Published in final edited form as:

J R Stat Soc Series B Stat Methodol. 2017 November ; 79(5): 1565–1582. doi:10.1111/rssb.12216.

Concordance-Assisted Learning for Estimating Optimal Individualized Treatment Regimes

Caiyun Fan¹, Wenbin Lu², Rui Song², and Yong Zhou^{3,4}

¹School of Business Information, Shanghai University of International Business and Economics Shanghai 201620, China

²Department of Statistics, North Carolina State University Raleigh, NC 27695, U.S.A.

³School of Statistics and Management, Shanghai University of Finance and Economics Shanghai 200433, China

⁴Academy of Mathematics and Systems Science, Chinese Academy of Sciences Beijing 100190, China

Abstract

In this article, we propose a new concordance-assisted learning for estimating optimal individualized treatment regimes. We first introduce a type of concordance function for prescribing treatment and propose a robust rank regression method for estimating the concordance function. We then find treatment regimes, up to a threshold, to maximize the concordance function, named prescriptive index. Finally, within the class of treatment regimes that maximize the concordance function, we find the optimal threshold to maximize the value function. We establish the convergence rate and asymptotic normality of the proposed estimator for parameters in the prescriptive index. An induced smoothing method is developed to estimate the asymptotic variance of the proposed estimator. We also establish the $n^{1/3}$ -consistency of the estimated optimal threshold and its limiting distribution. In addition, a doubly robust estimator of parameters in the prescriptive index is developed under a class of monotonic index models. The practical use and effectiveness of the proposed methodology are demonstrated by simulation studies and an application to an AIDS data.

Keywords

Concordance; Optimal treatment regime; Propensity score; Rank estimation; Value function

1 Introduction

An individualized treatment regime is a deterministic function of predictors, such as clinical and genetic factors of patients, which aims to account for patients' heterogeneity in response to treatment and maximize an expected clinical outcome of interest. Deriving optimal individualized treatment regimes has recently attracted a lot of attention for treating many complex diseases, such as cancer and AIDS.

There is a large and growing body of literature for estimating the optimal individualized treatment regimes with a single decision time point and multiple decision time points. The

latter is referred as optimal dynamic treatment regimes. Two main dynamic learning approaches based on backward induction were proposed for estimating optimal dynamic treatment regimes. One is the Q-learning (e.g. Watkins, 1989; Watkins and Dayan, 1992; Zhao et al., 2009, 2011; Qian and Murphy, 2011), which posits regression models for the outcome of interest and the so-called Q-functions. The other is the A-learning (e.g. Murphy, 2003; Blatt et al., 2004), which directly builds models for the contrast functions and uses doubly robust estimating equations to estimate the contrast functions by incorporating the estimated propensity score functions. Compared with the Q-learning, the A-learning is more robust to model misspecification. More recently, for deriving the optimal treatment regime with a single decision time point, Zhang et al. (2012b) formulated the problem in a missing data framework and proposed inverse propensity score weighted (IPSW) and augmented-IPSW (AIPSW) estimators for the mean potential outcome under a given treatment regime, i.e. the value function. The estimated optimal treatment regime is then obtained by maximizing the estimated value function within a class of prespecified regimes, such as linear decision rules. The value function-based optimization method was extended to estimate the optimal dynamic treatment regime in Zhang et al. (2013a). In addition, Zhao et al. (2012) and Zhang et al. (2012a) recast the estimation of the optimal treatment regime from a classification perspective and use machine learning tools, such as outcome-weighted support vector machine, to optimize the estimated value function. Other development includes the subgroup identification method of Foster et al. (2011), target population selection method of Zhao et al. (2013), and marker-guided treatment selection method of Matsouaka et al. (2014).

The value function-based estimation method proposed by Zhang et al. (2012b) is robust and appealing. It does not require to correctly specify the underlying model for the response and it finds the best treatment regime that maximizes the estimated value function in a class of interested regimes even when the true optimal regime is not contained in this class.

However, it also have some limitations. First, the convergence rates of the estimators for the parameters in the decision rules are slower than the standard $n^{1/2}$ -rate, and their asymptotic distributions are not normal. The inference of these estimators were not studied in Zhang et al. (2012b). Second, the optimization of the estimated value function can be challenging especially when the dimension of predictors is relatively large since it is very bumpy. Zhang et al. (2012b) proposed to use a genetic algorithm to search for the optimal.

In this work, we propose a new criterion for optimal treatment decision. First, we introduce a type of concordance function for prescribing treatment. Here, concordance for treatment prescription means that if one subject has bigger benefit of receiving a treatment compared with another subject, he or she is more likely to be assigned to this treatment by the regime, which is a natural requirement for a good treatment regime. Then, we propose a robust rank regression method for estimating the concordance function. The proposed rank estimator is a U-statistic of order 2 similar to the maximum rank correlation estimator that has been widely studied in econometrics literature (Han, 1987; Sherman, 1993; Cavanagh and Sherman, 1998; Chen, 2002; Abrevaya, 2003). Second, we find treatment regimes, up to a threshold, to maximize the concordance function, named prescriptive index. The optimization of the estimated concordance function can be done using a simplex algorithm, for example, the

optim function in $\mathbf{R}$, even with a relatively larger number of predictors. Finally, within the class of treatment regimes that maximize the concordance function, we find the optimal threshold to maximize the IPSW estimator of the value function. A similar method was used by Matsouaka et al. (2014) to estimate the optimal threshold in marker-guided treatment regimes. We establish the $n^{1/2}$ -consistency and asymptotic normality of the proposed estimator for parameters in the prescriptive index. An induced smoothing method is developed to estimate the asymptotic variance of the proposed estimator. We also establish the $n^{1/3}$ -consistency of the estimated optimal threshold and its limiting distribution. The asymptotic distribution of the estimated value function for the estimated optimal treatment regimes is also derived. In addition, a doubly robust estimator of parameters in the prescriptive index is developed under a class of monotonic index models.

The rest of the article is organized as follows. In Section 2, we introduce the notation, the concordance function for treatment prescription, and the concordance-assisted estimation methods of the optimal treatment regime. The asymptotic properties of the proposed estimators for the parameters in the optimal regime are also studied. Section 3 presents extensive simulation studies to demonstrate the performance of the proposed methods. An application to an AIDS clinical trial data is given in Section 4, followed by a discussion section. All of the technical proofs are provided in the Appendix.

The data that are analysed in the paper and the programs that were used to analyse them can be obtained from <http://www4.stat.ncsu.edu/~lu/programcodes.html>.

2 Our Estimation Method

2.1 Notation and IPSW Estimation

Let Y be the continuous response variable, $\mathbf{X}$ be the p -dimensional vector of covariates, and A , taking values in $\mathcal{A} = \{1, 0\}$, be the treatment indicator. It is assumed that larger value of Y implies better response. The observed data are $\{(Y_i, \mathbf{X}_i, A_i), i = 1, \dots, n\}$, which are independently and identically distributed across i . Let $Y^*(a)$ denote the potential outcome that would result if the subject was given treatment $a \in \mathcal{A}$. A treatment regime $d(\mathbf{x})$ is a deterministic function that maps $\mathbf{x} \in \mathcal{X}$ to $a \in \mathcal{A}$. An optimal treatment regime in class $\mathcal{D}$ is defined as $d^{opt} = \arg \max_{d \in \mathcal{D}} E\{Y^*(d(\mathbf{X}))\}$, where $\mathcal{D}$ is a class of treatment regimes of interest. For example, we may consider a class of linear decision rules $d(\mathbf{x}) = \mathbb{I}(\boldsymbol{\beta}' \mathbf{x} > c)$, where $\boldsymbol{\beta}$ is p -dimensional vector of parameters and c is a scalar. Here $E\{Y^*(d(\mathbf{X}))\}$ is called the value function of a given treatment regime d . To estimate the value function based on observed data, two assumptions are typically made: (i) (stable unit treatment value assumption) $Y = Y^*(1)A + Y^*(0)(1 - A)$; (ii) (no unmeasured confounders assumption) $A \perp \{Y^*(0), Y^*(1)\} | \mathbf{X}$. Based on these two assumptions, Zhang et al. (2012b) proposed an IPSW estimator for the value function, that is

$$\hat{V}_n(d) = \frac{1}{n} \sum_{i=1}^n \frac{Y_i I\{A_i = d(\mathbf{X}_i)\}}{A_i \pi(\mathbf{X}_i) + (1 - A_i)\{1 - \pi(\mathbf{X}_i)\}}, \quad (1)$$

where $\pi(X_j) = P(A_j = 1|X_j)$ is the propensity score. To search for the optimal treatment regime in a class of linear decision rules, they suggested to maximize $\hat{V}_n(d) \equiv \hat{V}_n(\beta, c)$ with respect to β and c under the constraint $\|(c, \beta')'\| = 1$, where $\|a\|$ is the Euclidean norm of a vector a .

2.2 Concordance-Assisted Learning

From now on, we consider linear decision rules $d(x) = I(\beta'x \geq c)$ for simplicity. Note that the value function $E\{Y^*(d(X))\} = E\{Y^*(1) - Y^*(0)\}d(X) + E\{Y^*(0)\}$. To maximize the value function, it is equivalent to maximize $E\{Y^*(1) - Y^*(0)\}d(X)$. Here $Y^*(1) - Y^*(0)$ is the gain of receiving treatment 1 against treatment 0 of a subject. The optimal treatment regime that maximizes the estimated value function tends to assign a subject to treatment 1 if his or her $Y^*(1) - Y^*(0) > 0$ and 0 otherwise. For an optimal treatment regime, it is also natural to require that for any two subjects i and j , if $Y_i^*(1) - Y_i^*(0) > Y_j^*(1) - Y_j^*(0)$, the regime should be more likely to assign subject i to treatment 1 compared with subject j , i.e. $\beta'X_i > \beta'X_j$ in terms of linear decision rules. This motivates us to propose a concordance-assisted learning (CAL) for estimating the optimal treatment regime in two steps. In the first step, we find β to maximize the concordance function, defined as

$$C(\beta) = E\{[Y_i^*(1) - Y_i^*(0)] - [Y_j^*(1) - Y_j^*(0)]I(\beta'X_i > \beta'X_j)\},$$

with the constraint $\|\beta\| = (\beta'\beta)^{1/2} = 1$. Let β^* denote the maximizer of $C(\beta)$. In the second step, we find c to maximize the value function $V(\beta^*, c) = E\{Y^*(I(\beta^{*'}X \geq c))\}$. Let c^* denote the maximizer. The optimal linear decision rule under the concordance-assisted learning is $d^{*,opt}(x) = I(\beta^{*'}X \geq c^*)$. Here, the index $\beta^{*'}X$ is named prescriptive index, that is, the larger the prescriptive index of a subject, the more benefit he or she tends to gain if assigned to treatment 1.

Note that the optimal treatment regime defined under the concordance-assisted learning may not maximize the value function in the class of linear decision rules, but it maximizes the concordance function. Moreover, among all linear decision rules that maximize the concordance function, it gives the maximal value function. Under the stable unit treatment value assumption and no unmeasured confounders assumption, it can be shown that

$$E\{[Y^*(1) - Y^*(0)]I(\beta'X \geq c)\} = E[D(X)I(\beta'X \geq c)] \quad \text{and}$$

$$E([Y_i^*(1) - Y_i^*(0)] - [Y_j^*(1) - Y_j^*(0)]I(\beta'X_i > \beta'X_j)) = E([D(X_i) - D(X_j)]I(\beta'X_i > \beta'X_j)),$$

where $D(X_j) = E(Y_j|A_j = 1, X_j) - E(Y_j|A_j = 0, X_j)$. For a class of monotonic index models with $D(X_i) = Q(\beta_0'X_i)$, where $Q(\cdot)$ is an unspecified strictly monotone increasing function and $\|\beta_0\| = 1$. Define $c_0 = Q^{-1}(0)$. Since $Q(\cdot)$ is a strictly monotone increasing function, it is easy to show that the maximizer of $E[Q(\beta_0'X_i)I(\beta'X \geq c)]$ is given by $\beta = \beta_0$ and $c = c_0$. In addition, following similar arguments by Cavanagh and Sherman (1998) for maximum

rank correlation estimator, the maximizer of the concordance function

$E[\{Q(\beta'_0 \mathbf{X}_i) - Q(\beta'_0 \mathbf{X}_j)\}I(\beta' \mathbf{X}_i > \beta' \mathbf{X}_j)]$ can also be shown to be $\beta^* = \beta_0$. Similarly, we have $c^* = c_0$. Therefore, the optimal treatment regime under the concordance-assisted learning coincides with the optimal treatment regime that maximizes the value function.

2.3 Estimation With Known Propensity Score

In this section we assume that the propensity score model $\pi(\mathbf{X})$ is known as in randomized clinical trials. Similar to the argument of A-learning (Murphy, 2003), we have

$$E\left[\frac{\{Y - \nu(\mathbf{X})\}\{A - \pi(\mathbf{X})\}}{\pi(\mathbf{X})\{1 - \pi(\mathbf{X})\}}|\mathbf{X}\right] = E(Y|\mathbf{X}, A=1) - E(Y|\mathbf{X}, A=0) = D(\mathbf{X}),$$

where $\nu(\mathbf{X})$ is an arbitrary function of $\mathbf{X}$. This motivates us to consider the following estimator of the concordance function

$$\begin{aligned}\hat{C}_n(\beta, \hat{\theta}) &= \frac{1}{n(n-1)} \sum_{i \neq j} \left[\frac{\{Y_i - \nu(\mathbf{X}_i, \hat{\theta})\}\{A_i - \pi(\mathbf{X}_i)\}}{\pi(\mathbf{X}_i)\{1 - \pi(\mathbf{X}_i)\}} - \frac{\{Y_j - \nu(\mathbf{X}_j, \hat{\theta})\}\{A_j - \pi(\mathbf{X}_j)\}}{\pi(\mathbf{X}_j)\{1 - \pi(\mathbf{X}_j)\}} \right] \\ &\quad \times I(\beta' \mathbf{X}_i > \beta' \mathbf{X}_j) \\ &\equiv \frac{1}{n(n-1)} \sum_{i \neq j} \Lambda_{ij}(\hat{\theta}) I(\beta' \mathbf{X}_i > \beta' \mathbf{X}_j),\end{aligned}\tag{2}$$

where $\nu(\mathbf{X}, \theta)$ is a posited parametric model for $\mu(\mathbf{X}) \equiv E(Y|\mathbf{X}, A=0)$, such as a constant model or linear model, and $\hat{\theta}$ is an estimator of θ . Define $\hat{\beta} = \operatorname{argmax}_{\|\beta\|=1} \hat{C}_n(\beta, \hat{\theta})$. In addition, an estimator of c^* is given by $\hat{c} = \arg \max_c \hat{V}_n(\hat{\beta}, c)$, where $\hat{V}_n(\hat{\beta}, c)$ is the IPSW estimator of the value function for the decision rule $d(\mathbf{x}) = I(\hat{\beta}' \mathbf{x} > c)$ as defined in (1). In our implementation, we consider a linear model for $\nu(\mathbf{X}, \theta)$, and $\hat{\theta}$ is the associated least squares estimator based on data from subjects with $A=0$. In addition, since $\hat{C}_n(\beta, \hat{\theta})$ is a U-statistic with order of 2, it is much less bumpy than the IPSW estimator of the value function, and its optimization can be directly obtained using the *optim* function in **R** even with relatively large p . Finally, since c is a scalar, the optimization of $\hat{V}_n(\hat{\beta}, c)$ can be simply done using the grid search.

Next we study the asymptotic properties of the estimators $\hat{\beta}$ and $\hat{c}$. For simplicity of presentation, we first introduce some notation. Define

$$\tau(\beta, \mathbf{X}_1, \mathbf{X}_2) = \{D(\mathbf{X}_1) - D(\mathbf{X}_2)\}I(\beta' \mathbf{X}_1 > \beta' \mathbf{X}_2) \quad \text{and}$$

$$\varrho(\beta, \mathbf{x}) = E\{\tau(\beta, \mathbf{x}, \mathbf{X})\} + E\{\tau(\beta, \mathbf{X}, \mathbf{x})\}.$$

Let $\nabla_{m\beta}(\beta, \mathbf{x})$ denote the m th partial derivative operator with respect to β , and define

$$|\nabla_m|\varrho(\beta, \mathbf{x}) = \sum_{i_1 + \dots + i_m = m} \left| \frac{\partial^m \varrho(\beta, \mathbf{x})}{\partial \beta_{i_1} \dots \partial \beta_{i_m}} \right|.$$

To establish the asymptotic results, we need the following conditions.

- (C1) The propensity score $\pi(\mathbf{x})$ is known and $0 < \pi(\mathbf{x}) < 1$ for all $\mathbf{x} \in \mathcal{X}$.
- (C2) The estimator $\hat{\theta}$ converges almost surely to a deterministic vector of parameters θ^* , and $\sqrt{n}(\hat{\theta} - \theta^*) = O_p(1)$.
- (C3) The concordance function $C(\beta) = E\{\tau(\beta, X_1, X_2)\}$ has a unique maximizer at $\beta = \beta^* = (\beta_1^*, \dots, \beta_p^*)'$ with $\|\beta^*\| = 1$.
- (C4) (i) The support of X is not contained in a proper linear subspace of $\mathbb{R}^p$; (ii) The density function of $\beta'X$ is everywhere positive for $\beta \in \mathcal{B}$, where $\mathcal{B}$ is a neighborhood of β^* ; (iii) $E\{D(X)\}^2 < \infty$ and $E\left(\left[\frac{\{Y - \nu(X, \theta^*)\}\{A - \pi(X)\}}{\pi(X)\{1 - \pi(X)\}}\right]^2\right) < \infty$.
- (C5) (i) The function $\varrho(\beta, \mathbf{x})$ is twice differentiable with respect to β ; (ii) There is an integrable function $T(\mathbf{x})$ such that for any $\mathbf{x} \in \mathcal{X}$ and β_1 and β_2 with $\|\beta_1\| = \|\beta_2\| = 1$, $\|\nabla_2 \varrho(\beta_1, \mathbf{x}) - \nabla_2 \varrho(\beta_2, \mathbf{x})\| < T(\mathbf{x})\|\beta_1 - \beta_2\|$; (iii) $E\{|\nabla_1 \varrho(\beta^*, X)|^2\} < \infty$ and $E\{|\nabla_2 \varrho(\beta^*, X)|\} < \infty$; (iv) $E\{\nabla_2 \varrho(\beta^*, X)\}$ is negative definite.
- (C6) The value function $V(\beta^*, c) = E\{Y^* (\beta^{*'} X - c)\}$ has a unique maximizer at $c = c^*$. In addition, there exist a neighborhood of c^* and a constant $M > 0$ such that $V(\beta^*, c) - V(\beta^*, c^*) = -M(c - c^*)^2$ for every c in this neighborhood.

Condition (C1) is assumed to simplify the theoretical arguments. It can be extended to the situation when the propensity score model is correctly specified, for example, using a logistic regression. The parameters in the propensity score model can be consistently estimated from data. The variation of the estimators then needs to be taken into account when deriving the asymptotic variance of $\hat{\beta}$. Condition (C2) usually holds for the least squares estimator under mild conditions. Conditions (C3) and (C4) are assumed to establish the consistency of $\hat{\beta}$. In particular, condition (C3) assumes the existence and uniqueness of population parameters that maximize the concordance function. For the class of monotonic index models, we have $C(\beta) = E[\{Q(\beta_0' X_i) - Q(\beta_0' X_j)\}I(\beta' X_i > \beta' X_j)]$. Following similar arguments by Cavanagh and Sherman (1998), condition (C2) can be shown to hold with $\beta^* = \beta_0$. Moreover, condition (ii) of (C4) is assumed to show the continuity of $C(\beta)$. It holds when there is a component of X , say X_j , that has an everywhere positive density conditional on the rest covariates and $\beta_j^* \neq 0$. Condition (iii) of (C4) is assumed to show the uniform convergence of $\hat{C}_n(\beta, \hat{\theta})$ to $C(\beta)$. Condition (C5) is assumed to ensure the asymptotic normality of $\hat{\beta}$. Conditions (C4)–(C5) are often used to establish the large sample properties of maximum rank correlation estimators (e.g. Sherman, 1993; Cavanagh and Sherman, 1998). Condition (C6) is assumed to establish the consistency and convergence rate of $\hat{c}$. For the class of monotonic index models, we have

$V(\boldsymbol{\beta}^*, c) = V(\boldsymbol{\beta}_0, c) = E\{Y_i^*(0)\} + E\{Q(\boldsymbol{\beta}_0' \mathbf{X}_i)I(\boldsymbol{\beta}_0' \mathbf{X}_i \geq c)\} = E\{Y_i^*(0)\} + \int_c^\infty Q(u)f(u)du$, where $f(\cdot)$ is the density function of $\boldsymbol{\beta}_0' \mathbf{X}_i$. Then, $V(\boldsymbol{\beta}^*, c)/c = -Q(c)f(c)$, which is less than 0 if $c > c_0 = Q^{-1}(0)$ and bigger than 0 otherwise. Therefore, $V(\boldsymbol{\beta}^*, c)$ has a unique maximizer with $c^* = c_0$. In addition, $\partial^2 V(\boldsymbol{\beta}^*, c)/\partial c^2 = -\dot{Q}(c)f(c) - Q(c)\dot{f}(c)$, where $\dot{a}(c)$ is the first derivative of $a(c)$. Then, $\partial^2 V(\boldsymbol{\beta}^*, c)/\partial c^2|_{c=c^*} = -\dot{Q}(c^*)f(c^*) < 0$. This implies that there exist a neighborhood of c^* and a constant $M > 0$ such that $V(\boldsymbol{\beta}^*, c) - V(\boldsymbol{\beta}^*, c^*) - M(c - c^*)^2$ for every c in this neighborhood. Condition (C6) holds.

Theorem 1—Under conditions (C1)–(C5), we have, as $n \rightarrow \infty$,

- i. $\|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\| \rightarrow 0$ almost surely;
- ii. $\sqrt{n}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*) \rightarrow N(0, \boldsymbol{\Sigma})$ in distribution, where $\boldsymbol{\Sigma} = \mathbf{V}^{-1} - (\mathbf{V}^{-1})' \mathbf{2V} = E\{\nabla_2 \varrho(\boldsymbol{\beta}^*, \mathbf{X})\}$, and $\mathbf{2V} = E\{\nabla_1 \varrho(\boldsymbol{\beta}^*, \mathbf{X})\nabla_1 \varrho(\boldsymbol{\beta}^*, \mathbf{X})'\}$.

Theorem 2—Under conditions (C1)–(C6), we have, as $n \rightarrow \infty$,

- i. $|\hat{c} - c^*| = O_p(n^{-1/3})$;
- ii. $n^{1/3}(\hat{c} - c^*)$ converges in distribution to $\arg\max_h G(h)$, where $G(h)$ is a two-sided Gaussian process defined in the Appendix.

Proofs of the above Theorems are delegated to the Appendix. Since $\hat{C}_n(\boldsymbol{\beta}, \hat{\boldsymbol{\theta}})$ is not a smooth function of $\boldsymbol{\beta}$, to estimate the asymptotic variance matrix of $\hat{\boldsymbol{\beta}}$, we derive an induced smoothing method similar to those studied in Brown and Wang (2005, 2007) and Pang et al. (2012). Define the smoothed concordance function $\tilde{C}_n(\boldsymbol{\beta}, \hat{\boldsymbol{\theta}}, \mathbf{H}) = E\{\hat{C}_n(\boldsymbol{\beta} + \mathbf{H}^{1/2}\mathbf{U}, \hat{\boldsymbol{\theta}})\}$, where $\mathbf{U}$ is a standard p -variate normal random vector, $\mathbf{H}$ is a $p \times p$ positive definite matrix with the order of $O(n^{-1})$, and the expectation is taken with respect to the distribution of $\mathbf{U}$. We have

$$\tilde{C}_n(\boldsymbol{\beta}, \hat{\boldsymbol{\theta}}, \mathbf{H}) = \frac{1}{n(n-1)} \sum_{i \neq j} \Lambda_{ij}(\hat{\boldsymbol{\theta}}) \left\{ 1 - \Phi \left(\frac{\mathbf{X}_{ij}' \boldsymbol{\beta}}{\sigma_{ij}^{\mathbf{H}}} \right) \right\},$$

where $\mathbf{X}_{ij} = \mathbf{X}_j - \mathbf{X}_i$, $\sigma_{ij}^{\mathbf{H}} = (\mathbf{X}_{ij}' \mathbf{H} \mathbf{X}_{ij})^{1/2}$, and $\Phi(\cdot)$ is the standard normal cumulative distribution function. Denote the first and second derivatives of $\tilde{C}_n(\boldsymbol{\beta}, \hat{\boldsymbol{\theta}}, \mathbf{H})$ with respect to $\boldsymbol{\beta}$ by

$$\tilde{S}_n(\boldsymbol{\beta}, \hat{\boldsymbol{\theta}}, \mathbf{H}) = \frac{\partial \tilde{C}_n(\boldsymbol{\beta}, \hat{\boldsymbol{\theta}}, \mathbf{H})}{\partial \boldsymbol{\beta}} = \frac{1}{n(n-1)} \sum_{i \neq j} \Lambda_{ij}(\hat{\boldsymbol{\theta}}) \phi \left(\frac{\mathbf{X}_{ij}' \boldsymbol{\beta}}{\sigma_{ij}^{\mathbf{H}}} \right) \left(\frac{\mathbf{X}_{ij}}{\sigma_{ij}^{\mathbf{H}}} \right),$$

$$\tilde{V}_n(\boldsymbol{\beta}, \hat{\boldsymbol{\theta}}, \mathbf{H}) = \frac{\partial^2 \tilde{C}_n(\boldsymbol{\beta}, \hat{\boldsymbol{\theta}}, \mathbf{H})}{\partial \boldsymbol{\beta} \partial \boldsymbol{\beta}'} = -\frac{1}{n(n-1)} \sum_{i \neq j} \Lambda_{ij}(\hat{\boldsymbol{\theta}}) \dot{\phi} \left(\frac{\mathbf{X}_{ij}' \boldsymbol{\beta}}{\sigma_{ij}^{\mathbf{H}}} \right) \left(\frac{\mathbf{X}_{ij}}{\sigma_{ij}^{\mathbf{H}}} \right)^{\otimes 2},$$

where $\phi(\cdot)$ is the density function of the standard normal distribution, $\dot{\phi}(\cdot)$ is the first derivative of $\phi(\cdot)$, and for a vector $\mathbf{v}$, $\mathbf{v}^{\otimes 2} = \mathbf{v}\mathbf{v}'$. Then, an estimator of Σ is given by

$$\hat{\Sigma}(\hat{\beta}, \hat{\theta}, \mathbf{H}) = \tilde{V}_n^{-1}(\hat{\beta}, \hat{\theta}, \mathbf{H}) \tilde{\Delta}_n(\hat{\beta}, \hat{\theta}, \mathbf{H}) \{\tilde{V}_n^{-1}(\hat{\beta}, \hat{\theta}, \mathbf{H})\}',$$

where

$$\tilde{\Delta}_n(\hat{\beta}, \hat{\theta}, \mathbf{H}) = n \widehat{\text{var}}(\tilde{S}_n(\beta^*, \hat{\theta}, \mathbf{H})) = \frac{4}{n^3} \sum_{i=1}^n \left\{ \sum_j \Lambda_{ij}(\hat{\theta}) \phi\left(\frac{\mathbf{X}_{ij}' \hat{\beta}}{\sigma_{ij}^{\mathbf{H}}}\right) \left(\frac{\mathbf{X}_{ji}}{\sigma_{ij}^{\mathbf{H}}}\right) \right\}^{\otimes 2}.$$

We have the following Theorem.

Theorem 3—Under conditions (C1)–(C5), we have that, for any positive definite matrix $\mathbf{H}$, $\hat{\Sigma}(\hat{\beta}, \hat{\theta}, \mathbf{H}) \rightarrow \Sigma$ almost surely, as $n \rightarrow \infty$.

The proof of Theorem 3 can follow similar arguments of Zhang et al. (2013b) for a self-induced smoothing approach for transformation models. The details are omitted here. In our numerical implementation, we set $\mathbf{H} = n^{-1}I_p$ for its simplicity. From our numerical experience, the results are not sensitive to the choice of $\mathbf{H}$ as long as it is in the order of $O(n^{-1})$. In practice, one more iteration may help to slightly improve the accuracy of the estimated variance of $\hat{\beta}$. Specifically, we initially set $\mathbf{H}$ as $\mathbf{H}^{(0)} = n^{-1}I_p$ and compute $\hat{\Sigma}(\hat{\beta}, \hat{\theta}, \mathbf{H}^{(0)})$. Then, we update $\mathbf{H}$ as $\mathbf{H}^{(1)} = n^{-1}\hat{\Sigma}(\hat{\beta}, \hat{\theta}, \mathbf{H}^{(0)})$. Moreover, a new estimator $\tilde{\beta}$ can be defined as the maximizer of the smoothed concordance function $\tilde{C}_n(\beta, \hat{\theta}, \mathbf{H})$. It can be shown that $\hat{\beta}$ and $\tilde{\beta}$ have the same asymptotic distribution, while $\tilde{\beta}$ may have slightly smaller standard deviation in finite samples compared with $\hat{\beta}$.

Next, we establish the asymptotic distribution of the estimated value function, which is given by

$$\hat{V}_n(\beta, c) = \frac{1}{n} \sum_{i=1}^n \frac{Y_i I\{A_i = I(\beta' \mathbf{X}_i \geq c)\}}{A_i \pi(\mathbf{X}_i) + (1 - A_i)\{1 - \pi(\mathbf{X}_i)\}}.$$

Define $V(\beta, c) = E[Y^* \{I(\beta' \mathbf{X} \geq c)\}]$.

Theorem 4—Under conditions (C1)–(C6), we have, as $n \rightarrow \infty$,

$$\sqrt{n}\{\hat{V}_n(\hat{\beta}, \hat{c}) - V(\beta^*, c^*)\} \rightarrow N(0, \Sigma_V)$$

in distribution, where Σ_V is defined in the Appendix.

2.4 Doubly Robust Estimation

If the propensity score $\pi(\mathbf{X})$ is unknown as in observational studies, a parametric model, such as a logistic regression model, is usually assumed for the propensity score. Let $\pi(\mathbf{X}, \alpha)$

denote the posited propensity score model. The parameters α can be estimated from the observed data and let $\hat{\alpha}$ denote the estimator. If the propensity score model is misspecified, although $\pi(\mathbf{X}, \hat{\alpha})$ may not consistently estimate the true propensity score $\pi(\mathbf{X})$, it can be shown that $\hat{\alpha}$ converges almost surely to a deterministic vector of parameters α^* under mild conditions. Under such a situation, $\hat{C}_n(\beta, \hat{\theta})$ is not a consistent estimator of the concordance function. To improve the robustness of the estimator, $\hat{\beta}$, proposed in the previous section, we develop a doubly robust estimation method for the class of monotonic index models:

$$D(\mathbf{X}) = Q(\beta'_0 \mathbf{X}).$$

Recall that $\mu(\mathbf{X}) = E(Y | \mathbf{X}, A = 0)$. We have

$$E \left[\frac{\{Y - \mu(\mathbf{X})\} \{A - \pi(\mathbf{X}, \alpha)\}}{\pi(\mathbf{X}, \alpha) \{1 - \pi(\mathbf{X}, \alpha)\}} | \mathbf{X} \right] = E \left[\frac{AQ(\beta'_0 \mathbf{X}) \{A - \pi(\mathbf{X}, \alpha)\}}{\pi(\mathbf{X}, \alpha) \{1 - \pi(\mathbf{X}, \alpha)\}} | \mathbf{X} \right] = Q(\beta'_0 \mathbf{X}) \frac{\pi(\mathbf{X})}{\pi(\mathbf{X}, \alpha)}.$$

Define

$$\Lambda_{ij}^{DR}(\theta, \alpha) = \frac{\{Y_i - \nu(\mathbf{X}_i, \theta)\} \{A_i - \pi(\mathbf{X}_i, \alpha)\} A_j}{\pi(\mathbf{X}_i, \alpha) \{1 - \pi(\mathbf{X}_i, \alpha)\} \pi(\mathbf{X}_j, \alpha)} - \frac{\{Y_i - \nu(\mathbf{X}_j, \theta)\} \{A_j - \pi(\mathbf{X}_j, \alpha)\} A_i}{\pi(\mathbf{X}_j, \alpha) \{1 - \pi(\mathbf{X}_j, \alpha)\} \pi(\mathbf{X}_i, \alpha)},$$

where $\nu(\mathbf{X}, \theta)$ is a posited parametric model for $\mu(\mathbf{X})$. Then, if the propensity score model is correctly specified,

$$E\{\Lambda_{ij}^{DR}(\theta, \alpha) I(\beta' \mathbf{X}_i > \beta' \mathbf{X}_j)\} = E[\{Q(\beta'_0 \mathbf{X}_i) - Q(\beta'_0 \mathbf{X}_j)\} I(\beta' \mathbf{X}_i > \beta' \mathbf{X}_j)];$$

on the other hand, if the baseline mean model $\nu(\mathbf{X}, \theta)$ is correctly specified,

$$E\{\Lambda_{ij}^{DR}(\theta, \alpha) I(\beta' \mathbf{X}_i > \beta' \mathbf{X}_j)\} = E \left[\frac{\pi(\mathbf{X}_i) \pi(\mathbf{X}_j)}{\pi(\mathbf{X}_i, \alpha) \pi(\mathbf{X}_j, \alpha)} \{Q(\beta'_0 \mathbf{X}_i) - Q(\beta'_0 \mathbf{X}_j)\} I(\beta' \mathbf{X}_i > \beta' \mathbf{X}_j) \right].$$

Under either case, the maximizer of $E\{\Lambda_{ij}^{DR}(\theta, \alpha) I(\beta' \mathbf{X}_i > \beta' \mathbf{X}_j)\}$ is $\beta = \beta_0$. This motivates us to consider the following loss function

$$\begin{aligned} \hat{C}_n^{DR}(\beta, \hat{\theta}, \hat{\alpha}) = & \frac{1}{n(n-1)} \sum_{i \neq j} \left[\frac{\{Y_i - \nu(\mathbf{X}_i, \hat{\theta})\} \{A_i - \pi(\mathbf{X}_i, \hat{\alpha})\} A_j}{\pi(\mathbf{X}_i, \hat{\alpha}) \{1 - \pi(\mathbf{X}_i, \hat{\alpha})\} \pi(\mathbf{X}_j, \hat{\alpha})} \right. \\ & \left. - \frac{\{Y_j - \nu(\mathbf{X}_j, \hat{\theta})\} \{A_j - \pi(\mathbf{X}_j, \hat{\alpha})\} A_i}{\pi(\mathbf{X}_j, \hat{\alpha}) \{1 - \pi(\mathbf{X}_j, \hat{\alpha})\} \pi(\mathbf{X}_i, \hat{\alpha})} \right] I(\beta' \mathbf{X}_i > \beta' \mathbf{X}_j). \end{aligned} \quad (3)$$

Denote $\hat{\beta}^{DR} = \arg\max_{\|\beta\|=1} \hat{C}_n^{DR}(\beta, \hat{\theta}, \hat{\alpha})$. We have the following Theorem.

Theorem 5—Assume that either the propensity score model $\pi(\mathbf{X}, \boldsymbol{\alpha})$ or the baseline mean model $\nu(\mathbf{X}, \boldsymbol{\theta})$ is correctly specified and that $D(\mathbf{X}) = E\{Y^*(1) - Y^*(0)|\mathbf{X}\} = Q(\boldsymbol{\beta}'_0 \mathbf{X})$. Under conditions (C1')–(C4') given in the Appendix, we have, as $n \rightarrow \infty$,

- i. $\|\hat{\boldsymbol{\beta}}^{DR} - \boldsymbol{\beta}_0\| \rightarrow 0$ almost surely;
- ii. $\sqrt{n}(\hat{\boldsymbol{\beta}}^{DR} - \boldsymbol{\beta}_0) \rightarrow N(0, \boldsymbol{\Sigma}^{DR})$ in distribution, where the asymptotic variance matrix $\boldsymbol{\Sigma}^{DR}$ is defined in the Appendix.

Since the asymptotic variance matrix $\boldsymbol{\Sigma}^{DR}$ has a very complicated form, the direct estimation of $\boldsymbol{\Sigma}^{DR}$ may be difficult. Following the similar techniques in Jin et al. (2001), we derive a resampling method to estimate $\boldsymbol{\Sigma}^{DR}$. Specifically, we consider the perturbed loss function

$$\begin{aligned} \hat{C}_n^{DR,*}(\boldsymbol{\beta}, \hat{\boldsymbol{\theta}}^*, \hat{\boldsymbol{\alpha}}^*) = & \frac{1}{n(n-1)} \sum_{i \neq j} \xi_i \xi_j \left[\frac{\{Y_i - \nu(\mathbf{X}_i, \hat{\boldsymbol{\theta}}^*)\} \{A_i - \pi(\mathbf{X}_i, \hat{\boldsymbol{\alpha}}^*)\} A_j}{\pi(\mathbf{X}_i, \hat{\boldsymbol{\alpha}}^*) \{1 - \pi(\mathbf{X}_i, \hat{\boldsymbol{\alpha}}^*)\} \pi(\mathbf{X}_j, \hat{\boldsymbol{\alpha}}^*)} \right. \\ & \left. - \frac{\{Y_j - \nu(\mathbf{X}_j, \hat{\boldsymbol{\theta}}^*)\} \{A_j - \pi(\mathbf{X}_j, \hat{\boldsymbol{\alpha}}^*)\} A_i}{\pi(\mathbf{X}_j, \hat{\boldsymbol{\alpha}}^*) \{1 - \pi(\mathbf{X}_j, \hat{\boldsymbol{\alpha}}^*)\} \pi(\mathbf{X}_i, \hat{\boldsymbol{\alpha}}^*)} \right] I(\boldsymbol{\beta}' \mathbf{X}_i > \boldsymbol{\beta}' \mathbf{X}_j), \end{aligned}$$

where $\xi_1, \dots, \xi_n$ are independent and identically distributed exponential variables with mean of one, $\hat{\boldsymbol{\theta}}^* = \operatorname{argmin}_{\boldsymbol{\theta}} \sum_{i=1}^n \xi_i (1 - A_i) \{Y_i - \nu(\mathbf{X}_i, \boldsymbol{\theta})\}^2$, and $\hat{\boldsymbol{\alpha}}^*$ is the solution to the equation $\operatorname{argmin}_{\boldsymbol{\alpha}} \sum_{i=1}^n \xi_i (1 - A_i) \{A_i - \pi(\mathbf{X}_i, \boldsymbol{\alpha})\} = 0$. Denote $\hat{\boldsymbol{\beta}}^{DR,*} = \operatorname{argmax}_{\|\boldsymbol{\beta}\|=1} \hat{C}_n^{DR,*}(\boldsymbol{\beta}, \hat{\boldsymbol{\theta}}^*, \hat{\boldsymbol{\alpha}}^*)$. We can use the empirical variance matrix of $\hat{\boldsymbol{\beta}}^{DR,*}$ to estimate $\boldsymbol{\Sigma}^{DR}$.

3 Simulations

3.1 Simulations for Monotonic Index Models

In the first set of simulations, we consider a class of monotonic index models

$$Y = \mu(\mathbf{X}) + AD(\mathbf{X}) + \varepsilon, \quad (4)$$

where $\mathbf{X} = (X_1, X_2, X_3, X_4)'$, A is generated from $\operatorname{Bernoulli}(\pi(\mathbf{X}))$, and ε is generated from $N(0, 0.5^2)$. Four cases are studied:

- (I) $\mu(\mathbf{X}) = 1 + \gamma'_1 \mathbf{X}$, $D(\mathbf{X}) = 2\boldsymbol{\beta}'_0 \mathbf{X}$; (II) $\mu(\mathbf{X}) = 1 + \gamma'_1 \mathbf{X}$, $D(\mathbf{X}) = \exp(\boldsymbol{\beta}'_0 \mathbf{X}) - 1$; (III) $\mu(\mathbf{X}) = 1 + \sin(\gamma'_1 \mathbf{X}) + 0.5(\gamma'_2 \mathbf{X})^2$, $D(\mathbf{X}) = (\boldsymbol{\beta}'_0 \mathbf{X})^3$; and (IV) $\mu(\mathbf{X}) = 1 + (X_1 X_2 + 0.5 X_3^2)$, $D(\mathbf{X}) = (\boldsymbol{\beta}'_0 \mathbf{X})^3$. Here X_1, X_2, X_3 and X_4 are independent and identically distributed standard normal random variables. In each case, we set $\boldsymbol{\beta}_0 = (1, 1, -1, 1)'$, $\boldsymbol{\gamma}_1 = (1, -1, 1, 1)'$, and $\boldsymbol{\gamma}_2 = (1, 0, -1, 0)'$. For all the cases, the optimal treatment regime that maximizes the value function and the optimal treatment regime that is defined by our proposed concordance-assisted learning are the same, which is given by $d^{\text{opt}}(x) = I(\boldsymbol{\beta}'_0 x \geq c_0)$ with $c_0 = 0$ and $\boldsymbol{\beta}_0 = (0.5, 0.5, -0.5, 0.5)'$ after imposing the

constraint $\|\beta_0\| = 1$. For each case, we carry out 500 simulation runs with sample size $n = 200$. The optimization in the proposed method is done by the *optim* function in **R** with the default method “Nelder-Mead” for searching the maximizer.

For the propensity score model, we first consider randomized trials with $\pi(\mathbf{X}) = 0.5$. We compare the IPSW and AIPSW estimators of Zhang et al. (2012b), the estimator obtained based on a linear regression model for both the baseline covariate effect and treatment-covariates interaction, denoted by LR, the proposed concordance-assisted learning estimator, denoted by CAL, and its doubly robust variant, denoted by CAL-DR. In all the estimators, we assume that the propensity score is known and is set as 0.5. For the AIPSW estimator, a linear model was fitted for the augmented term as done in Zhang et al. (2012b). Under case I, the fitted linear model is correctly specified, while under cases II – IV, it is not. For the CAL and CAL-DR estimators, we consider a linear model for $v(\mathbf{X}, \boldsymbol{\theta})$, where $\boldsymbol{\theta}$ is estimated based on data from subjects with $A = 0$. For the IPSW, AIPSW and LR estimators, we normalize $\hat{\beta}$ to have norm one for comparisons with the CAL and CAL-DR estimators and adjust $\hat{c}$ accordingly by $\hat{c}/\|\hat{\beta}\|$. To assess the performance of the estimators, for $\hat{\beta}$, we report the mean and standard deviation of the estimators, the mean of the estimated standard errors, and the empirical coverage probability of 95% Wald-type confidence intervals; for $\hat{c}$, we report the mean and standard deviation of the estimators. Here, the standard error of the CAL estimator for β is estimated by the proposed induced smoothing method and that of the CAL-DR estimator is estimated by the proposed resampling method. To evaluate the accuracy of the estimated optimal treatment regime $\hat{d}^{opt}(\mathbf{x}) = I(\hat{\beta}'\mathbf{x} \geq \hat{c})$, we report the mean and standard deviation of the percentages of making correct decisions (PCD), defined as

$1 - n^{-1} \sum_{i=1}^n |I(\hat{\beta}'\mathbf{X}_i \geq \hat{c}) - I(\beta_0'\mathbf{X}_i \geq c_0)|$. Furthermore, we report the mean and standard deviation of the value functions and concordance functions for the estimated optimal treatment regime, which are obtained via simulations. Specifically, we generate data for $N = 10,000$ subjects from model (4) and obtain the value function and the concordance function for $\hat{d}^{opt}(\mathbf{x})$ by

$$\hat{V} = \frac{1}{N} \sum_{i=1}^N \left\{ \mu(\mathbf{X}_i) + \hat{d}^{opt}(\mathbf{X}_i) D(\mathbf{X}_i) \right\} \quad \text{and}$$

$$\hat{C} = \frac{1}{N(N-1)} \sum_{i \neq j} \{D(\mathbf{X}_i) - D(\mathbf{X}_j)\} I(\hat{\beta}'\mathbf{X}_i > \hat{\beta}'\mathbf{X}_j),$$

respectively. Similarly, we can compute the value function and the concordance function for the true optimal treatment regime, $d^{opt}(\mathbf{x})$, denoted by V_0 and C_0 , respectively. In all the tables, we report the mean of the estimators (denoted as Est.), standard deviation of the estimators (SD), the mean of the estimated standard errors (SE) and the empirical coverage probability (CP%) of Wald-type 95% confidence interval.

The simulation results for Cases I and II are summarized in Table 1 and those for Cases III and IV are given in the Supplementary Appendix for saving space. From the simulation

results, we make the following observations. First, our proposed CAL and CAL-DR estimators for β are nearly unbiased under all cases and have smaller biases than the IPSW and AIPSW estimators. In addition, the standard deviations of the CAL and CAL-DR estimators for β are much smaller than those of the IPSW and AIPSW estimators, indicating the big efficiency gain of the concordance-assisted learning for estimating the optimal treatment regime. One possible reason is that the CAL and CAL-DR estimators have faster convergence rate than the IPSW and AIPSW estimators. Second, the mean of estimated standard errors of the CAL and CAL-DR estimators for β is close to the standard deviation of the estimators and the empirical coverage probability of 95% confidence intervals is close to the nominal level under all cases. Third, the CAL-DR estimator for β has smaller standard deviations than the CAL estimator, indicating that the CAL-DR estimator may be more efficient than the CAL estimator. Fourth, the estimated value function and concordance function of the optimal treatment regimes obtained by the CAL and CAL-DR estimators are all close to their true values and are larger than those obtained by the IPSW estimator. Fifth, the PCDs of the optimal treatment regimes obtained by the CAL and CAL-DR estimators range from 0.875 to 0.925, which are much higher than those obtained by the IPSW estimator. Sixth, for case I, where the fitted linear model was correctly specified, the AIPSW estimators of β have much smaller standard deviations than the IPSW estimators, showing the efficiency gain of the AIPSW estimators as studied in the literature. However, for cases II – IV, where the fitted linear model was misspecified, the AIPSW and IPSW estimators have comparable performance in terms of the value function and PCD. Finally, the LR estimator has the best performance under case I as expected since the linear regression model is correctly specified, but it may have worse performance than the CAL and CAL-DR estimators when the linear model is misspecified. For example, in case II, the LR estimators of β have larger biases and standard deviations, and the optimal treatment regimes obtained by the LR estimators give smaller value function, PCD and concordance function.

For comparison, we also implemented the method of Zhao et al. (2013). For saving space, the results are provided in the Supplementary Appendix. We made the following observations. First, the method of Zhao et al. (2013) depends on the choice of ξ . In practice, a data-adaptive way is needed for choosing the optimal ξ . Second, for case I, where the linear models are correctly specified, the method of Zhao et al. (2013) with the best choice of ξ and the proposed methods have comparable performance in terms of the value ($\hat{V}$) and PCD. However, for cases II – IV, where fitted linear models are misspecified, the estimated optimal treatment rules obtained by the proposed methods give larger values and PCD than those of Zhao et al. (2013). Third, the estimated optimal treatment rules obtained by the proposed methods give larger concordance values ($\hat{C}$) with much smaller standard deviations than those of Zhao et al. (2013) for all cases. In summary, the proposed methods showed very competitive performance.

In addition, we compare the proposed methods and the method of Matsouaka et al. (2014) to examine the inference of the estimated value function. The results are provided in the Supplementary Appendix. All the methods have proper coverage probabilities close to the nominal level, and the estimators of the value functions have comparable standard deviations.

We also conducted simulations with $p = 10$ covariates, generated from the standard normal distribution. The results are provided in the Supplementary Appendix. It can be seen that the proposal methods give larger concordance, value and PCD compared with the IPSW method and require much shorter computational time on average. This demonstrates the ability of the proposed methods for handling relatively large number of covariates.

Next we consider simulations when the propensity score $\pi(\mathbf{X})$ is unknown as in observation studies and the posited model is misspecified. The simulation results are given in the Supplementary Appendix. We observe that the IPSW and CAL estimators have relatively large biases under all cases as expected; the CAL-DR estimator is nearly unbiased under Cases I and II and has much smaller biases than the IPSW and CAL estimators under Cases III and IV; the mean of the estimated standard errors of the CAL-DR estimators is close to the standard deviation of the estimators and the empirical coverage probability of 95% confidence intervals is close to the nominal level. These findings show that the CAL-DR estimator has the double robustness property under Cases I and II when the propensity score model is misspecified and has superior performance compared with the IPSW and CAL estimators under Cases III and IV when both the baseline mean model and the propensity score model are misspecified.

3.2 Simulations for General Models

In this Section, we consider several models where a monotonic index model for $D(\mathbf{X})$ is violated and the true optimal treatment regime may or may not be defined by a linear decision rule. We first consider model (4) with the following setting:

(V) $\mu(\mathbf{X}) = 1 + \sin(\gamma_1' \mathbf{X}) + 0.5(\gamma_2' \mathbf{X})^2$, $D(\mathbf{X}) = |X_2^2 + X_1 X_2 + X_3 X_4| \text{sign}(X_2 - X_1^2 + X_3 - X_4^2)$, where γ_1 and γ_2 are defined the same as before. Here, we only consider randomized studies with $\pi(\mathbf{X}) = 0.5$. In this case, the true optimal treatment regime is given by

$d^{\text{opt}}(\mathbf{X}) = I(X_2 - X_1^2 + X_3 - X_4^2 > 0)$, which is not a linear decision rule. We compare the mean and standard deviation of the IPSW, AIPSW, LR, CAL and CAL-DR estimators and the value function and concordance function of the estimated optimal treatment regime obtained by different methods. For all the methods, we search the optimal treatment regime in the class of linear decision rules. The simulation results are summarized in Table 2. From the simulation results, it can be seen that the CAL and CAL-DR methods give similar estimates of β and c , which may have relatively big difference from the estimates obtained by the IPSW and AIPSW methods; the concordance function of the estimated optimal treatment regimes obtained by the CAL and CAL-DR methods are larger than those obtained by the IPSW and AIPSW methods; the value function of the estimated optimal treatment regimes obtained by the CAL and CAL-DR methods are comparable or slightly larger than those obtained by the IPSW, AIPSW and LR methods and are closer to the value function of the true optimal treatment regime $d^{\text{opt}}(\mathbf{X}) = I(X_2 - X_1^2 + X_3 - X_4^2 > 0)$.

In addition, we consider the model

$$Y = \exp\{\mu(\mathbf{X}) + A Q(\mathbf{X})\} + \varepsilon,$$

where $\varepsilon \sim N(0, 0.5^2)$, $A \sim \text{Bernoulli}(0.5)$ and $\mathbf{X}$ is generated the same as before. The following case is studied: $(\text{VI}) \mu(\mathbf{X}) = 0.1\{1 + (X_1 X_2 + 0.5 X_3^2)\}$, $Q(\mathbf{X}) = \beta_0' \mathbf{X} - 1$, where β_0 is defined the same as before. In this case, the true optimal treatment regime is given by $d^{\text{opt}}(\mathbf{X}) = I(\beta_0' \mathbf{X} > 1)$, which is a linear decision rule. Here, the contrast function $D(\mathbf{X}) = \exp\{\mu(\mathbf{X}) + Q(\mathbf{X})\} - \exp\{\mu(\mathbf{X})\}$, which is not a monotonic index model. For all the methods, we search the optimal treatment regime in the class of linear decision rules. Thus, the IPSW and AIPSW estimators are consistent but the CAL and CAL-DR are not. The simulation results are also summarized in Table 2. From the simulation results, we observe that the IPSW and AIPSW estimators of β have relatively small biases under all cases as expected; compared with the IPSW and AIPSW estimators, the CAL and CAL-DR estimators have comparable biases under Cases VI but with smaller standard deviations; the concordance function of the estimated optimal treatment regimes obtained by the CAL and CAL-DR methods are larger than those obtained by the IPSW and AIPSW methods; the value function of the estimated optimal treatment regimes obtained by the CAL and CAL-DR methods are slightly larger than those obtained by the IPSW and AIPSW methods and are closer to the value function of the true optimal treatment regime $d^{\text{opt}}(\mathbf{X}) = I(\beta_0' \mathbf{X} > 1)$. The LR method shows worst performance under this case, yielding much smaller value function and PCD than other methods.

In summary, based on the above simulations the CAL and CAL-DR methods show competitive performance even when the monotonic index model for $D(\mathbf{X})$ is violated.

4 Real Data Analysis

We demonstrate the proposed method by an application to the data from AIDS Clinical Trials Group Protocol 175 (ACTG175), which consists of 2139 HIV-infected subjects. The enrolled subjects were randomized to four different treatment groups: zidovudine (ZDV) monotherapy, ZDV+didanosine (ddI), ZDV+zalcitabine, and ddI monotherapy. Here, we focus on the subset of patients receiving the treatment ZDV+ddI or ZDV+zalcitabine. Treatment indicator $A = 0$ denotes for the treatment ZDV+zalcitabine (524 subjects), while $A = 1$ denotes for the treatment ZDV+ddI (522 subjects). The CD4 count (cells/mm³) at 20 $\pm$ 5 weeks post-baseline is chosen to be the continuous response variable Y . Twelve covariates are considered, including five continuous variables: age (years), weight (kg), Karnofsky score (scale of 0 – 100), CD4 count (cells/mm³) at baseline, and CD8 count (cells/mm³) at baseline; and seven binary variables: hemophilia (0=no, 1=yes), homosexual activity (0=no, 1=yes), history of intravenous drug use (0=no, 1=yes), race (0=white, 1=non-white), gender (0=female, 1=male), antiretroviral history (0=naive, 1=experienced), and symptomatic status (0=asymptomatic, 1=symptomatic). Here, the propensity score is known and set as $\pi(\mathbf{X}) \equiv 0.5$. We consider both the CAL and CAL-DR estimators. In our analysis, the five continuous covariates are normalized to have mean zero and norm one. The CAL and CAL-DR estimators of β and their standard errors are given in Table 3. It can be seen that the CAL and CAL-DR methods yield comparable estimates of β while the CAL-DR estimator generally has smaller standard errors compared with the CAL estimator. In addition, from the P -values reported in Table 3, age is the only significant covariate at the level of 0.05 based on both methods. We refit the CAL and CAL-DR estimators with age as

the only covariate. The two methods yield the same estimated optimal treatment regime, which is given by $I(\text{age} > 37.5)$. The results suggest that ZDV+zalcitabine ($A=0$) is more favorable to young patients with AIDS, but ZDV+ddI ($A=1$) for old patients. A similar finding was also observed in Lu et al. (2013).

Furthermore, to compare the CAL and CAL-DR estimators with the IPSW estimator, we randomly divide the dataset into half training and half testing sets for 200 times. For each random split, we compute the CAL, CAL-DR and IPSW estimators based on the training set with age as the only covariate. Then, we compute the estimated value function $\hat{V}_n$ defined in (1) for the estimated optimal treatment regimes obtained by the CAL, CAL-DR and IPSW estimators. The averages of $\hat{V}_n$ are 395.1 (12.2) for the CAL and CAL-DR estimators, and 386.5 (18.9) for the IPSW estimators, respectively, where the quantities in the parenthesis are the corresponding standard deviations. These results demonstrate that the optimal treatment regimes obtained by the CAL and CAL-DR methods give larger value functions on average but with smaller standard deviations than that obtained by the IPSW method.

5 Discussion

We propose new methods for estimating the optimal treatment regime based on concordance-assisted learning. The estimators of the parameters in the prescriptive index for treatment decision have standard $n^{1/2}$ -rate, and their asymptotic distributions and inference are studied. It is worth noting that although the proposed concordance-assisted learning has appealing interpretation and works well under monotonic index models, the proposed procedure will not always recommend patients with positive $D(\mathbf{X}) = E(Y|A=1, \mathbf{X}) - E(Y|A=0, \mathbf{X})$ to receive treatment 1.

In this paper, we only consider a class of linear decision rules for simplicity. However, our method can be extended to incorporate a more general class of decision rules as $d(\mathbf{X}) = I\{g(\mathbf{X}) \geq c\}$, where $g(\cdot)$ is a nonparametric function of $\mathbf{X}$. The corresponding concordance function is defined as

$$C(g) = E[\{Y_i^*(1) - Y_i^*(0)\} - \{Y_j^*(1) - Y_j^*(0)\} I\{g(\mathbf{X}_i) > g(\mathbf{X}_j)\}].$$

and the associated optimal treatment regime is given by $d^{*,opt}(\mathbf{X}) = I\{g^*(\mathbf{X}) \geq c^*\}$, where $g^* = \arg\max_g C(g)$ and $c^* = \arg\max_c E\{Y^*[I\{g^*(\mathbf{X}) \geq c\}]\}$.

In addition, it is interesting to generalize the proposed methods to multiple-treatment setup. For example, consider a study with three treatments, denoted as 0, 1 and 2. Assume

monotonic index models $D_1(\mathbf{X}) \equiv E\{Y^*(1) - Y^*(0)|\mathbf{X}\} = Q_1(\beta_1' \mathbf{X})$ and

$D_2(\mathbf{X}) \equiv E\{Y^*(2) - Y^*(0)|\mathbf{X}\} = Q_2(\beta_2' \mathbf{X})$, where $Q_1(\cdot)$ and $Q_2(\cdot)$ are two strictly increasing functions. We can use the proposed methods to estimate β_1 and β_2 based on comparison of two treatments at a time: 1 vs. 0 and 2 vs. 0, respectively. The resulting

estimated optimal treatment rules are denoted as $\hat{d}_{10}^{opt}(\mathbf{x})$ and $\hat{d}_{20}^{opt}(\mathbf{x})$. Here, if $\hat{d}_{10}^{opt}(\mathbf{x})=1$, a subject with covariates $\mathbf{X}=\mathbf{x}$ is given treatment 1, and 0 otherwise. The rule $\hat{d}_{20}^{opt}(\mathbf{x})$ is

similarly defined. Therefore, when $\hat{d}_{10}^{\text{opt}}(x) = \hat{d}_{20}^{\text{opt}}(x) = 0$, treatment 0 is given; when $\hat{d}_{10}^{\text{opt}}(x) = 1$ and $\hat{d}_{20}^{\text{opt}}(x) = 0$, treatment 1 is given; when $\hat{d}_{10}^{\text{opt}}(x) = 0$ and $\hat{d}_{20}^{\text{opt}}(x) = 1$, treatment 2 is given; when $\hat{d}_{10}^{\text{opt}}(x) = \hat{d}_{20}^{\text{opt}}(x) = 1$, we will apply the concordance assisted learning based on the comparison between treatments 1 and 2, and obtain the resulting estimated optimal treatment rule $\hat{d}_{21}^{\text{opt}}(x)$, which assigns treatment 2 if $\hat{d}_{21}^{\text{opt}}(x) = 1$ and 1 otherwise. However, generalization to multiple category treatment is not straightforward other than using the above ad-hoc pairwise comparison strategy and its properties need to be investigated. Moreover, the proposed methods can also be extended to incorporate multiple decision time points. The related discussions are given in the Supplementary Appendix.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The authors are grateful to two referees, an Associate Editor and the Editor for their thoughtful and suggestive comments, which have helped to greatly improve on an earlier manuscript. The work was partially supported by National Institutes of Health grants P01 CA142538.

Appendix

Here, we only give the proofs for Theorems 1 and 2. Those for Theorems 4 and 5 are given in the Supplementary Appendix.

Appendix A: Proof of Theorem 1

To establish the consistency of $\hat{\beta}$, similar to Cavanagh and Sherman (1998), we need to show: (i) $C(\beta)$ has a unique maximizer at β^* ; (ii) $\sup_{\|\beta\|=1} |\hat{C}_n(\beta, \hat{\theta}) - C(\beta)| = o_p(1)$; (iii) $C(\beta)$ is continuous. Here, (i) is assumed by condition (C3), which holds for the class of monotonic index models. Under condition (C1), we have $E\{\Lambda_{12}(\theta)|X_1, X_2\} = D(X_1) - D(X_2)$ for any fixed θ . Define $f_{ij}(\beta, \theta) = \Lambda_{ij}(\theta)I(\beta' X_i > \beta' X_j) - C(\beta)$. Then, $\hat{C}_n(\beta, \theta) - C(\beta) = U_n f_{ij}(\beta, \theta)$, where U_n denotes the random measure putting mass $1/(n(n-1))$ on each pair of data. Therefore, $U_n f_{ij}(\beta, \theta)$ is a zero-mean U-process of order 2. In addition, by condition (iii) of (C4), it can be shown that the class $\{f_{12}(\beta, \theta) : \|\beta\| = 1, \theta - \theta^* = O_p(n^{-1/2})\}$ is Euclidean with a square-integrable envelope. Thus, (ii) holds. Finally, condition (ii) of (C4) implies that $P(\beta' X_1 = \beta' X_2) = 0$ for $\beta \in \mathcal{B}$. Then, $\tau(\beta_m, X_1, X_2) \rightarrow \tau(\beta, X_1, X_2)$ in probability, where $\{\beta_m\}$ is a sequence of elements of $\mathcal{B}$ converging to β as m goes to infinity. Applying the dominated convergence theorem, we have $C(\beta_m) \rightarrow C(\beta)$ as $m \rightarrow \infty$. Thus, (iii) is proved. The consistency of $\hat{\beta}$ then follows Amemiya (1985, pp.106–107).

Next, we establish the limiting distribution of $\hat{\beta}$. We have

$$\hat{C}_n(\beta, \hat{\theta}) - \hat{C}_n(\beta^*, \hat{\theta}) = \hat{C}_n(\beta, \hat{\theta}) - \hat{C}_n(\beta, \theta^*) - \{\hat{C}_n(\beta^*, \hat{\theta}) - \hat{C}_n(\beta^*, \theta^*)\} + \{\hat{C}_n(\beta, \theta^*) - \hat{C}_n(\beta^*, \theta^*)\},$$

where $\hat{C}_n(\boldsymbol{\beta}, \boldsymbol{\theta}^*)$ is a U -statistics of order two. Following similar arguments in Sherman (1993), by condition (C5), we have, uniformly over any $o_p(1)$ neighborhood of $\boldsymbol{\beta}^*$,

$$\hat{C}_n(\boldsymbol{\beta}, \boldsymbol{\theta}^*) - \hat{C}_n(\boldsymbol{\beta}^*, \boldsymbol{\theta}^*) = \frac{1}{2}(\boldsymbol{\beta} - \boldsymbol{\beta}^*)' V (\boldsymbol{\beta} - \boldsymbol{\beta}^*) + \frac{1}{\sqrt{n}}(\boldsymbol{\beta} - \boldsymbol{\beta}^*)' W_n + o_p(\|\boldsymbol{\beta} - \boldsymbol{\beta}^*\|^2) + o_p(1/n), \quad (\text{A.1})$$

where $W_n = n^{-1/2} \sum_{i=1}^n \nabla_1 \varrho(\boldsymbol{\beta}^*, \mathbf{X}_i) \rightarrow N(0, \Delta)$ in distribution and $2V = E\{\nabla_2 \varrho(\boldsymbol{\beta}^*, \mathbf{X})\}$.

In addition, we have $E\{\Lambda_{ij}(\boldsymbol{\theta}) / \boldsymbol{\theta} | \mathbf{X}_i, \mathbf{X}_j\} = E\{\Lambda_{ij}(\boldsymbol{\theta}) / \boldsymbol{\theta} \boldsymbol{\theta}' | \mathbf{X}_i, \mathbf{X}_j\} = 0$ for any $\boldsymbol{\theta}$. First applying second order Taylor expansion with respect to $\boldsymbol{\theta}$ and then following similar arguments to derive (A.1), we have, uniformly over any $o_p(1)$ neighborhood of $\boldsymbol{\beta}^*$,

$$\begin{aligned} \hat{C}_n(\boldsymbol{\beta}, \hat{\boldsymbol{\theta}}) - \hat{C}_n(\boldsymbol{\beta}, \boldsymbol{\theta}^*) - \{\hat{C}_n(\boldsymbol{\beta}^*, \hat{\boldsymbol{\theta}}) - \hat{C}_n(\boldsymbol{\beta}^*, \boldsymbol{\theta}^*)\} \\ = \{o_p(\|\boldsymbol{\beta} - \boldsymbol{\beta}^*\|^2) + o_p(1/n)\} O_p(\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\| + \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|^2) + o_p(\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|^2) \\ = o_p(\|\boldsymbol{\beta} - \boldsymbol{\beta}^*\|^2) + o_p(1/n). \end{aligned} \quad (\text{A.2})$$

Combining (A.1) and (A.2), we have

$$0 \leq \hat{C}_n(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\theta}}) - \hat{C}_n(\boldsymbol{\beta}^*, \hat{\boldsymbol{\theta}}) = \frac{1}{2}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*)' V (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*) + \frac{1}{\sqrt{n}}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*)' W_n + o_p(\|\boldsymbol{\beta} - \boldsymbol{\beta}^*\|^2) + o_p(1/n),$$

which implies that $\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^* = O_p(1/\sqrt{n})$ since V is negative definite. Therefore,

$$\hat{C}_n(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\theta}}) - \hat{C}_n(\boldsymbol{\beta}^*, \hat{\boldsymbol{\theta}}) = \frac{1}{2}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*)' V (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*) + \frac{1}{\sqrt{n}}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*)' W_n + o_p(1/n).$$

By Theorem 2 of Sherman (1993), we have

$$\sqrt{n}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*) = V^{-1} W_n + o_p(1) \rightarrow N(0, \Sigma), \quad (\text{A.3})$$

in distribution, where $\Sigma = V^{-1} \Delta V^{-1}$ with $\Delta = E\{\nabla_1 \varrho(\boldsymbol{\beta}^*, \mathbf{X}) \nabla_1 \varrho(\boldsymbol{\beta}^*, \mathbf{X})'\}$.

Appendix B: Proof of Theorem 2

Define $\tilde{Y} = \frac{Y}{\pi(\mathbf{X})A + \{1 - \pi(\mathbf{X})\}(1 - A)}$. Note that $V(\boldsymbol{\beta}, c) = E[\tilde{Y} I\{A = I(\boldsymbol{\beta}' \mathbf{X} - c)\}]$. By condition (C6), we have $V(\boldsymbol{\beta}^*, c) - V(\boldsymbol{\beta}^*, c^*) = -M(c - c^*)^2$. In addition, define

$$\mathcal{F}_1 = \left\{ \tilde{Y} I \{ A = I(\beta' X \geq c) \} : |c - c^*| \leq \eta, \sqrt{n}(\beta - \beta^*) = O_p(1) \right\},$$

where η is a positive constant. Under condition (C1), it can be easily shown that $\mathcal{F}_1$ is a Donsker class. Moreover, by Theorem 11.2 of Kosorok (2007), we have

$$E \left[\sqrt{n} \sup_{|c - c^*| < \eta} |\hat{V}_n(\hat{\beta}, c) - V(\hat{\beta}, c) - \{\hat{V}_n(\hat{\beta}, c^*) - V(\hat{\beta}, c^*)\}| \right] \leq \int_0^\eta \sqrt{\log N_{[]}(\varepsilon F, \mathcal{F}_1, L_2(P))} d\varepsilon \leq K\eta^{1/2}$$

for all n large enough and sufficiently small η , where $F = |\tilde{Y}|$ is a square-integrable envelope function, $N_{[]}(\varepsilon F, \mathcal{F}_1, L_2(P))$ is the bracket number of $\mathcal{F}_1$ with $L_2(P)$ norm metric, and K is a positive constant. Therefore,

$$\begin{aligned} & E \left[\sqrt{n} \sup_{|c - c^*| < \eta} |\hat{V}_n(\hat{\beta}, c) - \hat{V}_n(\hat{\beta}, c^*) - \{V(\beta^*, c) - V(\beta^*, c^*)\}| \right] \\ &= E \left[\sqrt{n} \sup_{|c - c^*| < \eta} |\hat{V}_n(\hat{\beta}, c) - V(\hat{\beta}, c) - \{\hat{V}_n(\hat{\beta}, c^*) - V(\hat{\beta}, c^*)\}| \right. \\ & \quad \left. + \sqrt{n} \sup_{|c - c^*| < \eta} |V(\hat{\beta}, c) - V(\beta^*, c) - \{V(\hat{\beta}, c^*) - V(\beta^*, c^*)\}| \right] \\ & \leq K\eta^{1/2} + \eta O(1) = O(1)\phi(\eta), \end{aligned}$$

where $\phi(\eta) = \eta^{1/2} + \eta$. Note that $\phi(\eta)/\eta^\alpha$ is decreasing in η for any $\alpha \in [1, 2)$. Set $r_n = n^{1/3}$, then $r_n^2 \phi(1/r_n) = n^{1/2} + n^{1/3} = O(n^{1/2})$. By Theorem 14.4 of Kosorok (2007), we obtain $n^{1/3}(\hat{c} - c^*) = O_p(1)$.

Next, we establish the limiting distribution of $\hat{h}_n = n^{1/3}(\hat{c} - c^*)$. We have that h_n is the argmax of the process $n^{2/3}\{\hat{V}_n(\hat{\beta}, n^{-1/3}h + c^*) - \hat{V}_n(\hat{\beta}, c^*)\}$ that is indexed by h . In addition,

$$\begin{aligned} & n^{2/3}\{\hat{V}_n(\hat{\beta}, n^{-1/3}h + c^*) - \hat{V}_n(\hat{\beta}, c^*)\} \\ &= n^{2/3}\{\hat{V}_n(\hat{\beta}, n^{-1/3}h + c^*) - V(\hat{\beta}, n^{-1/3}h + c^*)\} - n^{2/3}\{\hat{V}_n(\hat{\beta}, c^*) - V(\hat{\beta}, c^*)\} \\ & \quad + n^{2/3}\{V(\hat{\beta}, n^{-1/3}h + c^*) - V(\hat{\beta}, c^*)\}. \end{aligned}$$

We have

$$\begin{aligned} & n^{2/3}\{V(\hat{\beta}, n^{-1/3}h + c^*) - V(\hat{\beta}, c^*)\} \\ &= n^{2/3}\{V(\beta^*, n^{-1/3}h + c^*) - V(\beta^*, c^*)\} + n^{2/3}\left[\{V(\hat{\beta}, n^{-1/3}h + c^*) - V(\beta^*, n^{-1/3}h + c^*)\} \right. \\ & \quad \left. - \{V(\hat{\beta}, c^*) - V(\beta^*, c^*)\}\right] \\ &= \frac{1}{2}vh^2 + o_p(1) + O_p(n^{1/3}\|\hat{\beta} - \beta^*\|) = \frac{1}{2}vh^2 + o_p(1), \end{aligned}$$

where $v = 2V(\beta^*, c^*)/c^2$. In addition, define

$$\mathcal{F}_{2n} = \left\{ n^{1/6} [\tilde{Y}I\{A=I(\beta'X \geq n^{-1/3}h+c^*)\} - \tilde{Y}I\{A=I(\beta'X_i \geq c^*)\}]: \sqrt{n}(\beta - \beta^*)=O_p(1) \right\},$$

which can be shown to be a Donsker class. In addition, we have

$$\begin{aligned} & E \left(n^{1/3} [\tilde{Y}I\{A=I(\beta^{*'}X \geq n^{-1/3}h+c^*)\} - \tilde{Y}I\{A=I(\beta^{*'}X_i \geq c^*)\}]^2 \right) \\ &= E \left[n^{1/3} \tilde{Y}^2 \{I(c^* \leq \beta^{*'}X < n^{-1/3}h+c^*)I(h>0) + I(n^{-1/3}h+c^* \leq \beta^{*'}X < c^*)I(h<0)\} \right] \\ &\rightarrow q^*|h|, \text{ almost surely,} \end{aligned}$$

as $n \rightarrow \infty$, where q^* is a positive constant. Therefore, $n^{2/3} \{ \hat{V}_n(\hat{\beta}, n^{-1/3}h + c^*) - V(\hat{\beta}, n^{-1/3}h + c^*) \} - n^{2/3} \{ \hat{V}_n(\hat{\beta}, c^*) - V(\hat{\beta}, c^*) \}$ converges weakly to a two-sided Gaussian process $\sqrt{q^*} \mathcal{Z}(h)$, where $\mathcal{Z}(h)$ is a standard two-sided Brownian motion process. By the Argmax theorem of Kosorok (2007), $\hat{h}_n$ converges in distribution to $\arg\max_h G(h)$, where $G(h) = \sqrt{q^*} \mathcal{Z}(h) + vh^2/2$.

References

- Abrevaya J. Pairwise-difference rank estimation of the transformation model. *Journal of Business & Economic Statistics*. 2003; 21:437–47.
- Amemiya, T. *Advanced econometrics*. Harvard university press; 1985.
- Blatt D, Murphy S, Zhu J. A-learning for approximate planning. Unpublished Manuscript. 2004
- Brown B, Wang Y-G. Standard errors and covariance matrices for smoothed rank estimators. *Biometrika*. 2005; 92:149–158.
- Brown B, Wang Y-G. Induced smoothing for rank regression with censored survival times. *Statistics in medicine*. 2007; 26:828–836. [PubMed: 16646005]
- Cavanagh C, Sherman RP. Rank estimators for monotonic index models. *Journal of Econometrics*. 1998; 84:351–381.
- Chen S. Rank estimation of transformation models. *Econometrica*. 2002; 70:1683–1697.
- Foster JC, Taylor JM, Ruberg SJ. Subgroup identification from randomized clinical trial data. *Statistics in medicine*. 2011; 30:2867–2880. [PubMed: 21815180]
- Han AK. Non-parametric analysis of a generalized regression model: the maximum rank correlation estimator. *Journal of Econometrics*. 1987; 35:303–316.
- Jin Z, Ying Z, Wei LJ. A simple resampling method by perturbing the minimand. *Biometrika*. 2001; 88:381–390.
- Kosorok, MR. *Introduction to empirical processes and semiparametric inference*. Springer; 2007.
- Lu W, Zhang H, Zeng D. Variable selection for optimal treatment decision. *Statistical methods in medical research*. 2013; 22:493–504. [PubMed: 22116341]
- Matsouaka RA, Li J, Cai T. Evaluating marker-guided treatment selection strategies. *Biometrics*. 2014; 70:489–499. [PubMed: 24779731]
- Murphy SA. Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 2003; 65:331–355.
- Pang L, Lu W, Wang HJ. Variance estimation in censored quantile regression via induced smoothing. *Computational statistics & data analysis*. 2012; 56:785–796. [PubMed: 22547899]
- Qian M, Murphy SA. Performance guarantees for individualized treatment rules. *Annals of statistics*. 2011; 39(2):1180. [PubMed: 21666835]
- Sherman RP. The limiting distribution of the maximum rank correlation estimator. *Econometrica*. 1993; 61:123–137.

- Watkins CJ, Dayan P. Q-learning. *Machine learning*. 1992; 8:279–292.
- Watkins, CJCH. PhD thesis. University of Cambridge England; 1989. Learning from delayed rewards.
- Zhang B, Tsiatis AA, Davidian M, Zhang M, Laber E. Estimating optimal treatment regimes from a classification perspective. *Stat*. 2012a; 1:103–114. [PubMed: 23645940]
- Zhang B, Tsiatis AA, Laber EB, Davidian M. A robust method for estimating optimal treatment regimes. *Biometrics*. 2012b; 68:1010–1018. [PubMed: 22550953]
- Zhang B, Tsiatis AA, Laber EB, Davidian M. Robust estimation of optimal dynamic treatment regimes for sequential treatment decisions. *Biometrika*. 2013a; 100:681–694.
- Zhang J, Jin Z, Shao Y, Ying Z. Statistical inference on transformation models: a self-induced smoothing approach. *arXiv preprint arXiv:1302.6651*. 2013b
- Zhao L, Tian L, Cai, Tianxiand Claggett B, Wei LJ. Effectively selecting a target population for a future comparative study. *Journal of the American Statistical Association*. 2013; 108:527–539. [PubMed: 24058223]
- Zhao Y, Kosorok MR, Zeng D. Reinforcement learning design for cancer clinical trials. *Statistics in medicine*. 2009; 28:3294–3315. [PubMed: 19750510]
- Zhao Y, Zeng D, Rush AJ, Kosorok MR. Estimating individualized treatment rules using outcome weighted learning. *Journal of the American Statistical Association*. 2012; 107:1106–1118. [PubMed: 23630406]
- Zhao Y, Zeng D, Socinski MA, Kosorok MR. Reinforcement learning strategies for clinical trials in nonsmall cell lung cancer. *Biometrics*. 2011; 67:1422–1433. [PubMed: 21385164]

Table 1

Simulation results for Cases I and II. The true optimal regime is $d^{\text{opt}}(x) = I(\beta'_0 x \geq c_0)$ with $\beta_0 = (0.5, 0.5, -0.5, 0.5)'$, $c_0 = 0$.

		$\hat{\beta}_1$	$\hat{\beta}_2$	$\hat{\beta}_3$	$\hat{\beta}_4$	$\hat{e}$	$\hat{V}$	PCD	$\hat{C}$
Case I									
$V_0 = 2.595, C_0 = 2.258$									
IPSW	Est.	0.456	0.591	-0.561	0.495	0.063	2.473	0.883	2.131
	SD	0.202	0.193	0.153	0.152	0.237	0.104	0.053	0.125
AIPSW	Est.	0.575	0.568	-0.574	0.572	0.002	2.585	0.962	2.244
	SD	0.058	0.060	0.062	0.087	0.080	0.034	0.020	0.021
LR	Est.	0.500	0.500	-0.500	0.499	0.000	2.597	0.989	2.257
	SD	0.016	0.016	0.016	0.015	0.018	0.032	0.008	0.017
CAL	Est.	0.498	0.502	-0.495	0.497	0.040	2.543	0.920	2.249
	SD	0.048	0.043	0.042	0.046	0.234	0.066	0.055	0.019
	SE	0.051	0.051	0.050	0.050	-	-	-	-
	CP	95.0	96.6	96.2	94.8	-	-	-	-
CAL-DR	Est.	0.501	0.501	-0.497	0.499	0.031	2.545	0.925	2.256
	SD	0.022	0.021	0.021	0.019	0.222	0.009	0.056	0.017
	SE	0.023	0.023	0.021	0.024	-	-	-	-
	CP	95.4	97.2	93.4	98.0	-	-	-	-
Case II									
$V_0 = 7.723, C_0 = 6.248$									
IPSW	Est.	0.376	0.525	-0.530	0.453	-0.137	7.558	0.818	5.699
	SD	0.311	0.284	0.245	0.215	0.334	0.259	0.088	0.877
AIPSW	Est.	0.422	0.403	-0.408	0.432	0.160	7.471	0.773	5.209
	SD	0.291	0.325	0.324	0.392	0.416	0.629	0.127	1.405
LR	Est.	0.481	0.479	-0.486	0.480	0.472	7.663	0.811	6.125
	SD	0.135	0.134	0.136	0.136	0.114	0.538	0.042	0.555
CAL	Est.	0.498	0.491	-0.503	0.496	-0.097	7.667	0.875	6.226
	SD	0.054	0.057	0.057	0.057	0.456	0.179	0.090	0.526
	SE	0.064	0.062	0.063	0.062	-	-	-	-
	CP	97.2	95.0	95.6	95.6	-	-	-	-
CAL-DR	Est.	0.498	0.497	-0.504	0.498	-0.078	7.668	0.875	6.244

	$\hat{\varrho}$	PCD	$\hat{\lambda}$	$\hat{\varrho}$	$\hat{\beta}_4$	$\hat{\beta}_3$	$\hat{\beta}_2$	$\hat{\beta}_1$
SD	0.025	0.025	0.185	0.455	0.023	0.025	0.025	0.025
SE	0.026	0.026	-	-	0.026	0.023	0.026	0.026
CP	95.8	96.2	-	-	96.2	92.4	96.2	95.8
	-	-	-	-	-	-	-	-

[†]Est., mean of estimators; SD, standard deviation of estimators; SE, mean of estimated standard errors; CP, empirical coverage probability of 95% confidence interval (%).

Table 2

Simulation results for Cases V and VI.

		$\hat{\beta}_1$	$\hat{\beta}_2$	$\hat{\beta}_3$	$\hat{\beta}_4$	ϵ	$\hat{\nu}$	PCD	$\hat{C}$
Case V									
$V_0 = 3.297$									
IPSW	Est.	0.088	0.447	0.209	0.023	0.596	2.877	0.712	0.273
	SD	0.500	0.468	0.350	0.446	0.457	0.131	0.100	0.184
AIPSW	Est.	0.057	0.516	0.260	0.011	0.640	2.912	0.739	0.310
	SD	0.532	0.461	0.338	0.409	0.413	0.123	0.100	0.179
LR	Est.	0.080	0.570	0.254	0.022	-0.641	2.902	0.734	0.331
	SD	0.457	0.315	0.436	0.326	0.536	0.113	0.086	0.156
CAL	Est.	0.091	0.575	0.267	0.018	0.836	2.915	0.735	0.342
	SD	0.448	0.282	0.398	0.390	0.759	0.122	0.106	0.135
CAL-DR	Est.	0.085	0.585	0.258	0.023	0.848	2.920	0.738	0.342
	SD	0.448	0.287	0.410	0.365	0.753	0.119	0.101	0.140
Case VI									
$V_0 = 4.366$									
IPSW	Est.	0.444	0.486	-0.445	0.468	0.527	4.259	0.890	3.135
	SD	0.198	0.196	0.194	0.194	0.189	0.447	0.051	0.453
AIPSW	Est.	0.448	0.426	-0.423	0.434	0.355	4.204	0.845	2.972
	SD	0.412	0.241	0.243	0.287	0.229	0.498	0.124	0.732
LR	Est.	0.483	0.485	-0.477	0.461	0.306	4.175	0.685	3.195
	SD	0.151	0.142	0.161	0.154	0.095	0.421	0.040	0.427
CAL	Est.	0.498	0.495	-0.508	0.480	0.548	4.338	0.935	3.255
	SD	0.075	0.071	0.064	0.066	0.205	0.420	0.033	0.416
CAL-DR	Est.	0.506	0.501	-0.508	0.475	0.531	4.341	0.937	3.262
	SD	0.047	0.045	0.053	0.047	0.217	0.422	0.045	0.416

Table 3

Estimated Optimal Treatment Regimes for ACTG175 Data.

		$\hat{\beta}_1$	$\hat{\beta}_2$	$\hat{\beta}_3$	$\hat{\beta}_4$	$\hat{\beta}_5$	$\hat{\beta}_6$	$\hat{\beta}_7$	$\hat{\beta}_8$	$\hat{\beta}_9$	$\hat{\beta}_{10}$	$\hat{\beta}_{11}$	$\hat{\beta}_{12}$
CAL	Est.	0.929	-0.182	-0.160	-0.252	-0.076	-0.057	-0.059	0.001	-0.033	0.037	-0.019	-0.021
	SE	0.470	0.353	0.957	0.471	0.303	0.083	0.044	0.048	0.029	0.052	0.025	0.042
	PV	0.048	0.608	0.868	0.594	0.802	0.494	0.182	0.982	0.253	0.473	0.454	0.614
CAL-DR	Est.	0.946	-0.093	-0.143	-0.202	-0.166	-0.051	-0.052	0.001	-0.030	0.021	-0.017	-0.017
	SE	0.111	0.211	0.207	0.258	0.203	0.053	0.028	0.028	0.02	0.032	0.016	0.021
	PV	0.000	0.658	0.490	0.433	0.413	0.330	0.067	0.960	0.130	0.516	0.279	0.410