Journal of the American Statistical Association

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/uasa20

Estimation and Validation of Ratio-based Conditional Average Treatment Effects Using Observational Data

Steve Yadlowsky , Fabio Pellegrini , Federica Lionetto , Stefan Braune & Lu Tian

To cite this article: Steve Yadlowsky , Fabio Pellegrini , Federica Lionetto , Stefan Braune & Lu Tian (2020): Estimation and Validation of Ratio-based Conditional Average Treatment Effects Using Observational Data, Journal of the American Statistical Association, DOI: 10.1080/01621459.2020.1772080

To link to this article: https://doi.org/10.1080/01621459.2020.1772080

JOURNAL OF THE AMERICAN STATISTICAL ASSOCIATION
2020, VOL. 00, NO. 0, 1–18: Precision Medicine
https://doi.org/10.1080/01621459.2020.1772080

Estimation and Validation of Ratio-based Conditional Average Treatment Effects Using Observational Data

Steve Yadlowsky[a], Fabio Pellegrini[b], Federica Lionetto[c], Stefan Braune[d], and Lu Tian[e]

[a]Department of Electrical Engineering, Stanford University, Stanford, CA; [b]Biogen International GmbH, Baar, Switzerland; [c]PwC Data & Analytics, Zurich, Switzerland; [d]NeuroTransData, Neuburg an der Donau, Germany; [e]Department of Biomedical Data Science, Stanford University, Palo Alto, CA

ABSTRACT

While sample sizes in randomized clinical trials are large enough to estimate the average treatment effect well, they are often insufficient for estimation of treatment-covariate interactions critical to studying data-driven precision medicine. Observational data from real world practice may play an important role in alleviating this problem. One common approach in trials is to predict the outcome of interest with separate regression models in each treatment arm, and estimate the treatment effect based on the contrast of the predictions. Unfortunately, this simple approach may induce spurious treatment-covariate interaction in observational studies when the regression model is misspecified. Motivated by the need of modeling the number of relapses in multiple sclerosis (MS) patients, where the ratio of relapse rates is a natural choice of the treatment effect, we propose to estimate the conditional average treatment effect (CATE) as the ratio of expected potential outcomes, and derive a doubly robust estimator of this CATE in a semiparametric model of treatment-covariate interactions. We also provide a validation procedure to check the quality of the estimator on an independent sample. We conduct simulations to demonstrate the finite sample performance of the proposed methods, and illustrate their advantages on real data by examining the treatment effect of dimethyl fumarate compared to teriflunomide in MS patients. Supplementary materials for this article are available online.

ARTICLE HISTORY
Received June 2019
Accepted May 2020

KEYWORDS
Conditional average treatment effect; Doubly robust estimation; Heterogeneous treatment effect; Observational study; Precision medicine

1. Introduction

Recently, interest in recommending tailored preventative interventions or treatments to patients in clinical practice has prompted investigating the conditional average treatment effect (CATE) from data. Knowledge of the CATE as the contrast between the expected outcome under different interventions conditional on covariate levels would allow clinicians to understand how much a patient would benefit from a particular intervention based on their covariates. The primary statistical objective is to estimate these CATEs by examining treatment-covariate interactions (Tian et al. 2014).

The small sample size of trials is one of the biggest obstacles in such analyses. Most randomized clinical trials are designed to study the average treatment effect (ATE), rather than the CATE. Furthermore, to verify CATE estimates or high value subgroup of patients for whom the treatment is most effective, researchers use sample splitting (Athey and Imbens 2016) or, ideally, independent external validation (Basu et al. 2017) to account for the exploratory nature and the overfitting tendency of relevant statistical analyses, which further shrinks the available sample size. One important alternative is to use observational data from real world practice. Observational data often contain more samples, have broader target patient populations, and if collected from clinical practice, better represent realistic clinical conditions. However, patients receiving the treatment of interest and those receiving alternatives may be systematically different in observational data, which introduces new challenges in data analysis (Imbens and Rubin 2015).

Let $Y^{(1)}$ be the outcome if the unit were given treatment $r = 1$, and $Y^{(0)}$ be the outcome if given treatment $r = 0$. In this work, we study the estimation and validation of the ratio-based CATE,

$$D(z) = \frac{\mathrm{E}[Y^{(1)}|Z = z]}{\mathrm{E}[Y^{(0)}|Z = z]},$$

which targets the ratio of expected outcomes under different interventions. Our motivation is the study of differences in the effect of treatments on relapse in multiple sclerosis (MS) patients, using observational data from the NeuroTransData (NTD) registry (described below). For repeated events such as MS relapses, it is natural to look at the ratio of relapse rates under different potential treatments, as it is a relative measure of effectiveness.

We provide a general framework for the estimation and validation of such a ratio-based CATE score. We develop a doubly robust method for estimating the treatment contrast $D(z)$ in the semiparametric model where $D(z) = \exp(\delta^\top z)$, but the conditional means, $\mu_r(z) = \mathrm{E}[Y^{(r)} \mid Z = z]$, and propensity score, $\pi_r(z) = P(R = r \mid Z = z)$, are in nonparametric models. This method has advantageous statistical properties, such as Neyman orthogonality (Chernozhukov

CONTACT Steve Yadlowsky syadlows@stanford.edu Department of Electrical Engineering, Stanford University, Stanford, CA 94305.
Supplementary materials for this article are available online. Please go to www.tandfonline.com/r/JASA.

et al. 2018), which, along with our simulation and experimental evidence, justify the use of machine learning methods for nuisance parameter estimation in the observational data. We provide another method for estimating the treatment contrast by adjusting per-arm regression analyses for confounding, and demonstrate via simulation that it often performs as well in practice. Both methods have the appealing property that when there is no treatment effect ratio heterogeneity, they will infer that $D(z)$ is constant, even if the regression model is misspecified. Finally, we provide a method for validating the effectiveness of the learned CATE score in a statistically independent validation sample, and use this to compare the scores learned via different methods on the MS observational data.

1.1. A Motivating Example

The NTD MS registry includes about 25,000 patients with MS, which represents about 15% of all MS patients in Germany. It includes demographic, clinical history, patient related outcomes and clinical variables captured in real time during clinical visits. The focus of the analysis is to estimate the CATE of teriflunomide (TERI, $n = 1050$) compared with dimethyl fumarate (DMF, $n = 1741$) and stratify the patient population for tailored treatment recommendation. The outcome of the primary interest is the number of relapses per unit time; the average treatment effect is measured by the ratio of the expected relapse rate under TERI versus that under DMF. The ratio is important, because it better contextualizes the value of the treatment for these patients. Preventing 1 relapse per decade is more noticeable when the baseline is 1 relapse every 5 years than 1 relapse every year.

Three questions arise when trying to analyze CATEs in the NTD registry:

1. Can we use the ratio rather than the difference in expected outcomes to measure the treatment effect?
2. How do we adjust for differences in baseline covariates (confounders) between two treatment groups, and how do these affect estimation of the CATE?
3. How do we validate our CATE model based on observational data and will the resulting method provide a useful measure of the quality of the estimated CATE score?

The purpose of this article is to answer these questions and provide a statistical methodology for estimation and validation of the ratio-based CATE, based on the described semiparametric model.

1.2. Related Approaches

The estimation of the absolute difference in expectations of potential outcomes as the CATE, that is, $\mathrm{E}(Y^{(1)} - Y^{(0)} \mid Z = z)$, using observational data has been studied extensively in the literature (Green and Kern 2012; Xie, Brand, and Jann 2012; Lu et al. 2018; Wager and Athey 2018; Athey, Tibshirani, and Wager 2019; Künzel et al. 2019; Nie and Wager 2019). Recently, Powers et al. (2018) and Wendling et al. (2018) compared a number of approaches for learning this function. The basic idea is to either estimate the CATE based on separately estimated $\mathrm{E}(Y^{(r)} \mid Z = z), r = 0, 1$, or learn directly using modified outcomes. However, in some settings, the absolute difference in potential outcomes is not the best measure of treatment effect. For example, if the ratio-based CATE, $D(z)$, is constant for all z, but $\mathrm{E}[Y^{(0)}|Z = z]$ varies with z, then there will appear to be significant treatment effect heterogeneity measured by the absolute difference $\mathrm{E}[Y^{(1)} - Y^{(0)}|Z = z]$, which may not be of particular interest. In this work, we focus on estimating the ratio-based CATE, that is, the ratio of the conditional expectation of the potential outcomes given the baseline covariates. It is natural to consider applying the methodology for absolute differences to the log-transformed outcomes. However, for counting processes where ratio-based measures are most natural, there is a nonzero probability that the observed outcome is zero, making $\mathrm{E}[\log(Y^{(r)}) \mid Z = z]$ infinite. On the other hand, our contrast $D(z)$ is well-defined whenever $\mu_0(z) > 0$.

In the context of binary outcomes, the ratio-based CATE (or ratio-based ATE) is known as the risk ratio. Robins and Rotnitzky (2001) showed that in a generalized linear model, the identity link function (corresponding to the absolute risk difference) and the log link function (corresponding to the risk ratio) are the only link functions that admit doubly robust estimators of the constant treatment effect. Given the extensive literature on CATE estimation with absolute differences, our work on ratio-based CATE estimation is a timely contribution. Dukes and Vansteelandt (2018) studied G-estimation of the risk ratio in the parametric setting, and demonstrated certain double robustness of the method. Van der Laan and Rose (2011) showed that the related targeted maximum likelihood estimator for the risk ratio also satisfies double robustness. In this article, we demonstrate that the contrast regression posed is doubly robust in the discussed semiparametric model and has the Neyman orthogonality property, which justifies the use of a wide class of nonparametric and machine learning methods for fitting the outcome and propensity score models in our method. Furthermore, the methodology and estimating equation used in Van der Laan and Rose (2011) is optimized for binary outcomes, whereas our method is optimized for count data.

The outcome weighted learning (OWL) is another class of methods for developing precision medicine (Zhao et al. 2012, 2014; Chen et al. 2017; Zhou et al. 2017). OWL methods find a decision boundary in the covariate space to classify patients into those with a treatment benefit $\{z \mid D(z) < 1\}$ or harm $\{z \mid D(z) > 1\}$, if Y represents undesirable events. The OWL method and its variations convert the original task into a binary classification problem and directly target the decision boundary, bypassing the need to estimate the CATE (Zhang et al. 2012). In contrast, the regression approach above attempts to directly estimate the magnitude of the benefit, and then identify the high value subgroup of patients accordingly (Cai et al. 2010; Foster, Taylor, and Ruberg 2011; Zhao et al. 2013). The OWL approach avoids the more difficult task, but also fails to yield information about the size of the treatment benefit for individual patient. A good estimator of the CATE ensures a good ATE within the subgroup consisting of patients with the largest CATEs, thus having priority of receiving the treatment. In addition, we may also directly recommend the

treatment to patients whose estimated CATE outweighs the associated cost, which can also be patient-dependent. Therefore, in this article, we focus on the more general question of directly estimating the CATE rather than a binary recommendation rule.

For data from a randomized clinical trial, Zhao et al. (2013) proposed the following approach to estimate and validate CATE. The method consists of two main steps:

1. In the training set:
   - Fit separate regression models $\widehat{\mu}_r(z) \approx \mathrm{E}[Y^{(r)}|R = r, Z = z]$ for the potential outcomes in the treatment ($r = 1$) and control ($r = 0$) arms.
   - Estimate the CATE by $\widehat{D}(z) = \widehat{\mu}_1(z) - \widehat{\mu}_0(z)$.
2. In the validation set:
   - Estimate $AD(c) = \mathrm{E}\left\{Y^{(1)} - Y^{(0)} \mid \widehat{D}(Z) \geq c\right\}$, the ATE for a subgroup of patients $\{z \mid \widehat{D}(z) \geq c\}$, and denote the resulting estimator by $\widehat{AD}(c)$.
   - Draw the validation curve $q \mapsto \widehat{AD}\{\widehat{H}^{-1}(1-q)\}$, where $q \in [0, 1)$ and $\widehat{H}(\cdot)$ is the empirical cumulative distribution function of $\widehat{D}(Z)$. This curve graphically represents the relationship between the proportion of patients q in the subgroup with the CATE score above $\widehat{H}^{-1}(1-q)$ and the estimated ATE in that subgroup.
   - Observe the slope of $\widehat{AD}\{\widehat{H}^{-1}(1-q)\}$, which reflects the quality of the scoring system $\widehat{D}(z)$ in ranking the patients according to their estimated CATE, $\widehat{D}(z)$.

In this article, we extend the sequence of training and validation steps by Zhao et al. (2013) to the ratio-based CATE with observational data.

2. Method

The standard regression model for the number of relapses in terms of the baseline covariates is the Poisson or negative binomial regression by treatment arm:

$$\mathrm{E}(Y^{(r)} \mid Z, R = r) = \exp(\beta_r^\top \widetilde{Z}), \quad r = 0, 1, \tag{1}$$

where $\widetilde{Z} = (1, Z^\top)^\top$ is a $d + 1$ dimensional covariate vector, R is the binary indicator of the treatment received, and $Y^{(r)}$ is the potential number of relapses if the patient received the treatment $r \in \{0, 1\}$ or the ratio of the number relapses to an exposure time. We only observe $Y^{(r)}$ when $R = r$, but we are interested in $\mathrm{E}[Y^{(r)} \mid Z = z]$, the average outcome over all individuals with covariate z, if they had been prescribed treatment r. This creates a causal-missing data problem. In this work, we make the unconfoundedness assumption (Imbens and Rubin 2015) that identifies the relapse rate:

$$\{Y^{(1)}, Y^{(0)}\} \perp\!\!\!\perp R \mid Z. \tag{2}$$

This implies that $\mathrm{E}(Y \mid Z, R = r) = \mathrm{E}(Y^{(r)} \mid Z, R = r) = \mathrm{E}(Y^{(r)} \mid Z)$.

Our goal is to model the effect of the treatment on the relapse rate. We assume that the follow up time is the same for all patients. Otherwise, if unconfoundedness holds with respect to the follow up time, we may replace the outcome by the number of relapses divided by the exposure time. Under the above model, the ratio of the expected relapse rates

$$D(z) = \frac{\mathrm{E}(Y^{(1)}|Z = z)}{\mathrm{E}(Y^{(0)}|Z = z)} \tag{3}$$

is a natural measure of the CATE for the relapse rate that is insensitive to differences in exposure time between patients. We can estimate β_r by applying standard Poisson or negative binomial regression methods to the observed data in each arm to get the estimator $\widehat{\beta}_r$. With the estimated regression coefficients $\widehat{\beta}_0$ and $\widehat{\beta}_1$, let

$$\widehat{\mu}_r(z) = \exp\{\widehat{\beta}_r^\top \widetilde{z}\}$$

be an estimator of the relapse rate under treatment $r \in \{0, 1\}$. A simple estimate of the CATE under model (1) is

$$\widehat{D}(z) = \frac{\widehat{\mu}_1(z)}{\widehat{\mu}_0(z)} = \exp\{(\widehat{\beta}_1 - \widehat{\beta}_0)^\top \widetilde{z}\}. \tag{4}$$

2.1. Confounding Effect on Estimating the CATE

If the regression models for Y given Z and the treatment assignment R are correctly specified, then $\widehat{\beta}_r$ will converge to the true regression coefficient as the sample size increases regardless of the underlying distribution of Z in each arm of the study. In practice, these statistical models may be misspecified and only serve as working models approximating the true relationship between outcomes and covariates. In such a case, the estimated regression coefficients may converge to limits that introduce spurious predicted treatment heterogeneity.

The following toy example illustrates this phenomenon. Assume that $Y^{(r)} \mid Z = z$ follows a Poisson distribution with a rate of z^2 in both arms; therefore, $D(z) = 1$. To introduce confounding, assume that $Z \mid R = r \sim N\{(r - 0.5), 1\}, r = 0, 1$. When fitting a misspecified Poisson regression model

$$\mathrm{E}(Y|Z = z, R = r) = \exp(\beta_r^\top \widetilde{z})$$

in two arms separately, the regression coefficient of Z is $\beta_1 = (-0.5, 0.8)^\top$ in arm $r = 1$ and $\beta_0 = (-0.5, -0.8)^\top$ in arm $r = 0$. This is not surprising, because 70% of the Z_i in arm 1 are positive, where the quadratic function is increasing, inducing a positive association between Y and Z, and 70% of the Z_i in arm 0 are negative in arm 0, where the quadratic function in decreasing, inducing a negative association. Therefore, the estimated CATE score $\widehat{D}(z) = \exp\{(\widehat{\beta}_1 - \widehat{\beta}_0)^\top \widetilde{z}\} \approx \exp(1.6z)$ would suggest that the between group rate ratio increases with the value of z, while in fact it is a constant.

This simple example shows that $\widehat{D}(z)$ estimated from misspecified regression models may create spurious treatment-covariate interactions in observational studies. Therefore, the construction of the CATE score $\widehat{D}(z)$ should adjust for the covariates imbalance to avoid falsely demonstrating treatment effect heterogeneity due to confounding.

2.2. Training

In this section, we propose methodologies for estimating $D(z)$ using observational training data. First, we show how the parameters of the semiparametric model specifying that $D(z) = \exp(\delta_0^\top \widetilde{z})$ are identified, and use this to develop a doubly robust method that can be used with machine learning estimates of the nuisance parameters and provide conditions that allow valid statistical inference. Then, we provide an approach based on fitting separate regressions for each treatment arm after adjusting for confounding effect. This allows interpretation of the regression models in each arm , in the same way as fitting regression models to each arm of a randomized clinical trial. While the estimate of δ_0 from the latter method may be biased, we show that it is consistent when there is no treatment heterogeneity.

2.2.1. Contrast Regression Approach

While motivated by the Poisson regression model, the CATE model

$$D(z) = \exp(\delta_0^\top \widetilde{z})$$

arises from a more general semiparametric model that we will study in this section. Specifically, we model the conditional expectation of the potential outcomes $\mu_r(z)$ with the semiparametric regression model

$$\mathrm{E}(Y^{(r)}|Z = z) = \exp\left(r\delta_0^\top \widetilde{z}\right)\mu_0(z), \tag{5}$$

where $\mu_0(z)$ is a unknown, measurable, and nonnegative function in some nonparametric function class. We assume that the propensity score $\pi_r(z)$ is unknown, but is also in a nonparametric function class. This model represents the class of distributions for which $D(z)$ depends on z through $\delta_0^\top z$, and thus our goal is to estimate δ_0. We provide a doubly robust estimator of δ_0 and discuss assumptions under which this estimator is $\sqrt{n}$-consistent.

If $Y^{(1)}$ and $Y^{(0)}$ were both observed, then δ_0 is the solution to

$$\mathrm{E}\left[w(Z,\delta)\tilde{Z}\left\{Y^{(1)} - \exp(\delta^\top \tilde{Z})Y^{(0)}\right\}\right] = 0, \tag{6}$$

because applying the tower property of conditional expectations to (6) gives the equivalent estimating equation $\mathrm{E}[w(Z,\delta)\widetilde{Z}\{\mu_1(Z) - \exp(\delta^\top \widetilde{Z})\mu_0(Z)\}] = 0$, where $w(z,\delta) > 0$ is a given weight function. The solution is unique in any compact set Ω containing δ_0, as long as $\widetilde{Z}$ does not belong to a d or lower dimensional hyperplane and $w(z,\delta)$ is bounded above and below, for all $\delta \in \Omega$ and z. Like many causal inference and missing data problems, there are a variety of ways to develop estimating equations that are equivalent to (6) under condition (2), involving imputation with the mean functions $\mu_r(z)$ or inverse probability weighting using the propensity $\pi_r(z)$. Because the nuisance parameters $\mu_r(z)$ and $\pi_r(z)$ are rarely known in practice, operationalizing these estimators depends on estimating the nuisance parameters.

Therefore, we follow the approach advocated by Robins and Rotnitzky (2001) for developing doubly robust approaches for semiparametric models. Our estimator is closely related to the generalized linear model with the logarithmic link function presented in their paper, and the doubly robust estimator of the semiparametric risk ratio model presented in Van der Laan and Rose (2011). Specifically, for any candidate nuisance parameters $\mu : \mathbf{R}^d \to \mathbf{R}$ for the baseline mean (that will hopefully approximate μ_0) and $\pi : \mathbf{R}^d \to [0,1]$ for the propensity score (that will approximate $\pi_1(z)$) and parameters $\delta \in \Omega \subset \mathbf{R}^{d+1}$, we consider the estimating function

$$m(G;\delta,\mu,\pi) = \widetilde{Z}\frac{\{1-\pi(Z)\}RY - \pi(Z)(1-R)Y\exp(\delta^\top \widetilde{Z})}{e^{\delta^\top \widetilde{Z}}\pi(Z) + (1-\pi(Z))} - \widetilde{Z}\mu(Z)\exp(\delta^\top \widetilde{Z})\frac{R-\pi(Z)}{e^{\delta^\top \widetilde{Z}}\pi(Z) + (1-\pi(Z))},$$

where $G = (Y, R, Z^\top)^\top$.

If the propensity score is known, then substituting π_1 for π gives the population estimating function

$$\mathrm{E}\left[m(G;\delta,\mu,\pi_1)\right] = \mathrm{E}\left[\widetilde{Z}w_1(Z;\delta,\pi_1)\left\{\mu_1(Z) - \mu_0(Z)e^{\delta^\top \widetilde{Z}}\right\}\right],$$

with the weight function

$$w_1(z;\delta,\pi) = \frac{\pi(z)(1-\pi(z))}{e^{\delta^\top \widetilde{z}}\pi(z) + 1 - \pi(z)},$$

and δ_0 is a root of the corresponding estimating equation, for any bounded choice of μ. In Section 1 of the supplementary materials, we show that this weight function is optimal in minimizing the variance of the resulting estimator when $Y^{(r)} \mid Z = z$ follows a Poisson distribution with rate $\mu_r(z)$ and with the correct propensity score model $\pi_1(z)$.

On the other hand, by rewriting the estimating function as

$$m(G;\delta,\mu,\pi) = \widetilde{Z}R\left\{Y - \mu(Z)\exp(\delta^\top \widetilde{Z})\right\}\frac{1-\pi(Z)}{e^{\delta^\top \widetilde{Z}}\pi(Z) + (1-\pi(Z))} + \widetilde{Z}(1-R)\left[Y - \mu(Z)\right]\frac{\exp(\delta^\top \widetilde{Z})\pi(Z)}{e^{\delta^\top \widetilde{Z}}\pi(Z) + (1-\pi(Z))},$$

we observe that if $\mu_0(z) = \mathrm{E}(Y^{(r)}|Z = z)$ is known, then substituting μ_0 for μ gives the population estimating equation

$$\mathrm{E}\left[m(G;\delta,\mu_0,\pi)\right] = \mathrm{E}\left[\widetilde{Z}w_1(Z;\delta,\pi)\left\{\mu_1(Z) - \mu_0(Z)\exp(\delta^\top \widetilde{Z})\right\}\right] = 0$$

for which δ_0 is still a root, regardless of the choice of $\pi(z)$ used. In practice, neither the propensity score nor the conditional expectation $\mu_0(z)$ is known, which motivates the plug-in estimating equation

$$S_n(\delta) = n^{-1}\sum_{i=1}^{n} m(G_i;\delta,\widehat{\mu}_0,\widehat{\pi}_1) = 0,$$

where $G_i = (Y_i, R_i, Z_i^\top)^\top$, $\widehat{\mu}_0(z)$ is an estimator for $\mu_0(z)$ and $\widehat{\pi}_1(z)$ is an estimator for the propensity score $\pi_1(z)$. If either $\widehat{\mu}_0(z)$ or $\widehat{\pi}_1(z)$ is consistent, then the solution of the estimating equation is a consistent estimator of δ_0.

In Appendix A, we prove that this estimating equation satisfies the Neyman orthogonality condition (Chernozhukov et al. 2018). Therefore, using cross-fitting with this estimating equation allows for general use of machine learning estimators of the nuisance parameters, while still providing accurate confidence interval coverage. To compute the cross-fitting estimator $\widehat{\delta}_0$ of

δ_0, divide data into K nonoverlapping parts of approximately equal sizes indexed by $\mathcal{I}_k, k = 1, \ldots, K$. Construct initial regression estimates $\widehat{\mu}_r^{(-k)}(z)$ of $\mu_r(z)$ without using observations in $\mathcal{I}_k$, and construct estimates $\widehat{\pi}_1^{(-k)}(z)$ of the propensity score $\pi_1(z)$ without using observations in $\mathcal{I}_k$ likewise. Then, estimate δ_0 by searching for the root of the estimating equation

$$S_n^{(cf)}(\delta) = n^{-1} \sum_{k=1}^{K} \sum_{i \in \mathcal{I}_k} m(G_i; \delta, \widehat{\mu}_0^{(-k)}, \widehat{\pi}_1^{(-k)}).$$

The estimator $\widehat{\delta}$ is $\sqrt{n}$-consistent and asymptotically normal under the following sufficient assumptions (Theorem 1).

Assumption 1. (a) $Z \in \mathcal{Z}$, a bounded subset of $\mathbf{R}^d$ and the eigenvalues of $\mathrm{E}[ZZ^\top]$ are between $\lambda_{\min} > 0$ and $\lambda_{\max}$, (b) $\mu_1(z)$ and $\mu_0(z)$ are strictly positive and bounded on $\mathcal{Z}$, (c) there exists $\epsilon_\pi > 0$ such that $\epsilon_\pi \leq \pi_1(z) \leq 1 - \epsilon_\pi$, (d) there exists $\sigma_L > 0$ and σ_U such that $\mathrm{var}(Y^{(r)} \mid z) \in [\sigma_L, \sigma_U], z \in \mathcal{Z}$, (e) for some $q > 2, \mathrm{E}[|Y^{(r)}|^q \mid Z = z] \leq C < \infty$; (f) δ_0 is an interior point of a compact set $\Omega \in \mathbf{R}^{d+1}$; and (g) $(\mu_0, \pi_1) \in \mathcal{T}$, which is a set of measurable functions and there exist positive constants ϵ_π and ϵ_μ such that $\epsilon_\pi \leq \pi(z) \leq 1 - \epsilon_\pi$, and $\epsilon_\mu \leq \mu(z) \leq 1/\epsilon_\mu$ for any $(\mu, \pi) \in \mathcal{T}$.

Assumption 2. There exists $n_0, \epsilon_\pi > 0$, and $\epsilon_\mu > 0$ such that for all $n > n_0$, (a) $\epsilon_\pi \leq \widehat{\pi}(z) \leq 1 - \epsilon_\pi$; (b) $\epsilon_\mu \leq \widehat{\mu}_0(z) \leq 1/\epsilon_\mu$; (c) $\|\widehat{\mu}_0(z) - \mu_0(z)\|_{P,2} + \|\widehat{\pi}_1(z) - \pi_1(z)\|_{P,2} = o_P(n^{-1/4})$.

To specify the asymptotic distribution of the estimator $\widehat{\delta}$ under aforementioned assumptions, let

$$\widehat{w}^{(-k)}(\delta, Z_i) = \frac{e^{\delta^\top \widetilde{Z}_i} \widehat{\pi}_1^{(-k)}(Z_i) \widehat{\pi}_0^{(-k)}(Z_i)}{\left[e^{\delta^\top \widetilde{Z}_i} \widehat{\pi}_1^{(-k)}(Z_i) + \widehat{\pi}_0^{(-k)}(Z_i)\right]^2},$$

$$\begin{aligned}
\widehat{A}(\delta) &= n^{-1} \sum_{k=1}^{K} \sum_{i \in \mathcal{I}_k} \widetilde{Z}_i \widetilde{Z}_i^\top \widehat{w}^{(-k)}(\delta, Z_i) \\
&\quad \left\{ Y_i + \frac{\widehat{\mu}_0^{(-k)}(Z_i)}{\widehat{\pi}_1^{(-k)}(Z_i)} (R_i - \widehat{\pi}_1^{(-k)}(Z_i)) \right\} \\
&= n^{-1} \sum_{k=1}^{K} \sum_{i \in \mathcal{I}_k} \widetilde{Z}_i \widetilde{Z}_i^\top \widehat{w}^{(-k)}(\delta, Z_i) \\
&\quad \left[R_i Y_i + \frac{\widehat{\pi}_0^{(-k)}(Z_i)}{\widehat{\pi}_1^{(-k)}(Z_i)} \widehat{\mu}_0^{(-k)} R_i \right. \\
&\quad \left. + (1 - R_i)\{Y_i - \widehat{\mu}_0^{(-k)}(Z_i)\} \right],
\end{aligned}$$

and

$$\widehat{B}(\delta) = n^{-1} \sum_{k=1}^{K} \sum_{i \in \mathcal{I}_k} \widetilde{Z}_i \widetilde{Z}_i^\top \left(R_i \frac{\left[Y_i - e^{\delta^\top \widetilde{Z}_i} \widehat{\mu}_0^{(-k)}(Z_i)\right]}{e^{\delta^\top \widetilde{Z}_i} \widehat{\pi}_1^{(-k)}(Z_i) + \widehat{\pi}_0^{(-k)}(Z_i)} \widehat{\pi}_0^{(-k)}(Z_i) \right.$$

$$\left. - (1 - R_i) \frac{\left[Y_i - \widehat{\mu}_0^{(-k)}(Z_i)\right] e^{\delta^\top \widetilde{Z}_i}}{e^{\delta^\top \widetilde{Z}_i} \widehat{\pi}_1^{(-k)}(Z_i) + \widehat{\pi}_0^{(-k)}(Z_i)} \widehat{\pi}_1^{(-k)}(Z_i) \right)^2.$$

Then, applying Theorem 3.3 from Chernozhukov et al. (2018) under Assumptions 1 and 2 gives the following result:

Theorem 1. Let $\widehat{\delta}$ solve $S_n^{(cf)}(\delta) = 0$. Under Assumptions 1 and 2, $\sqrt{n}(\widehat{\delta} - \delta_0)$ converges weakly to a mean zero Gaussian distribution, whose variance can be consistently estimated by

$$\widehat{A}(\widehat{\delta})^{-1} \widehat{B}(\widehat{\delta}) \widehat{A}(\widehat{\delta})^{-1}.$$

See Appendix A for proof. Assumption 1 provides important regularity conditions that ensure finite, estimable parameters and nuisance parameters, and a non-degenerate asymptotic variance. Assumption 2 requires certain convergence rate for $\widehat{\mu}_0(\cdot)$ and $\widehat{\pi}_1(\cdot)$ in estimating $\mu_0(\cdot)$ and $\pi_1(\cdot)$, respectively. Under appropriate smoothness conditions for $\mu_0(z)$ and $\pi_1(z)$, there are multiple nonparametric estimators that achieve the required accuracy; see Chernozhukov et al. (2018) for a review of these estimators and their connection to cross-fitting estimators. The proposed estimating equation can be solved via Newton–Raphson method. Although we cannot guarantee that the derivative matrix $\widehat{A}(\delta)$ is positive definite in finite samples, its limit is positive definite, with a consistent estimate of either the propensity score or the main effect $\mu_0(z)$. We find good numerical convergence in practice, when the sample size is adequately large.

Remark 1. Constructing estimating equations $S_n^{(cf)}(\delta)$ based on different random partitions of the data and averaging the resulting solutions as the final estimator reduces the Monte Carlo variation due to randomly splitting data into K parts. Chernozhukov et al. (2018) showed that this estimator is asymptotically equivalent to $\widehat{\delta}$ analyzed above.

Remark 2. By comparing the conditional means to the baseline of $\mu_1(z)$, the semiparametric regression model (5) is equivalent to

$$\mathrm{E}\left[Y^{(r)} \mid Z = z\right] = \exp(-r\delta^\top \tilde{z}) \mu_1(z),$$

and a similar analysis to the above gives a set of symmetric estimating equations in terms of nuisance parameters μ_1 and π_1. In Section 2 of the supplementary materials, we provide procedure based on combining these symmetric estimating equations.

2.2.2. Two Regressions Approach

Returning to the misspecified regression from Section 2.1, recall that we were interested in ensuring that a CATE estimator will not introduce spurious heterogeneity due to confounding. If the data are from an RCT, so that $R \perp\!\!\!\perp \{Z, Y^{(r)}\}$, then

$$\begin{aligned}
&\mathrm{E}\left[\widetilde{Z}_i \left\{\mu_r(Z_i) - \exp(\beta^\top \widetilde{Z}_i)\right\} \mid R = r\right] \\
&= E\left[\widetilde{Z}_i \left\{\mu_r(Z_i) - \exp(\beta^\top \widetilde{Z}_i)\right\}\right] \triangleq s_r(\beta).
\end{aligned}$$

Now, assume that there is no heterogeneity, so that

$$D(z) = \exp(d_0), \tag{7}$$

holds. Even though the regression may be misspecified, the solutions β_r^* of the estimating equations $s_r(\beta) = 0$ in the two arms

will satisfy $\beta_1^* - \beta_0^* = (d_0, 0, \ldots, 0)^\top$, and thus, the estimated CATE score converges to $D(z) = \exp\{(\beta_1^* - \beta_0^*)^\top \widetilde{z}\} = \exp(d_0)$ in probability. This correctly suggests that there is no treatment effect heterogeneity even under misspecified regression models.

We propose an approach to correct for the confounding in observational data so that our estimator of the CATE will provide results as if estimated from a randomized trial, which avoids the spurious treatment effect heterogeneity from misspecified regression models. Specifically, this approach "recovers" fitting a simple regression model in both arms of an RCT, while viewing the regression model as a working model approximating the association of interest, and fitting the regression model as if the potential outcomes and covariates are observed in the entire cohort. Generally, this approach can produce biased estimate, when $\delta_0 \neq 0$ but model (5) is correctly specified. However, in simulation and on the real data from the NTD registry, we find that it performs well compared to other methods.

Constructing empirical versions of the estimating equations $s_r(\beta) = 0$ is not possible from the observed data, because $Y^{(r)}$ is only observed when $R = r$. Under the unconfoundedness assumption (2), we can apply methods developed to adjust for confounding when estimating the ATE to construct appropriate empirical estimating equations. To this end, let $\widehat{\pi}_r(z)$ be an estimator for the propensity score $\pi_r(z) = P(R = r|Z = z)$, and

$$\widehat{W}_i(r) = r\frac{R_i}{\widehat{\pi}_1(Z_i)} + (1 - r)\frac{1 - R_i}{\widehat{\pi}_0(Z_i)}.$$

Then, we can use the doubly robust estimating equation

$$S_r(\beta) = n^{-1}\sum_{i=1}^{n} \widetilde{Z}_i \left\{\widetilde{\mu}_r(Z_i) - \exp(\beta_r^\top \widetilde{Z}_i)\right\} = 0,$$

where $\widetilde{\mu}_r(z)$ is a special estimator of $\mu_r(z)$ constructed via the following steps:

1. Construct an initial nonparametric (or otherwise more flexible parametric or semiparametric) prediction for $Y_i^{(r)}$ given $Z_i = z$ via the estimated conditional expectation $\mathrm{E}(Y^{(r)}|Z = z)$, denoted by $\widehat{\mu}_r(z)$;
2. Solve the weighted estimating equations

$$n^{-1}\sum_{i=1}^{n} \widehat{W}_i(r)\widetilde{Z}_i \left(Y_i - \exp\left[\alpha_r \times \log\{\widehat{\mu}_r(Z_i)\} + \gamma_r^\top \widetilde{Z}_i\right]\right) = 0, \quad (8)$$

$$r = 0, 1; \quad (9)$$

   and denote the roots by $(\widehat{\alpha}_r, \widehat{\gamma}_r^\top)^\top, r = 0, 1$.
3. Let $\widetilde{\mu}_r(z) = \exp\left\{\widehat{\alpha}_r \times \log(\widehat{\mu}_r(Z_i)) + \widehat{\gamma}_r^\top \widetilde{Z}_i\right\}$ be the "calibrated" outcome predictions used in the estimating equation $S_r(\beta) = 0$.

This estimator is a doubly robust estimator: if either $\widehat{\mu}_r(\cdot)$ is a consistent estimator of $\mu_r(\cdot)$ or $\widehat{\pi}_r(\cdot)$ is a consistent estimator of $\pi_r(\cdot)$, then the solution to the augmented estimating equation converges to β_r, the solution of $s_r(\beta) = 0$ under (2) and mild regularity conditions (Bang and Robins 2005). The key observation is that if the propensity score is consistently estimated, Equation (9) ensures that

$$\mathrm{E}\left[\widetilde{Z}\{Y^{(r)} - \widetilde{\mu}_r(Z)\} \mid \widetilde{\mu}_r(\cdot)\right] = o_p(1).$$

If we suspect that the Poisson regression (1) is misspecified, the initial prediction rule should be based a more flexible model than the Poisson regression that better approximates the true model. For example, we may fit a regression model

$$\mathrm{E}(Y|Z, R = r) = \exp\{\eta_r^\top B(Z)\}, r = 0, 1,$$

where $B(z)$ is a rich set of basis functions capturing the complex nonlinear relationship between Y and Z. $\widehat{\mu}_r(z) = \exp\{\widehat{\eta}_r^\top B(z)\}$ can then be the initial prediction rule, where $\widehat{\eta}_r$ is the estimated regression coefficient. Alternatively, we may employ machine learning methods such as random forest or boosting to generate $\widehat{\mu}_r(z)$ (Friedman, Hastie, and Tibshirani 2000; Breiman 2001). We can improve the performance of the estimator by cross-fitting, which removes the dependence between $\widehat{\mu}_r(Z_i)$ and (Y_i, Z_i) induced by potential overfitting in constructing $\widehat{\mu}_r(z)$. Specifically, $\widehat{\beta}_r$ is the solution to the estimating equation $S_r^{(cf)}(\beta) = 0$, where

$$S_r^{(cf)}(\beta) = n^{-1}\sum_{k=1}^{K}\sum_{i\in\mathcal{I}_k} \widetilde{Z}_i \Bigg\{ \exp\left[\widehat{\alpha}_r \times \log\{\widehat{\mu}_r^{(-k)}(Z_i)\} + \widehat{\gamma}_r^\top \widetilde{Z}_i\right] - \exp(\beta_r^\top \widetilde{Z}_i)\Bigg\},$$

$\widehat{\gamma}_r$ and $\widehat{\alpha}_r$ are the roots of the estimating equation

$$n^{-1}\sum_{k=1}^{K}\sum_{i\in\mathcal{I}_k} \widehat{W}_i^{(-k)}(r)\widetilde{Z}_i \Bigg(Y_i - \exp\left[\alpha_r \times \log\{\widehat{\mu}_r^{(-k)}(Z_i)\} + \gamma_r^\top \widetilde{Z}_i\right]\Bigg) = 0,$$

data are divided into K nonoverlapping parts of approximately equal sizes indexed by $\mathcal{I}_k, k = 1, \ldots, K$, $\widehat{\mu}_r^{(-k)}(z)$ and $\widehat{\pi}_r^{(-k)}(z)$ are constructed using observations not in $\mathcal{I}_k$, and $\widehat{W}_i^{(-k)}(r)$ is the analog of $\widehat{W}_i(r)$ with $\widehat{\pi}_r^{(-k)}(z)$ plugged in. The estimated CATE score is thus

$$\widehat{D}_1(z) = \exp\left\{(\widehat{\beta}_1 - \widehat{\beta}_0)^\top \widetilde{z}\right\}.$$

Remark 3. One natural question is that if $\widetilde{\mu}_r(z)$, a high quality prediction rule for $Y_i^{(r)}|Z_i = z$, is already available, why do we need to reconstruct an estimator $\exp\{\widehat{\beta}_r^\top \widetilde{z}\}$ under a misspecified regression model? The initial prediction rule may be a complex function of z; therefore, it is not as transparent as that based on a simple regression model for clinical interpretation and practical use. We can view the regression-based CATE score as a "projection" of the initial prediction $\widetilde{\mu}_r(z)$ to a simpler functional space. This is in the same spirit of simplifying the estimated CATE by a classification tree (Foster, Taylor, and Ruberg 2011; Loh, He, and Man 2015).

2.3. Validation

For estimating and validating CATE models, Zhao et al. (2013) considered the absolute difference, $\mathrm{E}(Y^{(1)} - Y^{(0)} \mid Z = z)$ and $\mathrm{E}(Y^{(1)} - Y^{(0)}|D(Z) \geq c)$. In such a case, the ATE in $\{z \mid D(z) \geq c\}$, is a monotone increasing function of c. We generalize the method by Zhao et al. (2013) to address confounding due to

differences in baseline covariates between two treatment groups, and validate CATE estimators of the ratio of expected potential outcomes.

Let the ratio of average treatment effects among the subgroup of patients with the highest true CATE $\{z : D(z) \geq c\}$ be

$$AD_{\text{true}}(c) = \frac{\mathrm{E}\{Y^{(1)} \mid D(z) \geq c\}}{\mathrm{E}\{Y^{(0)} \mid D(z) \geq c\}}.$$

Because this subgroup is selected based on the true treatment effect ratios, $AD_{\text{true}}(c)$ is monotone increasing in c. Instead, if we order patients by the estimated CATE score $\widehat{D}(z)$ and let

$$AD(c) = \frac{\mathrm{E}\{Y^{(1)} \mid \widehat{D}(z) \geq c\}}{\mathrm{E}\{Y^{(0)} \mid \widehat{D}(z) \geq c\}},$$

then we would expect that when $\widehat{D}(z)$ is a good estimate of $D(z)$, $AD(c)$ will be monotone increasing; the trend of $AD(c)$ is a natural measure of the quality of the CATE score. To estimate the ATE in the subgroup of patients $\{z \mid \widehat{D}(z) \geq c\}$, we show how to adjust for confounding between two treatment arms using propensity score, regression or doubly robust estimators (Bang and Robins 2005; Kang and Schafer 2007).

The following theorem ensures that for CATE measured by the ratio of $\mu_r(z)$, as in the NTD example, $AD(c)$ is also monotone increasing.

Theorem 2. For nonnegative potential outcomes $Y^{(r)}, r = 0, 1$, Let

$$D(z) = \frac{\mathrm{E}(Y^{(1)} \mid Z = z)}{\mathrm{E}(Y^{(0)} \mid Z = z)}$$

and

$$AD(c) = \frac{\mathrm{E}(Y^{(1)} \mid D(z) \geq c)}{\mathrm{E}(Y^{(0)} \mid D(z) \geq c)}.$$

If all involved expectations are finite and $0 < D(Z) < \infty$ for almost every Z, $AD(c)$ is monotone increasing in c, and $AD(c) \geq c$ for any c.

See Section 3 of the supplementary materials for detailed proof. By Theorem 2, if we measure treatment effects by the ratio, we still can evaluate the quality of the constructed CATE scoring system by examining the "slope" of the curve $\widehat{AD}(c)$. Because, $AD(c) \geq c$, for any c, if $D(z)$ is the true CATE, the ATE in the subgroup consisting of patients with promising CATEs tends to be promising as well.

Remark 4. The monotonicity of $AD(c)$ depends on the metric used to measure the treatment effect. For example, $AD(c)$ is not necessarily monotone increasing if the treatment effect is measured by odds ratio (OR) for binary outcomes. In Section 4 of the supplementary materials, we provide a simple example where the largest marginal OR is not in the subgroup of patients with highest conditional OR. Generally, the ATE in a subgroup of patients with the largest CATE may not be large if the treatment effect is measured by a contrast other than the ratio or difference.

To estimate $AD(c)$ using observational data, we need to account for potential imbalances in baseline covariates between two arms, since the treatment assignment is not randomized. There are various ways to estimate the ATE in an observational study and all involve certain model assumptions. To construct a doubly robust approach, suppose that the validation set consists of m, independent identically distributed copies of (Y^V, R^V, Z^V), $\left\{(Y_i^V, R_i^V, Z_i^V), i = 1, \ldots, m\right\}$, where the superscript V indicates membership in the validation set. Then, estimate $AD(c)$ as follows: first estimate $\mu_r(z)$ by $\widehat{\mu}_{rc}(z)$ in the subgroup of patients $\{z^V \mid \hat{D}(z^V) \geq c\}$; and then estimate $\mathrm{E}(Y^{(r)} \mid \widehat{D}(Z^V) \geq c)$ by

$$\widehat{\mu}_r(c) = m_c^{-1} \sum_{\widehat{D}(Z_i^V) \geq c} \left[\widehat{\mu}_{cr}(Z_i^V) + \widehat{W}_i^V(r, c)\left\{Y_i^V - \widehat{\mu}_{cr}(Z_i^V)\right\}\right],$$

where

$$\widehat{W}_i^V(r, c) = r\frac{R_i^V}{\widehat{\pi}_{c1}(Z_i^V)} + (1 - r)\frac{1 - R_i^V}{\widehat{\pi}_{c0}(Z_i^V)},$$

$\widehat{\pi}_{cr}(z)$ is the estimator for $\pi_r(z)$ in the subgroup $\{z^V \mid \widehat{D}(z^V) \geq c\}$ from the validation set, and m_c is the subgroup size. Finally, let $\widehat{AD}(c)$ be the simple plug-in estimator $\widehat{\mu}_1(c)/\widehat{\mu}_0(c)$. This estimator is consistent for $AD(c)$, if either the propensity score or the main effect $\mu_r(z)$ is consistently estimated within the subgroup. One advantage of this approach is that it provides estimates of $\mu_r(z), r = 0, 1$, which allows interpretation of the treatment effect.

Remark 5. When the outcome of interest is time to a clinical event, such as death or relapse, the same method can be used to approximate and validate CATE, where the treatment effect is defined via the ratio of restricted mean time lost within a given time window, that is,

$$D(z) = \frac{\mathrm{E}\{\tau - (T^{(1)} \wedge \tau) | Z = z\}}{\mathrm{E}\{\tau - (T^{(0)} \wedge \tau) \mid Z = z\}}.$$

Here, $\tau > 0$ is a chosen constant (Uno et al. 2014) and $T^{(j)}$ is the event time of interest under treatment j. When the event rate is low, this ratio is similar to the hazard ratio and can be approximated by a multiplicative model (5). The detailed extension can be found in Section 5 of the supplementary materials.

3. Numerical Simulation

In this section, we conduct a simulation study to investigate the finite sample performance of the proposed method. The simulation design and discussion of the results are detailed below. Overall, these simulations show that when the ratio-based CATE is well-approximated by $D(z) = \exp(\delta^\top \widetilde{z})$ and the propensity score is correctly specified, the contrast regression and two regression approaches perform well. The contrast regression outperforms the two regression approach when the log-transformed CATE is well approximated by a linear function of z, but the Poisson model for the baseline rate $\mu_0(z)$ is misspecified. When the log-transformed CATE is highly nonlinear, the increased flexibility of boosting and other machine learning methods is advantageous; the ratio of the predictions from boosting outperforms the proposed method.

In the simulations, the covariate $Z \in \mathbf{R}^{10}$ is generated from a multivariate Gaussian, where the first 5 components are independent and the last 5 components are correlated with a common correlation coefficient of 0.5 but independent of the

first 5 components. The marginal distributions of the entries of Z are standard Gaussians. To ensure that the propensity score is bounded away from 0 and 1, any Z_i greater than 2 (or less than -2) is replaced by 2 (or -2). The treatment assignment $R \mid Z = z$ is generated from a Bernoulli distribution with a probability of $\pi_1(z) = \left\{1 + \exp(z_1 + 0.5z_2 - 0.5z_6)\right\}^{-1}$. With a slight abuse of notation, z_i stands for the ith component of the covariate vector rather than the covariate vector of the ith patient in this section.

We simulate a random follow up time $F^{(r)} \mid Z = z$ from a uniform distribution $U[0, 0.75]$ (discussion of incorporating follow up times is provided in the supplementary materials), and simulate the potential outcomes $Y^{(r)} \mid F^{(r)} = f, Z = z$ from a Poisson distribution Pois$\{\mu_r(z)f\}$, using the mean functions $\mu_r(z)$ described below, to illustrate a number of settings. The four different settings are:

1. Well-specified contrast

$$\begin{aligned} D(z) &= \exp\{-0.1 + 0.25(z_1 + z_6)\}, \\ \mu_1(z) &= \exp\{0.85 + 0.25(z_1 + z_6) + 1.5(|z_1| - |z_6|)\} \\ \mu_0(z) &= \exp\{0.95 + 1.5(|z_1| - |z_6|)\}; \end{aligned}$$

2. Well-specified Poisson

$$\begin{aligned} D(z) &= \exp\{0.375 + 0.125z_1 + 0.05z_2 - 0.25z_6\} \\ \mu_1(z) &= \exp\{0.925 + 0.125z_1 + 0.30z_2 + 0.25z_6\} \\ \mu_0(z) &= \exp\{0.550 + 0.25z_2 + 0.50z_6\}; \end{aligned}$$

3. Mild contrast misspecification

$$\begin{aligned} D(z) &= \exp\{\ 0.75 + 0.125z_1 + 0.05|z_2 + 0.5| - 0.25z_6\} \\ \mu_1(z) &= \exp\Big\{\ 0.50 + 0.125z_1 + 0.30|z_2 + 0.5| + 0.25z_6 \\ &\qquad + 0.5(|z_1| + |z_6|)\Big\} \\ \mu_0(z) &= \exp\Big\{ - 0.25 + 0.25|z_2 + 0.5| + 0.50z_6 \\ &\qquad + 0.5(|z_1| + |z_6|)\Big\}; \end{aligned}$$

4. Large contrast misspecification

$$\begin{aligned} D(z) &= \exp\Big\{0.915 - \ 0.25|z_1 + z_6 + 1| - 0.6|z_2 \\ &\qquad + 0.5| - 0.25z_6\Big\} \\ \mu_1(z) &= \exp\Big\{1.235 - 0.125|z_1 + z_6 + 1| - 0.3|z_2 \\ &\qquad + 0.5| - 0.125z_6 + 0.5(|z_1| + |z_6|)\Big\} \\ \mu_0(z) &= \exp\Big\{0.320 + 0.125|z_1 + z_6 + 1| + 0.3|z_2 + 0.5| \\ &\qquad + 0.125z_6 + 0.5(|z_1| + |z_6|)\Big\}. \end{aligned}$$

The proposed contrast regression is the most valuable for the well-specified contrast setting, where the Poisson regression is misspecified, but the underlying log-transformed CATE is still a linear combination of baseline covariates satisfying model (5). We expect that the naïve regression approach should work especially well in the well-specified Poisson setting, but that the propensity adjusted two regression and contrast regression approaches should also perform reasonably well. In the remaining two settings, neither the Poisson model for $\mu_r(z)$ nor the semiparametric model (5) for the CATE is correctly specified. While the linear approximation is reasonably good in the third setting, the log-transformed CATE is highly nonlinear in the fourth setting.

For each simulated dataset, we construct the CATE score using the following six methods:

1. contrast regression targeting $D(z)$ directly with the doubly robust adjustment,
2. two regression with the proposed doubly robust adjustment (the boosting-based estimation of $\mu_r(z)$ serves as the initial predictor),
3. naïve regression with a Poisson model in each arm,
4. boosting with a regression tree of depth 2 as the a base learner to estimate $\mu_r(z)$ in each arm separately and taking the ratio of two estimators,
5. modified outcome (MO) boosting regression according to Wendling et al. (2018),
6. Bayesian additive regression tree (BART) to estimate $\mu_r(z)$ and the CATE score taking the difference of two estimators (Lu et al. 2018).

The propensity score, when used, is always estimated by fitting a standard logistic regression model. We calculated the true $AD\{\widehat{H}^{-1}(1-q)\}$ and the validation curve, $q \mapsto AD\{\widehat{H}^{-1}(q)\}$, based on the constructed CATE scores. The steeper the slope of the curve, the better the performance of the CATE score. We have also directly calculated the correlation coefficients between the estimated CATE score and the true CATE (after log-transformation). After repeating this process 200 times, we summarize the performance of each method based on the median of the resulting validation curves (Figure 1), where the validation curve of the true CATE serves as the benchmark. Figure 2 shows the distribution of the correlation coefficients between the estimated CATE score and the truth. As expected, the contrast regression outperforms the two regression and other approaches in ranking the magnitude of CATE in the well-specified contrast setting. In most other cases, the CATE estimated by the two proposed methods have similar concordance with the true CATE. In the simulation for the well-specified Poisson model, the naïve regression performs the best, however the boosting and two proposed methods are only slightly inferior to the naïve regression. In the large contrast misspecification setting, where the log-transformed CATE is highly nonlinear, the ratio of the predictions from boosting outperforms the proposed methods, suggesting that the increased flexibility of the boosting approach (or other machine learning method) is advantageous.

In addition, for settings 1 and 2, model (5) is correctly specified and $\widehat{\delta}$ from the contrast regression is a consistent estimator of δ_0. We have also examined the empirical bias and the coverage level of 95% confidence intervals in estimating δ_0 based on 400 replications. The results are summarized in Table 1. The proposed contrast regression estimator is almost unbiased and the empirical coverage level of the constructed 95% confidence interval is close to the nominal level.

We also designed some simulations for survival outcomes, which can be found in Section 7 of the supplementary materials.

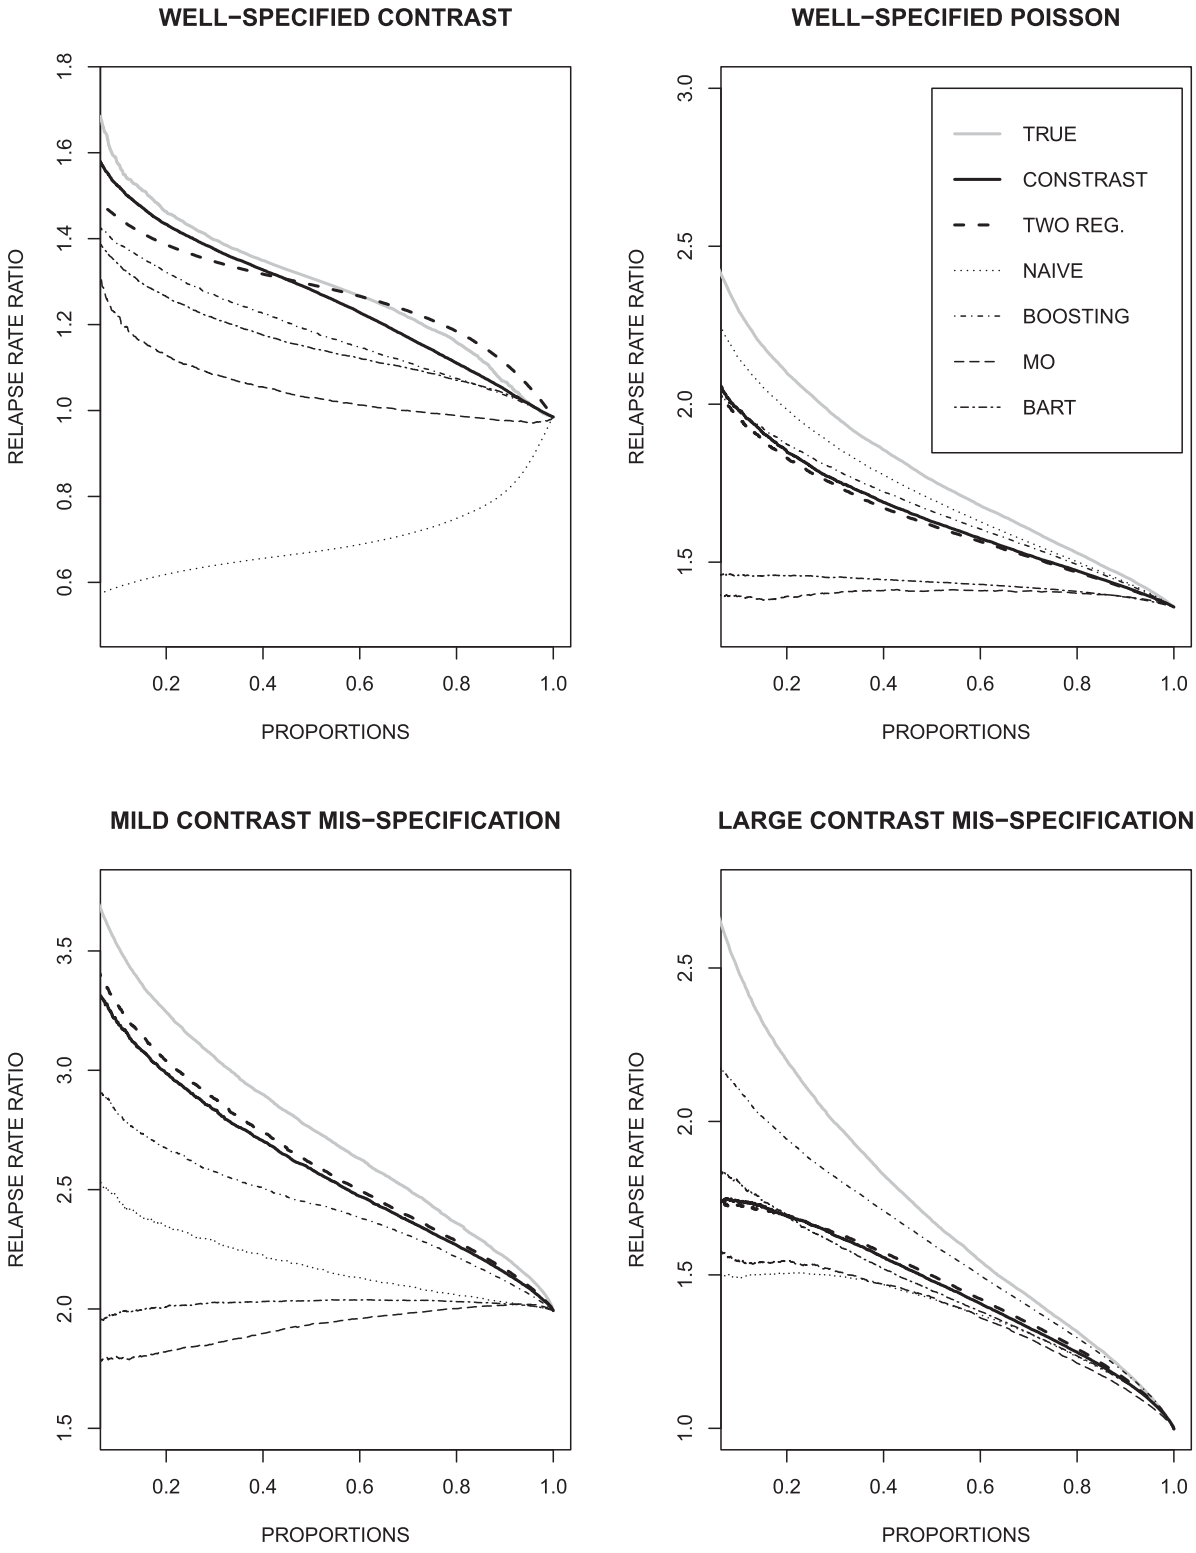

Figure 1. The ATE in subgroups of patients identified by different CATE scores in four simulation settings including: the true CATE, contrast regression, two regression, naïve regression, boosting, modified outcome boosting, and Bayesian additive regression trees (BART).

In general, the methods provided here adapt well to survival outcomes with censoring.

4. Treatments for Multiple Sclerosis

We return to our motivating example–measuring treatment effect heterogeneity between the TERI and DMF drugs for MS. As discussed in Section 1.1, one of the primary endpoints of interest is the relapse rate of severe symptoms of MS. The NTD registry records observational data of MS patients, including their treatments, relapses, and covariates over time. Hypothesized heterogeneity may be due to different drug treatment pathways, leading to different effectiveness across individuals. Here, we provide an in-depth description of the data, the results

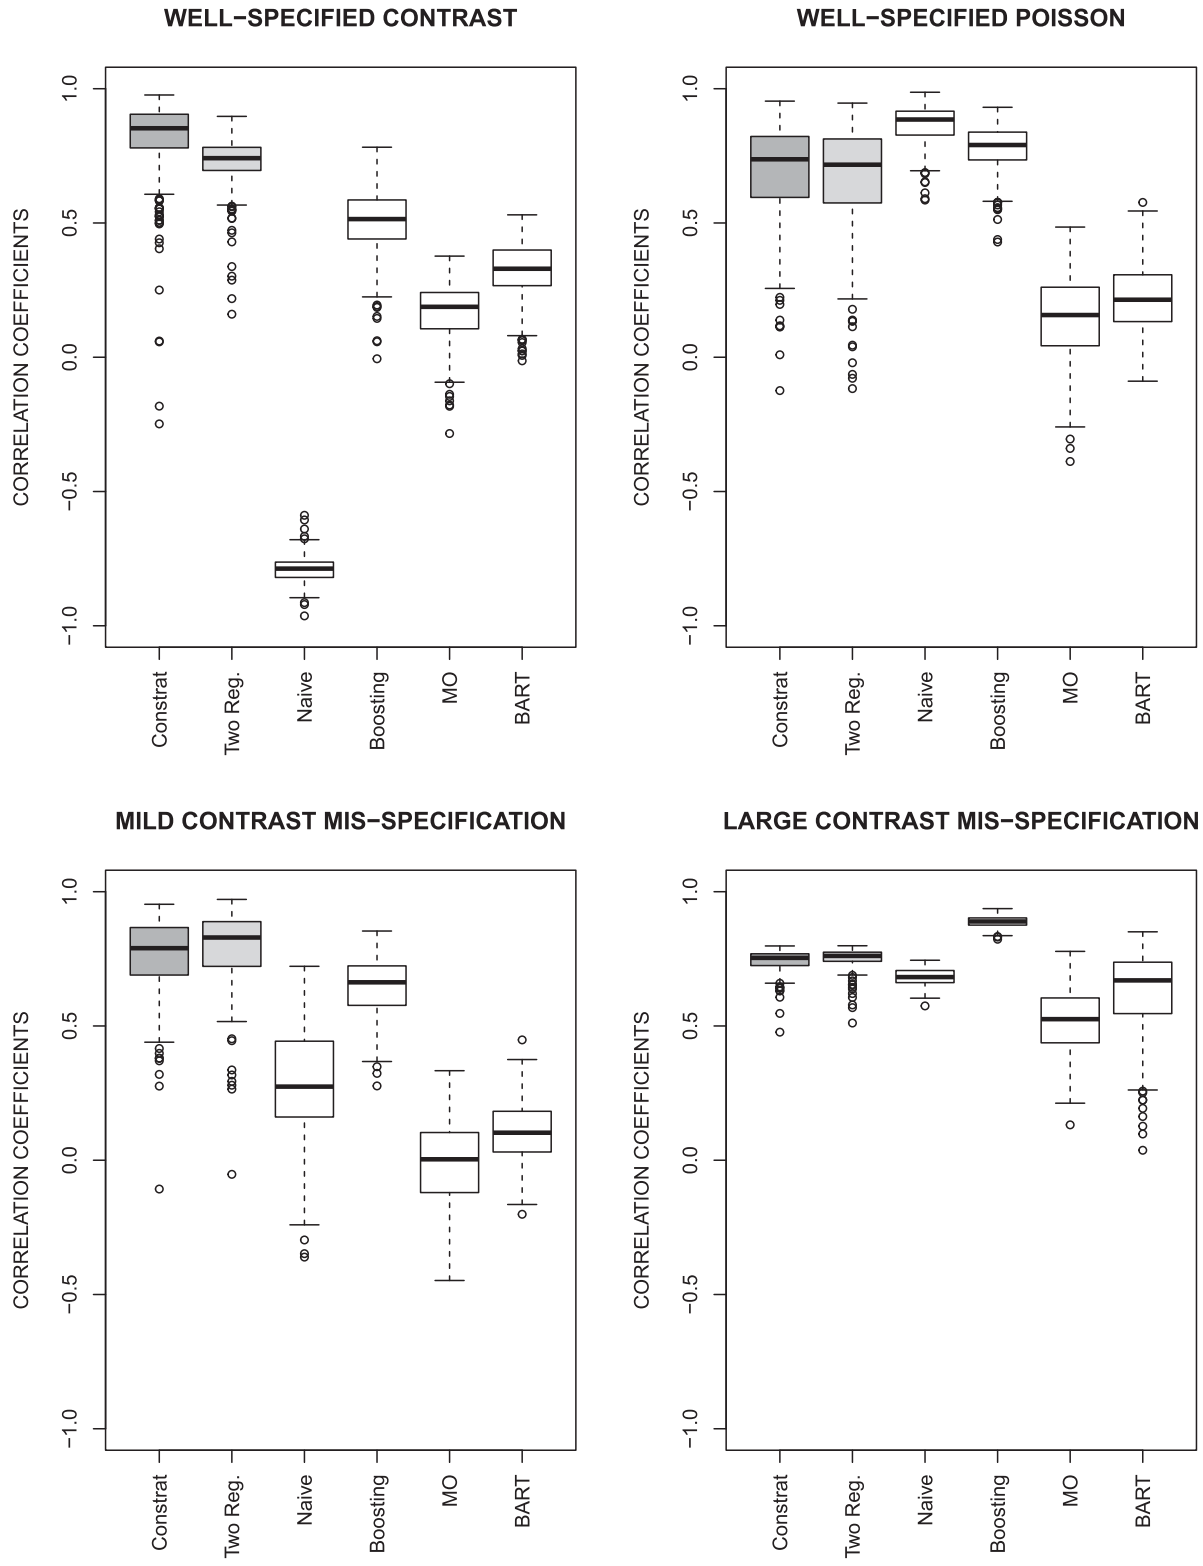

Figure 2. The distribution of correlation coefficients between the estimated and true CATE in four simulation settings; there are six methods considered from left to right: contrast regression (dark gray), two regression (light gray), naïve regression, boosting, modified outcome boosting, and Bayesian additive regression trees (BART).

of applying the proposed analyses on these observational data, and some implications of using the proposed methods to measure the CATE. Additional analysis and results for the time to relapse can be found in Section 8 of the supplementary materials.

4.1. Experimental Design

The NTD registry captured 1050 MS patients receiving TERI and 1741 patients receiving DMF between January 1, 2009 and July 1, 2018. Covariates of interest include age, number of prior treatments, MS duration, prior usage of glatiramer acetate (GA),

Table 1. Empirical bias and coverage level of the 95% confidence interval in estimating δ_0 based on 400 replications under the well-specified contrast and well-specified Poisson simulation settings.

	Well-specified Poisson			Well-specified contrast		
Covariates	Coefficients	Bias	Coverage	Coefficients	Bias	Coverage
Intercept	0.375	0.005	93.0%	−0.10	0.000	92.5%
Z_1	0.125	−0.021	93.8%	0.25	−0.013	91.3%
Z_2	0.050	−0.009	95.0%	0	−0.007	92.0%
Z_3	0	−0.003	94.5%	0	−0.007	95.0%
Z_4	0	0.003	93.0%	0	0.003	94.5%
Z_5	0	0.005	91.8%	0	0.003	93.5%
Z_6	−0.25	0.004	93.8%	0.25	−0.009	94.0%
Z_7	0	0.005	92.5%	0	0.006	95.0%
Z_8	0	0.000	93.5%	0	−0.001	92.8%
Z_9	0	−0.001	94.0%	0	−0.002	95.3%
Z_{10}	0	−0.004	93.0%	0	−0.003	93.5%

prior usage of interferon (IFN), number of relapses in one year and in two years prior to the index therapy, baseline Expanded Disability Status Scale (EDSS), and baseline pyramidal EDSS score. The data contain few missing values, thanks to processes to manage the definition of minimum datasets, mandatory data entry fields, and positive missing data confirmation. More details on the data source and management are available in Section 9 of the supplementary materials.

We implemented the standard regression, and the two proposed methods to construct the CATE scores approximating the "individualized" relapse rate ratio. To implement the proposed procedure, we estimated the baseline relapse rates $\mu_r(z)$ using boosting with the Poisson likelihood. The base learners are depth 2 regression trees, and the number of trees is selected via 5-fold cross-validation. The propensity score is constructed based on the standard logistic regression model. The proposed CATE score is based on the average of three replicates of 7-fold cross-fitting. One advantage of the contrast regression is that the standard errors can be estimated using the formulas provided in Section 2.2.1, and so we include these in our results.

We use repeated cross-validation to compare and evaluate the performance of the CATE scores objectively using the validation curves described in Section 2.3. To this end, we considered four CATE scores: a score based on predicted relapse rates using boosting method, a score based on naïve Poisson regression, and two scores based on our new proposal with the boosting-based prediction as the initial prediction. In each iteration, the data are split into a training set (67%) used to fit the CATE score and a testing set (33%) used to construct the validation curve. After repeating this process 50 times, we report the twice median validation curve for each CATE score in the training (left) and test (right) sets. We also use the estimated CATE score to split the patients in the testing set into two group of equal sizes. Then, we estimate the ratio of average relapse rates in two subgroups separately.

4.2. Results

Table 2 summarizes the distribution of covariates by treatment arm. The patients receiving TERI are different from patients receiving DMF in several key ways. For example, the patients receiving TERI tend to be older (45 vs. 40), have a longer disease duration (8.1 years vs. 6.6 years) and have higher EDSS scores (2.03 vs. 1.84) than those receiving DMF.

Table 2. Baseline characteristics of RRMS patients at the initiation of therapy with DMF and TERI: mean (standard deviation) for continuous covariate and number (proportion) for binary covariate.

Variable	TERI ($n = 1050$)	DMF ($n = 1741$)	p-value
Exposure time (year)	2.11(1.71)	2.17(1.72)	0.603
Age	44.86(10.20)	39.91(10.74)	0.0000
# prior treatments	0.97(0.93)	0.96(0.98)	0.4703
MS duration (year)	8.11(7.64)	6.57(6.60)	0.0000
GA	821(78.2%)	1327(76.2%)	0.246
IFN	502(47.8%)	886(50.9%)	0.118
# relapses (prior year)	0.42(0.60)	0.46(0.65)	0.2032
# relapses (prior 2 years)	0.64(0.84)	0.71(0.90)	0.095
EDSS	2.03(1.51)	1.84(1.50)	0.0006
Pyramidal EDSS	0.92(1.10)	0.77(1.04)	0.0000

Table 3. The estimated weights in constructed CATE scores (TERI vs. DMF).

	Ratio of relapse rate		
	Naïve reg.	Two reg.	Contrast reg.
Intercept	0.692	0.476	0.670 (0.711)[a]
Age	0.013	0.013	0.017 (0.013)
# prior treatments	−0.303	0.011	−0.088 (0.195)
MS duration (years)	0.022	0.045	0.028 (0.028)
GA	−0.584	0.517	−0.700 (0.349)
IFN	−0.304	−0.024	−0.185 (0.318)
# relapses (prior year)	−0.258	−0.661	−0.811 (0.271)
# relapses (prior two years)	0.191	0.360	0.444 (0.201)
EDSS	−0.046	−0.247	−0.233 (0.114)
Pyramidal EDSS	0.006	0.027	−0.004 (0.160)

[a]The estimated standard error of the weight.

In the entire cohort, the estimated ratio of the relapse rates (TERI vs. DMF) is 1.270 (95% confidence interval (CI): 1.121, 1.439; $p < 0.001$) using Poisson regression alone. After adjusting for confounding using the doubly robust procedure, the annual relapse rate is 0.308 for TERI and 0.237 for DMF; the estimated relapse rate ratio is 1.299 (95% CI: 1.018, 1.658).

Table 3 summarizes the estimated weights in the log-transformed CATE score for the naïve regression, two regression, and contrast regression approaches, as well as the estimated standard errors from the contrast regression. Based on contrast regression, GA, the number of relapses in the year prior to the therapy, the number of relapses in two years prior to the therapy, and baseline EDSS have a statistically significant impact on the treatment effectiveness at the 0.05 level. These weights suggest, for example, that patients experiencing more relapses in the previous two years and has a lower EDSS score tend to benefit more from DMF, even when looking at the relative ratio of treatment benefit.

The composition of the CATE score based on the naïve Poisson regression is different from that based on two new proposals. For example, the weight of EDSS from the naïve approach is substantially smaller than those in new CATE scores. Figure 3 shows a scatterplot of these three CATE scores in the entire cohort, demonstrating a positive correlation but also ample differences between the naïve and new CATE scores. On the other hand, the two new CATE scores are highly concordant. Comparing the cross-validation performance in Figure 4, the two proposed CATE scores that adjust for imbalance in baseline covariates appear to have a similarly superior performance to

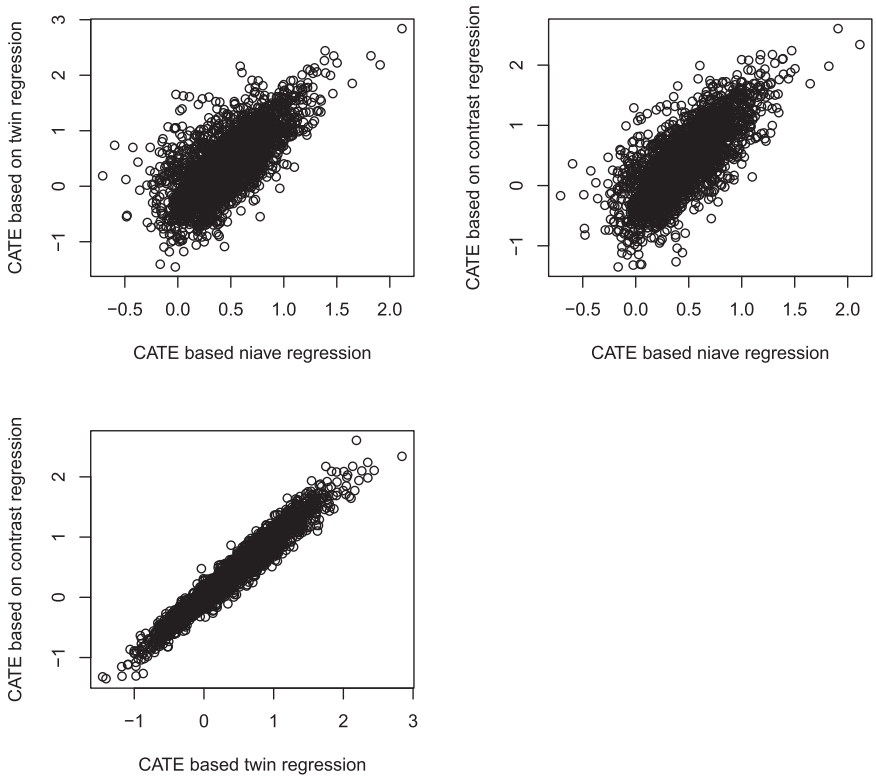

Figure 3. The log-transformed CATE scores based on the standard regression approach and the proposed doubly robust adjustment method: the CATE for the ratio of relapse rates.

naïve regression and boosting in the testing set, and both suggest a moderate treatment effect heterogeneity due to the monotone shape (with nonmonotone noise) of the validation curve.

Using the CATE score from contrast regression to split the patients in the testing set into two equal groups, the median ratio of the relapse rates (TERI vs. DMF) is 1.673 in 50% patients that would benefit most from DMF and 1.113 in remaining 50% patients. Using the CATE score from the two regression approach gives median ratios of 1.723 and 1.089, respectively. The distribution of the estimated ratios across different cross-validation replicates are summarized in Figure 5.

As a cautionary note, this observed difference in treatment effect may not be adequately stable due to the limited sample size in the testing set (on average, there are only 465 patients in each of the two subgroups). However, the results still exhibit signals for the presence of treatment effect heterogeneity captured by two proposed approaches.

A important observation is that the estimated treatment effect heterogeneity does not alter the recommendation of the treatment, since DMF appears to be superior to TERI in most if not all of the patients in terms of reducing relapse rate, although the relative benefit may vary in different subgroups.

5. Discussion

We show that estimation and validation of the ratio-based CATE benefits from many of the same approaches such as doubly robust estimation and semiparametric modeling that work well for the difference-based CATE. We also extend the regression approach by Zhao et al. (2013) to develop a precision medicine strategy from observational data. There are three important messages learned in this practice. First, the metric for the treatment effect has an important impact on the estimation and validation of the CATE. The treatment effects measured by the absolute difference and relative ratio both depend on the outcomes distribution in the control arm in simple ways, such that that the group of patients with large CATEs also have a large ATE. This is not necessarily true for treatment effect measured by odds or hazard ratio, where the ATE and CATE do not always align. Second, by borrowing appropriate techniques developed for estimating ATE in causal inference to adjust the standard estimation procedure, we eliminate the spurious heterogeneity caused by the imbalance in covariates in regression modeling for treatment covariates interactions. Lastly, we proposed a set of methods for estimating the ratio-based CATE, which may result in very different conclusions in comparison with most current methods, which target difference-based CATE.

We note that in this work, we have assumed that training and validation sets follow the same distribution. If the distribution of the validation set or the target population is different from that of the training set, the proposed estimation procedures need to be modified to adjust the distribution of covariates of the patients in treatment arm r of the training set to match that of the target population. Otherwise, the same CATE score may define a different subgroup of patients in the target population, that is, $\{z \mid \widehat{D}(z) \geq c\}$ may be different from $\{z^V \mid \widehat{D}(z^V) \geq c\}$, so that the ATE observed in the high value subgroup $\{z \mid \widehat{D}(z) \geq c\}$ may not be reproducible. Furthermore, the validity of all of our results depends on making the unconfoundedness assumption for causal effects. This is the most important for CATE validation, which would benefit greatly from being an

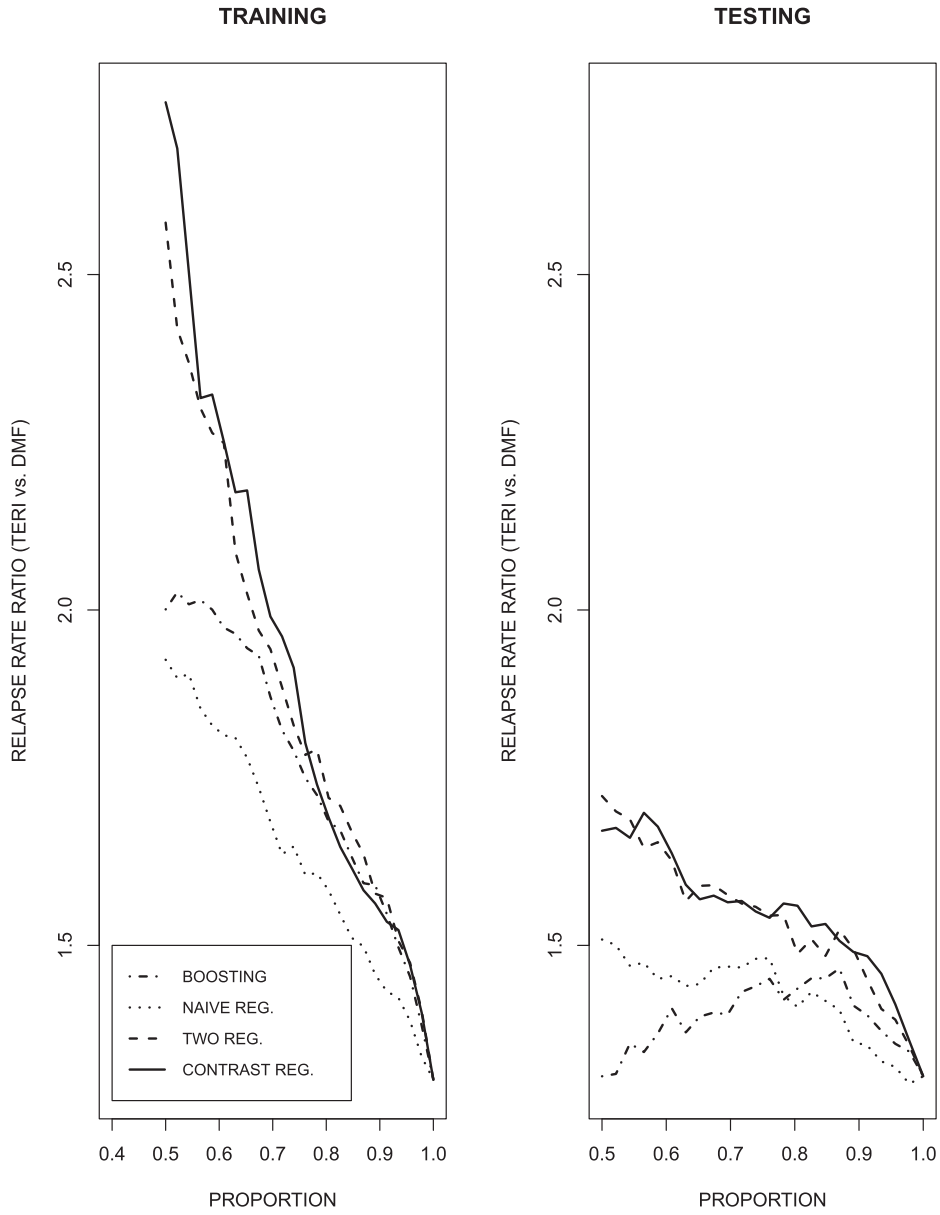

Figure 4. The ATE (relapse rate ratio of TERI vs. DMF) in subgroups of patients based on the CATE scores constructed in the training set (two proposed methods, naïve regression and boosting) in the NTD registry.

external randomized clinical trial, so that we can estimate the ATE in the identified high value subgroup without systematic biases from unmeasured confounding.

Qi et al. (2020) discussed the optimal treatment recommendation in the presence of $K > 2$ treatments. The proposed two regression approach can be used to approximate $\mu_k(z) = \mathrm{E}(Y^{(r)} \mid Z = z), k = 0, 1, \ldots, K$, and select treatments accordingly. The contrast regression can directly estimate $D_{ij}(z) = \mu_i(z)/\mu_j(z)$ based on the limiting estimating equation:

$$\mathrm{E}\Bigg[w(Z,\delta)\widetilde{Z}\Bigg\{ \prod_{k\neq i} \pi_k(Z) I(R=i) Y - \prod_{k\neq j} \pi_k(Z) I(R=j) Y \exp(\delta^\top \widetilde{Z}) \Bigg\}\Bigg] = 0$$

if $D_{ij}(z) = \exp(\delta_{ij}^\top \widetilde{z})$, where $\pi_k(z) = P(R = k \mid Z = z)$ and $w(z, \delta_{ij})$ is a weight function. Appropriate doubly robust augmentation based on $\{(I(R = r) - \pi_r(Z), r = 1, \ldots, K\}$ may further improve efficiency. However, the resulting estimators do not necessarily have the property that $D_{ij}(z) = D_{il}(z)D_{lj}(z)$, which warrants further research.

Appendix A: Proof of Theorem 1

Proof. To prove Theorem 1, it is sufficient to verify that the problem and assumptions satisfy those of Theorem 3.3 in Chernozhukov et al. (2018), which we repeat here for the reader's convenience. Let $c_0 > 0$, $c_1 > 0$, $a > 1$, $v > 0$, $s > 0$, and $q > 2$ be finite constants, and let $\{\delta_n\}_{n\geq 1}$, $\{\Delta_n\}_{n\geq 1}$, and $\{\tau_n\}_{n\geq 1}$ be some sequence of positive constants converging to 0. Define the following assumptions (Chernozhukov et al. 2018, Assumptions 3.3 and 3.4).

Assumption 3. For all $n \geq 3$ and $P \in \mathcal{P}_n$, (a) $\mathrm{E}\{m(G; \delta_0, \mu_0, \pi_1)\} = 0$, and Ω contains a ball of radius $c_1 n^{-1/2} \log n$ centered at δ_0; (b) the

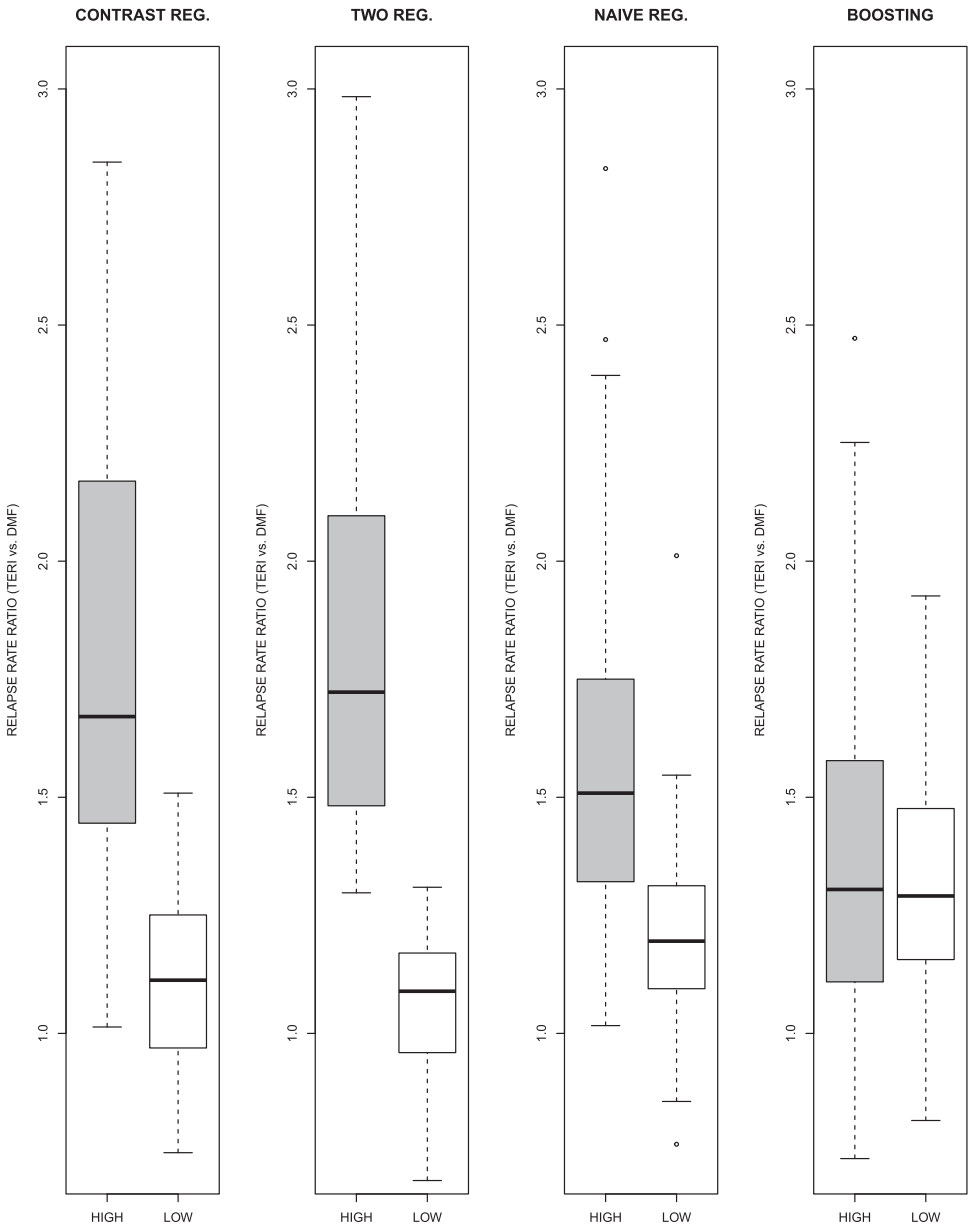

Figure 5. The cross-validated ATE (relapse rate ratio of TERI vs. DMF) of subgroups of patients identified by different CATE scores (two proposed methods, naïve regression and boosting) in the NTD registry.

map $(\delta, \mu, \pi) \to \mathrm{E}\{m(G; \delta, \mu, \pi)\}$ is twice continuously Gateaux-differentiable on $\Omega \times \mathcal{T}$; (c) for all $\delta \in \Omega$, $2\|\mathrm{E}\{m(G; \delta, \mu_0, \pi_1)\}\| \geq \|J_0(\delta - \delta_0)\| \wedge c_0$, where J_0 is the Jacobian matrix of $\delta \mapsto \mathrm{E}[m(G; \delta, \mu_0, \pi_1]$ at δ_0; (d) the score $m(g; \delta, \mu, \pi)$ obeys the Neyman orthogonality condition

$$\frac{\mathrm{d}}{\mathrm{d}r}\mathrm{E}\left[m(G, \delta_0, \mu_0 + r(\bar{\mu} - \mu_0), \pi_1 + r(\bar{\pi} - \pi_1))\right]\Big|_{r=0} = 0,$$

for any $(\bar{\mu}, \bar{\pi}) \in \mathcal{T}$.

Assumption 4. Let K be a fixed integer. For all $n \geq 3$ and $P \in \mathcal{P}_n$, the following conditions hold:

(a) Given a random subset I of $\{1, \ldots, n\}$ of size n/K, the nuisance parameter estimators $(\widehat{\mu}_0^{-k}, \widehat{\pi}_1^{-k})_{1 \leq k \leq K}$ belong to the realization set $\mathcal{T}_n$ with probability $1 - \Delta_n$, where $\mathcal{T}_n$ contains (μ_0, π_1) and is constrained by the conditions below;

(b) $\mathcal{F}_{1,(\mu,\pi)} = \{m_j(g; \delta, \mu, \pi) \mid j = 1, \ldots, d+1, \delta_0 \in \Omega\}$ is suitably measurable and its uniform covering entropy obeys

$$\sup_Q \log N(\epsilon \|F_{1,(\mu,\pi)}\|_{Q,2}, \mathcal{F}_{1,(\mu,\pi)}, \|\cdot\|_{Q,2}) \leq \nu \log(a/\epsilon_N)$$

for $\epsilon_N \in (0, 1]$, where $F_{1,(\mu,\pi)}$ is a measurable envelope for $\mathcal{F}_{1,(\mu,\pi)}$ that satisfies $\|F_{1,(\mu,\pi)}\|_{P,q} \leq c_1$;

(c) $r_n = \sup_{(\mu,\pi)\in\mathcal{T}_n, \delta_0 \in \Omega} \|\mathrm{E}\{m(G; \delta, \mu, \pi)\} - \mathrm{E}\{m(G; \delta_0, \mu_0, \pi_0)\}\| \leq \delta_n \tau_n$;

(d) $r_n' \log^{1/2}(1/r_n') \leq \delta_n$, where

$$r_n' = \sup_{(\mu,\pi)\in\mathcal{T}_n, \|\delta - \delta_0\| \leq \tau_n} \left(\mathrm{E}\left\{\|m(G; \delta, \mu, \pi) - m(G; \delta_0, \mu_0, \pi_0)\|^2\right\}\right)^{1/2};$$

(e) $\lambda_n \le \delta_n n^{-1/2}$, where

$$\lambda_n = \sup_{r\in(0,1),(\mu,\pi)\in\mathcal{T}_n,\|\delta-\delta_0\|\le\tau_n} \|\partial_r^2 \mathrm{E}\big[m\{G;\delta_0 + r(\delta-\delta_0), \mu_0 + r(\mu-\mu_0), \pi_1 + r(\pi-\pi_1)\}\big]\|.$$

(f) all eigenvalues of the matrix $\mathrm{E}\Big[m(G;\delta_0,\mu_0,\pi_1)m^\top(G;\delta_0,\mu_0,\pi_1)\Big]$ are bounded below by a positive constant.

Theorem 3 (Chernozhukov et al. (2018, Theorem 3.3)). Suppose that Assumptions 3 and 4 hold. In addition, suppose that $\delta_n \ge n^{-1/2+1/q}\log(n)$ and that $n^{-1/2}\log(n) \le \tau_n \le \delta_n$ for all $n \ge 1$ and a constant $q > 2$. Then, the DML2 estimator $\widehat{\delta}$ concentrates in a $1/\sqrt{n}$ neighborhood of δ_0, and are approximately linear and centered Gaussian:

$$\sqrt{n}\sigma^{-1}\left(\widehat{\delta}-\delta\right) = \frac{1}{\sqrt{n}}\sum_{i=1}^{n}\bar{\psi}(G_i) + O_p(\rho_n) \xrightarrow{d} N(0,I),$$

uniformly over $P \in \mathcal{P}_n$, where the size of the remainder term obeys

$$\rho_n = n^{-1/2+1/q}\log(n) + r_n'\log^{1/2}(1/r_n') + n^{1/2}\lambda_n + n^{1/2}\lambda_n',$$

$\bar{\psi}(\cdot) = -\sigma^{-1}J_0^{-1}m(\cdot,\delta_0,\mu_0,\pi_1)$ is the influence function, and the approximate variance is

$$\Sigma = J_0^{-1}\mathrm{E}\Big[m(G,\delta_0,\mu_0,\pi_1)m(G,\delta_0,\mu_0,\pi_1)^\top\Big]J_0^{-\top}.$$

We proceed by verifying the assumptions of Theorem 3. Let $\mathcal{P}$ be a set of measures satisfying Assumption 1 and $\mathcal{T}$ be a measurable subset of the pairs of functions (π,μ) such that for each $Q \in \mathcal{P}$, $\mu \in L_2(Q)$, $\pi \in L_\infty(Q)$, and $\epsilon_\pi \le \pi(z) \le 1-\epsilon_\pi$ Q-almost everywhere.

We proved Assumption 3(a) in the main text; see equation (6), by using the assumption in (5): $\mu_1(z) = \mu_0(z)\exp(\delta_0^\top\widetilde{z})$, and thus

$$\mathrm{E}[m(G;\delta_0,\mu_0)] = \mathrm{E}[w(Z,\delta_0)(\mu_1(Z) - \mu_0(Z)\exp(\delta_0^\top\widetilde{Z}))] = 0.$$

We will frequently use the fact that for $\delta \in \Omega$, $Z \in \mathcal{Z}$, and $\pi(z)$ satisfying $0 < \epsilon_\pi \le \pi(z) \le 1-\epsilon_\pi$,

$$\sup_{z\in\mathcal{Z},\delta\in\Omega}\{\exp(\delta^\top\widetilde{z})\pi(z) + 1 - \pi(z)\}^{-1} \le C_0 \tag{A.1}$$

for a constant C_0. Applying this with $\pi = \pi_1$ by Assumption 1(c), along with Assumption 1(e) ensures that $\partial\mathrm{E}[m(G;\delta,\mu_0,\pi_1) \mid Z = z]/\partial\delta$ and the second derivative have an integrable envelope function, and therefore, $\mathrm{E}[m(G;\delta,\mu_0,\pi_1)]$ is differentiable with respect to δ with Jacobian

$$J_0(\delta) = \mathrm{E}\left(\widetilde{Z}\widetilde{Z}^\top\frac{e^{\delta^\top\widetilde{Z}_i}\pi_1(Z)\pi_0(Z)\{e^{\delta_0^\top\widetilde{Z}}\pi_1(Z)+\pi_0(Z)\}\mu_0(Z)}{\left[e^{\delta^\top\widetilde{Z}}\pi_1(Z)+\pi_0(Z)\right]^2}\right),$$

which is continuous in δ and positive definite with its smallest eigenvalue uniformly bounded away from zero for $\delta \in \Omega$. We can choose a small open ball centered at δ_0, $\mathcal{N}$, such that for any $\delta \in \mathcal{N}$, $|J_0(\delta) - J_0(\delta_0)|_{ij} < \epsilon$, for all components $1 \le i,j \le d+1$, where ϵ is a small constant to be specified later. By the intermediate value theorem, for any $\delta \in \mathcal{N}$, there exists $\bar{\delta} \in \mathcal{N}$ such that

$$\begin{aligned}\|\mathrm{E}\{m(G;\delta,\mu_0,\pi_1)\}\| &= \|J_0(\bar{\delta})(\delta-\delta_0)\| \ge \|J_0(\delta_0)(\delta-\delta_0)\| \\ &\quad - \|\{J_0(\bar{\delta}) - J_0(\delta_0)\}(\delta-\delta_0)\| \\ &\ge \|J_0(\delta_0)(\delta-\delta_0)\| - \epsilon(d+1)\|\delta-\delta_0\| \\ &\ge \|J_0(\delta_0)(\delta-\delta_0)\| - \|J_0(\delta_0)(\delta-\delta_0)\|/2,\end{aligned}$$

if $\epsilon \le \lambda_0/2(d+1)$, where λ_0 is the smallest eigenvalue of

$$J_0 = J_0(\delta_0) = \mathrm{E}\left(\widetilde{Z}\widetilde{Z}^\top\frac{e^{\delta_0^\top\widetilde{Z}_i}\pi_1(Z)\pi_0(Z)\mu_0(Z)}{\left[e^{\delta_0^\top\widetilde{Z}}\pi_1(Z)+\pi_0(Z)\right]}\right).$$

Therefore, $2\|\mathrm{E}\{m(G;\delta,\mu_0,\pi_1)\}\| \ge \|J_0(\delta-\delta_0)\|$, for any δ within the ball. For δ outside $\mathcal{N}$, let $c_0 = \inf_{\delta\in\Omega-\mathcal{N}}\|\mathrm{E}\{m(G;\delta,\mu_0,\pi_1)\}\|$. $c_0 > 0$ due to the uniform continuity of $\mathrm{E}\{m(G;\delta,\mu_0,\pi_1)\}$ in the compact set $\Omega - \mathcal{N}$. This verifies Assumption 3(c).

To verify the Gateaux-differentiability of m, note that for $\delta \in \Omega$, $(\mu,\pi) \in \mathcal{T}$, $(\mu + rd_\mu,\pi) \in \mathcal{T}$, and $(\mu,\pi + rd_\pi) \in \mathcal{T}$,

$$\begin{aligned}&\frac{1}{r}\left[\mathrm{E}\{m(G;\delta,\mu+rd_\mu,\pi)\} - \mathrm{E}\{m(G;\delta,\mu,\pi)\}\right] \\ &\quad = -\mathrm{E}\left\{\frac{\{\pi_1(Z)(1-\pi(Z)) - \pi_0(Z)\pi(Z)\}e^{\delta^\top\widetilde{Z}}}{e^{\delta^\top\widetilde{Z}}\pi(Z)+1-\pi(Z)}d_\mu(Z)\right\},\end{aligned}$$

and

$$\begin{aligned}&\frac{1}{r}\left[\mathrm{E}\{m(G;\delta,\mu,\pi+rd_\pi)\} - \mathrm{E}\{m(G;\delta,\mu,\pi)\}\right] \\ &= -\mathrm{E}\left\{\frac{\begin{bmatrix}\pi_1(Z)\{\mu_1(Z) - \mu(Z)e^{\delta^\top\widetilde{Z}}\} \\ +\pi_0(Z)\{\mu_0(Z)-\mu(Z)\}\end{bmatrix}e^{\delta^\top\widetilde{Z}}}{\left(e^{\delta^\top\widetilde{Z}}\pi(Z)+1-\pi(Z)\right)\left(e^{\delta^\top\widetilde{Z}}\{\pi(Z)+rd_\pi(Z)\}+1-\pi(Z)-rd_\pi(Z)\right)}d_\pi(Z)\right\}.\end{aligned}$$

By dominated convergence theorem, we may exchange the $\lim_{r\to0}$ with the expectation, and the Gateuaux derivative with respect to μ along the direction of d_μ exists. Similarly, the Gateuaux derivative with respect to π along the direction of d_π also exists, and is

$$-\mathrm{E}\left\{\frac{\left[\pi_1(Z)\{\mu_1(Z)-\mu(Z)e^{\delta^\top\widetilde{Z}}\}+\pi_0(Z)\{\mu_0(Z)-\mu(Z)\}\right]e^{\delta^\top\widetilde{Z}}}{\left\{e^{\delta^\top\widetilde{Z}}\pi(Z)+1-\pi(Z)\right\}^2}d_\pi(Z)\right\}.$$

The smoothness of the numerator, and boundedness of the denominator similarly allow for second-order differentiability.

To examine the orthogonality condition, that is, Assumption 3(d), let $(\bar{\mu},\bar{\pi}) \in \mathcal{T}$, $d_\mu(z) = \bar{\mu}(z) - \mu_0(z)$, $d_\pi(z) = \bar{\pi}(z) - \pi_1(z)$, and

$$\bar{m}(\mu,\pi,f) = \mathrm{E}\Big[Z\Big\{\frac{(1-\pi(Z))R(Y-\mu(Z)e^{f(Z)})}{e^{f(Z)}\pi(Z)+1-\pi(Z)} - \frac{\pi(Z)(1-R)(Y-\mu(Z))e^{f(Z)}}{e^{f(Z)}\pi(Z)+1-\pi(Z)}\Big\}\Big].$$

Then,

$$\begin{aligned}g_0(r) &= \bar{m}(\mu_0 + rd_\mu,\pi_1 + rd_\pi,f_0) \\ &= \mathrm{E}\left(Z\frac{(\pi_0(Z) - rd_\pi(Z))R(Y_1 - (\mu_0(Z)+rd_\mu(Z))e^{f_0(Z)})}{e^{f_0(Z)}(\pi_1(Z)+rd_\pi(Z)) + 1 - \pi_1(Z) - rd_\pi(Z)}\right) \\ &\quad - \mathrm{E}\left(Z\frac{(\pi_1(Z)+rd_\pi(Z))(1-R)(Y_0 - (\mu_0(Z)+rd_\mu(Z)))e^{f_0(Z)}}{e^{f_0(Z)}(\pi_1(Z)+rd_\pi(Z)) + 1 - \pi_1(Z) - rd_\pi(Z)}\right) \\ &= \mathrm{E}\left(Z\frac{r^2e^{f_0(Z)}d_\pi(Z)d_\mu(Z)}{e^{f_0(Z)}\{\pi_1(Z)+rd_\pi(Z)\}+\pi_0(Z)-rd_\pi(Z)}\right),\end{aligned}$$

where $f_0(z) = \delta_0^\top z$. Let $\bar{\mu}(z) = \mu_0(z) + d_\mu(z)$ and $\bar{\pi}(z) = \pi_1(z) + d_\pi(z)$. Using a similar argument as above, because (A.1) is bounded for $\pi = \pi_1$ by Assumption 1(c) and $\mu_1(z)$ is integrable by Assumption 1(e), the dominated convergence theorem yields

$$\frac{dg_0(r)}{dr} = 2r\mathrm{E}\left(Z\frac{e^{f_0(Z)}d_\pi(Z)d_\mu(Z)}{e^{f_0(Z)}(\pi_1(Z)+rd_\pi(Z))+1-\pi_1(Z)-rd_\pi(Z)}\right) - r^2\mathrm{E}\left(Z\frac{e^{f_0(Z)}(e^{f_0(Z)}-1)d_\pi(Z)^2d_\mu(Z)}{\{e^{f_0(Z)}(\pi_1(Z)+rd_\pi(Z))+1-\pi_1(Z)-rd_\pi(Z)\}^2}\right).$$

Therefore,

$$\left.\frac{dg_0(r)}{dr}\right|_{r=0} = 0,$$

which verifies Assumption 3(d).

Assumption 2 implies that there exists sequences $\log(n)n^{-1/4} \le a_n = o(1)$, and $\Delta_n' = o(1)$, such that

$$\|\widehat{\pi}(\cdot)-\pi_1(\cdot)\|_{P,2} + \|\widehat{\mu}(\cdot)-\mu_0(\cdot)\|_{P,2} \le a_n n^{-1/4},$$

with probability $1-\Delta_n'/2$. Note that a_n can be chosen such that these hold when $\widehat{\pi}$ and $\widehat{\mu}$ are estimated using only $(1-K^{-1})n$ (as opposed to n) samples. Let

$$\mathcal{T}_n = \Big\{(\pi,\mu) \mid \pi,\mu \text{ are measureable}, \pi(\cdot)\in[\epsilon_\pi, 1-\epsilon_\pi], \mu(\cdot) \in [\epsilon_\mu, \epsilon_\mu^{-1}], \text{ and } \|\pi(\cdot)-\pi_1(\cdot)\|_{P,2} + \|\mu(\cdot)-\mu_0(\cdot)\|_{P,2} \le a_n n^{-1/4}\Big\},$$

Then, $P(\{(\widehat{\pi}^{(-k)},\widehat{\mu}^{(-k)})\in\mathcal{T}_n\}_{k=1}^K) \ge 1-K\Delta_n'$. Let $\Delta_n = K\Delta_n'$, and Assumption 4(a) is satisfied.

For $Q\in\mathcal{P}$, and $(\mu,\pi)\in\mathcal{T}$,

$$\begin{aligned}
&\|m(G;\bar{\delta},\pi,\mu) - m(G;\delta_0,\pi,\mu)\|_{Q,2}\\
=&\left\|\widetilde{Z}\left(\pi(1-R)Y^{(0)} + \mu(R-\pi)\right)\left(\frac{\exp(\bar{\delta}^\top\widetilde{Z})}{e^{\bar{\delta}^\top\widetilde{Z}}\pi+(1-\pi)} - \frac{\exp(\delta_0^\top\widetilde{Z})}{e^{\delta_0^\top\widetilde{Z}}\pi+(1-\pi)}\right)\right\|_{Q,2}\\
\le&\left\|\widetilde{Z}\{\pi\pi_0\mu_0 + \mu(\pi_1-\pi)\}\right\|_{Q,2}\left\|(1-\pi)\frac{\exp(\bar{\delta}_0^\top\widetilde{Z}) - \exp(\delta^\top\widetilde{Z})}{(e^{\bar{\delta}^\top\widetilde{Z}}\pi+(1-\pi))(e^{\delta_0^\top\widetilde{Z}}\pi+(1-\pi))}\right\|_{Q,2}\\
\le& C\left\|\exp(\bar{\delta}_0^\top\widetilde{z}) - \exp(\delta^\top\widetilde{z})\right\|_\infty \le CL_{\mathrm{rad}}\|\bar{\delta}-\delta_0\|_\infty,
\end{aligned}$$

where we suppressed Z in functions such as $\pi(Z)$, $\mu(Z)$, etc. to simplify notation, L_{rad} is the Lipschitz constant of $t\mapsto\exp(t)$ over $|t|\le\sup_{\delta\in\Omega,z\in\mathcal{Z}}|\delta^\top\widetilde{z}|$. Therefore, $m(G;\delta,\mu,\pi)$ is Lipschitz in δ. For all $(\mu,\pi)\in\mathcal{T}$, (A.1) and the fact that $\mu\in L_2(Q)$ imply that there exists a squared-integrable envelope function $F_{1,(\mu,\pi)}$. This, the Lipschitz constraint, and the bound $\log N(\epsilon_N,\Omega,\|\cdot\|_\infty) \le \widetilde{\nu}\log(\widetilde{a}/\epsilon_N)$ on the parameter space imply that $\sup_Q \log N(\epsilon_N\|F_{1,(\mu,\pi)}\|_{Q,2},\mathcal{F}_{1,\mu,\pi},\|\cdot\|_{Q,2}) \le \nu\log(a/\epsilon_N)$. Thus, Assumption 4(b) is verified.

$$\begin{aligned}
r_n &= \|\mathrm{E}\{m(G;\delta,\mu,\pi)\} - \mathrm{E}\{m(G;\delta,\mu_0,\pi_1)\}\|\\
&= \left\|\mathrm{E}\left[\widetilde{Z}\frac{(\pi_1-\pi)\exp(\delta^\top\widetilde{Z})(\pi_1(\mu_1-\mu\exp(\delta^\top\widetilde{Z})) + (1-\pi)(\mu_0-\mu))}{\left(e^{\delta^\top\widetilde{Z}}\pi+(1-\pi)\right)\left(e^{\delta^\top\widetilde{Z}}\pi_1+(1-\pi_1)\right)}\right]\right\|\\
&\le \widetilde{C}_1\|\pi-\pi_1\|_{P,2},
\end{aligned}$$

for a finite constant $\widetilde{C}_1$. Therefore, by Assumption 4(c), we can choose $\tau_n = (a_n^{3/8}n^{-1/4})$ and $\delta_n = \sqrt{a_n}$ to satisfy $r_n \le \widetilde{C}_1\|\pi-\pi_1\|_{P,2} \le \delta_n\tau_n$ for adequately large n, using the definition of $\mathcal{T}_n$. Next,

$$\begin{aligned}
&\left\{\mathrm{E}\|m(G;\delta,\mu,\pi) - m(G;\delta_0,\mu_0,\pi_1)\|^2\right\}^{1/2}\\
\le&\sqrt{3}\left\{\mathrm{E}\|m(G;\delta,\mu,\pi) - m(G;\delta,\mu_0,\pi)\|^2\right\}^{1/2}\\
&+\sqrt{3}\left\{\mathrm{E}\|m(G;\delta,\mu_0,\pi) - m(G;\delta,\mu_0,\pi_1)\|^2\right\}^{1/2}\\
&+\sqrt{3}\left\{\mathrm{E}\|m(G;\delta,\mu_0,\pi_1) - m(G;\delta_0,\mu_0,\pi_1)\|^2\right\}^{1/2}\\
=&\sqrt{3}\left(\mathrm{E}\left[\|\widetilde{Z}\|^2\frac{\exp(2\delta^\top\widetilde{Z})\{\pi_1(1-\pi)^2+\pi_0\pi^2\}}{\left(e^{\delta^\top\widetilde{Z}}\pi+(1-\pi)\right)^2}(\mu-\mu_0)^2\right]\right)^{1/2}\\
&+\sqrt{3}\left(\mathrm{E}\left[\|\widetilde{Z}\|^2\frac{\exp(2\delta^\top\widetilde{Z})\left\{\pi_1\left(Y^{(1)}-\mu_0\exp(\delta^\top\widetilde{Z})\right)^2 + \pi_0\left(Y^{(0)}-\mu_0\right)^2\right\}}{\left(e^{\delta^\top\widetilde{Z}}\pi+(1-\pi)\right)^2\left(e^{\delta^\top\widetilde{Z}}\pi_1+\pi_0\right)^2}(\pi_1-\pi)^2\right]\right)^{1/2}\\
&+\sqrt{3}\left(\mathrm{E}\left[\|\widetilde{Z}\|^2\frac{\pi_1^2\pi_0^2\left\{\pi_1\left(Y^{(1)}-\mu_1\right)^2+\pi_0\left(Y^{(0)}-\mu_0\right)^2 + (\pi_1\mu_1+\pi_0\mu_0)^2/\pi_1\right\}}{\left(e^{\delta^\top\widetilde{Z}}\pi+(1-\pi)\right)^2\left(e^{\delta^\top\widetilde{Z}}\pi_1+\pi_0\right)^2}(e^{\delta^\top\widetilde{Z}}-e^{\delta_0^\top\widetilde{Z}})^2\right]\right)^{1/2}\\
\le&\widetilde{C}_2\left(\|\pi-\pi_1\|_{P,2} + \|\mu-\mu_0\|_{P,2} + \|\delta-\delta_0\|_2\right),
\end{aligned}$$

where $(\mu,\pi)\in\mathcal{T}_n$, $\|\delta-\delta_0\|_2\le\tau_n$, $\widetilde{C}_2$ is a finite constant that depends on constants such as σ_U^2, ϵ_π, ϵ_μ in Assumption 1. Therefore,

$$\begin{aligned}
r_n' &= \sup_{(\mu,\pi)\in\mathcal{T}_n,\|\delta-\delta_0\|_2\le\tau_n}\left\{\mathrm{E}\|m(G;\delta,\mu,\pi) - m(G;\delta_0,\mu_0,\pi_1)\|^2\right\}^{1/2}\\
&\le \widetilde{C}_2\left(a_n n^{-1/4} + \tau_n\right).
\end{aligned}$$

Thus, $r_n'\log^{1/2}(1/r_n') \le \widetilde{C}_2 a_n^{3/8}n^{-1/4}\sqrt{\log(n)} \le \delta_n$ and, thus, Assumption 4(d) is satisfied.

Let $d_f(z) = \bar{f}(z) - f_0(z) = (\bar{\delta}-\delta_0)^\top\widetilde{z}$. Define

$$
\begin{aligned}
k(r) &= \tilde{m}(\mu_0 + rd_\mu, \pi_1 + rd_\pi, f + rd_f) \\
&= \mathrm{E}\left(Z \frac{\left[\begin{array}{l} (\pi_0 - rd_\pi) R (Y_1 - (\mu_0 + rd_\mu) e^{f_0 + rd_f}) \\ \quad -(\pi_1 + rd_\pi)(1-R)(Y_0 - (\mu_0 + rd_\mu)) e^{f_0 + rd_f} \end{array} \right]}{e^{f_0 + rd_f}(\pi_1 + rd_\pi) + \pi_0 - rd_\pi} \right) \\
&= \mathrm{E}\left(Z e^{f_0} \frac{\left[(1 - e^{rd_f}) \pi_0 \pi_1 \mu_0 - r(1 - e^{rd_f}) \pi_1 \mu_0 d_\pi + r^2 e^{rd_f} d_\mu d_\pi \right]}{e^{f_0 + rd_f}(\pi_1 + rd_\pi) + \pi_0 - rd_\pi} \right) \\
&= \mathrm{E}\left\{ Z e^{f_0(Z)} \frac{h_1(Z,r)}{h_2(Z,r)} \right\},
\end{aligned}
$$

where $h_1(z,r) = (1 - e^{rd_f(z)})\pi_0(z)\pi_1(z)\mu_0(z) - r(1 - e^{rd_f(z)})\pi_1(z)\mu_0(z)d_\pi(z) + r^2 e^{rd_f(z)} d_\mu(z) d_\pi(z)$, and $h_2(z,r) = e^{f_0(z) + rd_f(z)}\{\pi_1(z) + rd_\pi(z)\} + \pi_0(z) - rd_\pi(z)$. Similar as above, by the dominated convergence theorem, we have

$$
\begin{aligned}
\frac{d^2 k(r)}{dr^2} =& \mathrm{E}\left(Z e^{f_0(Z)} \left[\frac{\partial^2 h_1(Z,r)/\partial r^2}{h_2(Z,r)} - 2 \frac{\partial h_1(Z,r)/\partial r \partial h_2(Z,r)/\partial r}{h_2^2(Z,r)} \right.\right. \\
& \left.\left. - \frac{h_1(Z,r)\partial^2 h_2(Z,r)/\partial r^2}{h_2^2(Z,r)} + 2 \frac{h_1(Z,r)\{\partial h_2(Z,r)/\partial r\}^2}{h_2^3(Z,r)} \right] \right),
\end{aligned}
$$

where

$$
\begin{aligned}
\frac{\partial h_1(z,r)}{\partial r} =& -e^{rd_f} \pi_0 \pi_1 \mu_0 d_f - (1 - e^{rd_f}) \pi_1 \mu_0 d_\pi + R_{11}(r,z;\kappa) \\
\frac{\partial^2 h_1(z,r)}{\partial r^2} =& -e^{rd_f} \pi_0 \pi_1 \mu_0 d_f^2 + R_{12}(r,z;\kappa) \\
\frac{\partial h_2(r)}{\partial r} =& (e^{f_0 + rd_f} - 1) d_\pi + e^{f_0 + rd_f} \pi_1 d_f + R_{21}(r,z;\kappa) \\
\frac{\partial^2 h_1(r)}{\partial r^2} =& e^{f_0 + rd_f} \pi_1 d_f^2 + R_{22}(r,z;\kappa),
\end{aligned}
$$

$\kappa = (d_\pi, d_\mu, d_f)^\top$, and $R_{ij}(r,z;\kappa)$ is a function of r, z satisfying that

$$
\sup_{(r,z) \in [0,1] \times \mathcal{Z}} \frac{|R_{ij}(r,z;\kappa)|}{|d_\pi(z) d_\mu(z)| + |d_f(z) d_\pi(z)|} \leq \widetilde{C}_3
$$

for a constant $\widetilde{C}_3$. Therefore, after careful regrouping,

$$
\begin{aligned}
\lambda_n =& \sup_{r \in (0,1), (\bar{\mu}, \bar{\pi}) \in \mathcal{T}_n, |\bar{\delta} - \delta_0| \leq \tau_n} \left\| \frac{d^2 k(r)}{dr^2} \right\| \\
\leq& \widetilde{C}_4 (\|d_f\|_{P,2}^2 + \|d_f\|_{P,2} \|d_\pi\|_{P,2} + \|d_\pi\|_{P,2} \|d_\mu\|_{P,2} + \|d_\pi\|_{P,2}^2), \\
\leq& \widetilde{C}_5 (\tau_n^2 + \tau_n a_n n^{-1/4} + a_n^2 n^{-1/2}) \leq \sqrt{a_n} n^{-1/2} = \delta_n n^{-1/2},
\end{aligned}
$$

where $\widetilde{C}_i$ are finite constants. Thus, Assumption 4(d) is verified.

Lastly, to verify Assumption 4(e), note that

$$
\begin{aligned}
& \mathrm{E}\left[m(G; \delta_0, \mu_0, \pi_1) m^\top (G; \delta_0, \mu_0, \pi_1) \right] \\
&= \mathrm{E}\left\{ \widetilde{Z}\widetilde{Z}^\top \pi_0(Z)\pi_1(Z) \frac{\pi_0(Z) \mathrm{var}(Y^{(1)} \mid Z) + \pi_1(Z) \mathrm{var}(Y^{(0)} \mid Z) e^{2\delta_0^\top \widetilde{Z}}}{\left\{ e^{\delta_0^\top \widetilde{Z}} \pi_1(Z) + \pi_0(Z) \right\}^2} \right\},
\end{aligned}
$$

which is non-degenerate, because Assumption 1 ensures that $\mathrm{var}(Y^{(r)} \mid Z = z) \geq \sigma_L > 0$.

Note that our choice of δ_n and τ_n satisfies $\log(n)/\sqrt{n} \leq a_n^{3/8} n^{-1/4} = \tau_n \leq \sqrt{a_n} = \delta_n$, and $\delta_n = \sqrt{a_n} \geq \log(n)^{-1/2} \geq n^{-1/2+1/q} \log(n)$ for any constant $q > 2$. Therefore, all assumptions of Theorem 3 are verified. Applying this theorem completes the proof. □

Supplementary Materials

The supplementary materials provide additional results for showing optimality of the weights for the contrast regression under Poisson distributed outcomes, details regarding constructing symmetric contrast regression estimating equations, and proofs of Theorem 2. They also provide extensions for estimating and validating relative treatment effects for survival outcomes, along with experimental results for these settings.

Acknowledgments

The authors thank Hongseok Namkoong and the anonymous reviewers for helpful comments.

Funding

SY is partially supported by the Stanford Graduate Fellowship and NHLBI award R01HL144555-01. Dr. Tian's research is partially supported by R01HL089778-05 and 1UL1TR003142 from National Institutes of Health, USA.

References

Athey, S., and Imbens, G. (2016), "Recursive Partitioning for Heterogeneous Causal Effects," *Proceedings of the National Academy of Sciences of the United States of America*, 113, 7353–7360. [1]

Athey, S., Tibshirani, J., and Wager, S. (2019), "Generalized Random Forests," *The Annals of Statistics*, 47, 1148–1178. [2]

Bang, H., and Robins, J. M. (2005), "Doubly Robust Estimation in Missing Data and Causal Inference Models," *Biometrics*, 61, 962–973. [6,7]

Basu, S., Sussman, J. B., Rigdon, J., Steimle, L., Denton, B. T., and Hayward, R. A. (2017), "Benefit and Harm of Intensive Blood Pressure Treatment: Derivation and Validation of Risk Models Using Data From the SPRINT and ACCORD Trials," *PLoS Medicine*, 14, e1002410. [1]

Breiman, L. (2001), "Random Forests," *Machine Learning*, 45, 5–32. [6]

Cai, T., Tian, L., Wong, P. H., and Wei, L. (2010), "Analysis of Randomized Comparative Clinical Trial Data for Personalized Treatment Selections," *Biostatistics*, 12, 270–282. [2]

Chen, S., Tian, L., Cai, T., and Yu, M. (2017), "A General Statistical Framework for Subgroup Identification and Comparative Treatment Scoring," *Biometrics*, 73, 1199–1209. [2]

Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., and Robins, J. (2018), "Double/Debiased Machine Learning for Treatment and Structural Parameters," *The Econometrics Journal*, 31, C1–C68. [2,4,5,13,15]

Dukes, O., and Vansteelandt, S. (2018), "A Note on G-Estimation of Causal Risk Ratios," *American Journal of Epidemiology*, 187, 1079–1084. [2]

Foster, J. C., Taylor, J. M., and Ruberg, S. J. (2011), "Subgroup Identification From Randomized Clinical Trial Data," *Statistics in Medicine*, 30, 2867–2880. [2,6]

Friedman, J., Hastie, T., and Tibshirani, R. (2000), "Additive Logistic Regression: A Statistical View of Boosting," *The Annals of Statistics*, 28, 337–407. [6]

Green, D. P., and Kern, H. L. (2012), "Modeling Heterogeneous Treatment Effects in Survey Experiments With Bayesian Additive Regression Trees," *Public Opinion Quarterly*, 76, 491–511. [2]

Imbens, G. W., and Rubin, D. B. (2015), *Causal Inference in Statistics, Social, and Biomedical Sciences*, New York: Cambridge University Press. [1,3]

Kang, J. D., and Schafer, J. L. (2007), "Demystifying Double Robustness: A Comparison of Alternative Strategies for Estimating a Population Mean from Incomplete Data," *Statistical Science*, 22, 523–539. [7]

Künzel, S. R., Sekhon, J. S., Bickel, P. J., and Yu, B. (2019), "Metalearners for Estimating Heterogeneous Treatment Effects Using Machine Learning," *Proceedings of the National Academy of Sciences of the United States of America*, 116, 4156–4165. [2]

Loh, W.-Y., He, X., and Man, M. (2015), "A Regression Tree Approach to Identifying Subgroups With Differential Treatment Effects," *Statistics in Medicine*, 34, 1818–1833. [6]

Lu, M., Sadiq, S., Feaster, D. J., and Ishwaran, H. (2018), "Estimating Individual Treatment Effect in Observational Data Using Random Forest Methods," *Journal of Computational and Graphical Statistics*, 27, 209–219. [2,8]

Nie, X., and Wager, S. (2019), "Quasi-Oracle Estimation of Heterogeneous Treatment Effects," arXiv no. 1712.04912. [2]

Powers, S., Qian, J., Jung, K., Schuler, A., Shah, N. H., Hastie, T., and Tibshirani, R. (2018), "Some Methods for Heterogeneous Treatment Effect Estimation in High Dimensions," *Statistics in Medicine*, 37, 1767–1787. [2]

Qi, Z., Liu, D., Fu, H., and Liu, Y. (2020), "Multi-Armed Angle-Based Direct Learning for Estimating Optimal Individualized Treatment Rules With Various Outcomes," *Journal of the American Statistical Association*, 15, 1–33. [13]

Robins, J. M., and Rotnitzky, A. (2001), "Comment on: Inference for Semiparametric Models: Some Questions and an Answer," *Statistica Sinica*, 11, 920–936. [2,4]

Tian, L., Alizadeh, A. A., Gentles, A. J., and Tibshirani, R. (2014), "A Simple Method for Estimating Interactions Between a Treatment and a Large Number of Covariates," *Journal of the American Statistical Association*, 109, 1517–1532. [1]

Uno, H., Claggett, B., Tian, L., Inoue, E., Gallo, P., Miyata, T., Schrag, D., Takeuchi, M., Uyama, Y., and Zhao, L. (2014), "Moving Beyond the Hazard Ratio in Quantifying the Between-Group Difference in Survival Analysis," *Journal of Clinical Oncology*, 32, 2380. [7]

Van der Laan, M. J., and Rose, S. (2011), *Targeted Learning: Causal Inference for Observational and Experimental Data*, New York: Springer-Verlag. [2,4]

Wager, S., and Athey, S. (2018), "Estimation and Inference of Heterogeneous Treatment Effects Using Random Forests," *Journal of the American Statistical Association*, 113, 1228–1242. [2]

Wendling, T., Jung, K., Callahan, A., Schuler, A., Shah, N., and Gallego, B. (2018), "Comparing Methods for Estimation of Heterogeneous Treatment Effects Using Observational Data From Health Care Databases," *Statistics in Medicine*, 37, 3309–3324. [2,8]

Xie, Y., Brand, J. E., and Jann, B. (2012), "Estimating Heterogeneous Treatment Effects With Observational Data," *Sociological Methodology*, 42, 314–347. [2]

Zhang, B., Tsiatis, A. A., Davidian, M., Zhang, M., and Laber, E. (2012), "Estimating Optimal Treatment Regimes From a Classification Perspective," *Stat*, 1, 103–114. [2]

Zhao, L., Tian, L., Cai, T., Claggett, B., and Wei, L.-J. (2013), "Effectively Selecting a Target Population for a Future Comparative Study," *Journal of the American Statistical Association*, 108, 527–539. [2,3,6,12]

Zhao, Y., Zeng, D., Laber, E. B., Song, R., Yuan, M., and Kosorok, M. R. (2014), "Doubly Robust Learning for Estimating Individualized Treatment With Censored Data," *Biometrika*, 102, 151–168. [2]

Zhao, Y., Zeng, D., Rush, A. J., and Kosorok, M. R. (2012), "Estimating Individualized Treatment Rules Using Outcome Weighted Learning," *Journal of the American Statistical Association*, 107, 1106–1118. [2]

Zhou, X., Mayer-Hamblett, N., Khan, U., and Kosorok, M. R. (2017), "Residual Weighted Learning for Estimating Individualized Treatment Rules," *Journal of the American Statistical Association*, 112, 169–187. [2]