

Estimation and Optimization of Composite Outcomes

Daniel J. Lockett

*Genospace
Boston, MA 02108, USA*

DANIEL.LUCKETT.WORK@GMAIL.COM

Eric B. Laber

*Department of Statistical Science
Duke University
Durham, NC 27708, USA*

ERIC.LABER@DUKE.EDU

Siyeon Kim

*Department of Biostatistics
University of North Carolina at Chapel Hill
Chapel Hill, NC 27607, USA*

SIYEONKK@LIVE.UNC.EDU

Michael R. Kosorok

*Departments of Biostatistics and Statistics & Operations Research
University of North Carolina at Chapel Hill
Chapel Hill, NC 27599, USA*

KOSOROK@UNC.EDU

Editor: George Konidaris

Abstract

There is tremendous interest in precision medicine as a means to improve patient outcomes by tailoring treatment to individual characteristics. An individualized treatment rule formalizes precision medicine as a map from patient information to a recommended treatment. A treatment rule is defined to be optimal if it maximizes the mean of a scalar outcome in a population of interest, e.g., symptom reduction. However, clinical and intervention scientists often seek to balance multiple and possibly competing outcomes, e.g., symptom reduction and the risk of an adverse event. One approach to precision medicine in this setting is to elicit a composite outcome which balances all competing outcomes; unfortunately, eliciting a composite outcome directly from patients is difficult without a high-quality instrument, and an expert-derived composite outcome may not account for heterogeneity in patient preferences. We propose a new paradigm for the study of precision medicine using observational data that relies solely on the assumption that clinicians are approximately (i.e., imperfectly) making decisions to maximize individual patient utility. Estimated composite outcomes are subsequently used to construct an estimator of an individualized treatment rule which maximizes the mean of patient-specific composite outcomes. The estimated composite outcomes and estimated optimal individualized treatment rule provide new insights into patient preference heterogeneity, clinician behavior, and the value of precision medicine in a given domain. We derive inference procedures for the proposed estimators under mild conditions and demonstrate their finite sample performance through a suite of simulation experiments and an illustrative application to data from a study of bipolar depression.

Keywords: Individualized treatment rules, Inverse reinforcement learning, Precision medicine, Utility functions

1. Introduction

Precision medicine is an approach to healthcare that involves tailoring treatment based on individual patient characteristics (Hamburg and Collins, 2010; Collins and Varmus, 2015). Accounting for heterogeneity by tailoring treatment has the potential to improve patient outcomes in many therapeutic areas. An individualized treatment rule formalizes precision medicine as a map from the space of patient covariates into the space of allowable treatments (Murphy, 2003; Robins, 2004). Almost all methods for estimating individualized treatment rules have been designed to optimize a scalar outcome (exceptions will be discussed shortly). However, in practice, clinical decision making often requires balancing trade-offs between multiple outcomes. For example, clinicians treating patients with bipolar disorder must manage both depression and mania. Antidepressants may help correct depressive episodes but may also induce manic episodes (Sachs et al., 2007; Ghaemi, 2008; Goldberg, 2008; Wu et al., 2015). We propose a novel framework for using observational data to estimate a composite outcome and the corresponding optimal individualized treatment rule.

The estimation of optimal individualized treatment rules has been studied extensively, leading to a wide range of estimators designed to suit an array of data structures and data-generating processes (Kosorok and Laber, 2019; Tsiatis et al., 2020). These estimators include: regression-based methods like Q-learning (Murphy, 2005; Qian and Murphy, 2011; Schulte et al., 2014; Laber et al., 2014a), A-learning (Murphy, 2003; Robins, 2004; Blatt et al., 2004; Moodie et al., 2007; Wallace and Moodie, 2015), and regret regression (Henderson et al., 2010); direct-search methods (Rubin and van der Laan, 2012; Zhang et al., 2012b; Zhao et al., 2012; Zhang et al., 2013; Zhou et al., 2017) based on inverse probability weighting (Robins, 1999; Murphy et al., 2001; van der Laan and Petersen, 2007; Robins et al., 2008); and hybrid methods (Taylor et al., 2015; Zhang et al., 2018). The preceding methods require specification of a single scalar outcome that will be used to define an optimal regime; were individual patient utilities known, then they could be used as the outcome in any of these methods. However, in general such utilities are not known though they can be elicited provided a high-quality instrument is available (Butler et al., 2018); in the absence of such an instrument, preference elicitation is difficult to apply. A method for constructing a composite utility that is best predicted using a non-parametric machine learning model is proposed by (Benkeser et al., 2020); however, they do not consider heterogeneous utilities or the construction of precision medicine strategies.¹

We propose a new paradigm for estimating optimal individualized treatment rules from observational data without eliciting patient preferences. The key premise is that clinicians are attempting to act optimally with respect to each patient’s utility and thus the observed treatment decisions contain information about individual patient utilities. This idea is similar to that introduced by Wallace et al. (2018) (see also Wallace et al., 2016); however, we provide an estimator for the probability that a patient is treated optimally, rather than assuming that all patients are treated optimally. We construct estimators of individual patient utilities which do not require that clinicians are acting optimally, only that they approximately follow an optimal policy. This approach allows us to describe the goals of the decision maker and how these goals vary across patients, determine what makes a patient

1. Nevertheless, one could imagine how these flexible estimators could be integrated into our framework to reduce dependence on parametric models. See concluding remarks for additional discussion.

more or less likely to be treated optimally under standard care, and estimate the decision rule which optimizes patient-specific composite outcomes. We develop this approach in the context of a single-stage, binary decision in the presence of two outcomes. An extension to the setting with more than two outcomes is discussed in the Appendix.

Other methods for estimating optimal treatment rules in presence of multiple outcomes include using an expert-derived composite outcome for all patients (Thall et al., 2002, 2007; Murray et al., 2016; Moser et al., 2020). However, this does not account for differences in the utility function across patients and in some cases it may not be possible to elicit a high-quality composite outcome from an expert. Alternatively, multiple outcomes can be incorporated using set-valued treatment regimes (Laber et al., 2014b; Lizotte and Laber, 2016; Wu, 2016), constrained optimization (Linn et al., 2015; Laber et al., 2018; Wang et al., 2018), or inverse preference elicitation (Lizotte et al., 2012). Schnell et al. (2017) extend methods for estimating the benefiting subgroup to the case of multiple outcomes using the concept of admissibility (see also Schnell et al., 2016). However, none of these approaches provide a method for estimating an individual patient’s utility.

This work is closely related to inverse reinforcement learning (Kalman, 1964; Ng et al., 2000; Abbeel and Ng, 2004; Ratliff et al., 2006), which involves studying decisions made by an expert and constructing the utility function that is optimized by the expert’s decisions. Inverse reinforcement learning has been successfully applied in navigation (Ziebart et al., 2008) and human locomotion (Mombaur et al., 2009). Inverse reinforcement learning methods assume that decisions are made in a single environment. However, in the context of precision medicine, both the utility function and the probability of optimal treatment may vary across patients. Our approach is a version of inverse reinforcement learning with multiple environments.

This work is also related to the notion of stated and revealed preferences in the health economics literature.² Viewed through this lens, our work might be characterized as using clinical decisions as a kind of surrogate for patient revealed preferences thereby avoiding the need for the elicitation of stated preferences using specialized instruments. This is advantageous as the construction of high-quality instruments is difficult and collection of preference information is not routine in many areas (Carlsson and Martinsson, 2003; Ryan et al., 2007; de Bekker-Grob et al., 2012; Soekhai et al., 2019); though see Butler et al. (2018) for an illustrative application when such an instrument is available. Challenges associated with preference elicitation for precision medicine are discussed in Laber et al. (2014b), Lizotte and Laber (2016).

In Section 2, we introduce a pseudo-likelihood method to estimate patient utility functions from observational data. In Section 3, we state a number of theoretical results pertaining to the proposed method, including consistency and inference for the maximum pseudo-likelihood estimators. Section 4 presents a series of simulation experiments used to evaluate the finite sample performance of the proposed methods. Section 5 presents an illustrative application using data from the STEP-BD bipolar disorder study. Conclusions and a discussion of future research are given in Section 6. Proofs are given in the appendix along with additional simulation results and a discussion of an extension to more than two outcomes.

2. We gratefully acknowledge an anonymous referee for identifying this connection.

2. Pseudo-likelihood Estimation of Utility Functions

Assume the available data are $(\mathbf{X}_i, A_i, Y_i, Z_i)$, $i = 1, \dots, n$, which comprise n independent and identically distributed copies of (\mathbf{X}, A, Y, Z) , where $\mathbf{X} \in \mathcal{X} \subseteq \mathbb{R}^p$ are patient covariates, $A \in \mathcal{A} = \{-1, 1\}$ is a binary treatment, and Y and Z are two real-valued outcomes for which higher values are more desirable. The extension to scenarios with more than two outcomes is discussed in the Appendix. An individualized treatment rule is a function $d : \mathcal{X} \rightarrow \mathcal{A}$ such that, under d , a patient presenting with covariates $\mathbf{X} = \mathbf{x}$ will be assigned to treatment $d(\mathbf{x})$. Let $Y^*(a)$ denote the potential outcome under treatment $a \in \mathcal{A}$, and for any regime d , define $Y^*(d) = \sum_{a \in \mathcal{A}} Y^*(a) 1\{d(\mathbf{X}) = a\}$. An optimal regime for the outcome Y , say d_Y^{opt} , satisfies $\mathbb{E}Y^*(d_Y^{\text{opt}}) \geq \mathbb{E}Y^*(d)$ for any other regime d . The optimal regime for the outcome Z , say d_Z^{opt} , is defined analogously. In order to identify these optimal regimes, and subsequently to identify the optimal regime across the class of utility functions introduced below, we make the following assumptions.

Assumption 1 *Consistency*, $Y = Y^*(A)$ and $Z = Z^*(A)$.

Assumption 2 *Positivity*, $\Pr(A = a | \mathbf{X} = \mathbf{x}) \geq c > 0$ for some constant c and all pairs $(\mathbf{x}, a) \in \mathcal{X} \times \mathcal{A}$.

Assumption 3 *Ignorability*, $\{Y^*(-1), Y^*(1)\} \perp A | \mathbf{X}$ and $\{Z^*(-1), Z^*(1)\} \perp A | \mathbf{X}$.

In addition we assume that there is no interference between units nor are the multiple versions of treatment (Rubin, 1980). These assumptions are standard in causal inference (Robins, 2004; Hernan and Robins, 2010). Assumption 3 is not empirically verifiable in observational studies (Rosenbaum and Rubin, 1983; Rosenbaum, 1984).

Define $Q_Y(\mathbf{x}, a) = \mathbb{E}(Y | \mathbf{X} = \mathbf{x}, A = a)$. Then, under the preceding assumptions, it can be shown that $d_Y^{\text{opt}}(\mathbf{x}) = \arg \max_{a \in \mathcal{A}} Q_Y(\mathbf{x}, a)$ (Zhang et al., 2012b; Qian and Murphy, 2011). Similarly, it follows that $d_Z^{\text{opt}}(\mathbf{x}) = \arg \max_{a \in \mathcal{A}} Q_Z(\mathbf{x}, a)$ where $Q_Z(\mathbf{x}, a) = \mathbb{E}(Z | \mathbf{X} = \mathbf{x}, A = a)$. In general, $d_Y^{\text{opt}}(\mathbf{x})$ need not equal $d_Z^{\text{opt}}(\mathbf{x})$; therefore, if both Y and Z are clinically relevant, neither d_Y^{opt} nor d_Z^{opt} may be acceptable. We assume that there exists an unknown and possibly covariate-dependent utility $U = u(Y, Z)$, where $u : \mathbb{R}^2 \rightarrow \mathbb{R}$ measures the “goodness” of the outcome pair (y, z) . The optimal regime with respect to U , say d_U^{opt} , satisfies $\mathbb{E}U^*(d_U^{\text{opt}}) = \mathbb{E}u\left\{Y^*\left(d_U^{\text{opt}}\right), Z^*\left(d_U^{\text{opt}}\right)\right\} \geq \mathbb{E}u\{Y^*(d), Z^*(d)\} = \mathbb{E}U^*(d)$ for any other regime d . The goal is to use the observed data to estimate the utility and subsequently d_U^{opt} . Define $Q_U(\mathbf{x}, a) = \mathbb{E}(U | \mathbf{X} = \mathbf{x}, A = a)$. For the class of utility functions we consider below, $Q_U(\mathbf{x}, a)$ is a (possibly covariate-dependent) convex combination of $Q_Y(\mathbf{x}, a)$ and $Q_Z(\mathbf{x}, a)$ and is therefore identifiable under the stated causal assumptions and furthermore $d_U^{\text{opt}}(\mathbf{x}) = \arg \max_{a \in \mathcal{A}} Q_U(\mathbf{x}, a)$.

We assume that clinicians act with the goal of optimizing each patient’s utility and that their success in identifying the optimal treatment depends on individual patient characteristics. Therefore, we assume that the clinicians are approximately, i.e., imperfectly, assigning treatment according to $d_U^{\text{opt}}(\mathbf{x})$. If the clinician were always able to correctly identify the

optimal treatment and assign $A = d_U^{\text{opt}}(\mathbf{X})$ for each patient, there would be no need to estimate the optimal treatment policy (Wallace et al., 2016). Instead, we assume that the decisions of the clinician are imperfect and that $\Pr \left\{ A = d_U^{\text{opt}}(\mathbf{x}) | \mathbf{X} = \mathbf{x} \right\} = \text{expit}(\mathbf{x}^\top \beta)$ where β is an unknown parameter. We show in Section 2.2 that the model is identifiable under mild conditions; e.g., these exclude the possibility of a malevolent clinician that is systematically assigning poor treatments. We implicitly assume throughout that \mathbf{X} may contain higher order terms, interactions, or basis functions constructed from patient covariates.

2.1 Fixed Utility

We begin by assuming that the utility function is constant across patients and takes the form $u(y, z; \omega) = \omega y + (1 - \omega)z$ for some $\omega \in [0, 1]$. Lemma 1 of Butler et al. (2018) states that, for a broad class of utility functions, the optimal individualized treatment rule is equivalent to the optimal rule for a utility function of this form. Define $Q_\omega(\mathbf{x}, a) = \omega Q_Y(\mathbf{x}, a) + (1 - \omega)Q_Z(\mathbf{x}, a)$ and define $d_\omega^{\text{opt}}(\mathbf{x}) = \arg \max_{a \in \mathcal{A}} Q_\omega(\mathbf{x}, a)$. Let $\hat{Q}_{Y,n}$ and $\hat{Q}_{Z,n}$ denote estimators of Q_Y and Q_Z obtained from regression models fit to the observed data (Qian and Murphy, 2011). For a fixed value of ω , let $\hat{Q}_{\omega,n}(\mathbf{x}, a) = \omega \hat{Q}_{Y,n}(\mathbf{x}, a) + (1 - \omega) \hat{Q}_{Z,n}(\mathbf{x}, a)$ and subsequently let $\hat{d}_{\omega,n}(\mathbf{x}) = \arg \max_{a \in \mathcal{A}} \hat{Q}_{\omega,n}(\mathbf{x}, a)$ be the plug-in estimator of $d_\omega^{\text{opt}}(\mathbf{x})$. Given $\hat{Q}_{Y,n}$ and $\hat{Q}_{Z,n}$, $\hat{d}_{\omega,n}(\mathbf{x})$ can be computed for each $\omega \in [0, 1]$.

The joint distribution of (\mathbf{X}, A, Y, Z) is

$$\begin{aligned} f(\mathbf{X}, A, Y, Z) &= f(Y, Z | \mathbf{X}, A) f(A | \mathbf{X}) f(\mathbf{X}) \\ &= f(Y, Z | \mathbf{X}, A) f(\mathbf{X}) \frac{\exp \left[\mathbf{X}^\top \beta 1 \left\{ A = d_\omega^{\text{opt}}(\mathbf{X}) \right\} \right]}{1 + \exp(\mathbf{X}^\top \beta)}. \end{aligned}$$

Assuming that $f(Y, Z | \mathbf{X}, A)$ and $f(\mathbf{X})$ do not depend on ω or β , the likelihood for (ω, β) is

$$\mathcal{L}_n(\omega, \beta) \propto \prod_{i=1}^n \frac{\exp \left[\mathbf{X}_i^\top \beta 1 \left\{ A_i = d_\omega^{\text{opt}}(\mathbf{X}_i) \right\} \right]}{1 + \exp(\mathbf{X}_i^\top \beta)}, \quad (1)$$

which depends on the unknown function d_ω^{opt} . Plugging in $\hat{d}_{\omega,n}$ for d_ω^{opt} into (7) yields the pseudo-likelihood

$$\hat{\mathcal{L}}_n(\omega, \beta) \propto \prod_{i=1}^n \frac{\exp \left[\mathbf{X}_i^\top \beta 1 \left\{ A_i = \hat{d}_{\omega,n}(\mathbf{X}_i) \right\} \right]}{1 + \exp(\mathbf{X}_i^\top \beta)}. \quad (2)$$

If we let $\hat{\omega}_n$ and $\hat{\beta}_n$ denote the maximum pseudo-likelihood estimators obtained by maximizing (2), then an estimator of the utility function is $\hat{u}_n(y, z) = u(y, z; \hat{\omega}_n) = \hat{\omega}_n y + (1 - \hat{\omega}_n)z$ and $\text{expit}(\mathbf{x}^\top \hat{\beta}_n)$ is an estimator of the probability that a patient presenting with covariates \mathbf{x} would be treated optimally under standard care. An estimator of the optimal policy at \mathbf{x} is $\hat{d}_{\hat{\omega}_n,n}(\mathbf{x}) = \arg \max_{a \in \mathcal{A}} \hat{Q}_{\hat{\omega}_n,n}(\mathbf{x}, a)$.

Because the pseudo-likelihood given in (2) is non-smooth in ω , standard gradient-based optimization algorithms cannot be used. However, for a given ω , it is straightforward to

compute the profile estimator $\widehat{\beta}_n(\omega) = \arg \max_{\beta \in \mathbb{R}^p} \widehat{\mathcal{L}}_n(\omega, \beta)$. We can compute the profile pseudo-likelihood estimator over a grid of values for ω and select the point on the grid yielding the largest pseudo-likelihood. The algorithm to construct $(\widehat{\omega}_n, \widehat{\beta}_n)$ is given in Algorithm 1 below. Step (3) can be accomplished using logistic regression. The theoretical

Algorithm 1: Pseudo-likelihood estimation of fixed utility function.

- 1 Set a grid $0 = \omega_0 < \omega_1 < \dots < \omega_K = 1$;
 - 2 **for** $m = 0, \dots, K$ **do**
 - 3 | compute $\widehat{\beta}_n(\omega_m) = \arg \max_{\beta \in \mathbb{R}^p} \widehat{\mathcal{L}}_n(\omega_m, \beta)$;
 - 4 **end**
 - 5 Select $\widehat{m}_n = \arg \max_{0 \leq m \leq K} \widehat{\mathcal{L}}_n \left\{ \omega_m, \widehat{\beta}_n(\omega_m) \right\}$;
 - 6 Set $(\widehat{\omega}_n, \widehat{\beta}_n) = \left\{ \omega_{\widehat{m}_n}, \widehat{\beta}_n(\omega_{\widehat{m}_n}) \right\}$;
-

properties of this estimator are discussed in Section 3.

2.2 Patient-specific Utility

Outcome preferences can vary widely across patients in some application domains, including schizophrenia (Kinter, 2009; Strauss et al., 2010) and pain management (Gan et al., 2004). To accommodate this setting, we assume that the utility function takes the form $u(y, z; \mathbf{x}, \omega) = \omega(\mathbf{x})y + \{1 - \omega(\mathbf{x})\}z$ where $\omega : \mathcal{X} \rightarrow [0, 1]$ is a smooth function. For illustration, we let $\omega(\mathbf{x}; \theta) = \text{expit}(\mathbf{x}^\top \theta)$ where θ is an unknown parameter. Misspecified utility models are discussed in the Appendix. Define $Q_\theta(\mathbf{x}, a) = \omega(\mathbf{x}; \theta)Q_Y(\mathbf{x}, a) + \{1 - \omega(\mathbf{x}; \theta)\}Q_Z(\mathbf{x}, a)$ and define $d_\theta^{\text{opt}}(\mathbf{x}) = \arg \max_{a \in \mathcal{A}} Q_\theta(\mathbf{x}, a)$. Let $\widehat{Q}_{Y,n}$ and $\widehat{Q}_{Z,n}$ denote estimators of Q_Y and Q_Z obtained from regression models fit to the observed data. For a fixed value of θ , let $\widehat{Q}_{\theta,n}(\mathbf{x}, a) = \omega(\mathbf{x}; \theta)\widehat{Q}_{Y,n}(\mathbf{x}, a) + \{1 - \omega(\mathbf{x}; \theta)\}\widehat{Q}_{Z,n}(\mathbf{x}, a)$ and subsequently let $\widehat{d}_{\theta,n}(\mathbf{x}) = \arg \max_{a \in \mathcal{A}} \widehat{Q}_{\theta,n}(\mathbf{x}, a)$ be the plug-in estimator of $d_\theta^{\text{opt}}(\mathbf{x})$. Assume that decisions are made according to the model $\Pr \left\{ A = d_\theta^{\text{opt}}(\mathbf{x}) | \mathbf{X} = \mathbf{x} \right\} = \text{expit}(\mathbf{x}^\top \beta)$. We compute the estimators $(\widehat{\theta}_n, \widehat{\beta}_n)$ of (θ, β) by maximizing the pseudo-likelihood

$$\widehat{\mathcal{L}}_n(\theta, \beta) \propto \prod_{i=1}^n \frac{\exp \left[\mathbf{X}_i^\top \beta \mathbf{1} \left\{ A_i = \widehat{d}_{\theta,n}(\mathbf{X}_i) \right\} \right]}{1 + \exp(\mathbf{X}_i^\top \beta)}. \quad (3)$$

An estimator for the utility function is $\widehat{u}_n(y, z; \mathbf{x}) = \omega(\mathbf{x}; \widehat{\theta}_n)y + \{1 - \omega(\mathbf{x}; \widehat{\theta}_n)\}z$ and an estimator for the optimal decision function is $\widehat{d}_{\widehat{\theta}_n,n}$. The model, as stated is not identifiable. However, we show below that it is identifiable under the following conditions.

Assumption 4 *The following conditions hold.*

1. $\beta \in \mathcal{B} \subset \mathbb{R}^p$ and $\theta \in \Theta \subset \mathbb{R}^q$, where \mathcal{B} and Θ are compact.
2. $\beta_0 \neq 0$.

3. \mathcal{X} is bounded ($\mathbf{X} \in \mathcal{X} \subset \mathbb{R}^p$ a.s.).
4. Let \mathcal{X}_S be the collection of subsets of \mathcal{X} consisting of sets of the form $\{\mathbf{x} \in \mathcal{X} : d_\theta(\mathbf{x}) \neq d_{\theta_0}(\mathbf{x})\}$ for $\theta \in \Theta \setminus \{\theta_0\}$, together with the complements of these sets. Then:
 - (a) For all $X_S \in \mathcal{X}_S$, $0 < \Pr(\mathbf{X} \in X_S) < 1$, and
 - (b) $E(\mathbf{X}\mathbf{X}^T | \mathbf{X} \in X_S)$ is full rank $\forall X_S \in \mathcal{X}_S$.

Theorem 5 (Identifiability) Under Assumption 4, (θ_0, β_0) is uniquely identified under the model given by $\mathcal{L}_n(\theta, \beta)$.

Remark 6 A less technical but sufficient condition is to assume that (β_0, θ_0) satisfies $\Pr\{A = d_{\theta_0}(X) | \mathbf{X}\} > 1/2$ almost surely, i.e., that clinical decisions are always better than a coin toss. A proof of sufficiency is given in the Appendix.

As before, the pseudo-likelihood given in (3) is non-smooth in θ and standard gradient-based optimization methods cannot be used. It is again straightforward to compute the profile pseudo-likelihood estimator $\hat{\beta}_n(\theta) = \arg \max_{\beta \in \mathbb{R}^p} \hat{\mathcal{L}}_n(\theta, \beta)$ for any $\theta \in \mathbb{R}^p$. However, because it is computationally infeasible to compute $\hat{\beta}_n(\theta)$ for all θ on a grid if θ is of moderate dimension, we generate a random walk through the parameter space using the Metropolis algorithm as implemented in the `metrop` function in the R package `mcmc` (Geyer and Johnson, 2017) and compute the profile pseudo-likelihood for each θ on the random walk. Let $\tilde{\mathcal{L}}_n(\theta) = \max_{\beta \in \mathbb{R}^p} \hat{\mathcal{L}}_n(\theta, \beta)$. We can compute $\tilde{\mathcal{L}}_n(\theta) = \hat{\mathcal{L}}_n\{\theta, \hat{\beta}_n(\theta)\}$ by estimating $\hat{\beta}_n(\theta)$ using logistic regression as described in Section 2.1. The algorithm to construct a random walk through the parameter space is given in Algorithm 2 below. After generating a

Algorithm 2: Pseudo-likelihood estimation of patient-dependent utility function

- 1 Set a chain length, B , fix $\Sigma \succcurlyeq 0$, and initialize θ^1 to a starting value in \mathbb{R}^p ;
 - 2 **for** $b = 2, \dots, B$ **do**
 - 3 Generate $\mathbf{e} \sim N(0, \Sigma)$;
 - 4 Set $\tilde{\theta}^{b+1} = \theta^b + \mathbf{e}$;
 - 5 Compute $p = \min\left\{\tilde{L}_n(\tilde{\theta}^{b+1}) / \tilde{L}_n(\tilde{\theta}^b), 1\right\}$;
 - 6 Generate $U \sim U(0, 1)$; if $U \leq p$, set $\theta^{b+1} = \tilde{\theta}^{b+1}$; otherwise, set $\theta^{b+1} = \theta^b$;
 - 7 **end**
-

chain $(\theta^1, \dots, \theta^B)$, we select the θ^k that leads to the largest value of $\tilde{\mathcal{L}}_n(\theta^k)$ as the maximum pseudo-likelihood estimator. Standard practice is to choose the variance of the proposal distribution, σ^2 , so that the acceptance proportion is between 0.25 and 0.5 (Geyer and Johnson, 2017).

3. Theoretical Results

Here we state a number of theoretical results pertaining to the proposed pseudo-likelihood estimation method for utility functions. We state results for a patient-specific utility func-

tion; the setting where the utility function is fixed is a special case. All proofs are deferred to the Appendix.

We assume that $\Pr \left\{ A = d_U^{\text{opt}}(\mathbf{x}) \mid \mathbf{X} = \mathbf{x} \right\} = \text{expit}(\mathbf{x}^\top \beta_0)$ and that the true utility function is $u(y, z; \mathbf{x}, \theta_0) = \omega(\mathbf{X}; \theta_0) y + \{1 - \omega(\mathbf{X}; \theta_0)\} z$, where $\omega(\mathbf{X}; \theta)$ has bounded continuous derivative on compact sets and $d_{\theta_0}^{\text{opt}}(\mathbf{X}) = d_\theta^{\text{opt}}(\mathbf{X})$ almost surely implies $\theta = \theta_0$, i.e., the model introduced in Section 2.2 is well-defined and correctly specified with true parameters $\beta_0 \in \mathbb{R}^p$ and $\theta_0 \in \mathbb{R}^d$. We further assume that the estimators $\widehat{Q}_{Y,n}(\mathbf{x}, a)$ and $\widehat{Q}_{Z,n}(\mathbf{x}, a)$ are pointwise consistent for all ordered pairs (\mathbf{x}, a) . Along with Assumptions 1-3, we implicitly assume that the densities $f(Y, Z \mid \mathbf{X}, A)$ and $f(\mathbf{X})$ exist. The following result states the consistency of the maximum pseudo-likelihood estimators for the utility function and the probability of optimal treatment. The proof involves verifying the conditions of Theorem 2.12 of Kosorok (2008).

Theorem 7 (Consistency with patient-specific utility) *Let the maximum pseudo-likelihood estimators be as in Section 2.2, $(\widehat{\theta}_n, \widehat{\beta}_n) = \arg \max_{\theta \in \mathbb{R}^p, \beta \in \mathcal{B}} \widehat{\mathcal{L}}_n(\theta, \beta)$. Assume that \mathcal{B} is a compact set with $\beta_0 \in \mathcal{B}$ and that $\|\mathbb{E} \mathbf{X}\| < \infty$. Then, $\|\widehat{\theta}_n - \theta_0\| \xrightarrow{P} 0$ and $\|\widehat{\beta}_n - \beta_0\| \xrightarrow{P} 0$ as $n \rightarrow \infty$, where $\|\cdot\|$ is the Euclidean norm.*

Let $V_\theta(d) = \mathbb{E} \{u(Y, Z; \mathbf{X}, \theta) \mid A = d(\mathbf{X})\}$ be the mean composite outcome in a population where decisions are made according to d . The following result establishes the consistency of the value of the estimated optimal policy. The proof uses general theory developed by Qian and Murphy (2011).

Theorem 8 (Value consistency with patient-specific utility) *Let $\widehat{\theta}_n$ be the maximum pseudo-likelihood estimator for θ and let $\widehat{d}_{\widehat{\theta}_n, n}$ be the associated estimated optimal policy. Then, under the given assumptions, $\left| V_{\widehat{\theta}_n}(\widehat{d}_{\widehat{\theta}_n, n}) - V_{\theta_0}(d_{\theta_0}^{\text{opt}}) \right| \xrightarrow{P} 0$ as $n \rightarrow \infty$.*

Next, we derive the convergence rate and asymptotic distribution of $(\widehat{\theta}_n, \widehat{\beta}_n)$. Assume that \mathcal{X} is a bounded subset of \mathbb{R}^p and let $\|\cdot\|_{\mathcal{X}}$ be the sup norm over \mathcal{X} , i.e., for $f: \mathcal{X} \rightarrow \mathbb{R}$, $\|f\|_{\mathcal{X}} = \sup_{\mathbf{x} \in \mathcal{X}} |f(\mathbf{x})|$. Let $\dot{\omega}(\mathbf{x}) = (\partial/\partial \theta)\omega(\mathbf{x}; \theta)$. Assume that $\|\dot{\omega}_{\theta_0}(\mathbf{x})\|_{\mathcal{X}} < \infty$ and that $\lim_{\theta \rightarrow \theta_0} \|\dot{\omega}(\mathbf{x}) - \dot{\omega}_{\theta_0}(\mathbf{x})\|_{\mathcal{X}} = 0$. Define $R_Y(\mathbf{x}) = Q_Y(\mathbf{x}, 1) - Q_Y(\mathbf{x}, -1)$ to be the treatment contrast for outcome Y at patient covariates $\mathbf{X} = \mathbf{x}$; define $R_Z(\mathbf{x}) = Q_Z(\mathbf{x}, 1) - Q_Z(\mathbf{x}, -1)$ analogously. Let $R_0(\mathbf{x}) = R_Y(\mathbf{x}) - R_Z(\mathbf{x})$ denote the difference in the treatment contrasts across the two outcomes. Similarly, define $\widehat{R}_{Y,n}(\mathbf{x}) = \widehat{Q}_{Y,n}(\mathbf{x}, 1) - \widehat{Q}_{Y,n}(\mathbf{x}, -1)$, $\widehat{R}_{Z,n}(\mathbf{x}) = \widehat{Q}_{Z,n}(\mathbf{x}, 1) - \widehat{Q}_{Z,n}(\mathbf{x}, -1)$, and $\widehat{R}_{0,n}(\mathbf{x}) = \widehat{R}_{Y,n}(\mathbf{x}) - \widehat{R}_{Z,n}(\mathbf{x})$. Define $D_\theta(\mathbf{x}) = \omega(\mathbf{x}; \theta) R_Y(\mathbf{x}) + \{1 - \omega(\mathbf{x}; \theta)\} R_Z(\mathbf{x})$ to be the convex combination of treatment contrasts dictated by $\omega(\mathbf{x}; \theta)$ and let $\widehat{D}_{\theta,n}(\mathbf{x}) = \omega(\mathbf{x}; \theta) \widehat{R}_{Y,n}(\mathbf{x}) + \{1 - \omega(\mathbf{x}; \theta)\} \widehat{R}_{Z,n}(\mathbf{x})$. Note that $d_\theta^{\text{opt}}(\mathbf{x}) = \text{sign} \{D_\theta(\mathbf{x})\}$ and $\widehat{d}_{\theta,n}(\mathbf{x}) = \text{sign} \{\widehat{D}_{\theta,n}(\mathbf{x})\}$. Further define

$$\begin{aligned} P_\beta(\mathbf{x}) &= \text{expit}(\mathbf{x}^\top \beta), \\ \psi_{i,A} &= \left[1 \left\{ A_i = d_{\beta_0}^{\text{opt}}(\mathbf{X}_i) \right\} - P_{\beta_0}(\mathbf{X}_i) \right] \mathbf{X}_i, \\ I_n(\beta) &= \mathbb{E}_n [P_\beta(\mathbf{X}) \{1 - P_\beta(\mathbf{X})\} \mathbf{X} \mathbf{X}^\top], \\ I_0 &= \mathbb{E} [P_{\beta_0}(\mathbf{X}) \{1 - P_{\beta_0}(\mathbf{X})\} \mathbf{X} \mathbf{X}^\top]. \end{aligned}$$

We use the following regularity conditions.

Assumption 9 *There exist independent and identically distributed influence vectors $\psi_{1,Y}, \psi_{2,Y}, \dots \in \mathbb{R}^{q_1}$, and $\psi_{1,Z}, \psi_{2,Z}, \dots \in \mathbb{R}^{q_2}$, and vector basis functions $\phi_Y(\mathbf{x})$ and $\phi_Z(\mathbf{x})$ such that both*

$$\left\| \sqrt{n} \left\{ \widehat{R}_{Y,n}(\mathbf{x}) - R_Y(\mathbf{x}) \right\} - \phi_Y(\mathbf{x})^\top n^{-1/2} \sum_{i=1}^n \psi_{i,Y} \right\|_{\mathcal{X}} = o_P(1)$$

and

$$\left\| \sqrt{n} \left\{ \widehat{R}_{Z,n}(\mathbf{x}) - R_Z(\mathbf{x}) \right\} - \phi_Z(\mathbf{x})^\top n^{-1/2} \sum_{i=1}^n \psi_{i,Z} \right\|_{\mathcal{X}} = o_P(1).$$

Let $Z_{Y,n} = n^{-1/2} \sum_{i=1}^n \psi_{i,Y}$, $Z_{Z,n} = n^{-1/2} \sum_{i=1}^n \psi_{i,Z}$, $Z_{A,n} = n^{-1/2} \sum_{i=1}^n \psi_{i,A}$, and $q = q_1 + q_2$. Furthermore, assume that $\|R_Y(\mathbf{x})\|_{\mathcal{X}}$, $\|R_Z(\mathbf{x})\|_{\mathcal{X}}$, $\|\phi_Y(\mathbf{x})\|_{\mathcal{X}}$, and $\|\phi_Z(\mathbf{x})\|_{\mathcal{X}}$ are bounded by some $M < \infty$. Let $\Sigma_0 = \mathbb{E} \left[\left\{ \left(\psi_{1,Y}^\top, \psi_{1,Z}^\top, \psi_{1,A}^\top \right)^\top \right\}^{\otimes 2} \right]$ be positive definite and finite, where $u^{\otimes 2} = uu^\top$.

Assumption 10 *The following conditions hold.*

1. *The random variable $D_{\theta_0}(\mathbf{X})$ has a continuous density function f in a neighborhood of 0 with $f_0 = f(0) \in (0, \infty)$;*
2. *The conditional distribution of \mathbf{X} given that $|D_{\theta_0}(\mathbf{X})| \leq \epsilon$ converges to a non-degenerate distribution as $\epsilon \downarrow 0$;*
3. *There exist $\delta_1, \delta_2 > 0$ such that*

$$\liminf_{\epsilon \downarrow 0} \Pr \left[|\mathbf{X}^\top \beta_0| \geq \delta_1, |\{R_Y(\mathbf{X}) - R_Z(\mathbf{X})\} \dot{\omega}_{\theta_0}(\mathbf{X})^\top t| \geq \delta_1 \mid |D_{\theta_0}(\mathbf{X})| \leq \epsilon \right] \geq \delta_2,$$

where S^d is the d -dimensional unit sphere.

Assumption 11 *Define, for $Z_Y \in \mathbb{R}^{q_1}$, $Z_Z \in \mathbb{R}^{q_2}$, and $U \in \mathbb{R}^d$,*

$$(Z_Y, Z_Z, U) \mapsto k_0(Z_Y, Z_Z, U) = \mathbb{E} \left[\mathbf{X} \{2P_{\beta_0}(\mathbf{X}) - 1\} \cdot |\omega(\mathbf{X}; \theta_0) R_Y(\mathbf{X}) \phi_Y(\mathbf{X})^\top Z_Y + \{1 - \omega(\mathbf{X}; \theta_0)\} R_Z(\mathbf{X}) \phi_Z(\mathbf{X})^\top Z_Z + R_0(\mathbf{X}) \dot{\omega}_{\theta_0}(\mathbf{X})^\top U \mid D_{\theta_0}(\mathbf{X}) = 0 \right]. \quad (4)$$

Assume that $U \mapsto \beta_0^\top k_0(Z_Y, Z_Z, U)$ has a unique, finite minimum over \mathbb{R}^d for all $(Z_Y^\top, Z_Z^\top)^\top \in \mathbb{R}^q$.

Remark 12 *Assumption 9 establishes a rate of convergence for the estimated Q -functions and is automatically satisfied if the Q -functions are estimated using linear or generalized linear models with or without interactions or higher order terms. Assumption 10 is needed to ensure that there is positive probability of patients with \mathbf{x} values near the boundary between where each treatment is optimal. Assumption 11 is standard in M -estimation.*

Let $(\widehat{\theta}_n, \widehat{\beta}_n)$ be the maximum pseudo-likelihood estimators given in Section 2.2. The following theorem states the asymptotic distribution of $(\widehat{\theta}_n, \widehat{\beta}_n)$.

Theorem 13 (Asymptotic distribution) *Under the given regularity conditions*

$$\sqrt{n} \begin{pmatrix} \widehat{\theta}_n - \theta_0 \\ \widehat{\beta}_n - \beta_0 \end{pmatrix} \rightsquigarrow \begin{pmatrix} U \\ I_0^{-1} \{Z_A - k_0(Z_Y, Z_Z, U)\} \end{pmatrix} \equiv \begin{pmatrix} U \\ B \end{pmatrix}, \quad (5)$$

where $(Z_Y^\top, Z_Z^\top, Z_A^\top)^\top \sim N(0, \Sigma_0)$, and $U = \arg \min_{u \in \mathbb{R}^d} \beta_0^\top k_0(Z_Y, Z_Z, u)$.

Let $\xrightarrow[Z^*]{P}$ denote convergence in probability over Z^* , as defined in Section 2.2.3 and Chapter 10 of Kosorok (2008). Theorem 14 below establishes the validity of a parametric bootstrap procedure for approximating the sampling distribution of $(\widehat{\theta}_n, \widehat{\beta}_n)$.

Theorem 14 (Parametric bootstrap) *Assume $\widehat{\Sigma}_n = \Sigma_0 + o_P(1)$ and $h_n = \widehat{v}_n n^{-1/5}$, where $\widehat{v}_n \xrightarrow{P} v_0 \in (0, \infty)$ and v_0 is the standard error of $D_{\theta_0}(X)$. Assume the regularity conditions given above hold. Let $Z^* \sim N(0, I^{r \times r})$, where $I^{r \times r}$ is an $r \times r$ identity matrix and $r = q + p$. Let $\widetilde{Z}_n = \widehat{\Sigma}_n^{1/2} Z^* = (\widetilde{Z}_Y^\top, \widetilde{Z}_Z^\top, \widetilde{Z}_A^\top)^\top$, where*

$$\widehat{\Sigma}_n^{1/2} = \begin{bmatrix} \widehat{\Sigma}_1^{1/2} & 0 \\ \widehat{\Sigma}_2 \widehat{\Sigma}_1^{-1/2} & (\widehat{\Sigma}_2 - \widehat{\Sigma}_{21} \widehat{\Sigma}_1^{-1} \widehat{\Sigma}_{12})^{1/2} \end{bmatrix},$$

$\widehat{\Sigma}_1$ is the top left $q \times q$ block of $\widehat{\Sigma}_n$ (corresponding to Z_Y and Z_Z), $\widehat{\Sigma}_2$ is the lower right $p \times p$ block, $\widehat{\Sigma}_{21}$ is the upper right $q \times p$ block, $\widehat{\Sigma}_{12} = \widehat{\Sigma}_{21}^\top$, and the matrix square roots are the symmetric square roots obtained from the associated Eigenvalue decompositions. Let

$$\widetilde{T}_n(\mathbf{X}, Z_Y, Z_Z) = \omega(\mathbf{X}; \widehat{\theta}_n) \widehat{R}_{Y,n}(\mathbf{X}) \phi_Y(\mathbf{X})^\top Z_Y + \{1 - \omega(\mathbf{X}; \widehat{\theta}_n)\} \widehat{R}_{Z,n}(\mathbf{X}) \phi_Z(\mathbf{X})^\top Z_Z$$

and define

$$\begin{aligned} \widetilde{k}_n(Z_Y, Z_Z, U) &= \mathbb{E}_n \left[\mathbf{X} \left\{ 2P_{\widehat{\beta}_n}(\mathbf{X}) - 1 \right\} \cdot \left| \widetilde{T}_n(\mathbf{X}, Z_Y, Z_Z) \right. \right. \\ &\quad \left. \left. + \left\{ \widehat{R}_{Y,n}(\mathbf{X}) - \widehat{R}_{Z,n}(\mathbf{X}) \right\} \dot{\omega}_{\widehat{\theta}_n}(\mathbf{X})^\top U \right| \cdot h_n^{-1} \phi_0 \left\{ \widehat{D}_{\widehat{\theta}_n, n}(\mathbf{X}) / h_n \right\} \right] \\ &\quad \times \left\{ \mathbb{E}_n \left[h_n^{-1} \phi_0 \left\{ \widehat{D}_{\widehat{\theta}_n, n}(\mathbf{X}) / h_n \right\} \right] \right\}^{-1}, \end{aligned}$$

where ϕ_0 is the standard normal density. Define $\widetilde{U}_n = \arg \min_{u \in \mathbb{R}^d} \widehat{\beta}_n^\top \widetilde{k}_n(\widetilde{Z}_Y, \widetilde{Z}_Z, u)$ and $\widetilde{B}_n = I_n (\widehat{\beta}_n)^{-1} \left\{ \widetilde{Z}_A - \widetilde{k}_n(\widetilde{Z}_Y, \widetilde{Z}_Z, \widetilde{U}_n) \right\}$. Then,

$$\begin{pmatrix} \widetilde{U}_n \\ \widetilde{B}_n \end{pmatrix} \xrightarrow[Z^*]{P} \begin{pmatrix} U \\ B \end{pmatrix}, \quad (6)$$

where $(U^\top, B^\top)^\top$ is as defined in Theorem 13.

If we fix a large number of bootstrap replications, B , then $(\tilde{U}_{n,b}, \tilde{B}_{n,b})$, $b = 1, \dots, B$ will provide an approximation to the sampling distribution of the maximum pseudo-likelihood estimators. In Sections 4 and 5, we demonstrate the use of the bootstrap to test for heterogeneity of patient preferences.

Remark 15 *In Theorem 13, it can be seen that when β_0 only involves an intercept, there is no relationship between β_0 and U , as the argmax of an objective function does not change under multiplication by a positive scalar. This relationship is more complex when β_0 includes covariate effects. Theorem 13 also indicates that the asymptotic behaviors of $\hat{\theta}$ and $\hat{\beta}$ are driven largely by what happens at the boundary where $D_{\theta_0}(X) = 0$.*

4. Simulation Experiments

4.1 Fixed Utility Simulations

To examine the finite sample performance of the proposed methods, we begin with the following simple generative model. Let $\mathbf{X} = (X_1, \dots, X_5)^\top$ be a vector of independent normal random variables with mean 0 and standard deviation 0.5. Let treatment be assigned according to $\Pr\{A = d_\omega^{\text{opt}}(\mathbf{x}) | \mathbf{X} = \mathbf{x}\} = \rho$, i.e., the probability that the clinician correctly identifies the optimal treatment is constant across patients. Let ϵ_Y and ϵ_Z be independent normal random variables with mean 0 and standard deviation 0.5 and let $Y = A(4X_1 - 2X_2 + 2) + \epsilon_Y$ and $Z = A(2X_1 - 4X_2 - 2) + \epsilon_Z$. We estimated Q_Y and Q_Z using linear models, implemented the proposed method for a variety of n , ω , and ρ values, and examined $\hat{\omega}_n$, $\hat{\rho}_n$, and $\hat{d}_{\hat{\omega}_n, n}$, across 500 Monte Carlo replications per scenario.

Table 1 contains mean estimates of ω and ρ across replications along with the associated standard deviation across replications, and estimated error rate defined as the proportion of subjects to whom the estimated optimal policy does not recommend the true optimal treatment; to better characterize sampling variability in the estimated error rate the last column displays the median along with the first and third quartiles of the sampling distribution of the estimated error rate.

The pseudo-likelihood method performs well at estimating both ω and ρ , with estimation improving with larger sample sizes as expected. Table 2 contains estimated values of the true optimal policy, a policy where the utility function is estimated (the proposed method), policies estimated to maximize the two outcomes individually (corresponding to fixing $\omega = 1$ and $\omega = 0$), and the standard of care. The value of the standard of care is the mean composite outcome under the generative model. For each policy, the value is estimated by generating a testing sample of size 500 with treatment assigned according to the policy and averaging utilities (calculated using the true ω) in the testing set. The standard deviation across replications is included in parentheses.

The column labeled “estimated ω ” refers to the proposed method. We see that the proposed method produces values which increase with n and generally come close to the true optimal policy. In all settings, the proposed method offers significant improvement over the standard of care. The proposed method also offers improvement over policies to maximize each individual outcome.

n	ω	ρ	$\hat{\omega}_n$	$\hat{\rho}_n$	Error rate	Median(25th, 75th)
100	0.25	0.60	0.34 (0.24)	0.61 (0.08)	0.12 (0.13)	0.07 (0.03, 0.13)
			0.25 (0.05)	0.80 (0.04)	0.03 (0.02)	0.02 (0.01, 0.03)
	0.75	0.60	0.66 (0.24)	0.61 (0.07)	0.12 (0.13)	0.07 (0.03, 0.14)
			0.75 (0.05)	0.80 (0.04)	0.03 (0.02)	0.02 (0.01, 0.03)
200	0.25	0.60	0.28 (0.16)	0.61 (0.04)	0.07 (0.08)	0.04 (0.02, 0.10)
			0.25 (0.02)	0.80 (0.03)	0.01 (0.01)	0.01 (0.01, 0.02)
	0.75	0.60	0.72 (0.16)	0.61 (0.04)	0.07 (0.09)	0.03 (0.01, 0.08)
			0.75 (0.03)	0.80 (0.03)	0.01 (0.01)	0.01 (0.01, 0.02)
300	0.25	0.60	0.26 (0.11)	0.61 (0.03)	0.05 (0.06)	0.03 (0.01, 0.06)
			0.25 (0.02)	0.80 (0.02)	0.01 (0.01)	0.01 (0.00, 0.01)
	0.75	0.60	0.74 (0.13)	0.61 (0.03)	0.06 (0.07)	0.03 (0.01, 0.08)
			0.75 (0.02)	0.80 (0.02)	0.01 (0.01)	0.01 (0.01, 0.01)
500	0.25	0.60	0.25 (0.08)	0.61 (0.02)	0.04 (0.04)	0.02 (0.01, 0.04)
			0.25 (0.01)	0.80 (0.02)	0.01 (0.01)	0.01 (0.00, 0.01)
	0.75	0.60	0.75 (0.08)	0.61 (0.02)	0.04 (0.04)	0.02 (0.01, 0.05)
			0.75 (0.01)	0.80 (0.02)	0.01 (0.01)	0.01 (0.00, 0.01)

Table 1: Estimation results for simulations where utility and probability of optimal treatment are fixed.

n	ω	ρ	Optimal	Estimated ω	Y only	Z only	Standard of care
100	0.25	0.60	1.90 (0.07)	1.70 (0.38)	0.38 (0.11)	1.76 (0.08)	0.37 (0.24)
			1.90 (0.07)	1.89 (0.07)	0.39 (0.12)	1.76 (0.08)	1.14 (0.21)
	0.75	0.60	1.90 (0.06)	1.71 (0.37)	1.76 (0.08)	0.39 (0.12)	0.37 (0.23)
			1.90 (0.06)	1.89 (0.07)	1.76 (0.08)	0.39 (0.12)	1.14 (0.21)
200	0.25	0.60	1.90 (0.07)	1.82 (0.21)	0.39 (0.11)	1.76 (0.07)	0.39 (0.16)
			1.90 (0.07)	1.90 (0.06)	0.39 (0.11)	1.76 (0.07)	1.14 (0.14)
	0.75	0.60	1.90 (0.07)	1.82 (0.22)	1.76 (0.07)	0.38 (0.11)	0.39 (0.17)
			1.90 (0.07)	1.90 (0.06)	1.76 (0.07)	0.38 (0.11)	1.14 (0.15)
300	0.25	0.60	1.90 (0.07)	1.86 (0.13)	0.39 (0.11)	1.76 (0.07)	0.38 (0.14)
			1.90 (0.07)	1.89 (0.06)	0.39 (0.11)	1.76 (0.07)	1.14 (0.12)
	0.75	0.60	1.90 (0.06)	1.85 (0.17)	1.77 (0.07)	0.38 (0.11)	0.38 (0.14)
			1.90 (0.06)	1.90 (0.07)	1.77 (0.07)	0.38 (0.11)	1.13 (0.12)
500	0.25	0.60	1.90 (0.06)	1.88 (0.10)	0.39 (0.10)	1.76 (0.07)	0.38 (0.10)
			1.90 (0.06)	1.90 (0.07)	0.39 (0.11)	1.76 (0.07)	1.14 (0.09)
	0.75	0.60	1.89 (0.07)	1.88 (0.08)	1.77 (0.07)	0.39 (0.11)	0.38 (0.11)
			1.89 (0.07)	1.90 (0.07)	1.77 (0.07)	0.39 (0.11)	1.14 (0.09)

Table 2: Value results for simulations where utility and probability of optimal treatment are fixed.

To further examine the performance of the proposed method, we allow the probability of optimal treatment to depend on patient covariates. Let $\Pr \left\{ A = d_{\omega}^{\text{opt}}(\mathbf{X}) \right\} = \text{expit}(0.5 +$

X_1). This corresponds to the case where $\beta = (0.5, 1, 0, \dots, 0)^\top$, where the first element of β is an intercept. Let \mathbf{X} , Y , and Z be generated as described above. In this generative model, the probability that a patient is treated optimally in standard care is larger for patients with positive values of X_1 and smaller for patients with negative values of X_1 . We applied the proposed method to 500 replications of this generative model for various n and ω . Table 3 contains mean estimates of ω , root mean squared error (RMSE) of $\hat{\beta}_n$, and the error rate along with its standard error and quartiles.

n	ω	$\hat{\omega}_n$	RMSE of $\hat{\beta}_n$	Error rate	Median(25th, 75th)
100	0.25	0.34 (0.23)	1.32 (0.50)	0.10 (0.14)	0.04 (0.02, 0.10)
	0.75	0.71 (0.21)	1.37 (0.48)	0.10 (0.11)	0.06 (0.03, 0.12)
200	0.25	0.27 (0.13)	0.81 (0.30)	0.04 (0.08)	0.02 (0.01, 0.03)
	0.75	0.75 (0.15)	0.85 (0.29)	0.07 (0.07)	0.04 (0.02, 0.10)
300	0.25	0.26 (0.09)	0.60 (0.21)	0.03 (0.05)	0.01 (0.01, 0.02)
	0.75	0.75 (0.10)	0.63 (0.22)	0.04 (0.05)	0.03 (0.01, 0.07)
500	0.25	0.25 (0.03)	0.44 (0.14)	0.01 (0.01)	0.01 (0.00, 0.01)
	0.75	0.76 (0.07)	0.46 (0.16)	0.03 (0.04)	0.02 (0.01, 0.04)

Table 3: Estimation results for simulations where utility is fixed and probability of optimal treatment is variable.

Estimation of the observational policy (as defined by β) improves with larger sample sizes. The probability that the estimated policy assigns the optimal treatment also increases with the sample size. The true value of ω does not affect estimation of ω or β .

Table 4 contains estimated values of the true optimal policy, a policy where the utility function is estimated (the proposed method), policies estimated to maximize each outcome individually, and the standard of care. Values are estimated from independent testing sets of size 500 as described above. The value under the standard of care is the mean composite outcome under the generative model.

n	ω	Optimal	Estimated ω	Y only	Z only	Standard of care
100	0.25	1.90 (0.07)	1.72 (0.40)	0.39 (0.12)	1.76 (0.07)	0.33 (0.23)
	0.75	1.90 (0.06)	1.75 (0.30)	1.76 (0.08)	0.39 (0.12)	0.56 (0.23)
200	0.25	1.90 (0.06)	1.85 (0.23)	0.37 (0.11)	1.76 (0.07)	0.34 (0.16)
	0.75	1.89 (0.06)	1.83 (0.18)	1.76 (0.07)	0.38 (0.11)	0.58 (0.16)
300	0.25	1.90 (0.06)	1.88 (0.16)	0.39 (0.10)	1.77 (0.07)	0.33 (0.14)
	0.75	1.90 (0.06)	1.87 (0.08)	1.76 (0.07)	0.40 (0.11)	0.57 (0.13)
500	0.25	1.90 (0.07)	1.89 (0.06)	0.39 (0.10)	1.76 (0.07)	0.33 (0.11)
	0.75	1.90 (0.07)	1.88 (0.07)	1.77 (0.07)	0.39 (0.11)	0.58 (0.10)

Table 4: Value results for simulations where utility is fixed and probability of optimal treatment is variable.

The proposed method (found in the column labeled “estimated ω ”) produces values that are close to the true optimal policy in large samples and a significant improvement

over standard of care in small to moderate samples. We note that value under the standard of care differs across ω . When ω is close to 1, the composite outcome places more weight on Y , for which the magnitude of the association with X_1 is larger. Because patients with larger values of X_1 are more likely to be treated optimally in this generative model, the standard of care produces larger composite outcomes when ω is closer to 1. Likewise, the mean composite outcome under policies to maximize each individual outcome varies with the true value of ω .

4.2 Patient-specific Utility Simulations

Next, we examine the case where the utility function is allowed to vary across patients. Let \mathbf{X} , Y , and Z be generated as above. Again, assume that $\Pr \left\{ A = d_{\theta}^{\text{opt}}(\mathbf{X}) \right\} = \text{expit}(0.5 + X_1)$, i.e., $\beta = (0.5, 1, 0, \dots, 0)^\top$. Consider the composite outcome $U = \omega(\mathbf{X}; \theta)Y + \{1 - \omega(\mathbf{X}, \theta)\} Z$, where $\omega(\mathbf{X}; \theta) = \text{expit}(1 - 0.5X_1)$, i.e., $\theta = (1, -0.5, 0, \dots, 0)^\top$, where the first element of θ is an intercept. We implemented the proposed method for various n and examined estimation of θ and β across 500 replications. Each replication is based on a simulated Markov chain of length 10,000 as described in Section 2.2. Results are summarized in Table 5.

n	RMSE of $\widehat{\theta}_n$	RMSE of $\widehat{\beta}_n$	Error rate	Median(25th, 75th)
100	0.85 (0.45)	1.28 (0.46)	0.10 (0.08)	0.06 (0.04, 0.13)
200	0.72 (0.31)	0.81 (0.27)	0.07 (0.06)	0.05 (0.04, 0.08)
300	0.63 (0.16)	0.63 (0.21)	0.05 (0.04)	0.04 (0.03, 0.06)
500	0.60 (0.09)	0.46 (0.15)	0.05 (0.02)	0.04 (0.03, 0.05)

Table 5: Estimation results for simulations where both utility and probability of optimal treatment are variable.

Larger sample sizes produce marginal decreases in the RMSE of $\widehat{\theta}_n$. The estimated policy assigns the true optimal treatment more than 80% of the time for all sample sizes and the error rate decreases as the sample size increases. Table 6 contains estimated values of the true optimal policy, the policy estimated using the proposed method, policies estimated to maximize each outcome individually, and standard of care.

n	Optimal	Estimated ω	Y only	Z only	Standard of care
100	1.74 (0.06)	1.65 (0.14)	1.66 (0.06)	1.41 (0.08)	0.50 (0.21)
200	1.74 (0.06)	1.69 (0.11)	1.67 (0.06)	1.41 (0.08)	0.49 (0.15)
300	1.74 (0.06)	1.71 (0.07)	1.66 (0.06)	1.41 (0.07)	0.50 (0.13)
500	1.74 (0.06)	1.71 (0.07)	1.66 (0.06)	1.41 (0.08)	0.50 (0.11)

Table 6: Value results for simulations where both utility and probability of optimal treatment are variable.

The proposed method produces policies that achieve significant improvement over the standard of care across sample sizes.

Finally, we examine the performance of the parametric bootstrap as described in Section 3. Let \mathbf{X} be a bivariate vector of normal random variables with mean 0, standard deviation 0.5, and correlation zero. Let Y and Z be generated as above and let $\beta = (2.5, 1, 0)^\top$ where the first element of β is an intercept. Let $\theta_{(1)}$ be the vector θ with the first element removed. We are interested in testing the null hypothesis $H_0 : \|\theta_{(1)}\| = 0$, which corresponds to a test for heterogeneity of patient preferences. The table below contains estimated power across 500 Monte Carlo replications under the null hypothesis, where the true value is $\theta = (1, 0, 0)^\top$, and two alternative hypotheses: $H_1 : \theta = (1, 4, 3)^\top$, and $H_2 : \theta = (1, 6, 6)^\top$. All tests were conducted at level $\alpha = 0.05$ and based on 1000 bootstrap samples. The last column in Table 7 shows the average agreement between the bootstrap and estimated optimal decision rule when $\theta_0 = (1, 0, 0)$; the results suggest the decision rule is stable at sample sizes and generative models considered.

n	Type 1 error	Power against H_1	Power against H_2	Stability
100	0.002	0.238	0.264	0.805
200	0.004	0.790	0.782	0.831
300	0.004	0.958	0.948	0.833
500	0.002	0.998	0.990	0.829

Table 7: Power of bootstrap test for homogeneity of utility function

The proposed bootstrap procedure produces type I error rates near nominal levels under the null and moderate power in large samples under alternative hypotheses.

5. Case Study: The STEP-BD Standard Care Pathway

The Systematic Treatment Enhancement Program for Bipolar Disorder (STEP-BD) was a landmark study of the effects of antidepressants in patients with bipolar disorder (Sachs et al., 2007). In addition to a randomized trial assessing outcomes for patients given an antidepressant or placebo, the STEP-BD study also included a large-scale observational study, the standard care pathway. As our method requires observational data on clinical decision making, we apply the proposed method to the observational data from the STEP-BD standard care pathway to estimate decision rules for the use of antidepressants in patients with bipolar disorder. (Clearly, as clinicians are not generally assigning treatment according to their best clinical judgment in a randomized clinical trial, the proposed method is not applicable to the randomized pathway of STEP-BD.)

Although bipolar disorder is characterized by alternating episodes of depression and mania, recurrent depression is the leading cause of impairment among patients with bipolar disorder (Judd et al., 2002). However, the use of antidepressants has not become standard care in bipolar disorder due to the risk of antidepressants inducing manic episodes in certain patients (Ghaemi, 2008; Goldberg, 2008). Thus, the clinical decision in the treatment of bipolar disorder is whether to prescribe antidepressants to a specific patient in order to balance trade-offs between symptoms of depression, symptoms of mania, and other side effects of treatment.

We use the SUM-D score for depression symptoms and the SUM-M score for mania symptoms as outcomes. We consider a patient treated if they took any one of ten antidepres-

sants that appear in the STEP-BD standard care pathway (Deseryl, Serzone, Citalopram, Escitalopram Oxalate, Prozac, Fluvoxamine, Paroxetine, Zoloft, Venlafaxine, or Bupropion). To generate candidate predictors for our model we made use of a complimentary randomized pathway in the STEP-BD trial. In this pathway, the patients are drawn from the same population, and the same variables are measured; however, treatment is randomly assigned so that there is no unmeasured confounding. Using step-wise variable selection to construct an outcome model from these data identified the following variables: mood elevation, anxiety, irritability, baseline SUM-M, and baseline SUM-D. We also used a step-wise logistic regression for the propensity score in the observational pathway to identify any additional potential confounders (Moodie et al., 2012). In addition to the variables in the outcome model, the logistic regression model identified race, insurance status, age, and substance abuse. The union of variables identified in through the randomized pathway and the propensity score were used in our models of the Q -functions and as tailoring variables in our treatment rules. Figure 1 contains box plots of SUM-D scores on the log scale by substance abuse and treatment. Figure 2 contains box plots of SUM-M scores on the log scale by substance abuse and treatment. For both outcomes, lower scores are more desirable. Figure 1 indicates that those without a history of substance abuse benefit

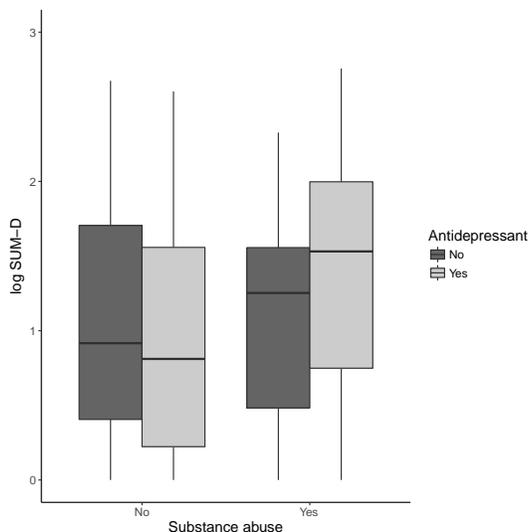


Figure 1: Box plots of log SUM-D by substance abuse and treatment.

from treatment with antidepressants. However, among those with a history of substance abuse, patients treated with antidepressants appear to have worse symptoms of depression. Figure 2 indicates that treatment has no effect on symptoms of mania among those without a history of substance abuse. However, among those with a history of substance abuse, it appears that treatment may be inducing manic episodes. Thus, a sensible treatment policy would be one that tends to prescribe antidepressants only to patients without a history of substance abuse.

We analyzed these data using the proposed method for optimizing composite outcomes. Results are summarized in Table 8 below. We estimated policies where both utility and probability of optimal treatment are fixed (fixed-fixed), where utility is fixed but probabil-

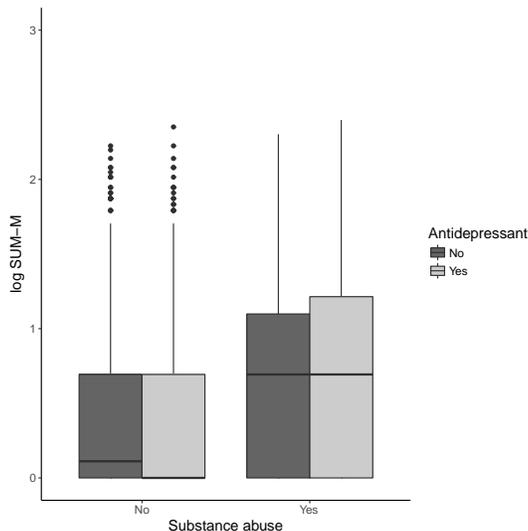


Figure 2: Box plots of log SUM-M by substance abuse and treatment.

ity of optimal treatment is assumed to vary between patients (fixed-variable), and where both utility and probability of optimal treatment are assumed to vary between patients (variable-variable). For both the fixed-variable policy and the variable-variable policy, we report $\mathbb{E}_n \left\{ \text{expit} \left(\mathbf{X}^\top \hat{\beta}_n \right) \right\}$ in place of $\hat{\rho}_n$ and for the variable-variable policy, we report $\mathbb{E}_n \left\{ \text{expit} \left(\mathbf{X}^\top \hat{\theta}_n \right) \right\}$ in place of $\hat{\omega}_n$. Thus, for parameters that are assumed to vary across patients, Table 8 contains the mean estimate in the sample. To evaluate each estimated policy, we used five-fold cross-validation of the inverse probability weighted estimator (IPWE) of the value for each outcome; i.e., for each fold, we used the training portion to estimate the optimal policy and propensity score, and we used the testing portion to compute the IPWE of the value; taking the average of the IPWE value estimates across folds yields the reported values. For both SUM-D and SUM-M, lower scores are preferred. Value is reported as the percent improvement over standard of care, calculated using the estimated utility function. Large percent improvements in value are preferred.

Policy	SUM-D	SUM-M	Value (% improvement)	$\hat{\omega}_n$	$\hat{\rho}_n$
fixed-fixed	2.336	0.857	1.8%	0.039	0.431
fixed-variable	2.324	0.838	3.9%	0.039	0.440
variable-variable	2.321	0.804	8.3%	0.334	0.448
standard of care	2.480	0.868	0.0%	.	.

Table 8: Results of analysis of STEP-BD data for SUM-D and SUM-M.

All estimated policies produce more desirable SUM-D scores and SUM-M scores compared to standard of care and improved value according to the estimated utility. Allowing the probability of optimal treatment to vary between patients leads to further improvements in value, as does allowing the utility function to vary between patients. All policies produce similar estimates for the probability of optimal treatment averaged across patients.

	Intercept	Age	Substance	Mood elevation	Insurance	Race
Estimate	2.427	-0.177	-1.666	-2.632	4.263	1.078
Standard error	3.039	0.953	2.746	2.238	3.241	3.274

Table 9: Estimates of $\hat{\theta}_n$ in the variable-variable policy

The resulting decision rules can be written as the sign of a linear combination of the covariates. As an example, the fixed-fixed policy assigns treatment with antidepressants when $0.032 - 0.001(\text{age}) - 0.646(\text{substance abuse}) - 0.007(\text{mood elevation}) + 0.007(\text{medical insurance}) + 0.129(\text{white})$ is non-negative. The negative coefficient for substance abuse means that a history of substance abuse indicates that a patient should not be prescribed antidepressants. Prior research has shown that patients with a history of substance abuse are more likely to abuse antidepressants (Evans and Sullivan, 2014). This may contribute to the poor outcomes experienced by STEP-BD patients with a history of substance abuse who were treated with antidepressants. Table 9 displays estimates and standard errors of the components of $\hat{\theta}_n$ in the variable-variable policy. A test for preference heterogeneity based on 1000 bootstrap samples generated according to Theorem 14 yielded a p-value < 0.001 .

As a secondary analysis, we use the SUM-D score and a side effect score as the outcomes. Eight side effects were recorded in the STEP-BD standard care pathway (tremors, dry mouth, sedation, constipation, diarrhea, headache, poor memory, sexual dysfunction, and increased appetite). Patients rated the severity of each side effect from 0 to 4 with larger values indicating more severe side effects. We took the mean score across side effects as the second outcome. Results are summarized in Table 10, reported analogously to those in Table 8.

Policy	SUM-D	Side effect score	Value (% improvement)	$\hat{\omega}_n$	$\hat{\rho}_n$
fixed-fixed	2.377	0.156	5.2%	0.601	0.462
fixed-variable	2.384	0.159	5.7%	0.100	0.472
variable-variable	2.430	0.161	6.1%	0.378	0.487
standard of care	2.480	0.172	0.0%	.	.

Table 10: Results of analysis of STEP-BD data for SUM-D and Side effect score.

	Intercept	Age	Substance	Mood	Irritable	Anxiety	Insurance	Race
Estimate	-3.125	-4.614	2.094	-0.609	2.594	0.332	-3.493	-2.563
Standard error	2.257	2.300	2.395	2.603	2.610	2.538	2.599	2.449

Table 11: Estimates of $\hat{\theta}_n$ in the variable-variable policy

Each estimated policy produces improved SUM-D scores and improved side effect scores compared to the standard of care. Each policy also produces improvement in value according to the estimated utility function. Again, allowing the utility function to vary between patients results in further improvements in value. Each policy produces similar estimates of the probability that patients are treated optimally in standard care. The variable-variable policy places more weight on SUM-D scores on average compared to the other policies.

Table 11 displays estimates and standard errors of coefficients in $\hat{\theta}_n$ in the variable-variable policy. The bootstrap procedure for testing the null hypothesis that patient preferences are homogeneous based on 1000 bootstrap samples yielded a p-value < 0.001 .

6. Discussion

The estimation of individualized treatment rules has been well-studied in the statistical literature. Existing methods have typically defined the optimal treatment rule as optimizing the mean of a fixed scalar outcome. However, clinical practice often requires consideration of multiple outcomes. Thus, there is a disconnect between existing statistical methods and current clinical practice. It is reasonable to assume that clinicians make treatment decisions for each patient with the goal of maximizing that patient’s utility. Therefore, it is natural to use observational data to estimate patient utilities from observed clinician decisions. This represents a new paradigm for the use of observational data in the context of precision medicine in that clinical decisions are viewed as a (noisy) surrogate for patient preferences and leveraged to improve the quality of a learned treatment rule and to generate new insights into heterogeneity in patient preferences.

The proposed methodology offers many opportunities for future research. In the present manuscript, we have considered only the simplest case— that of one decision time, two outcomes, and two possible treatments. Scenarios with more than two outcomes are discussed in the Appendix, and the simulation results there demonstrate that the proposed method performs well with three outcomes. Extensions to more than two treatments or multiple time points represent potential areas for future research. The proposed method requires positing a parametric model for the utility function. Model misspecification is discussed in the Appendix, and the simulation results there demonstrate that the proposed method performs reasonably well when important covariates are omitted from the model for the utility function. However, the use of semi- or non-parametric models is an important extension. A more technical direction for future work is a more nuanced study of the affect of boundary conditions on the resulting rate of convergence (see Assumption 10). Finally, while we have proposed our utility function estimator inside the framework of one-stage Q-learning, the pseudo-likelihood utility function estimator could be used alongside other existing one-stage optimal treatment policy estimators based on (augmented) inverse probability weighting (e.g., Zhao et al., 2012; Zhang et al., 2012a). There is a great future for the development of methods for optimizing composite outcomes in precision medicine and application of these methods in clinical studies.

Acknowledgments

We would like to acknowledge support for this project from the National Science Foundation (NSF grants DMS-1555141 and DMS-1513579) and the National Institutes of Health (NIH grants R01 DK108073 and P01 CA142538). We also gratefully acknowledge the National Institute of Mental Health for providing access to the STEP-BD data set. Finally, we would like to thank the Editor, Associate Editor, and two anonymous referees for their constructive feedback during the review process.

Appendix A: Proofs

Proof [Proof of Theorem 5] Consider $(\beta, \theta) \in \mathcal{B} \times \Theta$ and suppose

$$P_\beta(\mathbf{X})^{1\{Y=d_\theta(\mathbf{X})\}}(1-P_\beta(\mathbf{X}))^{1-1\{Y=d_\theta(\mathbf{X})\}} = P_{\beta_0}(\mathbf{X})^{1\{Y=d_{\theta_0}(\mathbf{X})\}}(1-P_{\beta_0}(\mathbf{X}))^{1-1\{Y=d_{\theta_0}(\mathbf{X})\}}, \quad \forall(\mathbf{X}, Y) \text{ a.s.} \quad (7)$$

Let $X_S = \{\mathbf{X} : d_\theta(\mathbf{X}) = d_{\theta_0}(\mathbf{X})\}$. For all $\mathbf{X} \in X_S$, we have that $P_\beta(\mathbf{X}) = P_{\beta_0}(\mathbf{X}) \Rightarrow \beta = \beta_0$ by 4 in Assumption 4 which implies identifiability of $P_\beta(\mathbf{X})$ on X_S .

Now suppose $\theta \neq \theta_0$, then since $\beta = \beta_0$, we have by (7), that $\forall \mathbf{X} \in X_S^c$, $P_\beta(\mathbf{X}) = 1 - P_{\beta_0}(\mathbf{X})$ which by applying 4 in Assumption 4 again, we obtain that $\beta = -\beta_0$. However, this is impossible by 2 in Assumption 4. Thus $\theta = \theta_0$, and we obtain that $(\beta, \theta) = (\beta_0, \theta_0)$. ■

Proof [Proof of Theorem 7] The log of the pseudo-likelihood is given by

$$\widehat{\ell}_n(\theta, \beta) = \mathbb{E}_n \left[\mathbf{X}^\top \beta 1 \left\{ A = \widehat{d}_{\theta, n}(\mathbf{X}) \right\} - \log \{1 + \exp(\mathbf{X}^\top \beta)\} \right].$$

Let $\widehat{m}(\cdot, \cdot; \theta, \beta) : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}$ be defined by $\widehat{m}(\mathbf{X}, A; \theta, \beta) = \mathbf{X}^\top \beta 1 \left\{ A = \widehat{d}_{\theta, n}(\mathbf{X}) \right\} - \log \{1 + \exp(\mathbf{X}^\top \beta)\}$ and consider the class of functions $\{\widehat{m}(\cdot, \cdot; \theta, \beta) : \theta \in \mathbb{R}^p, \beta \in \mathcal{B}\}$. The class $\{\log \{1 + \exp(\mathbf{X}^\top \beta)\} : \beta \in \mathcal{B}\}$ is contained in a VC class by Lemma 9.9 (viii) and (v) of Kosorok (2008). By Theorem 9.3 of Kosorok (2008), this is also a Glivenko–Cantelli (GC) class.

Let $u(\mathbf{X}, A; \theta) = \omega(\mathbf{X}; \theta) \left\{ \widehat{Q}_{Y, n}(\mathbf{X}, A) - \widehat{Q}_{Z, n}(\mathbf{X}, A) \right\} + \widehat{Q}_{Z, n}(\mathbf{X}, A)$, which lies in a VC class indexed by $\theta \in \mathbb{R}^p$ by Lemma 9.6 and Lemma 9.9 (viii), (vi), and (v) of Kosorok (2008). We have that

$$1 \left\{ A = \widehat{d}_{\theta, n}(\mathbf{X}) \right\} = 1(A = 1) 1 \{u(\mathbf{X}, 1; \theta) - u(\mathbf{X}, -1, \theta) \geq 0\} \\ + 1(A = -1) 1 \{u(\mathbf{X}, 1; \theta) - u(\mathbf{X}, -1, \theta) < 0\},$$

and it follows that $1 \left\{ A = \widehat{d}_{\theta, n}(\mathbf{X}) \right\}$ is contained in a GC class indexed by $\theta \in \mathbb{R}^p$. From Corollary 9.27 (ii) of Kosorok (2008) it follows that $\mathbf{X}^\top \beta 1 \left\{ A = \widehat{d}_{\theta, n}(\mathbf{X}) \right\}$ lies in a GC class indexed by $(\theta, \beta) \in \mathbb{R}^p \times \mathcal{B}$ as long as $\mathbf{X}^\top \beta$ is uniformly bounded by a function with finite mean, which holds as long as \mathcal{B} is compact and $\|\mathbb{E}\mathbf{X}\| < \infty$. It follows that

$$\sup_{(\theta, \beta) \in \mathbb{R}^p \times \mathcal{B}} \left| (\mathbb{E}_n - \mathbb{E}) \left[\mathbf{X}^\top \beta 1 \left\{ A = \widehat{d}_{\theta, n}(\mathbf{X}) \right\} - \log \{1 + \exp(\mathbf{X}^\top \beta)\} \right] \right| \xrightarrow{P} 0.$$

Next, define

$$\widehat{M}(\theta, \beta) = \mathbb{E} \{ \widehat{m}(\mathbf{X}, A; \theta, \beta) \} = \mathbb{E} \left(\mathbf{X}^\top \beta \mathbb{E} \left[1 \left\{ A = \widehat{d}_{\theta, n}(\mathbf{X}) \right\} \mid \mathbf{X} \right] \right) - \mathbb{E} \log \{1 + \exp(\mathbf{X}^\top \beta)\}$$

and note that $\widehat{M}(\theta, \beta)$ is continuous in β . The inside expectation of the first piece is

$$\begin{aligned} \mathbb{E} \left[1 \left\{ A = \widehat{d}_{\theta,n}(\mathbf{X}) \right\} \middle| \mathbf{X} \right] &= \text{expit}(\mathbf{X}^\top \beta_0) 1 \left\{ \widehat{d}_{\theta,n}(\mathbf{X}) = d_{\theta_0}^{\text{opt}}(\mathbf{X}) \right\} \\ &\quad + \{1 - \text{expit}(\mathbf{X}^\top \beta_0)\} 1 \left\{ \widehat{d}_{\theta,n}(\mathbf{X}) \neq d_{\theta_0}^{\text{opt}}(\mathbf{X}) \right\}, \end{aligned}$$

using the fact that $\Pr \left\{ A = d_{\theta_0}^{\text{opt}}(\mathbf{X}) \right\} = \text{expit}(\mathbf{X}^\top \beta_0)$. Define $a(\mathbf{X}) = Q_Y(\mathbf{X}, 1) - Q_Y(\mathbf{X}, -1) - Q_Z(\mathbf{X}, 1) + Q_Z(\mathbf{X}, -1)$ and $b(\mathbf{X}) = Q_Z(\mathbf{X}, 1) - Q_Z(\mathbf{X}, -1)$. Similarly, define $\widehat{a}(\mathbf{X}) = \widehat{Q}_{Y,n}(\mathbf{X}, 1) - \widehat{Q}_{Y,n}(\mathbf{X}, -1) - \widehat{Q}_{Z,n}(\mathbf{X}, 1) + \widehat{Q}_{Z,n}(\mathbf{X}, -1)$ and $\widehat{b}(\mathbf{X}) = \widehat{Q}_{Z,n}(\mathbf{X}, 1) - \widehat{Q}_{Z,n}(\mathbf{X}, -1)$. Then,

$$\begin{aligned} 1 \left\{ \widehat{d}_{\theta,n}(\mathbf{X}) = d_{\theta_0}^{\text{opt}}(\mathbf{X}) \right\} &= 1 \left[\left\{ \omega(\mathbf{X}; \theta) \widehat{a}(\mathbf{X}) + \widehat{b}(\mathbf{X}) \right\} \left\{ \omega(\mathbf{X}; \theta) a(\mathbf{X}) + b(\mathbf{X}) \right\} \geq 0 \right] \\ &= 1 \left[\omega(\mathbf{X}; \theta) \left\{ \omega(\mathbf{X}; \theta) a(\mathbf{X}) \widehat{a}(\mathbf{X}) + \widehat{a}(\mathbf{X}) b(\mathbf{X}) \right\} \right. \\ &\quad \left. + \omega(\mathbf{X}; \theta) a(\mathbf{X}) \widehat{b}(\mathbf{X}) + \widehat{b}(\mathbf{X}) b(\mathbf{X}) \geq 0 \right], \end{aligned}$$

and thus $\mathbb{E} \left[1 \left\{ A = \widehat{d}_{\theta,n}(\mathbf{X}) \right\} \middle| \mathbf{X} \right]$ is continuous in θ .

Let $m(\mathbf{X}, A; \theta, \beta) = \mathbf{X}^\top \beta 1 \left\{ A = d_{\theta_0}^{\text{opt}}(\mathbf{X}) \right\} - \log \{1 + \exp(\mathbf{X}^\top \beta)\}$. Because the model is identifiable and $\mathcal{L}_n(\theta, \beta)$ is a parametric log-likelihood, $\mathbb{E}m(\mathbf{X}, A; \theta, \beta)$ has unique maximizers at θ_0 and β_0 . Let $\widetilde{\theta}_n$ and $\widetilde{\beta}_n$ be the maximizers of $\mathbb{E}\widehat{m}(\mathbf{X}, A; \theta, \beta)$. Because $\mathbb{E} \left\{ \left| \widehat{d}_{\theta,n}(\mathbf{X}) - d_{\theta_0}^{\text{opt}}(\mathbf{X}) \right| \right\} \rightarrow 0$ in probability, for any $\theta \in \mathbb{R}^d$, $\mathbb{E} \left[1 \left\{ A = \widehat{d}_{\theta,n}(\mathbf{X}) \right\} \middle| \mathbf{X} \right] - \mathbb{E} \left[1 \left\{ A = d_{\theta_0}^{\text{opt}}(\mathbf{X}) \right\} \middle| \mathbf{X} \right] \rightarrow 0$ in probability, uniformly in θ over compact subsets of \mathbb{R}^d , which implies that both $\widetilde{\theta}_n \rightarrow \theta_0$ and $\widetilde{\beta}_n \rightarrow \beta_0$ in probability. The claim now follows from Lemma 14.3 and Theorem 2.12 of Kosorok (2008). \blacksquare

Proof [Proof of Theorem 8] Define $Q_{\theta_0}(\mathbf{x}, a)$ and $Q_{\widehat{\theta}_n}(\mathbf{x}, a)$ as defined in Section 2. Let $u(Y, Z; A, \mathbf{X}, \theta) = \omega(\mathbf{X}; \theta) Q_Y(\mathbf{X}, A) + \{1 - \omega(\mathbf{X}; \theta)\} Q_Z(\mathbf{X}, A)$. Under the given assumptions, for some constant $0 < c < \infty$,

$$\begin{aligned} &\left| V \left(\widehat{d}_{\widehat{\theta}_n,n} \right) - V \left(d_{\theta_0}^{\text{opt}} \right) \right| \\ &\leq c \left| \mathbb{E} \left\{ u(Y, Z; A, \mathbf{X}, \theta_0) - \widehat{Q}_{\widehat{\theta}_n,n}(\mathbf{X}, A) \right\}^2 - \mathbb{E} \left\{ u(Y, Z; A, \mathbf{X}, \theta_0) - Q_{\theta_0}(\mathbf{X}, A) \right\}^2 \right|^{1/2} \quad (8) \end{aligned}$$

by equation (3.1) of Qian and Murphy (2011) (see also Murphy, 2005). The right hand side of (8) converges in probability to 0 by the consistency of $\widehat{\theta}_n$, consistency of $\widehat{Q}_{Y,n}$ and $\widehat{Q}_{Z,n}$, and the continuous mapping theorem. The result follows. \blacksquare

Proof [Proof of Theorem 13] By definition of $(\widehat{\theta}_n, \widehat{\beta}_n)$,

$$\begin{aligned}
 0 &\leq \widehat{\ell}_n(\widehat{\theta}_n, \widehat{\beta}_n) - \widehat{\ell}_n(\theta_0, \beta_0) \\
 &= \sum_{i=1}^n \left[\mathbf{X}_i^\top \widehat{\beta}_n \mathbf{1} \left\{ A_i = \widehat{d}_{\widehat{\theta}_n, n}(\mathbf{X}_i) \right\} - \mathbf{X}_i^\top \beta_0 \mathbf{1} \left\{ A_i = \widehat{d}_{\theta_0, n}(\mathbf{X}_i) \right\} \right. \\
 &\quad \left. - \left(\widehat{\beta}_n - \beta_0 \right)^\top \mathbf{X}_i P_{\beta_0}(\mathbf{X}_i) \right] - \frac{1}{2} \sqrt{n} \left(\widehat{\beta}_n - \beta_0 \right)^\top I_n(\beta_*) \sqrt{n} \left(\widehat{\beta}_n - \beta_0 \right) \\
 &= \sqrt{n} \left(\widehat{\beta}_n - \beta_0 \right)^\top n^{-1/2} \sum_{i=1}^n \mathbf{X}_i \left[\mathbf{1} \left\{ A_i = \widehat{d}_{\widehat{\theta}_n, n}(\mathbf{X}_i) \right\} - P_{\beta_0}(\mathbf{X}_i) \right] \\
 &\quad - \frac{1}{2} \sqrt{n} \left(\widehat{\beta}_n - \beta_0 \right)^\top I_n(\beta_*) \sqrt{n} \left(\widehat{\beta}_n - \beta_0 \right) \\
 &\quad + \sum_{i=1}^n \mathbf{X}_i^\top \beta_0 \left[\mathbf{1} \left\{ A_i = \widehat{d}_{\widehat{\theta}_n, n}(\mathbf{X}_i) \right\} - \mathbf{1} \left\{ A_i = \widehat{d}_{\theta_0, n}(\mathbf{X}_i) \right\} \right],
 \end{aligned}$$

where β_* is a point between $\widehat{\beta}_n$ and β_0 . Using the definition of a maximizer and letting $\widehat{u}_n(\theta) = n^{-1/2} \sum_{i=1}^n \mathbf{X}_i \left[\mathbf{1} \left\{ A_i = \widehat{d}_{\theta, n}(\mathbf{X}_i) \right\} - P_{\beta_0}(\mathbf{X}_i) \right]$, we have that $\sqrt{n} \left(\widehat{\beta}_n - \beta_0 \right) = I_n(\beta_*)^{-1} \widehat{u}_n \left(\widehat{\theta}_n \right)$ by setting $\frac{\partial}{\partial \beta} \widehat{\ell}_n(\widehat{\theta}, \beta) |_{\beta} = 0$, since $I_n(\beta_*) \xrightarrow{P} I_0$ and I_0 is positive definite. Next, note that

$$\begin{aligned}
 \widehat{u}_n \left(\widehat{\theta}_n \right) &= n^{-1/2} \sum_{i=1}^n \mathbf{X}_i \left[\mathbf{1} \left\{ A_i = \widehat{d}_{\widehat{\theta}_n, n}(\mathbf{X}_i) \right\} - \mathbf{1} \left\{ A_i = d_{\theta_0}^{\text{opt}}(\mathbf{X}_i) \right\} \right] + Z_{A, n} \\
 &= \mathbb{G}_n \left(\mathbf{X} \left[\mathbf{1} \left\{ A = \widehat{d}_{\widehat{\theta}_n, n}(\mathbf{X}) \right\} - \mathbf{1} \left\{ A = d_{\theta_0}^{\text{opt}}(\mathbf{X}) \right\} \right] \right) + Z_{A, n} \\
 &\quad + \sqrt{n} \mathbb{E} \left(\mathbf{X} \left[\mathbf{1} \left\{ A = \widehat{d}_{\widehat{\theta}_n, n}(\mathbf{X}) \right\} - \mathbf{1} \left\{ A = d_{\theta_0}^{\text{opt}}(\mathbf{X}) \right\} \right] \right) \\
 &= Z_{A, n} + \sqrt{n} \mathbb{E} \left(\mathbf{X} \left[\mathbf{1} \left\{ A = \widehat{d}_{\widehat{\theta}_n, n}(\mathbf{X}) \right\} - \mathbf{1} \left\{ A = d_{\theta_0}^{\text{opt}}(\mathbf{X}) \right\} \right] \right) \{1 + o_P(1)\} + o_P(1),
 \end{aligned}$$

where $\mathbb{G}_n f = n^{1/2} (\mathbb{E}_n - \mathbb{E}) f(\mathbf{X})$ and $Z_{A, n} = n^{-1/2} \sum_{i=1}^n \left[\mathbf{1} \left\{ A_i = d_{\theta_0}^{\text{opt}}(\mathbf{X}_i) \right\} - P_{\beta_0}(\mathbf{X}_i) \right]$.

We also have $\mathbf{1} \left\{ A = \widehat{d}_{\widehat{\theta}_n, n}(\mathbf{X}) \right\} - \mathbf{1} \left\{ A = d_{\theta_0}^{\text{opt}}(\mathbf{X}) \right\} = - \left[2 \cdot \mathbf{1} \left\{ A = d_{\theta_0}^{\text{opt}}(\mathbf{X}) \right\} - 1 \right]$

$\cdot \mathbf{1} \left\{ \widehat{d}_{\widehat{\theta}_n, n}(\mathbf{X}) \neq d_{\theta_0}^{\text{opt}}(\mathbf{X}) \right\}$, which implies that

$$\begin{aligned}
 &\sqrt{n} \mathbb{E} \left(\mathbf{X} \left[\mathbf{1} \left\{ A = \widehat{d}_{\widehat{\theta}_n, n}(\mathbf{X}) \right\} - \mathbf{1} \left\{ A = d_{\theta_0}^{\text{opt}}(\mathbf{X}) \right\} \right] \right) \\
 &= -\sqrt{n} \mathbb{E} \left[\mathbf{X} \{2P_{\beta_0}(\mathbf{X}) - 1\} \mathbf{1} \left\{ \widehat{d}_{\widehat{\theta}_n, n}(\mathbf{X}) \neq d_{\theta_0}^{\text{opt}}(\mathbf{X}) \right\} \right] \\
 &= -\sqrt{n} \mathbb{E} \left\{ \mathbf{X} \{2P_{\beta_0}(\mathbf{X}) - 1\} \left(\mathbf{1} \left[0 \leq D_{\theta_0}(\mathbf{X}) < - \left\{ \widehat{D}_{\widehat{\theta}_n, n}(\mathbf{X}) - D_{\theta_0}(\mathbf{X}) \right\} \right] \right. \right. \\
 &\quad \left. \left. + \mathbf{1} \left[- \left\{ \widehat{D}_{\widehat{\theta}_n, n}(\mathbf{X}) - D_{\theta_0}(\mathbf{X}) \right\} \leq D_{\theta_0}(\mathbf{X}) < 0 \right] \right) \right\} \\
 &= -\sqrt{n} \mathbb{E} \left[\mathbf{X} \{2P_{\beta_0}(\mathbf{X}) - 1\} \left(\int_0^{-\left(\widehat{D}_{\widehat{\theta}_n} - D_{\theta_0} \right)} f_0(u) du + \int_{-\left(\widehat{D}_{\widehat{\theta}_n} - D_{\theta_0} \right)}^0 f_0(u) du \right) \right] | |D_{\theta_0}(\mathbf{X})| \leq \epsilon_n \\
 &\quad \times (1 + o_P(1))
 \end{aligned}$$

$$= -\mathbb{E} \left[\mathbf{X} \{2P_{\beta_0}(\mathbf{X}) - 1\} \cdot \left| \sqrt{n} \left\{ \widehat{D}_{\widehat{\theta}_n, n}(\mathbf{X}) - D_{\theta_0}(\mathbf{X}) \right\} \right| \middle| D_{\theta_0}(\mathbf{X}) = 0 \right] f_0 + o_P(1),$$

for some sequence ϵ_n converging down to zero slowly enough. The second to last equality holds by the fact that $\left\| \widehat{D}_{\widehat{\theta}_n, n}(\mathbf{x}) - D_{\theta_0}(\mathbf{x}) \right\|_{\mathcal{X}} = o_P(1)$ and Assumption 10 yields that the distribution of \mathbf{X} converges to a random variable independent of $D_{\theta_0}(X)$ conditional on $|D_{\theta_0}(X)| \leq \epsilon_n$, as n goes to ∞ .

Note that

$$\begin{aligned} \sqrt{n} \left\{ \widehat{D}_{\widehat{\theta}_n, n}(\mathbf{X}) - D_{\theta_0}(\mathbf{X}) \right\} &= \sqrt{n} \left[\omega(\mathbf{X}; \widehat{\theta}_n) \left\{ \widehat{R}_{Y, n}(\mathbf{X}) - R_Y(\mathbf{X}) \right\} \right. \\ &\quad \left. + \left\{ 1 - \omega(\mathbf{X}; \widehat{\theta}_n) \right\} \left\{ \widehat{R}_{Z, n}(\mathbf{X}) - R_Z(\mathbf{X}) \right\} \right] \\ &\quad + \sqrt{n} \left\{ \omega(\mathbf{X}; \widehat{\theta}_n) - \omega(\mathbf{X}; \theta_0) \right\} \left\{ R_Y(\mathbf{X}) - R_Z(\mathbf{X}) \right\} \\ &= \omega(\mathbf{X}; \theta_0) \phi_Y(\mathbf{X})^\top Z_{Y, n} + \left\{ 1 - \omega(\mathbf{X}; \theta_0) \right\} \phi_Z(\mathbf{X})^\top Z_{Z, n} \\ &\quad + \dot{\omega}_{\theta_0}(\mathbf{X}) \left\{ R_Y(\mathbf{X}) - R_Z(\mathbf{X}) \right\} \sqrt{n} \left(\widehat{\theta}_n - \theta_0 \right) \left\{ 1 + o_P(1) \right\} \\ &= O_P \left(1 + \sqrt{n} \left\| \widehat{\theta}_n - \theta_0 \right\| \right), \end{aligned}$$

thus, $\left\| \widehat{u}_n(\widehat{\theta}_n) \right\| = O_P \left(1 + \sqrt{n} \left\| \widehat{\theta}_n - \theta_0 \right\| \right)$. Letting $v_n(\widehat{\theta}_n, \beta_*) = n^{-1/2} \widehat{u}_n(\widehat{\theta}_n)^\top I_n^{-1}(\beta_*) \cdot \widehat{u}_n(\widehat{\theta}_n)$,

$$\begin{aligned} 0 &\leq n^{-1/2} \left\{ \widehat{\ell}_n(\widehat{\theta}_n, \widehat{\beta}_n) - \widehat{\ell}_n(\theta_0, \beta_0) \right\} \\ &= n^{-1/2} \sum_{i=1}^n \mathbf{X}_i^\top \beta_0 \left[1 \left\{ A_i = \widehat{d}_{\widehat{\theta}_n, n}(\mathbf{X}) \right\} - 1 \left\{ A_i = \widehat{d}_{\theta_0, n}(\mathbf{X}) \right\} \right] + v_n(\widehat{\theta}_n, \beta_*) / 2 \\ &= n^{1/2} \mathbb{E} \left(\mathbf{X}^\top \beta_0 \left[1 \left\{ A = \widehat{d}_{\widehat{\theta}_n, n}(\mathbf{X}) \right\} - 1 \left\{ A = \widehat{d}_{\theta_0, n}(\mathbf{X}) \right\} \right] \right) + o_P \left(1 + \sqrt{n} \left\| \widehat{\theta}_n - \theta_0 \right\| \right) \\ &= n^{1/2} \mathbb{E} \left(\mathbf{X}^\top \beta_0 \left[1 \left\{ A = \widehat{d}_{\widehat{\theta}_n, n}(\mathbf{X}) \right\} - 1 \left\{ A = d_{\theta_0}^{\text{opt}}(\mathbf{X}) \right\} \right] \right) \\ &\quad - n^{1/2} \mathbb{E} \left(\mathbf{X}^\top \beta_0 \left[1 \left\{ A = \widehat{d}_{\theta_0, n}(\mathbf{X}) \right\} - 1 \left\{ A = d_{\theta_0}^{\text{opt}}(\mathbf{X}) \right\} \right] \right) + o_P \left(1 + \sqrt{n} \left\| \widehat{\theta}_n - \theta_0 \right\| \right) \\ &= -\mathbb{E} \left[\mathbf{X}^\top \beta_0 \{2P_{\beta_0}(\mathbf{X}) - 1\} \cdot \left| \sqrt{n} \left\{ \widehat{D}_{\widehat{\theta}_n, n}(\mathbf{X}) - D_{\theta_0}(\mathbf{X}) \right\} \right| \middle| D_{\theta_0}(\mathbf{X}) = 0 \right] f_0 r_n \\ &\quad + O_P(1) + o_P \left(1 + \sqrt{n} \left\| \widehat{\theta}_n - \theta_0 \right\| \right) \\ &\leq -\mathbb{E} \left[\mathbf{X}^\top \beta_0 \{2P_{\beta_0}(\mathbf{X}) - 1\} \cdot \left\{ R_Y(\mathbf{X}) - R_Z(\mathbf{X}) \right\} \dot{\omega}_{\theta_0}(\mathbf{X})^\top \sqrt{n} \left(\widehat{\theta}_n - \theta_0 \right) \right] \middle| D_{\theta_0}(\mathbf{X}) = 0 \right] f_0 r_n \\ &\quad + O_P(1) + o_P \left(1 + \sqrt{n} \left\| \widehat{\theta}_n - \theta_0 \right\| \right) \\ &\leq -\delta_1^2 \left(\frac{\exp(\delta_1) - 1}{\exp(\delta_1) + 1} \right) \inf_{t \in S^d} \left\{ \Pr \left[\left| \mathbf{X}^\top \beta_0 \right| \geq \delta_1, \left| \left\{ R_Y(\mathbf{X}) - R_Z(\mathbf{X}) \right\} \dot{\omega}_{\theta_0}(\mathbf{X})^\top t \right| \geq \delta_1 \middle| D_{\theta_0}(\mathbf{X}) = 0 \right] \right\} \\ &\quad \times \sqrt{n} \left\| \widehat{\theta}_n - \theta_0 \right\| f_0 r_n + O_P(1) + o_P \left(1 + \sqrt{n} \left\| \widehat{\theta}_n - \theta_0 \right\| \right) \\ &\leq -\delta_2 \delta_1^2 \left(\frac{\exp(\delta_1) - 1}{\exp(\delta_1) + 1} \right) \sqrt{n} \left\| \widehat{\theta}_n - \theta_0 \right\| \left\{ 1 + o_P(1) \right\} + O_P(1) + o_P \left(1 + \sqrt{n} \left\| \widehat{\theta}_n - \theta_0 \right\| \right), \end{aligned}$$

where $r_n = 1 + o_P(1)$. This implies that $\sqrt{n} \|\hat{\theta}_n - \theta_0\| = O_P(1)$. From the above calculations, and the fact that the arg min or arg max of a function $t \mapsto g(t)$ does not change if we add a term that is invariant in t , we obtain that $\hat{\theta}_n = \arg \min_{t \in \mathbb{R}^d} M_n(t)$, where

$$M_n(t) = n^{-1/2} \sum_{i=1}^n \mathbf{X}_i^\top \beta_0 \left[1 \left\{ A_i = \hat{d}_{t,n}(\mathbf{X}) \right\} - 1 \left\{ A_i = d_{\theta_0}^{\text{opt}}(\mathbf{X}) \right\} \right] + v_n(t, \beta_*)/2.$$

Let $M(u) = \beta_0^\top k_0(Z_Y, Z_Z, u)$ and $U = \arg \min_{u \in \mathbb{R}^d} M(u)$. We will show that $M_n(\theta_0 + u/\sqrt{n}) \rightsquigarrow M(u)$ in $\ell^\infty(K)$ for any compact subset K of \mathbb{R}^d . Then, it will follow from the argmax Theorem (See chapter 14 of Kosorok, 2008) that $\tilde{U}_n \rightsquigarrow U$, where $\tilde{U}_n = \arg \min_{u \in \mathbb{R}^d} M_n(\theta_0 + u/\sqrt{n})$. Let $h_n(u) = \theta_0 + u/\sqrt{n}$. Similar arguments along with Assumptions 11 and 10, yield that, for any compact $K \subset \mathbb{R}^d$,

$$\begin{aligned} & \arg \min_{u \in K} M_n \{h_n(u)\} \\ &= \arg \min_{u \in K} n^{1/2} \mathbb{E} \left(\mathbf{X}^\top \beta_0 \left[1 \left\{ A = \hat{d}_{h_n(u),n}(\mathbf{X}) \right\} - 1 \left\{ A = d_{\theta_0}^{\text{opt}}(\mathbf{X}) \right\} \right] \right) + o_P(1) \\ &= \arg \min_{u \in K} n^{1/2} \mathbb{E} \left\{ \mathbf{X}^\top \beta_0 \{2P_{\beta_0}(\mathbf{X}) - 1\} \left(1 \left[- \left\{ \hat{D}_{h_n(u),n}(\mathbf{X}) - D_{\theta_0}(\mathbf{X}) \right\} \right. \right. \right. \\ &\quad \left. \left. \left. \leq D_{\theta_0}(\mathbf{X}) < 0 \right] + 1 \left[0 \leq D_{\theta_0}(\mathbf{X}) < - \left\{ \hat{D}_{h_n(u),n}(\mathbf{X}) - D_{\theta_0}(\mathbf{X}) \right\} \right] \right) \right\} + o_P(1) \\ &= \arg \min_{u \in K} \mathbb{E} \left[\mathbf{X}^\top \beta_0 \{2P_{\beta_0}(\mathbf{X}) - 1\} \left| \sqrt{n} \left\{ \hat{D}_{h_n(u),n}(\mathbf{X}) - D_{\theta_0}(\mathbf{X}) \right\} \right| \right. \\ &\quad \left. \left| D_{\theta_0}(\mathbf{X}) = 0 \right] f_0 + o_P(1), \end{aligned}$$

However,

$$\begin{aligned} \sqrt{n} \left\{ \hat{D}_{h_n(u),n}(\mathbf{X}) - D_{\theta_0}(\mathbf{X}) \right\} &\rightsquigarrow \omega(\mathbf{X}; \theta_0) \phi_Y(\mathbf{X})^\top Z_Y + \{1 - \omega(\mathbf{X}; \theta_0)\} \phi_Z(\mathbf{X})^\top Z_Z \\ &\quad + R_0(\mathbf{X}) \dot{\omega}_{\theta_0}(\mathbf{X})^\top u \end{aligned}$$

uniformly over \mathcal{X} when \mathbf{X} has its conditional distribution given $D_{\theta_0}(\mathbf{X}) = 0$. By applying continuous mapping theorem, this implies that $M_n(\theta_0 + u/\sqrt{n}) \rightsquigarrow M(u)$ in $\ell^\infty(K)$ as desired and thus $\tilde{U}_n \rightsquigarrow U$. It is straightforward to verify the remaining conclusions of the theorem using previous arguments. \blacksquare

Proof [Proof of Theorem 14] Using the assumptions, the fact that both $\sqrt{n} (\hat{\theta}_n - \theta_0) = O_P(1)$ and $\sqrt{n} (\hat{\beta}_n - \beta_0) = O_P(1)$, and standard arguments, we obtain that, for any compact $K_1 \subset \mathbb{R}^q$, $\sup_{(Z_Y^\top, Z_Z^\top)^\top \in K_1} \left\| \tilde{T}_n \left\{ \mathbf{x}, \tilde{Z}_Y(Z_Y, Z_Z), \tilde{Z}_Z(Z_Y, Z_Z) \right\} - T_0(\mathbf{x}, Z_Y, Z_Z) \right\|_{\mathcal{X}} = o_P(1)$, where

$$\begin{pmatrix} \tilde{Z}_Y(Z_Y, Z_Z) \\ \tilde{Z}_Z(Z_Y, Z_Z) \end{pmatrix} = \hat{\Sigma}_1^{1/2} \Sigma_1^{-1/2} \begin{pmatrix} Z_Y \\ Z_Z \end{pmatrix},$$

Σ_1 is the upper left $q \times q$ block of Σ_0 , and also $T_0(\mathbf{x}, Z_Y, Z_Z) = \omega(\mathbf{x}; \theta_0) \phi_Y(\mathbf{x})^\top Z_Y + \{1 - \omega(\mathbf{x}; \theta_0)\} \phi_Z(\mathbf{x})^\top Z_Z$. Furthermore,

$$\begin{aligned} & \left\| \left\{ \widehat{R}_{Y,n}(\mathbf{x}) - \widehat{R}_{Z,n}(\mathbf{x}) \right\} \dot{\omega}_{\widehat{\theta}_n}(\mathbf{x})^\top u - \{R_Y(\mathbf{x}) - R_Z(\mathbf{x})\} \dot{\omega}_{\theta_0}(\mathbf{x})^\top u \right\|_{\mathcal{X}} \\ & \leq \left\| \left\{ \widehat{R}_{Y,n}(\mathbf{x}) - \widehat{R}_{Z,n}(\mathbf{x}) \right\} \dot{\omega}_{\widehat{\theta}_n}(\mathbf{x}) - \{R_Y(\mathbf{x}) - R_Z(\mathbf{x})\} \dot{\omega}_{\theta_0}(\mathbf{x}) \right\|_{\mathcal{X}} \cdot \|u\| \\ & = O_P\left(n^{-1/2}\right) \|u\|, \end{aligned}$$

$\left\| \widehat{D}_{\widehat{\theta}_n,n}(\mathbf{x}) - D_{\theta_0}(\mathbf{x}) \right\|_{\mathcal{X}} = O_P\left(n^{-1/2}\right)$, and $\left\| P_{\widehat{\beta}_n}(\mathbf{x}) - P_{\beta_0}(\mathbf{x}) \right\|_{\mathcal{X}} = O_P\left(n^{-1/2}\right)$. Thus,

$$\begin{aligned} & \sup_{(Z_Y^\top, Z_Z^\top)^\top \in K_1} \mathbb{E}_n \left[\|\mathbf{X}\| \cdot \left| \left\{ 2P_{\widehat{\beta}_n}(\mathbf{X}) - 1 \right\} \widetilde{T}_n \left\{ \mathbf{X}, \widetilde{Z}_Y(Z_Y, Z_Z), \widetilde{Z}_Z(Z_Y, Z_Z) \right\} \right| \frac{1}{h_n} \phi_0 \left\{ \frac{\widehat{D}_{\widehat{\theta}_n,n}(\mathbf{X})}{h_n} \right\} \right] \\ & \leq O_P(1) \mathbb{E}_n \left[\frac{1}{h_n} \phi_0 \left\{ \frac{\widehat{D}_{\widehat{\theta}_n,n}(\mathbf{X})}{h_n} \right\} \right]. \quad (9) \end{aligned}$$

However,

$$\begin{aligned} & \mathbb{E}_n \left(\frac{1}{h_n} \left[\phi_0 \left\{ \frac{\widehat{D}_{\widehat{\theta}_n,n}(\mathbf{X})}{h_n} \right\} - \phi_0 \left\{ \frac{D_{\theta_0}(\mathbf{X})}{h_n} \right\} \right] \right) \\ & = \mathbb{E}_n \left[\frac{1}{h_n^3} \int_0^1 \left\{ (1-s)D_{\theta_0}(\mathbf{X}) + s\widehat{D}_{\widehat{\theta}_n,n}(\mathbf{X}) \right\} \phi_0 \left\{ \frac{(1-s)D_{\theta_0}(\mathbf{X}) + s\widehat{D}_{\widehat{\theta}_n,n}(\mathbf{X})}{h_n} \right\} ds \right. \\ & \quad \left. \cdot \left\{ \widehat{D}_{\widehat{\theta}_n,n}(\mathbf{X}) - D_{\theta_0}(\mathbf{X}) \right\} \right] \\ & = O_P\left(\frac{1}{h_n^3 n^{1/2}}\right) \\ & \quad \cdot \mathbb{E}_n \left[\int_0^1 \left\{ (1-s)D_{\theta_0}(\mathbf{X}) + s\widehat{D}_{\widehat{\theta}_n,n}(\mathbf{X}) \right\} \phi_0 \left\{ \frac{(1-s)D_{\theta_0}(\mathbf{X}) + s\widehat{D}_{\widehat{\theta}_n,n}(\mathbf{X})}{h_n} \right\} ds \right] \\ & = O_P\left(\frac{1}{h_n^3 n^{1/2}}\right) O_P(h_n) = O_P\left(\frac{1}{h_n^2 n^{1/2}}\right) = o_P(1), \end{aligned}$$

since $|u\phi_0(u)| \leq (2\pi)^{-1/2}e^{-1} < \infty$. Now, since $\mathbb{E} [h_n^{-1}\phi_0\{D_{\theta_0}(\mathbf{X})/h_n\}] \xrightarrow{P} f_0$, we have that (9) is equal to $O_P(1)$. Thus, if $\|u_n\| \rightarrow \infty$,

$$\begin{aligned} \widehat{\beta}_n^\top \widetilde{k}_n \left(\widetilde{Z}_Y, \widetilde{Z}_Z, u_n \right) & \geq \mathbb{E}_n \left[\widehat{\beta}_n^\top \mathbf{X} \left\{ 2P_{\widehat{\beta}_n}(\mathbf{X}) - 1 \right\} \cdot \left| \left\{ \widehat{R}_{Y,n}(\mathbf{X}) - \widehat{R}_{Z,n}(\mathbf{X}) \right\} \dot{\omega}_{\widehat{\theta}_n}(\mathbf{X})^\top u_n \right| \right. \\ & \quad \left. \cdot \frac{1}{h_n} \phi_0 \left\{ \frac{\widehat{D}_{\widehat{\theta}_n,n}(\mathbf{X})}{h_n} \right\} \right] - O_P(1), \quad (10) \end{aligned}$$

where the $O_P(1)$ is uniform over K_1 . Thus, up to the $O_P(1)$ added on the right-hand side,

$$\begin{aligned}
 (10) &\geq \|u_n\| \\
 &\cdot \inf_{t \in S^d} \mathbb{E}_n \left[\widehat{\beta}_n^\top \mathbf{X} \left\{ 2P_{\widehat{\beta}_n}(\mathbf{X}) - 1 \right\} \left| \left\{ \widehat{R}_{Y,n}(\mathbf{X}) - \widehat{R}_{Z,n}(\mathbf{X}) \right\} \dot{\omega}_{\widehat{\theta}_n}(\mathbf{X})^\top t \right| \frac{1}{h_n} \phi_0 \left\{ \frac{\widehat{D}_{\widehat{\theta}_n,n}(\mathbf{X})}{h_n} \right\} \right] \\
 &\geq \|u_n\| \left(o_P(1) + \inf_{t \in S^d} \mathbb{E} \left[\beta_0^\top \mathbf{X} \left\{ 2P_{\beta_0}(\mathbf{X}) - 1 \right\} \left| \left\{ R_Y(\mathbf{X}) - R_Z(\mathbf{X}) \right\} \dot{\omega}_{\theta_0}(\mathbf{X})^\top t \right| \right. \right. \\
 &\quad \left. \left. \frac{1}{h_n} \phi_0 \left\{ \frac{D_{\theta_0}(\mathbf{X})}{h_n} \right\} \right] \right) \\
 &= \|u_n\| \\
 &\cdot \left(o_P(1) + \inf_{t \in S^d} \mathbb{E} \left[\beta_0^\top \mathbf{X} \left\{ 2P_{\beta_0}(\mathbf{X}) - 1 \right\} \left| \left\{ R_Y(\mathbf{X}) - R_Z(\mathbf{X}) \right\} \dot{\omega}_{\theta_0}(\mathbf{X})^\top t \right| \mathbb{1}_{D_{\theta_0}(\mathbf{X}) = 0} \right] f_0 \right) \\
 &\geq \|u_n\| \left[o_P(1) + \delta_2 \delta_1^2 \left\{ \frac{\exp(\delta_1) - 1}{\exp(\delta_1) + 1} \right\} \right],
 \end{aligned}$$

with the expectation over \mathbf{X} . Let $\widehat{U}_n(Z_Y, Z_Z) = \arg \min_{u \in \mathbb{R}^d} \widehat{\beta}_n^\top \widetilde{k}_n \left\{ \widetilde{Z}_Y(Z_Y, Z_Z), \widetilde{Z}_Z(Z_Y, Z_Z), u \right\}$, where, if the arg min set has more than one element, one can be chosen randomly or algorithmically. Since the $O_P(1)$ above is uniform over K_1 , we conclude that

$$\sup_{(Z_Y^\top, Z_Z^\top)^\top \in K_1} \left\| \widehat{U}_n(Z_Y, Z_Z) \right\| = O_P(1). \quad (11)$$

Now, let K_2 be any compact subset of \mathbb{R}^d . Previous and standard arguments give us that $\sup_{(Z_Y^\top, Z_Z^\top)^\top \in K_1} \sup_{u \in K_2} \left\| \widetilde{k}_n \left\{ \widetilde{Z}_Y(Z_Y, Z_Z), \widetilde{Z}_Z(Z_Y, Z_Z), u \right\} - k_0(Z_Y, Z_Z, u) \right\| = o_P(1)$. Thus, we also have that

$$\sup_{(Z_Y^\top, Z_Z^\top)^\top \in K_1} \sup_{u \in K_2} \left\| \widehat{\beta}_n^\top \widetilde{k}_n \left\{ \widetilde{Z}_Y(Z_Y, Z_Z), \widetilde{Z}_Z(Z_Y, Z_Z), u \right\} - \beta_0^\top k_0(Z_Y, Z_Z, u) \right\| = o_P(1). \quad (12)$$

Define $U_0(Z_Y, Z_Z) = \arg \max_{u \in \mathbb{R}^d} \beta_0^\top k_0(Z_Y, Z_Z, u)$. Previous arguments yield that

$$\sup_{(Z_Y^\top, Z_Z^\top)^\top \in K_1} \|U_0(Z_Y, Z_Z)\| = O(1). \quad (13)$$

By Assumption 11, the arg min for each $(Z_Y^\top, Z_Z^\top)^\top \in K_1$ is unique. Fix $\epsilon > 0$. By (11), there exists an $m_2 < \infty$ such that $\Pr \left(\sup_{(Z_Y^\top, Z_Z^\top)^\top \in K_1} \left\| \widehat{U}_n(Z_Y, Z_Z) \right\| < m_2 \right) \geq 1 - \epsilon$ for all n large enough. By (13), we can enlarge m_2 such that

$$\sup_{(Z_Y^\top, Z_Z^\top)^\top \in K_1} \|U_0(Z_Y, Z_Z)\| < m_2 < \infty.$$

We can also find an $\widehat{m}_1 < \infty$ such that $K_1 \subset K_{\widehat{m}_1}^q$ as defined in Corollary 17. It is straightforward to show that (1) and (3) in Corollary 17 are satisfied by $f(Z, u) = \beta_0^\top k_0(Z_Y, Z_Z, u)$, where $Z = (Z_Y^\top, Z_Z^\top)^\top$. Let $f_n(Z, u) = \widehat{\beta}_n^\top \widetilde{k}_n(Z_Y, Z_Z, u)$. Standard arguments and the given

assumptions yield that there exists a $w_1 < \infty$ such that $\sup_{Z \in K_{m_1}^q} \sup_{u \in K_{m_2}^d} |f_n(Z, u)| < w_1$ almost surely and

$$\sup_{Z_1, Z_2 \in K_{m_1}^q: \|Z_1 - Z_2\| < \delta} \|f_n(Z_1, u) - f_n(Z_2, u)\|_{K_{m_2}^d} < w_1 \delta$$

for all $\delta > 0$ and all $n \geq 1$ almost surely. Every subsequence in (12) has a further subsequence n'' on which the convergence in probability to zero can be replaced with almost sure convergence. Thus, (2) and (4) of Corollary 17 apply, using the fact that minimizing is equivalent to maximizing after a change in sign. Setting $\widehat{U}_n^*(Z_Y, Z_Z) = \arg \min_{u \in K_{m_2}^d} \widehat{\beta}_n^\top \widetilde{k}_n \left\{ \widetilde{Z}_Y(Z_Y, Z_Z), \widetilde{Z}_Z(Z_Y, Z_Z), u \right\}$, Corollary 17 now yields that

$$\sup_{(Z_Y^\top, Z_Z^\top)^\top \in K_{m_1}^q} \left\| \widehat{U}_{n''}^*(Z_Y, Z_Z) - U_0(Z_Y, Z_Z) \right\| \rightarrow 0$$

almost surely. Since this is true for every subsequence, we have that

$$\sup_{(Z_Y^\top, Z_Z^\top)^\top \in K_{m_1}^q} \left\| \widehat{U}_n^*(Z_Y, Z_Z) - U_0(Z_Y, Z_Z) \right\| \xrightarrow{P} 0$$

as $n \rightarrow \infty$. Note that, on K_2 , $\widehat{U}_n^*(Z_Y, Z_Z) = \widehat{U}_n(Z_Y, Z_Z)$ for all $(Z_Y^\top, Z_Z^\top)^\top \in K_{m_1}^q$. Hence,

$$\begin{aligned} & \limsup_{n \rightarrow \infty} \Pr \left\{ \sup_{(Z_Y^\top, Z_Z^\top)^\top \in K_{m_1}^q} \left\| \widehat{U}_n(Z_Y, Z_Z) - U_0(Z_Y, Z_Z) \right\| > \epsilon \right\} \\ & \leq \limsup_{n \rightarrow \infty} \left[\Pr \left\{ \widehat{U}_n(Z_Y, Z_Z) \in K_2, \sup_{(Z_Y^\top, Z_Z^\top)^\top \in K_{m_1}^q} \left\| \widehat{U}_n^*(Z_Y, Z_Z) - U_0(Z_Y, Z_Z) \right\| \geq \epsilon \right\} \right. \\ & \quad \left. + \Pr \left\{ \widehat{U}_n(Z_Y, Z_Z) \in K_2^c \right\} \right] \\ & \leq \epsilon. \end{aligned}$$

Since ϵ was arbitrary, we obtain that

$$\sup_{(Z_Y^\top, Z_Z^\top)^\top \in K_{m_1}^q} \left\| \widehat{U}_n(Z_Y, Z_Z) - U_0(Z_Y, Z_Z) \right\| = o_P(1).$$

Let $B_L(\mathbb{B})$ be the space of all Lipschitz continuous functions mapping $\mathbb{B} \rightarrow \mathbb{R}$ which are bounded by 1 and which have Lipschitz constant 1. Let \mathbb{E}_Z be expectation with respect to $Z_0^* = (Z_Y^{*\top}, Z_Z^{*\top}, Z_A^{*\top})^\top \sim N(0, \Sigma_0)$. Let $B_0(Z_0^*) = I_0^{-1} \left[Z_A^* - k_0 \left\{ Z_Y^*, Z_Z^*, \widetilde{U}_n(Z_0^*) \right\} \right]$ and let $f \in B_L(\mathbb{R}^{d+p})$. Then, using \widetilde{U}_n and \widetilde{B}_n as defined in the statement of the theorem,

$$\begin{aligned} & \left| \mathbb{E}_Z \left[f \left\{ \widetilde{U}_n(Z_0^*), \widetilde{B}_n(Z_0^*) \right\} - f \left\{ U_0(Z_0^*), B_0(Z_0^*) \right\} \right] \right| \\ & \leq \left| \mathbb{E}_Z \left[f \left\{ \widetilde{U}_n(Z_0^*), \widetilde{B}_n(Z_0^*) \right\} - f \left\{ U_0(Z_0^*), \widetilde{B}_n(Z_0^*) \right\} \right] \right| \\ & \quad + \left| \mathbb{E}_Z \left[f \left\{ U_0(Z_0^*), \widetilde{B}_n(Z_0^*) \right\} - f \left\{ U_0(Z_0^*), B_0(Z_0^*) \right\} \right] \right| \\ & = \left| \mathbb{E}_Z \left[f_1 \left\{ \widetilde{U}_n(Z_0^*) \right\} - f_1 \left\{ U_0(Z_0^*) \right\} \right] \right| + \left| \mathbb{E}_Z \left[f_2 \left\{ \widetilde{B}_n(Z_0^*) \right\} - f_2 \left\{ B_0(Z_0^*) \right\} \right] \right| \end{aligned}$$

for some other $f_1 \in B_L(\mathbb{R}^d)$ and $f_2 \in B_L(\mathbb{R}^p)$. Hence,

$$\begin{aligned} & \sup_{f \in B_L(\mathbb{R}^{d+p})} \left| \mathbb{E}_Z f \left\{ \tilde{U}_n(Z_0^*), \tilde{B}_n(Z_0^*) \right\} - \mathbb{E}_Z f \left\{ U_0(Z_0^*), B_0(Z_0^*) \right\} \right| \\ & \leq \sup_{f \in B_L(\mathbb{R}^d)} \left| \mathbb{E}_Z f \left\{ \tilde{U}_n(Z_0^*) \right\} - \mathbb{E}_Z f \left\{ U_0(Z_0^*) \right\} \right| \\ & \quad + \sup_{f \in B_L(\mathbb{R}^p)} \left| \mathbb{E}_Z f \left\{ \tilde{B}_n(Z_0^*) \right\} - \mathbb{E}_Z f \left\{ B_0(Z_0^*) \right\} \right| \\ & = A_n + B_n, \end{aligned}$$

where we define both $A_n = \sup_{f \in B_L(\mathbb{R}^d)} \left| \mathbb{E}_Z f \left\{ \tilde{U}_n(Z_0^*) \right\} - \mathbb{E}_Z f \left\{ U_0(Z_0^*) \right\} \right|$ and

$$B_n = \sup_{f \in B_L(\mathbb{R}^p)} \left| \mathbb{E}_Z f \left\{ \tilde{B}_n(Z_0^*) \right\} - \mathbb{E}_Z f \left\{ B_0(Z_0^*) \right\} \right|.$$

Fix some compact $K_1 \subset \mathbb{R}^q$ such that $\Pr \left\{ (Z_Y^{*\top}, Z_Z^{*\top})^\top \in K_1 \right\} \geq 1 - \epsilon$. Then,

$$\begin{aligned} & \sup_{f \in B_L(\mathbb{R}^d)} \left| \mathbb{E}_Z \left[f \left\{ \tilde{U}_n(Z_0^*) \right\} - \mathbb{E}_Z f \left\{ U_0(Z_0^*) \right\} \right] \right| \\ & \leq \sup_{f \in B_L(\mathbb{R}^d)} \left| \mathbb{E}_Z 1(Z_0^* \in K_1) f \left\{ \tilde{U}_n(Z_0^*) \right\} - \mathbb{E}_Z 1(Z_0^* \in K_1) f \left\{ U_0(Z_0^*) \right\} \right| \\ & \quad + 2\mathbb{E}_Z 1(Z_0^* \in K_1) \\ & = o_P(1) + 2\epsilon, \end{aligned}$$

which implies that $A_n = o_P(1)$ since ϵ was arbitrary. For $K_2 \subset \mathbb{R}^{q+p}$ such that $\Pr(Z_0^* \in K_2) \geq 1 - \epsilon$, previous arguments yield that

$$\sup_{Z_0^* \in K_2} \left\| \tilde{B}_n(Z_0^*) - B_0(Z_0^*) \right\| = o_P(1).$$

As before, we can argue that $B_n = o_P(1) + 2\epsilon$, which implies that $B_n = o_P(1)$ since ϵ was arbitrary. The result follows. \blacksquare

Theorem 16 *Let H be a compact set in a metric space with metric d and let \mathcal{F} be a compact subset of $C[H]$ with respect to the uniform norm, $\|\cdot\|_H$. For each $f \in \mathcal{F}$, let $u(f) = \arg \max_{h \in H} f(h)$, where, when the arg max is not unique, we select one element of the arg max set either randomly or algorithmically. Suppose also that there exists a closed $\mathcal{F}_1 \subset \mathcal{F}$ for which each $f \in \mathcal{F}_1$ has a unique maximum. Then,*

$$\lim_{\delta \downarrow 0} \sup_{f \in \mathcal{F}_1} \sup_{g \in \mathcal{F}: \|f-g\|_H < \delta} d\{u(f), u(g)\} = 0.$$

Proof [Proof of Theorem 16] Fix $\epsilon > 0$. For each $f \in \mathcal{F}_1$, there exists $\delta_f > 0$ such that

$$\sup_{h \in H \cap B_\epsilon\{u(f)\}^c} f(h) < f\{u(f)\} - 2\delta_f,$$

where $B_\epsilon(u)$ is the open d -ball of radius ϵ around u . This follows since the compactness of \mathcal{F} ensures that all $f \in \mathcal{F}$ are continuous. Let $g \in \mathcal{F}$ be such that $\|f - g\|_H < \delta_f$. Then, $f\{u(g)\} > g\{u(g)\} - \delta_f \geq g\{u(f)\} - \delta_f > f\{u(f)\} - 2\delta_f$, which implies that $d\{u(g), u(f)\} < \epsilon$. We have that $\cup_{f \in \mathcal{F}_1} \{g \in \mathcal{F} : \|g - f\|_H < \delta_f\}$ is an open cover of \mathcal{F}_1 . Since \mathcal{F}_1 is compact, there exists a set \mathcal{F}_1^ϵ such that \mathcal{F}_1^ϵ is finite and

$$\cup_{f \in \mathcal{F}_1^\epsilon} \{g \in \mathcal{F} : \|g - f\|_H < \delta_f\}$$

still covers \mathcal{F}_1 . Let $\{f_n\} \in \mathcal{F}_1$ and $\{g_n\} \in \mathcal{F}$ be sequences such that $\|f_n - g_n\| \rightarrow 0$. By compactness, every subsequence has a further subsequence n'' such that $f_{n''} \rightarrow f_0 \in \mathcal{F}_1$ and $g_{n''} \rightarrow g_0 \in \mathcal{F}$ so that both f_0 and g_0 are in a set of the form $\{g \in \mathcal{F} : \|g - f\|_H < \delta_f\}$ for some $f \in \mathcal{F}_1^\epsilon$. This implies that $d\{u(g_0), u(f_0)\} < \epsilon$. Since the subsequence was arbitrary, we have that $\limsup_{n \rightarrow \infty} d\{u(g_n), u(f_n)\} \leq \epsilon$. Since ϵ was arbitrary, we now have that $\limsup_{n \rightarrow \infty} d\{u(g_n), u(f_n)\} = 0$ for any sequences $\{f_n\} \in \mathcal{F}_1$ and $\{g_n\} \in \mathcal{F}$ such that $\|f_n - g_n\| \rightarrow 0$. This proves the result. \blacksquare

Corollary 17 *For $m_1 < \infty$, let $K_{m_1}^q = \{z \in \mathbb{R}^q : \|z\| \leq m_1\}$. Let $(z, u) \mapsto f(z, u)$ and $(z, u) \mapsto f_n(z, u)$ be a fixed function and a sequence of functions, respectively, from $K_{m_1}^q \times \mathbb{R}^d$ to \mathbb{R} . Suppose there exists $m_2 < \infty$ such that for each $z \in K_{m_1}^q$, $u(z) = \arg \max_{u \in \mathbb{R}^d} f(z, u) < m_2$ and is uniquely defined. Suppose also that $u_n(z) = \arg \max_{u \in \mathbb{R}^d} f_n(z, u) < m_2$ for all n large enough, where we allow the arg max to be non-unique, but we randomly or algorithmically select one element from the arg max set. Define $K_{m_2}^d$ similarly to $K_{m_1}^q$ and assume that*

1. $\sup_{z \in K_{m_1}^q} \sup_{u \in K_{m_2}^d} |f(z, u)| < \infty$
2. $\limsup_{n \rightarrow \infty} \sup_{z \in K_{m_1}^q} \sup_{u \in K_{m_2}^d} |f_n(z, u)| < \infty$
3. $\lim_{\delta \downarrow 0} \sup_{z_1, z_2 \in K_{m_1}^q : \|z_1 - z_2\| < \delta} \|f(z_1, \cdot) - f(z_2, \cdot)\|_{K_{m_2}^d} = 0$
4. $\lim_{\delta \downarrow 0} \sup_{z_1, z_2 \in K_{m_1}^q : \|z_1 - z_2\| < \delta} \|f_n(z_1, \cdot) - f_n(z_2, \cdot)\|_{K_{m_2}^d} = 0$

for all n large enough. Then, provided $\sup_{z \in K_{m_1}^q} \|f_n(z, \cdot) - f(z, \cdot)\|_{K_{m_2}^d} \rightarrow 0$,

$$\sup_{z \in K_{m_1}^q} \|u_n(z) - u(z)\| \rightarrow 0,$$

as $n \rightarrow \infty$.

Proof [Corollary 17] By the Arzelà–Ascoli Theorem, there exists a compact $K \subset C[H]$ for $H = K_{m_2}^d$, such that both $f(z, \cdot) \in K$ and $f_n(z, \cdot) \in K$, for all $z \in K_{m_1}^q$ and all n large enough. If we let $\mathcal{F}_1 = \{f(z, \cdot) : z \in K_{m_1}^q\}$, we can directly apply Theorem 16 to obtain that

$$\lim_{\delta \downarrow 0} \sup_{z \in K_{m_1}^q} \sup_{g \in K : \|g - f(z, \cdot)\|_H < \delta} \|u(g) - u\{f(z, \cdot)\}\|_H = 0.$$

Because $\sup_{z \in K_{m_1}^q} \|f_n(z, \cdot) - f(z, \cdot)\|_{K_{m_2}^d} < \delta$ for all n large enough, the result follows. \blacksquare

Proof [Remark 6.] We require estimation be restricted to parameters (β, θ) which satisfy $P_\beta \{A = d_\theta(\mathbf{X}) | \mathbf{X}\} > 1/2$ with probability one. Suppose towards a contradiction that such a set of parameters also satisfies

$$\begin{aligned} P_\beta(\mathbf{X})^{1\{A=d_\theta(\mathbf{X})\}} \{1 - P_\beta(\mathbf{X})\}^{1-1\{A=d_\theta(\mathbf{X})\}} \\ = P_{\beta_0}(\mathbf{X})^{1\{A=d_{\theta_0}(\mathbf{X})\}} \{1 - P_{\beta_0}(\mathbf{X})\}^{1-1\{A=d_{\theta_0}(\mathbf{X})\}} \quad \text{a.s.,} \end{aligned}$$

where $P_\beta(\mathbf{x}) = \text{expit}(\mathbf{x}^\top \beta)$. For any (\mathbf{x}, a) such that $a = d_\theta(\mathbf{x}) \neq d_{\theta_0}(\mathbf{x})$ and it follows that $P_\beta(\mathbf{x}) = P_{-\beta_0}(\mathbf{x}) < P_{\beta_0}(\mathbf{x}) = P_{-\beta}(\mathbf{x})$ which contradicts the restriction that the probability of recommending an optimal treatment exceeds $1/2$. Thus, it follows that $d_\theta(\mathbf{x}) = d_{\theta_0}(\mathbf{x})$ for almost all \mathbf{x} . This in turn implies that $P\beta(\mathbf{x}) = P_{\beta_0}(\mathbf{x})$ for almost all \mathbf{x} from which it follows that $\beta = \beta_0$ provided $\mathbb{E}\mathbf{X}\mathbf{X}^\top$ is full rank. \blacksquare

Appendix B: Three or More Outcomes

Assume the available data consist of $(\mathbf{X}_i, A_i, Y_{1,i}, \dots, Y_{K,i})$, $i = 1, \dots, n$, which comprise n independent and identically distributed copies of $(\mathbf{X}, A, Y_1, \dots, Y_K)$, where \mathbf{X} and A are as defined previously, and Y_1, \dots, Y_K are outcomes, with each outcome coded so that higher values are better. Assume there exists an unknown utility function $U = u(Y_1, \dots, Y_K)$, where $u : \mathbb{R}^K \rightarrow \mathbb{R}$, such that $u(y_1, \dots, y_K)$ quantifies the “goodness” of the outcome vector (y_1, \dots, y_K) . As before, let $U^*(d)$ be the potential utility under a treatment regime d . Let d_U^{opt} be the optimal regime for the utility defined by u , i.e., $\mathbb{E}U^*(d_U^{\text{opt}}) \geq \mathbb{E}U^*(d)$ for any regime d . The goal is to estimate the utility function and the associated optimal regime in the presence of more than two outcomes.

To begin, we assume that the utility function is constant across patients and takes the form $u(y_1, \dots, y_K; \omega) = \sum_{k=1}^{K-1} \omega_k y_k + \left(1 - \sum_{k=1}^{K-1} \omega_k\right) y_K$, where $\omega = (\omega_1, \dots, \omega_{K-1})$ is a vector of parameters with $\sum_{k=1}^{K-1} \omega_k \leq 1$ and $\omega_k \geq 0$ for $k = 1, \dots, K-1$. Thus, we assume that the utility function is a convex combination of the set of outcomes. Let d_ω^{opt} be the optimal regime for the utility defined by ω . Assume that there exists a true utility function defined by some $\omega_0 = (\omega_{1,0}, \dots, \omega_{K-1,0})$ such that observed decisions were made with the intent to maximize $U = u(y_1, \dots, y_K; \omega_0)$. Further assume that treatment decisions in the observed data follow $\Pr \{A = d_{\omega_0}^{\text{opt}}(\mathbf{x}) | \mathbf{X} = \mathbf{x}\} = \text{expit}(\mathbf{x}^\top \beta)$, where β is an unknown parameter.

Define $Q_{Y_k}(\mathbf{x}, a) = \mathbb{E}(Y_k | \mathbf{X} = \mathbf{x}, A = a)$, for $k = 1, \dots, K$. Define also

$$Q_\omega(\mathbf{x}, a) = \mathbb{E}\{u(Y_1, \dots, Y_K; \omega) | \mathbf{X} = \mathbf{x}, A = a\}$$

and note that $Q_\omega(\mathbf{x}, a) = \sum_{k=1}^{K-1} \omega_k Q_{Y_k}(\mathbf{x}, a) + \left(1 - \sum_{k=1}^{K-1} \omega_k\right) Q_{Y_K}(\mathbf{x}, a)$. The Q -functions for each outcome can be estimated from the observed data using regression models. Let $\widehat{Q}_{Y_k, n}(\mathbf{x}, a)$ denote an estimator for $Q_{Y_k}(\mathbf{x}, a)$. Then, an estimator for $Q_\omega(\mathbf{x}, a)$ is $\widehat{Q}_{\omega, n}(\mathbf{x}, a) = \sum_{k=1}^{K-1} \omega_k \widehat{Q}_{Y_k, n}(\mathbf{x}, a) + \left(1 - \sum_{k=1}^{K-1} \omega_k\right) \widehat{Q}_{Y_K, n}(\mathbf{x}, a)$. For any fixed ω , we can compute an

estimator for d_ω^{opt} as $\hat{d}_{\omega,n}(\mathbf{x}) = \arg \max_{a \in \mathcal{A}} \hat{Q}_{\omega,n}(\mathbf{x}, a)$. The pseudo-likelihood is

$$\hat{\mathcal{L}}_n(\omega, \beta) \propto \prod_{i=1}^n \frac{\exp \left[\mathbf{X}_i^\top \beta \mathbf{1} \left\{ A_i = \hat{d}_{\omega,n}(\mathbf{X}_i) \right\} \right]}{1 + \exp(\mathbf{X}_i^\top \beta)}, \quad (14)$$

for a vector β and a vector $\omega = (\omega_1, \dots, \omega_{K-1})$. For $K = 2$, this reduces to the formulation in Section 2. Estimators for β and $\omega_1, \dots, \omega_{K-1}$ can be obtained by maximizing the pseudo-likelihood in (14). Letting $\hat{\omega}_n = (\hat{\omega}_{1,n}, \dots, \hat{\omega}_{K-1,n})$ denote the maximum pseudo-likelihood estimator for ω , an estimator for the optimal regime is $\hat{d}_{\hat{\omega}_n,n}(\mathbf{x}) = \arg \max_{a \in \mathcal{A}} \hat{Q}_{\hat{\omega}_n,n}(\mathbf{x}, a)$.

When the number of outcomes is large, maximizing (2) using the grid search proposed in Section 2.1 is infeasible. However, we can use the Metropolis algorithm similar to that proposed for a patient-specific utility function. A patient-specific utility function can be accommodated by setting $u(y_1, \dots, y_K; \mathbf{x}, \theta) = \sum_{k=1}^{K-1} \text{expit}(\mathbf{x}^\top \theta_k) y_1 + \left\{ 1 - \sum_{k=1}^{K-1} \text{expit}(\mathbf{x}^\top \theta_k) \right\} y_K$ for a vector $\theta = (\theta_1^\top, \dots, \theta_{K-1}^\top)^\top$ and maximizing the pseudo-likelihood using the Metropolis algorithm.

To examine the performance of the proposed method in the presence of more than two outcomes, we use the following generative model. As before, let $\mathbf{X} = (X_1, \dots, X_5)^\top$ be independent normal random variables with mean 0 and standard deviation 0.5. Let ϵ_1, ϵ_2 , and ϵ_3 be independent normal random variables with mean 0 and standard deviation 0.5. Given treatment assignment, outcomes are generated according to $Y_1 = A(4X_1 - 2X_2 + 2) + \epsilon_1$, $Y_2 = A(2X_1 - 4X_2 - 2) + \epsilon_2$, and $Y_3 = 1 + A(X_1 + X_2 + 1)\epsilon_3$. For a fixed $\omega = (\omega_1, \omega_2)$ and fixed $\rho \in [0, 1]$, treatment assignment is made according to $\Pr \left\{ A = d_\omega^{\text{opt}}(\mathbf{x}) \mid \mathbf{X} = \mathbf{x} \right\} = \rho$.

We set $\omega_1 = 0.2$, $\omega_2 = 0.4$, and $\rho = 0.6, 0.8$. Table 12 contains parameter estimates averaged across 500 replications along with standard deviations (in parentheses) across replications. The error rate is the proportion of samples in a testing set that were assigned the optimal treatment by the estimated policy. Table 13 contains estimated values (cal-

n	ρ	$\hat{\omega}_{1,n}$	$\hat{\omega}_{2,n}$	$\hat{\rho}_n$	Error rate
100	0.60	0.21 (0.16)	0.34 (0.20)	0.63 (0.07)	0.15 (0.11)
	0.80	0.21 (0.07)	0.42 (0.09)	0.81 (0.04)	0.04 (0.03)
200	0.60	0.21 (0.13)	0.40 (0.17)	0.62 (0.04)	0.11 (0.09)
	0.80	0.21 (0.04)	0.41 (0.06)	0.80 (0.03)	0.03 (0.02)
300	0.60	0.21 (0.12)	0.39 (0.16)	0.62 (0.03)	0.09 (0.08)
	0.80	0.20 (0.03)	0.41 (0.04)	0.80 (0.02)	0.02 (0.01)
500	0.60	0.21 (0.09)	0.41 (0.12)	0.61 (0.02)	0.06 (0.05)
	0.80	0.20 (0.02)	0.40 (0.03)	0.80 (0.02)	0.01 (0.01)

Table 12: Estimation results for simulations where utility and probability of optimal treatment are fixed, with three outcomes.

culated by generating an independent testing set following the same generative model but with decisions made according to each policy) of the optimal policy, a policy where the utility function is estimated (the proposed method), policies estimated to maximize each outcome individually, and standard of care. The proposed method results in values close

n	ρ	Optimal	Estimated utility	Y_1 only	Y_2 only	Y_3 only	Standard of care
100	0.60	1.38 (0.04)	1.28 (0.15)	1.09 (0.06)	1.00 (0.06)	0.62 (0.10)	0.59 (0.14)
	0.80	1.39 (0.04)	1.37 (0.06)	1.09 (0.06)	1.00 (0.06)	0.62 (0.11)	0.99 (0.13)
200	0.60	1.38 (0.04)	1.32 (0.12)	1.09 (0.06)	1.00 (0.06)	0.62 (0.08)	0.60 (0.10)
	0.80	1.39 (0.04)	1.38 (0.05)	1.09 (0.05)	1.00 (0.06)	0.62 (0.09)	0.98 (0.09)
300	0.60	1.38 (0.04)	1.34 (0.10)	1.09 (0.06)	1.00 (0.06)	0.63 (0.08)	0.60 (0.08)
	0.80	1.38 (0.04)	1.39 (0.05)	1.10 (0.06)	1.00 (0.06)	0.63 (0.08)	0.99 (0.07)
500	0.60	1.39 (0.04)	1.36 (0.07)	1.09 (0.05)	1.00 (0.06)	0.62 (0.07)	0.60 (0.06)
	0.80	1.39 (0.04)	1.38 (0.05)	1.10 (0.06)	1.00 (0.06)	0.63 (0.07)	0.99 (0.06)

Table 13: Value results for simulations where utility and probability of optimal treatment are fixed, with three outcomes.

to the true optimal in large samples and larger than maximizing each individual outcome across sample sizes.

Appendix C: Misspecified Model for the Utility Function

In this section, we demonstrate that the proposed method achieves reasonable performance even in the presence of a misspecified model for the utility function. Let \mathbf{X} , Y , and Z be generated as above. Let the true underlying utility function be $u(y, z; \mathbf{x}, \theta) = \omega(\mathbf{x}; \theta)y + \{1 - \omega(\mathbf{x}; \theta)\}z$, where $\omega(\mathbf{x}; \theta) = \text{expit}(1 + x_1^2 + \mathbf{x}^\top \theta_0)$ with $\theta_0 = (-0.5, 0, 0, 1, 0.5)^\top$. The misspecified model fit to estimate the utility function contained only an intercept, X_1 , X_2 , X_3 , and X_4 . Therefore, these simulations represent the setting where one important covariate and a squared term for one covariate are omitted from the model for the utility function. Treatment was assigned according to $\Pr\{A = d_\omega^{\text{opt}}(\mathbf{x}) | \mathbf{X} = \mathbf{x}\} = \text{expit}(0.5 + x_1)$. Table 14 contains the estimated value when the model for the utility function is correctly specified and when the model is incorrectly specified, along with the value of the true optimal policy and the standard of care. The proposed method produces comparable results

n	Optimal	Correct	Misspecified	Standard of Care
100	1.86 (0.07)	1.61 (0.21)	1.64 (0.20)	0.59 (0.23)
200	1.85 (0.07)	1.68 (0.16)	1.69 (0.17)	0.57 (0.16)
300	1.86 (0.07)	1.72 (0.13)	1.74 (0.13)	0.57 (0.13)
500	1.86 (0.07)	1.77 (0.10)	1.76 (0.11)	0.58 (0.10)

Table 14: First simulation results with a misspecified model for the utility function.

regardless of whether the utility function is misspecified or not.

Table 15 contains results for the same model misspecification, but where the true utility function is $\omega(\mathbf{x}; \theta) = \text{expit}(1 + 4x_1^2 + \mathbf{x}^\top \theta_0)$ with $\theta_0 = (-0.5, 0, 0, 1, 4)^\top$, i.e., the coefficients for the components that are left out of the misspecified model are larger. When the coefficients of the components left out of the utility function model are larger, the proposed method produces better results when the model is correctly specified. However, even in the presence of model misspecification, the proposed method produces results that improve upon the standard of care.

n	Optimal	Correct	Misspecified	Standard of Care
100	2.11 (0.08)	1.63 (0.28)	1.64 (0.23)	0.69 (0.26)
200	2.11 (0.08)	1.76 (0.22)	1.68 (0.18)	0.67 (0.18)
300	2.10 (0.07)	1.84 (0.19)	1.70 (0.16)	0.68 (0.15)
500	2.10 (0.08)	1.88 (0.16)	1.73 (0.15)	0.67 (0.12)

Table 15: Second simulation results with a misspecified model for the utility function.

Appendix D: Misspecified Model for the Probability of Assigning the Optimal Treatment

Similar to the model with the misspecified utility function, the model with the misspecified probability of the optimal treatment resulted in the values that are comparable to the correct model. \mathbf{X}, Y and Z were generated as above. Let $\omega(\mathbf{x}; \theta) = \text{expit}(1 + 4x_1^2 + \mathbf{x}^T \theta_0)$, with $\theta_0 = (-0.5, 0, 0, 1, 0.5)$ and $\Pr\{A = d_\omega^{\text{opt}} | \mathbf{X} = \mathbf{x}\} = \text{expit}(0.5 + 4x_1^2 + x_1)$. For this simulation, the misspecified model for the probability of assigning the optimal treatment is assumed to not include $4X_1^2$. In Table 16, the estimated value of the model with the misspecified probability of assigning the optimal treatment is confirmed to be similar to the estimated value of the correct model.

n	Optimal	Correct	Misspecified	Standard of Care
100	2.01 (0.07)	1.90 (0.14)	1.89 (0.20)	1.29 (0.23)
200	2.01 (0.08)	1.96 (0.09)	1.95 (0.12)	1.29 (0.16)
300	2.01 (0.08)	1.97 (0.09)	1.97 (0.10)	1.29 (0.13)
500	2.00 (0.08)	1.98 (0.08)	1.98 (0.08)	1.30 (0.10)

Table 16: Value results with a misspecified model for the probability of assigning the optimal treatment.

In the next simulation, let $\omega(\mathbf{x}; \theta) = \text{expit}(1 + \mathbf{x}^T \theta_0)$ with $\theta_0 = (-0.5, 0, 0, 1, 0.5)$ and the true probability of assigning the optimal treatment is $\Pr\{A = d_\omega^{\text{opt}}(\mathbf{x}) | \mathbf{X} = \mathbf{x}\} = \text{expit}(0.5 + x_1)$ when the misspecified model is assumed as a constant. In Table 17, it is noticeable that the difference between the two estimated values is similar as in Table 16. For all settings above, the estimated value of the correct model is similar to the estimated value of the models with both the misspecified utility functions and the misspecified probability of optimal treatment.

n	Optimal	Correct	Misspecified	Standard of Care
100	1.76 (0.06)	1.66 (0.19)	1.67 (0.19)	0.51 (0.22)
200	1.76 (0.06)	1.71 (0.12)	1.71 (0.13)	0.49 (0.15)
300	1.76 (0.06)	1.72 (0.11)	1.74 (0.08)	0.50 (0.13)
500	1.76 (0.06)	1.75 (0.07)	1.75 (0.06)	0.50 (0.10)

Table 17: Value results with a misspecified model for the utility function.

Appendix E: Relationship between the probability of assigning the optimal treatment and the variance of the estimator of the utility

n	ρ		0.2	0.3	0.35	0.4
	ω					
100	0.25		0.25 (0.05)	0.25 (0.09)	0.27 (0.15)	0.33 (0.24)
	0.75		0.75 (0.05)	0.75 (0.10)	0.74 (0.14)	0.66 (0.24)
200	0.25		0.25 (0.03)	0.25 (0.05)	0.25 (0.08)	0.28 (0.16)
	0.75		0.75 (0.03)	0.75 (0.05)	0.75 (0.09)	0.73 (0.13)
300	0.25		0.25 (0.02)	0.25 (0.04)	0.25 (0.07)	0.26 (0.12)
	0.75		0.75 (0.02)	0.75 (0.04)	0.75 (0.07)	0.73 (0.13)
500	0.25		0.25 (0.01)	0.25 (0.02)	0.25 (0.04)	0.25 (0.08)
	0.75		0.75 (0.01)	0.75 (0.03)	0.75 (0.05)	0.76 (0.08)

Table 18: Estimates of ω across different ρ s

n	ρ		0.6	0.65	0.7	0.8
	ω					
100	0.25		0.34 (0.24)	0.27 (0.14)	0.25 (0.10)	0.25 (0.05)
	0.75		0.66 (0.24)	0.73 (0.15)	0.74 (0.10)	0.75 (0.05)
200	0.25		0.28 (0.16)	0.25 (0.08)	0.25 (0.05)	0.25 (0.02)
	0.75		0.72 (0.16)	0.75 (0.08)	0.75 (0.05)	0.75 (0.03)
300	0.25		0.26 (0.11)	0.25 (0.07)	0.25 (0.04)	0.25 (0.02)
	0.75		0.74 (0.13)	0.75 (0.07)	0.75 (0.04)	0.75 (0.02)
500	0.25		0.25 (0.08)	0.25 (0.04)	0.25 (0.03)	0.25 (0.01)
	0.75		0.75 (0.08)	0.76 (0.05)	0.75 (0.02)	0.75 (0.01)

Table 19: Estimates of ω across different ρ s

In this section, we explore the relationship between the probability of assigning the optimal treatment and the variance of the estimator of the utility, which was mentioned in Remark 15. We empirically examine how the standard error of the estimator of the utility changes as $\Pr\{A = d_\omega^{\text{opt}}(\mathbf{x}) | \mathbf{X} = \mathbf{x}\}$ varies. For simplicity, we focus on the fixed utility setting described in 4.1. Let \mathbf{X}, Y and Z be generated as in 4.1. For each ρ s, we estimate $\hat{\omega}$ when $\omega = 0.25$ and $\omega = 0.75$.

Table 18 and Table 19 display $\hat{\omega}$ and $s.e.(\hat{\omega})$ across values of n , ω , and ρ . As ρ deviates from 0.5 (so that β is not close to 0), $\hat{\omega}$ is closer to the true ω , and the standard error of $\hat{\omega}$ decreases. Moreover, if ω_1 and ω_2 satisfy $\omega_1 = 1 - \omega_2$, their estimates and standard errors are similar as $\text{logit}(\omega_1) = -\text{logit}(\omega_2)$.

Appendix F: Performance of Metropolis optimizer in the algorithm with the patient-specific utility.

In this section, we examine the performance of the Metropolis optimizer that is used in 2.2. We randomly select 10000 points in the unit ball around θ_0 , and use the point that yields the

n	Absolute difference-MH	Absolute difference-Unit ball	Distance-MH	Distance-Unit ball
100	1.49 (0.71)	1.66 (0.30)	0.85 (0.45)	0.86 (0.12)
200	1.30 (0.49)	1.60 (0.31)	0.72 (0.31)	0.84 (0.13)
300	1.19 (0.20)	1.63 (0.29)	0.63 (0.16)	0.86 (0.13)
500	1.15 (0.12)	1.63 (0.32)	0.60 (0.09)	0.85 (0.13)

Table 20: Absolute difference and the Euclidean distance of the estimates of the patient utility and θ_0

best likelihood as the reference for the comparison. We define $\hat{\theta}$ by Metropolis algorithm as $\hat{\theta}_{\text{MH}}$, and a best point from a unit ball around θ_0 as $\hat{\theta}_{\text{Uball}}$. We obtain the distance between θ_0 and $\hat{\theta}_{\text{MH}}$ and the distance between θ_0 and $\hat{\theta}_{\text{Uball}}$ for each of 500 replications.

In Table 20, absolute difference-MH was calculated as the mean of $|\hat{\theta}_{\text{MH}} - \theta_0|$, and absolute difference-Unit ball was calculated as the mean of $|\hat{\theta}_{\text{Uball}} - \theta_0|$ for 500 replications. Similarly, Distance-MH was calculated as the mean of $\sqrt{\sum_{j=1}^6 (\hat{\theta}_{\text{MH},j} - \theta_{0,j})^2}$, and Distance-Unit ball was calculated as the mean of $\sqrt{\sum_{j=1}^6 (\hat{\theta}_{\text{Uball},j} - \theta_{0,j})^2}$ for 500 replications. The standard errors of these distances are in the parentheses.

By Table 20, it is noticeable that both absolute difference and the Euclidean distance of $\hat{\theta}_{\text{MH}}$ are smaller, and we could conclude that it is reasonable to use Metropolis algorithm when estimating a patient-specific utility.

References

- Pieter Abbeel and Andrew Y Ng. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the Twenty-first International Conference on Machine Learning*. ACM, 2004.
- David Benkeser, Andrew Mertens, John M Colford, Alan Hubbard, Benjamin F Arnold, Aryeh Stein, and Mark J van der Laan. A machine learning-based approach for estimating and testing associations with multivariate outcomes. *The international journal of biostatistics*, 1(ahead-of-print), 2020.
- D Blatt, SA Murphy, and J Zhu. A-learning for approximate planning. *Ann Arbor*, 1001: 48109–2122, 2004.
- Emily L Butler, Eric B Laber, Sonia M Davis, and Michael R Kosorok. Incorporating patient preferences into estimation of optimal individualized treatment rules. *Biometrics*, 74(1): 18–26, 2018.
- Fredrik Carlsson and Peter Martinsson. Design techniques for stated preference methods in health economics. *Health economics*, 12(4):281–294, 2003.
- Francis S Collins and Harold Varmus. A new initiative on precision medicine. *New England Journal of Medicine*, 372(9):793–795, 2015.

- Esther W de Bekker-Grob, Mandy Ryan, and Karen Gerard. Discrete choice experiments in health economics: a review of the literature. *Health economics*, 21(2):145–172, 2012.
- Elizabeth A Evans and Maria A Sullivan. Abuse and misuse of antidepressants. *Substance Abuse and Rehabilitation*, 5:107–120, 2014.
- TJ Gan, DA Lubarsky, EM Flood, T Thanh, J Mauskopf, T Mayne, and C Chen. Patient preferences for acute pain treatment. *British Journal of Anaesthesia*, 92(5):681–688, 2004.
- Charles J. Geyer and Leif T. Johnson. *mcmc: Markov Chain Monte Carlo*, 2017. URL <https://CRAN.R-project.org/package=mcmc>. R package version 0.9-5.
- S Nassir Ghaemi. Why antidepressants are not antidepressants: STEP-BD, STAR*D, and the return of neurotic depression. *Bipolar Disorders*, 10(8):957–968, 2008.
- Joseph F Goldberg. The modern pharmacopoeia: A return to depressive realism? *Bipolar Disorders*, 10(8):969–972, 2008.
- Margaret A Hamburg and Francis S Collins. The path to personalized medicine. *New England Journal of Medicine*, 2010(363):301–304, 2010.
- Robin Henderson, Phil Ansell, and Deyadeen Alshibani. Regret-regression for optimal dynamic treatment regimes. *Biometrics*, 66(4):1192–1201, 2010.
- Miguel A Hernan and James M Robins. *Causal Inference*. CRC Boca Raton, FL, 2010.
- Lewis L Judd, Hagop S Akiskal, Pamela J Schettler, Jean Endicott, Jack Maser, David A Solomon, Andrew C Leon, John A Rice, and Martin B Keller. The long-term natural history of the weekly symptomatic status of bipolar I disorder. *Archives of General Psychiatry*, 59(6):530–537, 2002.
- Rudolf Emil Kalman. When is a linear control system optimal? *Journal of Basic Engineering*, 86(1):51–60, 1964.
- Elizabeth Timberlake Kinter. *Identifying Treatment Preferences of Patients with Schizophrenia in Germany: An Application of Patient-centered Care*. PhD thesis, Johns Hopkins University, 2009.
- Michael R Kosorok and Eric B Laber. Precision medicine. *Annual review of statistics and its application*, 6:263–286, 2019.
- Michael Rene Kosorok. *Introduction to Empirical Processes and Semiparametric Inference*. Springer, New York, 2008.
- Eric B Laber, Kristin A Linn, and Leonard A Stefanski. Interactive model building for Q-learning. *Biometrika*, 101(4):831–847, 2014a.
- Eric B Laber, Daniel J Lizotte, and Bradley Ferguson. Set-valued dynamic treatment regimes for competing outcomes. *Biometrics*, 70(1):53–61, 2014b.

- Eric B Laber, Fan Wu, Catherine Munera, Ilya Lipkovich, Salvatore Colucci, and Steve Ripa. Identifying optimal dosage regimes under safety constraints: An application to long term opioid treatment of chronic pain. *Statistics in medicine*, 37(9):1407–1418, 2018.
- Kristin A Linn, Eric B Laber, and Leonard A Stefanski. Estimation of dynamic treatment regimes for complex outcomes: Balancing benefits and risks. In *Adaptive Treatment Strategies in Practice: Planning Trials and Analyzing Data for Personalized Medicine*, pages 249–262. SIAM, 2015.
- Daniel J Lizotte and Eric B Laber. Multi-objective Markov decision processes for data-driven decision support. *Journal of Machine Learning Research*, 17(211):1–28, 2016.
- Daniel J Lizotte, Michael Bowling, and Susan A Murphy. Linear fitted-Q iteration with multiple reward functions. *Journal of Machine Learning Research*, 13:3253–3295, 2012.
- K Mombaur, A Truong, and J-P Laumond. Identifying the objectives of human path generation. *Computer Methods in Biomechanics and Biomedical Engineering*, 12:189–191, 2009.
- Erica EM Moodie, Thomas S Richardson, and David A Stephens. Demystifying optimal dynamic treatment regimes. *Biometrics*, 63(2):447–455, 2007.
- Erica EM Moodie, Bibhas Chakraborty, and Michael S Kramer. Q-learning for estimating optimal dynamic treatment rules from observational data. *Canadian Journal of Statistics*, 40(4):629–645, 2012.
- Elizabeth Charlotte Moser, Sarah E Hoffe, Jessica Frakes, Todd Anthony Aguilera, Mona Karim, Lauren Elizabeth Colbert, Shalini Moningi, Ching-Wei David Tzeng, Peter F Thall, Shubham Pant, et al. Adaptive dose optimization trial of stereotactic body radiation therapy (sbrt) with or without gc4419 (avasopasem manganese) in pancreatic cancer., 2020.
- Susan A Murphy. Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society, Series B*, 65(2):331–355, 2003.
- Susan A Murphy. A generalization error for Q-learning. *Journal of Machine Learning Research*, 6:1073–1097, 2005.
- Susan A Murphy, Mark J van der Laan, James M Robins, and Conduct Problems Prevention Research Group. Marginal mean models for dynamic regimes. *Journal of the American Statistical Association*, 96(456):1410–1423, 2001.
- Thomas A Murray, Peter F Thall, and Ying Yuan. Utility-based designs for randomized comparative trials with categorical outcomes. *Statistics in medicine*, 35(24):4285–4305, 2016.
- Andrew Y Ng, Stuart J Russell, et al. Algorithms for inverse reinforcement learning. In *ICML*, pages 663–670, 2000.

- Min Qian and Susan A Murphy. Performance guarantees for individualized treatment rules. *Annals of Statistics*, 39(2):1180, 2011.
- Nathan D Ratliff, J Andrew Bagnell, and Martin A Zinkevich. Maximum margin planning. In *Proceedings of the 23rd International Conference on Machine Learning*, pages 729–736. ACM, 2006.
- James Robins, Liliana Orellana, and Andrea Rotnitzky. Estimation and extrapolation of optimal treatment and testing strategies. *Statistics in Medicine*, 27(23):4678–4721, 2008.
- James M Robins. Testing and estimation of direct effects by reparameterizing directed acyclic graphs with structural nested models. *Computation, causation, and discovery*, pages 349–405, 1999.
- James M Robins. Optimal structural nested models for optimal sequential decisions. In *Proceedings of the Second Seattle Symposium in Biostatistics*, pages 189–326. Springer, 2004.
- Paul R Rosenbaum. From association to causation in observational studies: The role of tests of strongly ignorable treatment assignment. *Journal of the American Statistical Association*, 79(385):41–48, 1984.
- Paul R Rosenbaum and Donald B Rubin. Assessing sensitivity to an unobserved binary covariate in an observational study with binary outcome. *Journal of the Royal Statistical Society, Series B*, 45(2):212–218, 1983.
- Daniel B Rubin and Mark J van der Laan. Statistical issues and limitations in personalized medicine research with clinical trials. *The International Journal of Biostatistics*, 8(1), 2012.
- Donald B Rubin. Randomization analysis of experimental data: The fisher randomization test comment. *Journal of the American statistical association*, 75(371):591–593, 1980.
- Mandy Ryan, Karen Gerard, and Mabel Amaya-Amaya. *Using discrete choice experiments to value health and health care*, volume 11. Springer Science & Business Media, 2007.
- Gary S Sachs, Andrew A Nierenberg, Joseph R Calabrese, Lauren B Marangell, Stephen R Wisniewski, Laszlo Gyulai, Edward S Friedman, Charles L Bowden, Mark D Fossey, Michael J Ostacher, et al. Effectiveness of adjunctive antidepressant treatment for bipolar depression. *New England Journal of Medicine*, 356(17):1711–1722, 2007.
- Patrick Schnell, Qi Tang, Peter Müller, and Bradley P Carlin. Subgroup inference for multiple treatments and multiple endpoints in an Alzheimer’s disease treatment trial. *The Annals of Applied Statistics*, 11(2):949–966, 2017.
- Patrick M Schnell, Qi Tang, Walter W Offen, and Bradley P Carlin. A Bayesian credible subgroups approach to identifying patient subgroups with positive treatment effects. *Biometrics*, 72(4):1026–1036, 2016.

- P.J. Schulte, A.A. Tsiatis, E.B. Laber, and M. Davidian. Q- and A-learning methods for estimating optimal dynamic treatment regimes. *Statistical Science*, 29(4):640–661, 2014.
- Vikas Soekhai, Esther W de Bekker-Grob, Alan R Ellis, and Caroline M Vass. Discrete choice experiments in health economics: past, present and future. *Pharmacoeconomics*, 37(2):201–226, 2019.
- Gregory P Strauss, Benjamin M Robinson, James A Waltz, Michael J Frank, Zuzana Kasanova, Ellen S Herbener, and James M Gold. Patients with schizophrenia demonstrate inconsistent preference judgments for affective and nonaffective stimuli. *Schizophrenia Bulletin*, 37(6):1295–1304, 2010.
- Jeremy MG Taylor, Wenting Cheng, and Jared C Foster. Reader reaction to “a robust method for estimating optimal treatment regimes” by zhang et al.(2012). *Biometrics*, 71(1):267–273, 2015.
- Peter F Thall, Hsi-Guang Sung, and Elihu H Estey. Selecting therapeutic strategies based on efficacy and death in multicourse clinical trials. *Journal of the American Statistical Association*, 97(457):29–39, 2002.
- Peter F Thall, Christopher Logothetis, Lance C Pagliaro, Sijin Wen, Melissa A Brown, Dallas Williams, and Randall E Millikan. Adaptive therapy for androgen-independent prostate cancer: A randomized selection trial of four regimens. *Journal of the National Cancer Institute*, 99(21):1613–1622, 2007.
- A.A. Tsiatis, M. Davidian, S. Holloway, and E.B. Laber. *Dynamic Treatment Regimes*. CRC Press, 2020.
- Mark J van der Laan and Maya L Petersen. Causal effect models for realistic individualized treatment and intention to treat rules. *The International Journal of Biostatistics*, 3(1), 2007.
- Michael P Wallace and Erica EM Moodie. Doubly-robust dynamic treatment regimen estimation via weighted least squares. *Biometrics*, 71(3):636–644, 2015.
- Michael P Wallace, Erica EM Moodie, and David A Stephens. Personalized dose finding using outcome weighted learning comment. *Journal of the American Statistical Association*, 111(516):1530–1534, 2016.
- Michael P Wallace, Erica EM Moodie, and David A Stephens. Reward ignorant modeling of dynamic treatment regimes. *Biometrical Journal*, 2018.
- Yuanjia Wang, Haoda Fu, and Donglin Zeng. Learning optimal personalized treatment rules in consideration of benefit and risk: with an application to treating type 2 diabetes patients with insulin therapies. *Journal of the American Statistical Association*, 113(521):1–13, 2018.
- Fan Wu, Eric B Laber, Ilya A Lipkovich, and Emanuel Severus. Who will benefit from antidepressants in the acute treatment of bipolar depression? a reanalysis of the STEP-BD

- study by Sachs et al. 2007, using q-learning. *International Journal of Bipolar Disorders*, 3(1):7, 2015.
- Tianshuang Wu. *Set Valued Dynamic Treatment Regimes*. PhD thesis, University of Michigan, 2016.
- Baqun Zhang, Anastasios A Tsiatis, Marie Davidian, Min Zhang, and Eric B Laber. Estimating optimal treatment regimes from a classification perspective. *Stat*, 1(1):103–114, 2012a.
- Baqun Zhang, Anastasios A Tsiatis, Eric B Laber, and Marie Davidian. A robust method for estimating optimal treatment regimes. *Biometrics*, 68(4):1010–1018, 2012b.
- Baqun Zhang, Anastasios A Tsiatis, Eric B Laber, and Marie Davidian. Robust estimation of optimal dynamic treatment regimes for sequential treatment decisions. *Biometrika*, 100(3):681–694, 2013.
- Yichi Zhang, Eric B Laber, Marie Davidian, and Anastasios A Tsiatis. Interpretable dynamic treatment regimes. *Journal of the American Statistical Association*, 113(524):1541–1549, 2018.
- Yingqi Zhao, Donglin Zeng, A John Rush, and Michael R Kosorok. Estimating individualized treatment rules using outcome weighted learning. *Journal of the American Statistical Association*, 107(499):1106–1118, 2012.
- Xin Zhou, Nicole Mayer-Hamblett, Umer Khan, and Michael R Kosorok. Residual weighted learning for estimating individualized treatment rules. *Journal of the American Statistical Association*, 112(517):169–187, 2017.
- Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, and Anind K Dey. Maximum entropy inverse reinforcement learning. In *AAAI*, volume 8, pages 1433–1438. Chicago, IL, USA, 2008.