80e-01)
attempted suicide	1.80e-05	T	(-1.10e+00)	(2.56e-01)
duration of current major depression episode (in units of 6 months)	4.44e-02	F	(1.92e-02, -9.41e-04)	(2.78e-02, 5.59e-04)
# of Cumulative Illness Rating (CIRS) categories	2.05e-02	F	(3.40e-01, -4.51e-02)	(1.42e-01, 1.65e-02)
CIRS severity score	8.09e-02	F	(-6.21e-01, 3.65e-01)	(4.41e-01, 1.85e-01)
Have AxI Obsessive-Compulsive Disorder	5.20e-01	F	(-1.86e-01)	(2.89e-01)
Have AxI Panic Disorder	1.17e-01	F	(5.53e-01)	(3.53e-01)
Have AxI Post-Traumatic Stress Disorder	8.78e-01	F	(3.86e-02)	(2.51e-01)
Have AxI Agoraphobia Disorder	9.73e-01	F	(1.04e-02)	(3.12e-01)
Have AxI Social Anxiety Disorder	5.15e-01	F	(1.81e-01)	(2.78e-01)
Have AxI Generalized Anxiety Disorder	1.90e-01	F	(3.68e-01)	(2.81e-01)
Anxious Depression at screening	6.89e-01	F	(-9.17e-02)	(2.3e-01)
Atypical Depression at screening	4.15e-04	T	(9.79e-01)	(2.77e-01)
Melancholic Depression at screening	1.58e-01	F	(-3.48e-01)	(2.46e-01)
At level 1 exit				
Maximum side effects	6.46e-09	T	(5.53e-01, 8.10e-02, 1.85e+00, -2.62e-01, 1.39e+00, 5.88e-01)	(5.27e-01, 5.25e-01, 4.59e-01, 5.02e-01, 5.20e-01, 5.65e-01)
SF-36 Physical Health (in units of 5)	6.83e-09	T	(2.55e+00, -1.74e-01)	(4.17e-01, 2.91e-02)
SF-36 Mental Health (in units of 5)	1.36e-01	F	(5.55e-01, -5.48e-02)	(2.78e-01, 2.86e-02)
Quality of Life	4.84e-04	T	(-7.53e-02, 5.35e-04)	(3.12e-02, 3.46e-04)
Side effect burden	0.00e+00	T	(2.91e-01, 5.18e-01, -6.26e-02, 1.29e+00, 2.73e+00, -5.85e-01)	(2.00e-03, 2.41e-03, 2.60e-03, 2.50e-03, 1.96e-03, 1.21e-03)
Quick Inventory of Depressive Symptom (QIDS) to- tal score (self-reported)	7.61e-02	F	(2.82e-01, -1.03e-02)	(1.24e-01, 4.62e-03)
Hamilton rating scale for depression (17-item) score	4.83e-02	F	(1.16e-01, -1.88e-03)	(8.45e-02, 2.17e-03)
IDS (30 item) score	1.27e-01	F	(5.06e-02, -4.75e-04)	(4.47e-02, 6.72e-04)

Table 6: 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 *projected* TMLE. The standard error estimate $\hat{SE}(\beta_n)$ is computed using on the variance of the projected-IC and the functional delta method. Control false discovery rate at 0.05 with Benjamini-Hochberg procedure

	P	reject	β_n	$\hat{SE}(\beta_n)$	$\arg\min_v \ell_{OR}(v;\psi_0)$	$\arg\max_v \ell_{OR}(v;\psi_0)$
At screening						
race	2.8e-01	F	4.6e-01	4.3e-01	3.0e+00	2.0e+00
years of schooling completed	1.2e-01	F	2.7e+00	1.7e+00	2.7e+01	1.1e+01
employment status	1.9e-01	F	7.9e-01	6.0e-01	3.0e+00	2.0e+00
education level	9.2e-02	F	4.9e-01	2.9e-01	3.0e+00	2.0e+00
attempted suicide	1.8e-05	T	1.1e+00	2.6e-01	1.0e+00	0.0e+00
duration of current major depression episode (in units of 6 months)	3.3e-02	F	9.6e+00	4.5e+00	1.1e+02	1.0e+01
# of Cumulative Illness Rating (CIRS) categories	5.9e-03	T	3.1e+00	1.1e+00	1.2e+01	3.8e+00
CIRS severity score	2.7e-02	F	3.6e+00	1.6e+00	8.5e-01	4.0e+00
Have AxI Obsessive-Compulsive Disorder	5.2e-01	F	1.9e-01	2.9e-01	1.0e+00	0.0e+00
Have AxI Panic Disorder	1.2e-01	F	5.5e-01	3.5e-01	0.0e+00	1.0e+00
Have AxI Post-Traumatic Stress Disorder	8.8e-01	F	3.9e-02	2.5e-01	0.0e+00	1.0e+00
Have AxI Agoraphobia Disorder	9.7e-01	F	1.0e-02	3.1e-01	0.0e+00	1.0e+00
Have AxI Social Anxiety Disorder	5.2e-01	F	1.8e-01	2.8e-01	0.0e+00	1.0e+00
Have AxI Generalized Anxiety Disorder	1.9e-01	F	3.7e-01	2.8e-01	0.0e+00	1.0e+00
Anxious Depression at screening	6.9e-01	F	9.2e-02	2.3e-01	1.0e+00	0.0e+00
Atypical Depression at screening	4.2e-04	T	9.8e-01	2.8e-01	0.0e+00	1.0e+00
Melancholic Depression at screening	1.6e-01	F	3.5e-01	2.5e-01	1.0e+00	0.0e+00
At level 1 exit						
Maximum side effects	8.1e-09	T	2.1e+00	3.7e-01	4.0e+00	3.0e+00
SF-36 Physical Health (in units of 5)	9.0e-10	T	9.3e+00	1.5e+00	0.0e+00	7.3e+00
SF-36 Mental Health (in units of 5)	1.0e-01	F	2.6e+00	1.6e+00	1.2e+01	5.1e+00
Quality of Life	1.4e-04	T	2.6e+00	7.0e-01	7.0e+01	0.0e+00
Side effect burden	0.0e+00	T	3.3e+00	2.3e-03	6.0e+00	5.0e+00
Quick Inventory of Depressive Symptom (QIDS) to- tal score (self-reported)	2.7e-02	F	1.9e+00	8.7e-01	0.0e+00	1.4e+01
Hamilton rating scale for depression (17-item) score	2.0e-02	F	1.8e+00	7.6e-01	0.0e+00	3.1e+01
IDS (30 item) score	5.1e-02	F	1.3e+00	6.9e-01	0.0e+00	5.3e+01

7 Summary

In this work, we studied causal effect modification by a counterfactual modifier. The tools developed here are applicable in situations where the effect modifier of interest may be better cast as counterfactual variables. Examples of such situations include the study of time-varying effect modification, or the study of baseline effect modification with modifiers that are missing at random. We established the efficient influence function (EIC) for the corresponding marginal structural model parameters which provides the semiparametric efficiency bound for all the asymptotically linear estimators. This efficient influence function 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 treatment and censoring mechanisms, 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 projected influence function (projected IC), which equals the EIC in a model where the missingness mechanism (or more generally, pre-modifier intervention, A_0) is known. Though not fully efficient under the larger model, this projected IC is robust against misspecification of the outcome models or the treatment 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 two TMLE estimators and the IPW estimators are asymptotically linear, thereby allowing Central Limit Theorem based standard error estimates. Moreover, the full TMLE estimator using the EIC will be semiparametric efficient if all the components of the likelihood are consistently estimated.

8 Acknowledgements

We thank Josh Schwab for his invaluable help in the implementing the *projected* TMLE. This work was supported by NIH grants RC4MH092737 (P.I. Z. Luo) and R01 A1074345-07 (P.I. M. van der Laan).

9 Appendix

9.1 Proof of remark 1

The first equality in the remark follows from definition of $\Psi(P)$ and choice of $m_\psi(a_0', \underline{\mathbf{a}}_1, v)$, the third equality is trivial. We only show the second one.

$$\sum_{(a_0', \underline{\mathbf{a}}_1, v) \in \mathscr{A}' \times \mathscr{A} \times \mathscr{V}} h(a_0', \underline{\mathbf{a}}_1, v) \phi(a_0', \underline{\mathbf{a}}_1, v) \gamma_{a_0', v}(P) \left\{ \frac{\eta_{a_0', \underline{\mathbf{a}}_1, v}(P)}{\gamma_{a_0', v}(P)} - m_{\Psi(P)}(a_0', \underline{\mathbf{a}}_1, v) \right\}$$

$$= -E_P \sum_{(a_0', \underline{\mathbf{a}}_1, v) \in \mathscr{A}' \times \mathscr{A} \times \mathscr{V}} h(a_0', \underline{\mathbf{a}}_1, v) \phi(a_0', \underline{\mathbf{a}}_1, v) \left\{ Q^V(v \mid A_0 = a_0', W) Q_{t=0}^{a_0', \underline{\mathbf{a}}_1}(v, W) - \eta_{a_0', \underline{\mathbf{a}}_1, v}(P) \right\}$$

$$+ E_P \sum_{(a_0', \underline{\mathbf{a}}_1, v) \in \mathscr{A}' \times \mathscr{A} \times \mathscr{V}} h(a_0', \underline{\mathbf{a}}_1, v) \phi(a_0', \underline{\mathbf{a}}_1, v) Q^V(v \mid A_0 = a_0', W) \left\{ Q_{t=0}^{a_0', \underline{\mathbf{a}}_1}(v, W) - m_{\Psi(P)}(a_0', \underline{\mathbf{a}}_1, v) \right\}$$

The first-line in the right-hand-side of the above equation is zero by definition of $\eta_{a_0', \underline{\mathbf{a}}_1, v}(P)$.

9.2 Proof of lemma 1: Efficient influence function for $\Psi(P)$ under $\mathscr{M}$

In this appendix, we derive the efficient influence function 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

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

From definition of the maps Ψ, $\gamma_{a_0', v}$ and $\eta_{a_0', \underline{\mathbf{a}}_1, v}$, and our choice of working model $m_\psi(a_0', \underline{\mathbf{a}}_1, v)$, we know that at each $P_r(\alpha)$

$$0 = \sum_{(a_0', \underline{\mathbf{a}}_1, v) \in \mathscr{A}' \times \mathscr{A} \times \mathscr{V}} \gamma_{a_0', v}(P_r(\alpha)) h(a_0', \underline{\mathbf{a}}_1, v) \phi(a_0', \underline{\mathbf{a}}_1, v) \left[\frac{\eta_{a_0', \underline{\mathbf{a}}_1, v}(P_r(\alpha))}{\gamma_{a_0', v}(P_r(\alpha))} - m_{\Psi(P_r(\alpha))}\left(a_0, \underline{\mathbf{a}}', v\right) \right]. \tag{18}$$

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

$$\left.\frac{d\Psi(P_r(\alpha))}{d\alpha}\right|_{\alpha=0} \times \tag{19}$$

$$\left\{ \sum_{(a_0',\underline{\mathbf{a}}_1,v)\in\mathscr{A}'\times\mathscr{A}\times\mathscr{V}} \gamma_{a_0',v}(P) h(a_0',\underline{\mathbf{a}}_1,v)\phi(a_0',\underline{\mathbf{a}}_1,v) m_{\Psi(P)}(a_0',\underline{\mathbf{a}}_1,v)\left[1-m_{\Psi(P)}(a_0',\underline{\mathbf{a}}_1,v)\right]\phi(a_0',\underline{\mathbf{a}}_1,v)^\top \right\}$$

$$= \sum_{(a_0',\underline{\mathbf{a}}_1,v)\in\mathscr{A}'\times\mathscr{A}\times\mathscr{V}} h(a_0',\underline{\mathbf{a}}_1,v)\phi(a_0',\underline{\mathbf{a}}_1,v)\left\{ \left.\frac{d\eta_{a_0',\underline{\mathbf{a}}_1,v}(P_r(\alpha))}{d\alpha}\right|_{\alpha=0} - \left(\left.\frac{d\gamma_{a_0',v}(P_r(\alpha))}{d\alpha}\right|_{\alpha=0} \right) m_{\Psi(P)}(a_0',\underline{\mathbf{a}}_1,v) \right\}. \tag{20}$$

Next, we proceed to express $\left.\frac{d\gamma_{a_0',v}(P_r(\alpha))}{d\alpha}\right|_{\alpha=0}$ and $\left.\frac{d\eta_{a_0',\underline{\mathbf{a}}_1,v}(P_r(\alpha))}{d\alpha}\right|_{\alpha=0}$ as $E_P(D_{\gamma_{a_0',v}} \times r)$ and $E_P(D_{\rho_{a_0',\underline{\mathbf{a}}_1,v}} \times r)$, respectively, for some functions $D_{\gamma_{a_0',v}}$ and $D_{\eta_{a_0',\underline{\mathbf{a}}_1,v}}$ belonging to the Hilbert space $\mathscr{H}_P$.

For convenience of indexing, for a given vector ℓ_K, we shall use the short hand $h_t^{a_0',\underline{\mathbf{a}}_1,v} = (A_0 = a_0', V = v, \mathbf{L}_{t-1} = \ell_{t-1}, \mathbf{A}_t = \mathbf{a}_t)$, for $t = 1,\ldots,K$, and $h_t^{a_0',\underline{\mathbf{a}}_1,v} = (A_0 = a_0', V = v)$ for $t = 0$. From definition of $P_r(\alpha)$, it follows that

$$P_r(\alpha)(W) = P(W)\left(1+\alpha E_P(r \mid W)\right),$$

$$P_r(\alpha)(v \mid A_0 = a_0', W) = P(v \mid A_0 = a_0', W)\frac{1+\alpha E_P(r \mid V = v, A_0 = a_0', W)}{1+\alpha E_P(r \mid A_0 = a_0', W)},$$

and $P_r(\alpha)(l_t \mid W, H_t^{a_0',\underline{\mathbf{a}}_1,v}) = P(l_t \mid W, H_t^{a_0',\underline{\mathbf{a}}_1,v})\frac{1+\alpha E_P\left(r|l_t,H_t^{a_0',\underline{\mathbf{a}}_1,v}\right)}{1+\alpha E_P\left(r|H_t^{a_0',\underline{\mathbf{a}}_1,v}\right)}$.

Proposition 1. *For a given $v \in \mathscr{V}$, $\left.\frac{d\gamma_{a_0',v}(P_r(\alpha))}{d\alpha}\right|_{\alpha=0} = E_P(D_{\gamma_{a_0',v}} \times r)$, where*

$$D_{\gamma_{a_0',v}}(P) = \frac{I(A_0 = a_0')}{g^{A_0}(1 \mid W)}\left(I(V = v) - Q^V(v \mid A_0 = a_0', W)\right) + Q^V(v \mid A_0 = a_0', W) - \gamma_{a_0',v}(P). \tag{21}$$

Proof.

$$\left.\frac{d\gamma_{a_0',v}(P_r(\alpha))}{d\alpha}\right|_{\alpha=0} = \lim_{\alpha\to 0}\frac{\gamma_{a_0',v}(P_r(\alpha))-\gamma_{a_0',v}(P)}{\alpha} = \lim_{\alpha\to 0}\frac{1}{\alpha}\int_W \left\{ P(W)P(v\mid A_0=a_0',W)\alpha \right.$$
$$\left. \frac{E_P(r\mid v,A_0=a_0',W)-E_P(r\mid A_0=a_0',W)+E_P(r\mid W)+\alpha E_P(r\mid W)E_P(r\mid v,A_0=a_0',W)}{1+\alpha E_P(r\mid A_0=a_0',W)} \right\}$$
$$= \int_W P(W)P(vmidA_0=a_0',W)\left\{E_P(rI(V=v)\mid A_0=a_0',W)-E_P(r\mid A_0=a_0',W)+E_P(r\mid W)\right\}$$
$$= E_P\left\{\left(\frac{I(A_0=a_0')}{g^{A_0}(1\mid W)}\left(I(V=v)-P(V\mid A_0=a_0',W)\right)+Q^V(v\mid A_0=a_0',W)\right)\times r(O)\right\}$$
$$= E_P\left(D_{\gamma_v}(P)\times r(O)\right).$$

In obtaining the last equality, we note that centering the left factor of the integrand by $\gamma_{a_0',v}(P)$ does not change the expression because $E_P(\gamma_{a_0',v}(P)r(O)) = \gamma_{a_0',v}(P)E_P(r(O)) = 0$ by definition of r. It is straightforward to check that indeed $E_P D^*_{\gamma_{a_0',v}}(P) = 0$. Moreover, under our saturated model $D^*_{\gamma_{a_0',v}}(P)$ is in fact in the tangent space. Therefore, it is indeed the efficient influence curve. This concludes the proof of proposition 1. □

Proposition 2. *For a given* $\mathbf{a}\in\mathscr{A}$, $v\in\mathscr{V}$, $\left.\frac{d\eta_{a_0',\underline{\mathbf{a}}_1,v}(P_r(\alpha))}{d\alpha}\right|_{\alpha=0} = E_P(D_{\eta_{a_0',\underline{\mathbf{a}}_1,v}}\times r)$, *where*

$$\begin{aligned}
&D_{\eta_{a_0',\underline{\mathbf{a}}_1,v}}(P)\\
&\equiv \frac{I(A_0=a_0')}{g^{A_0}(1\mid W)}I(V=v)\sum_{t=1}^{K}C_t^{\underline{\mathbf{a}}_1}\left\{Q_{t+1}^{a_0',\underline{\mathbf{a}}_1}(\mathbf{L}_t,v,W)-Q_t^{a_0',\underline{\mathbf{a}}_1}(\mathbf{L}_{t-1},v,W)\right\}\\
&+\frac{I(A_0=a_0')}{g^{A_0}(1\mid W)}I(V=v)\left\{Q_{t=1}^{a_0',\underline{\mathbf{a}}_1}(L_0,v,W)-Q_{t=0}^{a_0',\underline{\mathbf{a}}_1}(v,W)\right\}\\
&+\frac{I(A_0=a_0')}{g^{A_0}(1\mid W)}Q_{t=0}^{a_0',\underline{\mathbf{a}}_1}(v,W)\left\{I(V=v)-Q^V(v\mid A_0=a_0',W)\right\}\\
&+Q^V(v\mid A_0=a_0',W)Q_{t=0}^{a_0',\underline{\mathbf{a}}_1}(v,W)-\eta_{a_0',\underline{\mathbf{a}}_1,v}(P). \qquad (22)
\end{aligned}$$

where with $C_t^{\underline{\mathbf{a}}_1} = \frac{I(\mathbf{A}_t=\mathbf{a}_t)}{\prod_{j=1}^t g^A(a_j\mid \mathbf{A}_{j-1}=\mathbf{a}_{j-1},\mathbf{L}_{j-1},V,A_0=a_0',W}$

Proof.

$$\left.\frac{d\eta_{a_0',\underline{\mathbf{a}}_1,v}(P_r(\alpha))}{d\alpha}\right|_{\alpha=0} = \lim_{\alpha\to 0}\frac{\eta_{a_0',\underline{\mathbf{a}}_1,v}(P_r(\alpha))-\eta_{a_0',\underline{\mathbf{a}}_1,v}(P)}{\alpha}$$

$$= \lim_{\alpha\to 0}\frac{1}{\alpha}\int_{w_1,l_0,\ell_K} y_K P(W)P(vmidA_0=a_0',W)\prod_{j=0}^{K}P(l_j\mid w_1,h_j^{\mathbf{a},v})\alpha$$

$$\times\left\{\frac{\sum_{t=0}^{K}E_P(r\mid l_t,w_1,h_t^{a_0',\underline{\mathbf{a}}_1,v})+E_P(r\mid v,A_0=a_0',W)+E_P(r\mid W)}{1+\alpha M_P'(w_1,l_0,\ell_K)}\right.$$

$$\left.\frac{-\sum_{t=0}^{K}E_P(r\mid w_1,h_t^{a_0',\underline{\mathbf{a}}_1,v})-E_P(r\mid A_0=a_0',W)+\alpha M_P(w_1,l_0,\ell_K)}{1+\alpha M_P'(w_1,l_0,\ell_K)}\right\}$$

$$= \int_{w_1,l_0,\ell_K} y_K P(w_1)P(vmidA_0=a_0',w_1)\prod_{j=0}^{K}P(l_j\mid w_1,h_j^{\mathbf{a},v})$$

$$\times\left\{\sum_{t=0}^{K}E_P(r\mid l_t,w_1,h_t^{a_0',\underline{\mathbf{a}}_1,v})-\sum_{t=0}^{K}E_P(r\mid w_1,h_t^{a_0',\underline{\mathbf{a}}_1,v})+E_P(r\mid v,A_0=a_0',w_1)\right.$$

$$\left.-E_P(r\mid A_0=a_0',w_1)+E_P(r\mid w_1)\right\}$$

$$= E_P\left\{\frac{I(A_0=a_0')}{g^{A_0}(1\mid W)}I(V=v)\sum_{t=0}^{K}\frac{I(\mathbf{A}_t=\mathbf{a}_t)}{\prod_{j=1}^{t}g^A(a_t\mid parents(a_t))}\right.$$

$$\left.\times\left[Q_{t+1}^{a_0',\underline{\mathbf{a}}_1}(\ell_t,v,w_1)-Q_t^{a_0',\underline{\mathbf{a}}_1}(\ell_{t-1},v,w_1)\right]\times r(O)\right\}$$

$$+E_P\left\{\frac{I(A_0=a_0')}{g^{A_0}(1\mid W)}Q_{t=0}^{a_0',\underline{\mathbf{a}}_1}(V=v,W)\left[I(V=v)-Q^V(v\mid A_0=a_0',W)\right]\times r(O)\right\}$$

$$+E_P\left\{Q^V(v\mid A_0=a_0',W)Q_{t=0}^{a_0',\underline{\mathbf{a}}_1}(V=v,W)\times r(O)\right\}$$

$$= E_P(D_{\eta_{a_0',\underline{\mathbf{a}}_1,v}}\times r)$$

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 2 □

Now, we derive the efficient influence function for Ψ at $P\in\mathscr{M}$. From proposition 1 and 2, after some simplifications, we conclude that the right-hand-side of (20) can be

written as

$$\sum_{(a_0',\underline{\mathbf{a}}_1,v)\in\mathscr{A}'\times\mathscr{A}\times\mathscr{V}} h(a_0',\underline{\mathbf{a}}_1,v)\phi(a_0',\underline{\mathbf{a}}_1,v)\left\{\left.\frac{d\eta_{a_0',\underline{\mathbf{a}}_1,v}(P_r(\alpha))}{d\alpha}\right|_{\alpha=0} - \left(\left.\frac{d\gamma_{a_0',v}(P_r(\alpha))}{d\alpha}\right|_{\alpha=0}\right) m_{\Psi(P)}(a_0',\underline{\mathbf{a}}_1,v)\right\}$$

$$= \sum_{(a_0',\underline{\mathbf{a}}_1,v)\in\mathscr{A}'\times\mathscr{A}\times\mathscr{V}} h(a_0',\underline{\mathbf{a}}_1,v)\phi(a_0',\underline{\mathbf{a}}_1,v)\left\{E_P(D_{\eta_{a_0',\underline{\mathbf{a}}_1,v}}(P)\times r) - E_P(D_{\gamma_v}(P)\times r)m_{\Psi(P)}(a_0',\underline{\mathbf{a}}_1,v)\right\}$$

$$= E_P\left\{\sum_{(a_0',\underline{\mathbf{a}}_1,v)\in\mathscr{A}'\times\mathscr{A}\times\mathscr{V}} \tilde{h}(a_0',\underline{\mathbf{a}}_1,v)\left[D_{\eta_{a_0',\underline{\mathbf{a}}_1,v}}(P) - D_{\gamma_v}(P)m_{\Psi(P)}(a_0',\underline{\mathbf{a}}_1,v)\right]\times r(O)\right\}$$

$$= E_P\left\{D^*(Q,g,\Psi(P))\times r\right\},$$

where

$$\begin{aligned}
D^*(Q,g,\psi) &= \frac{I(A_0=a_0')}{g^{A_0}(1\mid W)}\sum_{\mathbf{a}}\tilde{h}(a_0',\underline{\mathbf{a}}_1,V)\sum_{t=1}^{K} C_t^{\underline{\mathbf{a}}_1}\left\{Q_{t+1}^{a_0',\underline{\mathbf{a}}_1}(\mathbf{L}_t,V,W) - Q_t^{a_0',\underline{\mathbf{a}}_1}(\mathbf{L}_{t-1},V,W)\right\}\\
&+ \frac{I(A_0=a_0')}{g^{A_0}(1\mid W)}\sum_{\mathbf{a}}\tilde{h}(a_0',\underline{\mathbf{a}}_1,V)\left\{Q_{t=1}^{a_0',\underline{\mathbf{a}}_1}(L_0,V,W) - Q_{t=0}^{a_0',\underline{\mathbf{a}}_1}(V,W)\right\}\\
&+ \frac{I(A_0=a_0')}{g^{A_0}(1\mid W)}\sum_{(a_0',\underline{\mathbf{a}}_1,v)\in\mathscr{A}'\times\mathscr{A}\times\mathscr{V}}\left\{\tilde{h}(a_0',\underline{\mathbf{a}}_1,V)\left[Q_{t=0}^{a_0',\underline{\mathbf{a}}_1}(v,W) - m_\psi(a_0',\underline{\mathbf{a}}_1,v)\right]\right.\\
&\qquad\left.\times\left(I(V=v) - Q^V(v\mid A_0=a_0',W)\right)\right\}\\
&+ \sum_{(a_0',\underline{\mathbf{a}}_1,v)\in\mathscr{A}'\times\mathscr{A}\times\mathscr{V}}\tilde{h}(a_0',\underline{\mathbf{a}}_1,V)Q^V(v\mid A_0=a_0',W)\left\{Q_{t=0}^{a_0',\underline{\mathbf{a}}_1}(v,W) - m_\psi(a_0',\underline{\mathbf{a}}_1,v)\right\}.
\end{aligned}$$

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 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 first part of lemma 1, it suffices to see from (20) that the normalizing quantity is given by the inverse of

$$\frac{\partial U(\Psi(P),P)}{\partial\Psi(P)} = \sum_{(a_0',\underline{\mathbf{a}}_1,v)\in\mathscr{A}'\times\mathscr{A}\times\mathscr{V}}\gamma_{a_0',v}(P)h(a_0',\underline{\mathbf{a}}_1,v)\phi(a_0',\underline{\mathbf{a}}_1,v)m_{\Psi(P)}(a_0',\underline{\mathbf{a}}_1,v)\left[1-m_{\Psi(P)}(a_0',\underline{\mathbf{a}}_1,v)\right]\phi(a_0',\underline{\mathbf{a}}_1,v)^\top$$

$$= \int_W P(W)\sum_{v,\mathbf{a}} Q^V(v\mid A_0=a_0',W)h(a_0',\underline{\mathbf{a}}_1,v)\phi(a_0',\underline{\mathbf{a}}_1,v)m_{\Psi(P)}(a_0',\underline{\mathbf{a}}_1,v)\left[1-m_{\Psi(P)}(a_0',\underline{\mathbf{a}}_1,v)\right]\phi(a_0',\underline{\mathbf{a}}_1,v)^\top$$

$$= P\left\{\frac{I(A_0=a_0')}{g^{A_0}(1\mid W)}\sum_{\mathbf{a}}h(a_0',\underline{\mathbf{a}}_1,v)\phi(a_0',\underline{\mathbf{a}}_1,V)m_{\Psi(P)}(a_0',\underline{\mathbf{a}}_1,v)\left[1-m_{\Psi(P)}(a_0',\underline{\mathbf{a}}_1,v)\right]\phi(a_0',\underline{\mathbf{a}}_1,V)^\top\right\}$$

$$= M(\Psi(P),P)$$

To prove the robustness statement, 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 $P_0 D^*(Q_0, g(P), \psi_0)$ is $\int_{w_1} P_0(w_1) \sum_{(a_0', \underline{\mathbf{a}}_1, v) \in \mathscr{A}' \times \mathscr{A} \times \mathscr{V}} P_0(v \mid A_0 = a_0', w_1) \tilde{h}(a_0', \underline{\mathbf{a}}_1, v) \left\{ \frac{\eta_{a_0', \underline{\mathbf{a}}_1, v}(Q_0)}{\gamma_{a_0', v}(Q_0)} - m_{\psi_0}(a_0', \underline{\mathbf{a}}_1, v) \right\}$, which equals zero by definition of ψ_0.

Next, we consider the case $g(P) = g_0$. By telescoping the sums in (13) and applying definition of $Q_{t=0}^{a_0', \underline{\mathbf{a}}_1}(P_0)$, we obtain

$$
\begin{aligned}
& P_0 D^*(Q(P), g_0, \psi_0) \\
& = E_{P_0} \left\{ \sum_{(a_0', \underline{\mathbf{a}}_1, v) \in \mathscr{A}' \times \mathscr{A} \times \mathscr{V}} P_0(V = v \mid (A_0 = a_0', W) \tilde{h}(a_0', \underline{\mathbf{a}}_1, v) Q_{t=0}^{a_0', \underline{\mathbf{a}}_1}(P_0)(V = v, W) \right\} \\
& - E_{P_0} \left\{ \sum_{(a_0', \underline{\mathbf{a}}_1, v) \in \mathscr{A}' \times \mathscr{A} \times \mathscr{V}} P_0(V = v \mid (A_0 = a_0', W) \tilde{h}(a_0', \underline{\mathbf{a}}_1, v) m_{\psi_0}(a_0', \underline{\mathbf{a}}_1, v) \right\},
\end{aligned}
$$

which equals 0 by remark 1.

9.3 Proof of lemma 2

To see that D^{A_0} is a score, up to a normalizing matrix, on the model where g^{A_0} is known, we repeat the steps in section 9.2, but this time we only consider the the class of 1-dimensional submodels indexed by

$$
R_P = \left\{ r = r' - E_P(r' \mid A_0, W) + E_P(r' \mid W) - E_P(r') : r' \in \mathscr{H}_P \right\} \subset \mathscr{H}_P.
$$

It is straight forward to verify that $E_P(r \mid A_0 = a_0', W) = 0 = E_P(r \mid W)$ for such $r \in R_P$. Therefore, $D_{\gamma_{a_0', v}(P)}$ in this case becomes $\frac{I(A_0 = a_0')}{g^{A_0}(1 \mid W)} I(V = v) - \gamma_{a_0', v}(P)$ and

$$
\begin{aligned}
D_{\eta_{a_0', \underline{\mathbf{a}}_1, v}}(P) = & \frac{I(A_0 = a_0')}{g^{A_0}(1 \mid W)} I(V = v) \sum_{t=1}^{K} C_t^{\underline{\mathbf{a}}_1} \left\{ Q_{t+1}^{a_0', \underline{\mathbf{a}}_1}(\mathbf{L}_t, v, W) - Q_t^{a_0', \underline{\mathbf{a}}_1}(\mathbf{L}_{t-1}, v, W) \right\} \\
& + \frac{I(A_0 = a_0')}{g^{A_0}(1 \mid W)} I(V = v) \left\{ Q_{t=1}^{a_0', \underline{\mathbf{a}}_1}(L_0, v, W) - Q_{t=0}^{a_0', \underline{\mathbf{a}}_1}(v, W) \right\} \\
& + \frac{I(A_0 = a_0')}{g^{A_0}(1 \mid W)} I(V = v) Q_{t=0}^{a_0', \underline{\mathbf{a}}_1}(v, W) - \eta_{a_0', \underline{\mathbf{a}}_1, v}.
\end{aligned}
$$

The rest is straightforward from equation (18).

To see that $PD^{A_0}(Q,g,\Psi(P))=0$, it suffices to show that $PD_W^{A_0}(Q,g,\Psi(P))=0$:

$$
\begin{aligned}
&P\left\{\frac{I(A_0=a_0')}{g^{A_0}(1\mid W)}\sum_{\mathbf{a}}h(a_0',\underline{\mathbf{a}}_1,V)\phi(a_0',\underline{\mathbf{a}}_1,V)\left[Q_{t=1}^{a_0',\underline{\mathbf{a}}_1}(L_0,V,W)-m_{\Psi(P)}(a_0',\underline{\mathbf{a}}_1,V)\right]\right\}\\
&=P\left\{\frac{I(A_0=a_0')}{g^{A_0}(1\mid W)}\sum_{\mathbf{a}}h(a_0',\underline{\mathbf{a}}_1,V)\phi(a_0',\underline{\mathbf{a}}_1,V)\left\{Q_{t=1}^{a_0',\underline{\mathbf{a}}_1}(L_0,V,W)-Q_{t=0}^{a_0',\underline{\mathbf{a}}_1}(V,W)\right\}\right\}\\
&+P\left\{\frac{I(A_0=a_0')}{g^{A_0}(1\mid W)}\sum_{\mathbf{a}}h(a_0',\underline{\mathbf{a}}_1,V)\phi(a_0',\underline{\mathbf{a}}_1,V)\left[Q_{t=0}^{a_0',\underline{\mathbf{a}}_1}(V,W)-m_{\Psi(P)}(a_0',\underline{\mathbf{a}}_1,V)\right]\right\}.
\end{aligned}
$$

The first term in the right hand side of the equality is zero by definition of $Q_1^{a_0',\underline{\mathbf{a}}_1}$; the second term is zero by remark 1. There D^{A_0} is a valid estimating function.

To see its robustness properties under $g^{A_0}=g^{A_0}(P_0)$, first consider the case of $Q_t^{a_0',\underline{\mathbf{a}}_1}(P)=Q_t^{a_0',\underline{\mathbf{a}}_1}(P_0)$ for $t=1,\ldots,K$. Trivially, $P_0D_t^*(Q,g)=0$ for each $t\geq 1$ by definition of $Q_t^{a_0',\underline{\mathbf{a}}_1}$. On the term $D_W^{A_0}$, we have

$$
\begin{aligned}
&P_0\left\{\frac{I(A_0=a_0')}{g^{A_0}(P_0)(1\mid W)}\sum_{\mathbf{a}}\tilde{h}(a_0',\underline{\mathbf{a}}_1,V)\left[Q_{t=1}^{a_0',\underline{\mathbf{a}}_1}(P_0)(L_0,V,W)-m_{\psi_0}(a_0',\underline{\mathbf{a}}_1,V)\right]\right\}\\
&=P_0\left\{\sum_v Q^V(P_0)(v\mid A_0=a_0',W)\sum_{\mathbf{a}}\tilde{h}(a_0',\underline{\mathbf{a}}_1,V)\left[Q_{t=0}^{a_0',\underline{\mathbf{a}}_1}(P_0)(V=v,W)-m_{\psi_0}(a_0',\underline{\mathbf{a}}_1,V)\right]\right\}=0,
\end{aligned}
$$

per definition of $Q_{t=0}^{a_0',\underline{\mathbf{a}}_1}$ and ψ_0.

Next, we consider the case of $g(P)=g_0$. By telescoping the sums in (14) and applying definition of $Q_{t=0}^{a_0',\underline{\mathbf{a}}_1}$, we again obtain

$$
\begin{aligned}
P_0D^{A_0}(Q,g_0,\psi_0)=P_0&\left\{\sum_{(a_0',\underline{\mathbf{a}}_1,v)\in\mathscr{A}'\times\mathscr{A}\times\mathscr{V}}P_0(V=v\mid(A_0=a_0',W)\tilde{h}(a_0',\underline{\mathbf{a}}_1,v)Q_{t=0}^{a_0',\underline{\mathbf{a}}_1}(P_0)(V=v,W)\right\}\\
&-P_0\left\{\sum_{(a_0',\underline{\mathbf{a}}_1,v)\in\mathscr{A}'\times\mathscr{A}\times\mathscr{V}}P_0(V=v\mid(A_0=a_0',W)\tilde{h}(a_0',\underline{\mathbf{a}}_1,v)m_{\psi_0}(a_0',\underline{\mathbf{a}}_1,v)\right\}=0,
\end{aligned}
$$

This completes proof of lemma 2.

References

H. Bang and J. M. Robins. Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61:962–972, 2005.

P.J. Bickel, C.A.J. Klaassen, Y. Ritov, and J. Wellner. *Efficient and Adaptive Estimation for Semiparametric Models*. Springer-Verlag, 1997.

A. Ellis, S. Dusetzina, R. Hansen, B. Gaynes, J. Farley, and T. Stumer. Investigating differences in treatment effect estimates between propensity score matching and weighting: a demonstration using STAR*D trial data. *Pharmacoepidemiology and Drug Safety*, 22: 138–144, 2013.

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.

J. Pearl. *Causality: Models, Reasoning and Inference*. Cambridge University Press, New York, 2nd edition, 2009.

M. Petersen, S. Deeks, J. Martin, and M J van der Laan. History-adjusted marginal structural models for estimating time-varying effect modification. *American Journal of Epidemiology*, 166(9), 2007a.

M. Petersen, S. Deeks, and M J van der Laan. Individualized treatment rules: generating candidate clinical trials. *Statistics in Medicine*, 26(25), 2007b.

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`.

J. M. Robins. Marginal structural models. *Proceedings of the American Statistical Association. Section on Bayesian Statistical Science*, pages 1–10, 1997a.

J. M. Robins, M. A. Hernan, and B. Brumback. Marginal structural models and causal inference in epidemiology. *Epidemiology*, 11(5):550–560, 2000.

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.

M J van der Laan and M. Petersen. Statistical learning of origin-specific statically optimal individualized treatment rules. *Int. J. Biostat.*, 3(1), 2007.

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 and D.B. Rubin. Targeted maximum likelihood learning. *The International Journal of Biostatistics*, 2(1), 2006.

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, S. Dudoit, and M.J. van der Laan. Oracle inequalities for multi-fold cross-validation. *Statistics and Decisions*, 24(3):351–371, 2006.