Estimation and Validation of a Class of Conditional Average Treatment Effects Using Observational Data

Steve Yadlowsky, Fabio Pellegrini, Federica Lionetto, Stefan Braune, and Lu Tian [1]

December 2019

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 recommend interventions based on the contrast of the predicted outcomes. 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 patients, where the ratio of relapse rates is a natural choice of the treatment effect, we propose to estimate treatment-covariate interactions by coupling the standard regression approach with a doubly robust adjustment that mitigates finding spurious 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 the advantage of this approach on real data examining the treatment effect of dimethyl fumarate compared to teriflunomide in multiple sclerosis patients.

Key Words: Heterogeneous Treatment Effect; Conditional Average Treatment Effect, Doubly robust Estimation; Precision Medicine; Observational Study.

[1]Steve Yadlowsky is a Ph.D. student in the Department of Electrical Engineering, Stanford University, Stanford, CA 94305, USA (Email: syadlows@stanford.edu). Fabio Pellegrini is Distinguished Biostatistician, Senior Director, Global Head of RWE Innovations, ALIGN, Biogen International GmbH (Email: fabioi.pellegrini@biogen.com). Federica Lionetto is a Data Scientist at PwC Data & Analytics (Email: federica.lionetto@ch.pwc.com). Stefan Braune is Professor of Neurology and Head of real world evidence data management of NeuroTransData (Email: sbraune@neurotransdata.com). Lu Tian is a Professor, Department of Biomedical Data Science, Stanford University, Palo Alto, California 94305, U.S.A (Email: lutian@stanford.edu). This research was supported by grants from the National Institutes of Health. SY is partially supported by a Stanford Graduate Fellowship.

1. Introduction

Recently, interest in recommending tailored preventative interventions or treatments to patients in clinical practice has prompted investigating conditional average treatment effects (CATE) from data. Knowledge of CATE as the contrast between the expected outcome for different interventions conditional on different covariate levels would allow clinicians to identify whether a patient would benefit from a particular intervention based on their covariates. The primary statistical objective is to estimate these CATE by examining treatment-covariate interactions [Tian et al., 2014]. In this work, we study the estimation and validation of a special type of CATEs using observational data.

The estimation of the difference in expectations of potential outcomes as the CATE, i.e., $\mathrm{E}(Y^{(1)} - Y^{(0)} \mid Z = z)$, using observational data has been studied extensively in the literature [Green and Kern, 2012, Xie et al., 2012, Nie and Wager, 2019, Lu et al., 2018, Wager and Athey, 2018, Athey et al., 2019, Künzel et al., 2019]. Recently, Powers et al. [2018] and Wendling et al. [2018] compare a number of approaches for learning this function. The basic idea is either estimating the CATE based on separately estimated $\mathrm{E}(Y^{(r)} \mid Z = z), r = 0, 1$, or learning CATE directly based on 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

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

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 of the conditional expectation of the potential outcomes given the baseline covariates, rather than the difference thereof, motivated by the study of relapse in multiple sclerosis (MS) patients. We also provide a general framework for the estimation and validation of such a CATE score.

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 via the CATE, and then identify the high value subgroup of patients accordingly [Cai et al., 2010, Foster et al., 2011, Zhao et al., 2013]. The OWL approach avoids the more difficult task of estimating treatment effect for each patient, but also fails to yield information about the size of the treatment benefit for individual patients that may be important in cost-effectiveness analyses, which seek to balance the economic cost and risk of adverse events with treatment benefits. A good estimator of the CATE ensures a good average treatment effect (ATE) within the subgroup consisting of patients with the largest CATEs, thus assigning priority of receiving the treatment in the population. In the spirit of precision medicine, we may also directly recommend the treatment to individual patients whose estimated CATE outweighs the associated cost, which may also be patient-dependent. Therefore, in this paper, we focus on the more general question of directly estimating the CATE rather than a binary recommendation rule.

1.1. A motivating example

The NeuroTransData (NTD) multiple sclerosis (MS) registry includes about 25,000 patients with multiple sclerosis, 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 data contain few missing values, thanks to processes to manage the definition of minimum data sets, mandatory data entry fields, and positive missing data confirmation. Dynamic web-based data capturing, regular training of doctors and nurses, interactive chat forum for nurses and doctors, automated and manual feedback query system, daily automated analysis of data plausibility and correctness and annually on-site audit of procedures and source data by an external process quality certifier organization ensure the data is accuracy.

All data are pseudonymized and pooled to form the MS registry database. The codes are managed by the Institute for medical information processing, biometry and epidemiology (Institut für medizinische Informationsverarbeitung, Biometrie und Epidemiologie) at the

Ludwig Maximilian University in Munich, Germany, acting as an external trust center. This data acquisition and management protocol was approved by the ethical committee of the Bavarian Medical Board (Bayerische Landesärztekammer; June 14, 2012) and re-approved by the ethical committee of the Medical Board North-Rhine. (Ärztekammer Nordrhein, April 25, 2017). Compliance with European and German legislation is warranted including patient rights and informed consent requirements.

For this project, data were extracted from two groups of patients with relapsing remitting multiple sclerosis captured between Jan/1/2009 and July/1/2018, who received either Dimethyl fumarate (DMF, $n = 1741$) or Teriflunomide (TERI, $n = 1050$), respectively. The focus of the analysis is to estimate the CATE of TERI compared with DMF 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 (ATE) is measured by the ratio of the expected relapse rate under TERI versus that under DMF.

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 remainder of the introduction serves to detail the nuances of these questions, and the following sections discuss our approach to answer them.

1.2. The relapse rates ratio as the measure of CATE

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:

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

where $\widetilde{Z} = (1, Z^\top)^\top$ is a $d+1$ dimensional covariate vector, R is the binary indicator of the treatment received, and $(Y^{(r)}, F^{(r)})$ are the potential number of relapses and exposure time, when the patient receives the treatment $r \in \{0, 1\}$. Our goal is to model the effect of the treatment on the relapse rate. Under this regression model, the expected number of relapses per unit time for a group of patients with the same covariate $Z = z$ under treatment r, is

$$\mathrm{E}\left(\frac{Y^{(r)}}{F^{(r)}} \,\middle|\, Z = z\right) = \exp(\beta_r^\top \widetilde{z}),$$

where $\widetilde{z} = (1, z^\top)^\top$. 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 $r \in \{0, 1\}$. Noting that the ratio of the expected relapse rates

$$D(z) = \frac{\mathrm{E}(Y^{(1)}/F^{(1)} | Z = z)}{\mathrm{E}(Y^{(0)}/F^{(0)} | Z = z)} \tag{2}$$

is a natural measure of the CATE for the relapse rate that is insensitive to differences in exposure time between patients, a simple estimate of the CATE is

$$\widehat{D}(z) = \frac{\widehat{\mu}_1(z)}{\widehat{\mu}_0(z)} = \exp\{(\widehat{\beta}_1 - \widehat{\beta}_0)^\top \widetilde{z}\}, \tag{3}$$

under model (1.2).

We only observe $(Y^{(r)}, F^{(r)})$, when $R = r$, leading to a causal-missing data problem. We consider the following generalization of the unconfoundedness assumption [Imbens and Rubin, 2015] that identifies the relapse rate:

$$\{Y^{(1)}, F^{(1)}, Y^{(0)}, F^{(0)}\} \perp\!\!\!\perp R \mid Z. \tag{4}$$

This implies that $\mathrm{E}(Y \mid F, Z, R = r) = \mathrm{E}(Y^{(r)} \mid F^{(r)}, Z, R = r) = \mathrm{E}(Y^{(r)} \mid F^{(r)}, Z)$. Note that if the exposure times $F^{(r)} \geqslant \epsilon_F > 0$, all analyses only depend on $Y^{(r)}/F^{(r)}$. And so, in the rest of the paper, we assume that $F^{(r)} = 1$, without loss of generality. Consequently, $D(z)$ reduces to $\mathrm{E}(Y^{(1)} \mid Z = z)/\mathrm{E}(Y^{(0)} \mid Z = z)$.

1.3. Confounding effect on estimating CATE

If the regression models for Y given Z and the treatment assignment R are correctly specified, then the current regression method doesn't need to be altered in an observational study since $\widehat{\beta}_r$ is expected to converge to the true regression coefficient as the sample size increases. In practice, these statistical models may be mis-specified 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 $Y \mid Z = z$ follows a Poisson distribution with a rate of z^2 in both arms, i.e., there is no treatment effect for any $Z = z$. To mimic the confounding effect in an observational study, we assume that $Z \mid R = 1 \sim N(0.5, 1)$, and $Z \mid R = 0 \sim N(-0.5, 1)$. If we fit 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 positive in arm $R = 1$ and negative in arm $R = 0$. Specifically $\beta_1 = (-0.5, 0.8)^\top$ and $\beta_0 = (-0.5, -0.8)^\top$. This is not a surprise, 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 $\exp\{(\widehat{\beta}_1 - \widehat{\beta}_0)^\top \widetilde{z}\} \approx \exp(1.6z)$ 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 misspecified regression models may create spurious treatment covariate interactions in observational studies.

In general, assume that there is no treatment effect heterogeneity, so that

$$D(z) = \exp(d_0), \tag{5}$$

for all z, but (1.2) are misspecified, $\mathrm{E}(Y^{(r)} \mid Z = z) \neq \beta_r^\top \tilde{z}$. Under the Poisson regression

model, the estimator $\widehat{\beta}_r$ is a solution to the score equation:

$$\widetilde{S}_r(\beta) = n_r^{-1} \sum_{R_i=r} \widetilde{Z}_i \left\{ Y_i - \exp(\beta^\top \widetilde{Z}_i) \right\} = 0,$$

where $\{(Y_i, Z_i, R_i), i = 1, \cdots, n\}$ are observed training data, and n_r is the number of patients in arm $r = 0, 1$. If $\sum_{R_i=r} \widetilde{Z}_i \widetilde{Z}_i^\top$ is positive definite, then this solution is unique and the estimator is well-defined. Although the regression model is mis-specified, under mild regularity conditions, the resulting estimator $\widehat{\beta}_r$ will still converge to a limit β_r^* that is the root of the equation

$$\widetilde{s}_r(\beta) = \mathrm{E}\left[\widetilde{Z}_i \left\{ \mu_r(Z_i) - \exp(\beta^\top \widetilde{Z}_i) \right\} \,\middle|\, R_i = r \right] = 0,$$

where $\mu_r(z) = \mathrm{E}(Y^{(r)} \mid Z = z)$. If the data are from a randomized clinical trial, then $\widetilde{s}_r(\beta)$ becomes

$$s_r(\beta) = \mathrm{E}\left[\widetilde{Z}_i \left\{ \mu_r(Z_i) - \exp(\beta^\top \widetilde{Z}_i) \right\} \right].$$

Assumption (5) implies that $\beta_1^* - \beta_0^* = (d_0, 0, \cdots, 0)^\top$, and thus, the estimated CATE score converges to

$$D(z) = \exp\{(\beta_1^* - \beta_0^*)^\top \widetilde{z}\} = \exp(d_0).$$

This correctly suggests that there is no treatment effect heterogeneity even under mis-specified regression models. However, in an observational study, where the observed treatment assignment depends on the covariates, as the toy example above shows, in general,

$$\widetilde{s}_r(\beta) \neq s_r(\beta) \text{ and } \beta_1^* - \beta_0^* \neq (d_0, 0, \cdots, 0)^\top.$$

Consequently, $\widehat{D}(z)$ may indicate false dependence of the treatment effect on the covariates.

One lesson is that the construction of the CATE score, $\widehat{D}(z)$, still should adjust for the covariates imbalance, if the relevant regression models for outcomes could be misspecified.

1.4. Validation of CATE estimation models

For data from a randomized clinical trial, Zhao et al. [2013] proposed the following approach to estimate and validate heterogeneous treatment effects. The method consists of two main steps:

1. In the training set

   - assume separate regression models to link the potential outcomes $Y^{(r)}$ in control or treatment arm $R = r$, for $r = 0, 1$ (respectively) with the baseline covariates Z to estimate the expected outcome $\mu_r(z)$. Denote the prediction rules by $\widehat{\mu}_r(z)$.
   - estimate the CATE by the difference between $\widehat{\mu}_1(z)$ and $\widehat{\mu}_0(z)$, i.e., $\widehat{D}(z) = \widehat{\mu}_1(z) - \widehat{\mu}_0(z)$.

2. In the validation set,

   - estimate $AD(c) = \mathrm{E}\left\{Y^{(1)} - Y^{(0)} \mid \widehat{D}(Z) \geqslant c\right\}$, the ATE for a subgroup of patients

     $$\{z \mid \widehat{D}(z) \geqslant c\},$$

     and denote the resulting estimator by $\widehat{AD}(c)$.
   - consider the validation curve $\left(1 - q, \widehat{AD}\{\widehat{H}^{-1}(q)\}\right)$, where $q \in [0, 1)$ and $\widehat{H}(\cdot)$ is the empirical cumulative distribution function of $\widehat{D}(Z)$. Here, $\widehat{H}^{-1}(q)$ and $1 - q$ represent the cutoff value defining the subgroup and the proportion of patients in the subgroup, respectively. The validation curve graphically demonstrates the relationship between the size of the selected subgroup and the estimated ATE in that subgroup. It can be used to assist the evaluation of $\widehat{D}(z)$ and individualized treatment recommendation.
   - the slope of $\widehat{AD}\{\widehat{H}^{-1}(q)\}$ is a measure reflects the quality of the scoring system $\widehat{D}(z)$ in ranking the patients according to their estimated CATE, $\widehat{D}(z)$.

These two steps can be repeated multiple times in cross-validations for evaluating and comparing the performance of different scoring systems approximating the true CATE.

The small sample size of trials is one of the biggest obstacles in such analyses. Most randomized clinical trials are designed to study the ATE, rather than the CATE. Furthermore, to verify the 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 analysis, 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 statistical analysis [Imbens and Rubin, 2015].

Our goal is to 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 potential outcomes. If we select the subgroup of patients with the highest true CATE $\{z : D(z) \geqslant c\}$ and use the ratio to measure the treatment effect,

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

then we will show that $AD_{\text{true}}(c)$ is monotone in c. Consequently, if, instead, we group patients by an estimated CATE score $\widehat{D}(z)$ and

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

then 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 $\{z \mid \widehat{D}(z) \geqslant c\}$, one needs to adjust again for the potential imbalances between two treatment arms using propensity score, regression or doubly robust estimators [Bang and Robins, 2005, Kang and Schafer, 2007]. Finally, we will address these issues in details and discuss extensions to survival outcomes.

2. Method

In this section, we propose to extend the sequence of training and validation steps by Zhao et al. [2013] to observational studies.

2.1. Training

2.1.1. Two Regressions Approach

We first propose an approach to avoid the aforementioned spurious treatment effect heterogeneity with minimum model assumptions in estimating the CATE. To this end, we plan to quantify the association between the outcome and baseline covariates in two arms via a simple regression model, while (1) viewing the regression model as a working model approximating the association of interest; (2) fitting the regression model as if the potential outcomes and covariates are observed in the entire cohort. Specifically, we propose to construct the CATE score based on the equation $s_r(\beta) = 0$: the CATE score to be estimated is simply

$$\exp\left\{(\beta_1 - \beta_0)^\top \widetilde{z}\right\},$$

where β_r be the (unique) solution to the estimating equation $s_r(\beta) = 0$. This proposal is appealing since $D(z) = \exp(d_0)$ implies that $(\beta_1 - \beta_0)^\top \widetilde{z} = d_0$ is also a constant, i.e., there is no spurious treatment effect heterogeneity despite the fact that the regression model $\mu_r(z) = \exp(\beta_r^\top \widetilde{z})$ may be mis-specified.

It is tempting to estimate the proposed CATE score via solving empirical counterparts of estimating equations $s_r(\beta) = 0$. The challenge is that $Y^{(r)}$ is only observed when $R = r$, so we do not observe complete data to allow us to construct these equation empirically. There are two approaches for estimating β_r. If we can construct a consistent nonparametric estimator of $\mu_r(z)$ denoted by $\widehat{\mu}_r(z)$, then we may estimate β_r by solving a simple estimating equation

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

Alternatively, under the unconfoundedness assumption (4), we can adjust the distribution of covariates in each arm to approximate these equations by reweighitng. To be specific, instead of solving the estimating equation above, consider the inverse probability weighted (IPW) estimating equation

$$n^{-1}\sum_{i=1}^{n} \widehat{W}_i(r)\widetilde{Z}_i \left\{Y_i - \exp\left(\beta^\top \widetilde{Z}_i\right)\right\} = 0, r = 0, 1,$$

where

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

and $\widehat{\pi}_r(z)$ is an estimator for the propensity score $\pi_r(z) = P(R = r|Z = z)$. If unconfoundedness (4) holds, $\pi_r(z)$ is bounded away from 0 and 1, and $\widehat{\pi}_r(z)$ is uniformly consistent for estimating $\pi_r(z)$, then this IPW estimating function converges to $s_r(\beta)$ and its root converges to β_r in probability. The aforementioned two approaches have their own limitations and we propose to consider a more robust estimator by solving the following 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 non-parametric (or otherwise more flexible parametric or semi-parametric) 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. Solving weighted estimating equations

$$n^{-1}\sum_{i=1}^{n}\widehat{W}_i(r)\widetilde{Z}_i\left(Y_i - \exp\left[\alpha_r \times \log\{\widehat{\mu}_r(Z_i)\} + \gamma_r^\top \widetilde{Z}_i\right]\right) = 0, r = 0, 1; \tag{6}$$

   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 $\widetilde{\mu}_r(\cdot)$ 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 (4) and mild regularity conditions [Bang and Robins, 2005]. The key observation is that if the propensity score is consistently estimated, the equation (6) ensures that

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

If we are suspect that the Poisson regression (1.2) is mis-specified, 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, one may also employ machine learning methods such as random forest or boosting to generate $\widehat{\mu}_r(z)$ [Friedman et al., 2000, Breiman, 2001].

Furthermore, we propose to employ the cross-fitting procedure to improve the finite sample performance of the estimator by removing the correlation between $\widetilde{\mu}_r(Z_i)$ and Z_i induced by potential overfitting in constructing $\widetilde{\mu}_r(z)$. Specifically, we propose to estimate β_r by $\widehat{\beta}_r$, which is the root of the estimating equation:

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

where data are divided into K non-overlapping parts of approximately equal sizes indexed by $\mathcal{I}_k, k = 1, \cdots, K$, and $\widetilde{\mu}_r^{(-k)}(z)$ is constructed using observations not in $\mathcal{I}_k$. The estimated CATE score is thus

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

Remark 1. *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 [Loh et al., 2015, Foster et al., 2011].*

2.1.2. Contrast Regression Approach

While motivated by the Poisson regression model, the CATE model

$$D(z) = \exp(\delta_0^\top z)$$

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

$$\mathrm{E}(Y^{(r)}|Z=z) = \exp\left(r\delta_0^\top z\right)\mu_0(z), \tag{7}$$

where $\mu_0(z)$ is a unknown, measurable and non-negative function in some non-parametric function class. We assume $\pi_r(z)$ is unknown, but is also in a non-parametric 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. Under these assumptions, 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

$$\mathrm{E}\left[w(Z,\delta)\tilde{Z}\left\{Y^{(1)} - \exp(\delta^\top\tilde{Z})Y^{(0)}\right\}\right] = 0, \tag{8}$$

because applying the tower property of conditional expectations to (8) 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 $\widetilde{Z}$ doesn't 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 (8) under condition (4). For instance, because $\mathrm{E}[Y^{(r)} \mid Z = z] = \mu_r(z)$, applying the tower property of conditional expectations to (8) gives the equivalent estimating equation $\mathrm{E}[w(Z,\delta)\widetilde{Z}\{\mu_1(Z) - \exp(\delta^\top\widetilde{Z})\mu_0(Z)\}] = 0$. Equivalent inverse probability weighted estimators also exist that depend on $\pi_r(z)$. 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 semi-parametric 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 semi-parametric risk ratio model presented in Van der Laan and Rose [2011]. Specifically, for any candidate nuisance parameters $\mu : \mathbf{R}^d \to \mathbf{R}$ and

$\pi : \mathbf{R}^d \to [0,1]$ and parameters $\delta \in \Omega \subset \mathbf{R}^{d+1}$, we consider the estimating function

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

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 the Appendix, we show that this weight function is optimal under the assumption that $Y^{(r)} \mid Z = z$ follows a Poisson distribution with a rate of $\mu_r(z)$ in minimizing the variance of the resulting estimator under the correct propensity score model $\pi_1(z)$.

On the other hand, by re-writing 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 equations

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

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

In the Appendix, 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 non-overlapping 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. $\widehat{\pi}_0^{(-k)}(z) = 1 - \widehat{\pi}_1^{(-k)}(z)$. Then, estimate δ_0 by searching for the root of the estimating equation

$$
\begin{aligned}
S_n^{(cf)}(\delta) =& n^{-1} \sum_{k=1}^{K} \sum_{i \in \mathcal{I}_k} m(G_i; \delta, \widehat{\mu}_0^{(-k)}, \widehat{\pi}_1^{(-k)}) \\
=& n^{-1} \sum_{k=1}^{K} \sum_{i \in \mathcal{I}_k} \widetilde{Z}_i \left(R_i \frac{\left[Y_i - \exp(\delta^\top \widetilde{Z}_i) \widehat{\mu}_0^{(-k)}(Z_i) \right] \widehat{\pi}_0^{(-k)}(Z_i)}{e^{\delta^\top \widetilde{Z}_i} \widehat{\pi}_1^{(-k)}(Z_i) + \widehat{\pi}_0^{(-k)}(Z_i)} \right. \\
& \left. -(1 - R_i) \frac{\left[Y_i - \widehat{\mu}_0^{(-k)}(Z_i) \right] \exp(\delta^\top \widetilde{Z}_i) \widehat{\pi}_1^{(-k)}(Z_i)}{e^{\delta^\top \widetilde{Z}_i} \widehat{\pi}_1^{(-k)}(Z_i) + \widehat{\pi}_0^{(-k)}(Z_i)} \right) = 0.
\end{aligned}
$$

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 $E[ZZ^\top]$ 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 $\epsilon_\pi \leqslant \pi_1(z) \leqslant 1 - \epsilon_\pi$, (d) There exists $\sigma_L > 0$ and σ_U such that $var(Y^{(r)} \mid z) \in [\sigma_L, \sigma_U], z \in \mathcal{Z}$, (e) for some $q > 2$, $E[|Y^{(r)}|^q \mid Z = z] \leqslant C < \infty$; (f) δ_0 is a 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 $\epsilon_\pi \leqslant \pi(z) \leqslant 1 - \epsilon_\pi$, and $\epsilon_\mu \leqslant \mu(z) \leqslant 1/\epsilon_\mu$ 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) $\epsilon_\pi \leqslant$*

$\widehat{\pi}(z) \leqslant 1-\epsilon_\pi$; *(b)* $\epsilon_\mu \leqslant \widehat{\mu}_0(z) \leqslant 1/\epsilon_\mu$; *(c)* $\|\widehat{\mu}_0(z)-\mu_0(z)\|_{P,2} + \|\widehat{\pi}_1(z)-\pi_0(z)\|_{P,2} = o_P(n^{-1/4})$.

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

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

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

and

$$\begin{aligned}\widehat{B}(\delta) =& n^{-1}\sum_{k=1}^{K}\sum_{i\in\mathcal{I}_k}\widetilde{Z}_i\widetilde{Z}_i^\top\left(R_i\frac{\left[Y_i - e^{\delta^\top\widetilde{Z}_i}\widehat{\mu}_0^{(-k)}(Z_i)\right]}{e^{\delta^\top\widetilde{Z}_i}\widehat{\pi}_1^{(-k)}(Z_i) + \widehat{\pi}_0^{(-k)}(Z_i)}\widehat{\pi}_0^{(-k)}(Z_i)\right.\\ &\left.-(1-R_i)\frac{\left[Y_i - \widehat{\mu}_0^{(-k)}(Z_i)\right]e^{\delta^\top\widetilde{Z}_i}}{e^{\delta^\top\widetilde{Z}_i}\widehat{\pi}_1^{(-k)}(Z_i) + \widehat{\pi}_0^{(-k)}(Z_i)}\widehat{\pi}_1^{(-k)}(Z_i)\right)^2.\end{aligned}$$

Then, applying Theorem 3.3 from Chernozhukov et al. [2018] under the Assumptions 1 and 2 gives

Theorem 1. *Let* $\widehat{\delta}$ *solve* $S_n^{(cf)}(\delta) = 0$. *Under Assumptions 1 and 2,* $\sqrt{n}(\hat{\delta} - \delta_0)$ *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 2 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-fit semi-parametric estimators. The proposed estimating equation can be solved via Newton-Raphson method. Although we can't 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 2. *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 the random splitting data into K parts. Chernozhukov et al. [2018] shows that this estimator is asymptotically equivalent to $\widehat{\delta}$ analyzed above.*

Remark 3. *By comparing the conditional means to the baseline of $\mu_1(z)$, the semi-parametric regression model (7) is equivalent to*

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

and a similar analysis to the above gives a set of symmetric estimating equations in terms of nuisance parameters μ_1 and π_1. Averaging the resulting estimating equation with its counterpart based on μ_0 and π_1 yields two new equations:

$$\begin{aligned}\bar{S}_n(\delta) =& n^{-1}\sum_{i=1}^n \widetilde{Z}_i R_i \left[Y_i - \frac{e^{\delta^\top \widetilde{Z}_i}\widehat{\mu}_0(Z_i) + \widehat{\mu}_1(Z_i)}{2}\right] \frac{\widehat{\pi}_0(Z_i)}{e^{\delta^\top \widetilde{Z}_i}\widehat{\pi}_1(Z_i) + \widehat{\pi}_0(Z_i)} \\ & - \sum_{i=1}^n \widetilde{Z}_i(1-R_i)\left[Y_i e^{\delta^\top \widetilde{Z}_i} - \frac{e^{\delta^\top \widetilde{Z}_i}\widehat{\mu}_0(Z_i) + \widehat{\mu}_1(Z_i)}{2}\right] \frac{\widehat{\pi}_1(Z_i)}{e^{\delta^\top \widetilde{Z}_i}\widehat{\pi}_1(Z_i) + \widehat{\pi}_0(Z_i)} = 0,\end{aligned}$$

and

$$\begin{aligned}\bar{S}_n^{(cf)}(\delta) =& n^{-1}\sum_{k=1}^K \sum_{i\in\mathcal{I}_k} \Bigg(\widetilde{Z}_i R_i \frac{\left[Y_i - \{e^{\delta^\top \widetilde{Z}_i}\widehat{\mu}_0^{(-k)}(Z_i) + \widehat{\mu}_1^{(-k)}(Z_i)\}/2\right]}{e^{\delta^\top \widetilde{Z}_i}\widehat{\pi}_1^{(-k)}(Z_i) + \widehat{\pi}_0^{(-k)}(Z_i)} \widehat{\pi}_0^{(-k)}(Z_i) \\ & - \widetilde{Z}_i(1-R_i)\frac{\left[Y_i e^{\delta^\top \widetilde{Z}_i} - \{e^{\delta^\top \widetilde{Z}_i}\widehat{\mu}_0^{(-k)}(Z_i) + \widehat{\mu}_1^{(-k)}(Z_i)\}/2\right]}{e^{\delta^\top \widetilde{Z}_i}\widehat{\pi}_1^{(-k)}(Z_i) + \widehat{\pi}_0^{(-k)}(Z_i)} \widehat{\pi}_1^{(-k)}(Z_i) \Bigg) = 0.\end{aligned}$$

In practice, we recommend to solve $\bar{S}_n^{(cf)}(\delta) = 0$ for estimating δ_0 due to the symmetry between μ_1 and μ_0.

2.2. Validation

In the validation step, we rank the patients in the validation set according to the estimated CATE score $\widehat{D}(z)$ and estimate the ATE in the subgroup of patients with the most promising CATE scores, $\{z \mid \widehat{D}(z) \geqslant c\}$. Zhao et al. [2013] considered $D(z) = \mathrm{E}(Y^{(1)} - Y^{(0)} \mid Z = z)$ and $AD(c) = \mathrm{E}(Y^{(1)} - Y^{(0)} | D(Z) \geqslant c)$. In such a case, $AD(c)$, the ATE in $\{z \mid D(z) \geqslant c\}$, is a monotone increasing function of c. 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*

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

and

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

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

See Appendix 3 for detailed proof. Therefore, 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, $AD(c) \geqslant c$, 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. The general problem is that the ATE in a subgroup of patients with the largest CATE may not be large if the treatment effect is measured by a contrast that is more complex than the ratio or difference. In Appendix 3, we have provided a simple example where the marginal OR of a subgroup of patients with highest conditional OR is not the highest.*

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 random-

ized. 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), $\left\{(Y_i^V, R_i^V, Z_i^V), i = 1, \cdots, m\right\}$, 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 $\{z^V \mid \hat{D}(z^V) \geqslant c\}$; and then estimate $\mathrm{E}(Y^{(r)} \mid \widehat{D}(Z^V) \geqslant c)$ by

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

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 $\{z^V \mid \widehat{D}(z^V) \geqslant c\}$ 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. One advantage of this approach is that it not only provides an estimator for the ATE measured by the rate ratio but also the rates themselves, which facilitates our interpretation of the effect size.

2.3. Time-to-Event Outcomes

In precision medicine, one important type of clinically meaningful outcomes to model is time-to-event outcome. For example, among MS patients in the NTD registry, a useful endpoint to measure effectiveness is the time to first relapse. In this case, the observed training data consist of $\{(X_i, \Delta_i, R_i, Z_i), i = 1, \cdots, n\}$, where $X_i = T_i \wedge C_i$, $\Delta_i = I(T_i < C_i)$, T_i is the event time of interest and C_i is the censoring time. As previously, we will use $T_i^{(r)}$ and $C_i^{(r)}$ to refer to the potential event time and censoring time under treatment r. For time-to-event outcomes, the heterogeneous treatment effect often cannot be measured by a contrast of the conditional expectations $\mathrm{E}(T^{(r)}|Z = z)$, because this expectation may not be identifiable due to right censoring.

It is a popular practice to employ Cox proportional hazards model to analyze time-to-event outcomes. Let $\lambda_r(t \mid z)$ be the hazard function of $T_i^{(r)} \mid Z_i = z$, $r = 0, 1$. The Cox

proportional hazards model assumes that

$$\lambda_r(t \mid z) = \lambda_{r0}(t) \exp(\beta_r^\top z),$$

for $\lambda_{r0}(\cdot)$, an unspecified baseline hazard function in group r. In this case, $\exp\{(\beta_1 - \beta_0)^\top z\}$ is a natural measure of the CATE. If the proportional hazards assumption is violated, we can still define β_r as the root of the limit of the score equation of the Cox model. However, in this case, $\exp\{(\beta_1 - \beta_0)^\top z\}$ may no longer be a sensible surrogate for the CATE even if two arms have the same covariate distribution, as it may introduce spurious heterogeneity not present in real data. We provide an example in Appendix 3. The parameter closest to the mean survival time and still estimable is the restricted mean survival time, i.e., $E(T \wedge \tau)$. Then, the ATE (and CATE) are an appropriate contrast of restricted mean survival time.

In this section, we propose to use the ratio of restricted mean time lost (RMTL) due to the first relapse to measure the treatment effect, i.e.,

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

where $\tau > 0$ is a chosen constant [Uno et al., 2014], since it is expected to be concordant with the relapse rate ratio. When the relapse rate is low, this ratio is similar to the hazard ratio.

To approximate this CATE, we consider the score $\exp\{(\beta_1 - \beta_0)^\top \widetilde{z}\}$, where β_r is the root of the the equation

$$\mathrm{E}\left[\widetilde{Z}\left\{\mathrm{E}(\tau - T^{(r)} \wedge \tau \mid Z) - \exp(\beta_r^\top \widetilde{Z})\right\}\right] = 0.$$

While this regression model is not always natural for arbitrary choices of τ, it can still be a useful working model. For example, when $T^{(r)} \mid Z = z$ follows an exponential distribution with the failure rate of $\beta_r^\top \widetilde{z}$ and τ is large relative to $\beta_r^\top \widetilde{Z}$, this model is approximately correctly specified. Luckily, like the Poisson regression model discussed above, this approach is robust to mis-specification of the model in the sense that it doesn't introduce spurious treatment effect heterogeneity. Therefore, it remains a useful working model for estimating heterogeneous treatment effects measured by RMTL ratio.

Again, we propose two approaches for approximating $D(z)$. In the first approach, we

estimate β_r by solving the estimating equation based on cross-fitting:

$$n^{-1}\sum_{k=1}^{K}\sum_{i\in\mathcal{I}_k}\widetilde{Z}_i\left[\widetilde{\mu}_r^{(-k)}(Z_i,\tau)-\exp(\beta_r^\top\widetilde{Z}_i)\right]=0,$$

where $\widetilde{\mu}_r^{(-k)}(z;\tau)$ is an estimator of $E(\tau - T^{(r)}\wedge\tau \mid Z=z)$ constructed via the following step

1. Construct an initial estimator of $\mathrm{E}(\tau - T_i^{(r)}\wedge\tau|Z_i=z)$ based on observations not in $\mathcal{I}_k$. denote the estimator by $\widehat{\mu}_r^{(-k)}(z;\tau)$.

2. Solving weighted estimating equations

$$\sum_{i\notin\mathcal{I}_k}\widehat{L}_i(r)\widehat{W}_i^{(-k)}(r)\widetilde{Z}_i\left[(\tau-T_i\wedge\tau)-\exp\left\{\alpha_r\times\log(\widehat{\mu}_r^{(-k)}(Z_i,\tau))+\gamma_r^\top\widetilde{Z}_i\right\}\right]=0,$$

   and denote the roots by $(\widehat{\alpha}_r^{(-k)},\widehat{\gamma}_r^{(-k)\top})^\top, r=0,1.$

3. Let $\widetilde{\mu}_r^{(-k)}(z;\tau)=\exp\left\{\widehat{\alpha}_r^{(-k)}\times\log(\widehat{\mu}_r^{(-k)}(Z_i,\tau))+\widehat{\gamma}_r^{(-k)\top}\widetilde{Z}_i\right\}.$

In the second step of constructing $\widetilde{\mu}_r^{(-k)}(\cdot,\tau)$,

$$\widehat{L}_i(r)=\frac{I(T_i\wedge\tau<C_i)}{\widehat{K}_{Cr}(T_i\wedge\tau\mid Z_i)}$$

and

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

where $I(T_i\wedge\tau<C_i)=\Delta_i+(1-\Delta_i)I(X_i\geqslant\tau)$, and $\widehat{K}_{Cr}(\cdot\mid z)$ is an estimator of the survival function of the censoring time $C_i^{(r)}\mid Z_i=z$ under a regression model, such as the Cox model. The key assumption here is that the inverse probability weighting via $\widehat{L}_i(r)$ can correct the effect of right censoring, as long as $\widehat{K}_{Cr}(t\mid z)$ is a consistent estimator of the survival function of the censoring time $C^{(r)}$ given $Z=z$ and $P(C^{(r)}>\tau|Z=z)\geqslant\epsilon_\tau>0$. Let the resulting estimator of β_r be denoted by $\widehat{\beta}_r$ and $D(z)$ can be approximated by $\exp\{\widehat{\beta}_1-\widehat{\beta}_0)^\top z\}$.

Remark 5. *The initial estimator $\widehat{\mu}_r^{(-k)}(z;\tau)$ can be based on any sensible regression models characterizing the relationship between $T_i^{(r)}$ and Z_i. For example, one may employ the*

standard Cox model $\lambda_r(t \mid z) = \lambda_{r0}(t)\exp\{\eta_r^\top B(z)\}$. *Under this model,*

$$\tau - \int_0^\tau \exp\left[-\widehat{\Lambda}_{r0}(t)\exp\{\widehat{\eta}_r^\top B(z)\}\right] dt$$

can be used to approximate $E(\tau - T_i^{(r)} \wedge \tau | Z_i = z)$, *where* $\widehat{\Lambda}_{r0}(t)$ *is the Breslow estimator for the cumulative baseline hazard function and* $\widehat{\eta}_r$ *is the maximum partial likelihood estimator for* η_r *in arm* r *[Cheng et al., 1998]. Alternatively, we may assume a regression model tailored for restricted mean survival time [Tian et al., 2013]:*

$$E(T^{(r)} \wedge \tau | Z = z) = \tau \frac{\exp\{\eta_r^\top B(z)\}}{1 + \exp\{\eta_r^\top B(z)\}}$$

and estimate η_r *by solving the simple estimating equation*

$$\sum_{i=1}^n I(R_i = r)\widehat{L}_i(r)B(Z_i)\left\{T_i \wedge \tau - \tau \frac{\exp\{\eta_r^\top B(Z_i)\}}{1 + \exp\{\eta_r^\top B(Z_i)\}}\right\} = 0.$$

Under this model, $E(\tau - T_i^{(r)} \wedge \tau | Z_i = z)$ *can be approximated by*

$$\frac{\tau}{1 + \exp\{\widehat{\eta}_r^\top B(z)\}},$$

where $\widehat{\eta}_r$ *is the root to the estimating equation above.*

In the second approach, we may directly estimate δ_0 under the assumption that $D(z) = \exp(\delta_0^\top z)$. Specifically, we estimate δ_0 by solving the estimating equation by cross-fitting

$$n^{-1}\sum_{k=1}^K \sum_{i \in \mathcal{I}_k} \widehat{L}_i(R_i)\left(\widetilde{Z}_i R_i \frac{\left[(\tau - T_i \wedge \tau) - e^{\delta^\top \widetilde{Z}_i}\widehat{\mu}_0^{(-k)}(Z_i;\tau)\right]}{e^{\delta^\top \widetilde{Z}_i}\widehat{\pi}_1^{(-k)}(Z_i) + \widehat{\pi}_0^{(-k)}(Z_i)}\widehat{\pi}_0^{(-k)}(Z_i)\right.$$
$$\left. - \widetilde{Z}_i(1 - R_i)\frac{\left[(\tau - T_i \wedge \tau) - \widehat{\mu}_0^{(-k)}(Z_i;\tau)\right]e^{\delta^\top \widetilde{Z}_i}}{e^{\delta^\top \widetilde{Z}_i}\widehat{\pi}_1^{(-k)}(Z_i) + \widehat{\pi}_0^{(-k)}(Z_i)}\widehat{\pi}_1^{(-k)}(Z_i)\right) = 0,$$

or the counterpart of $\bar{S}_n^{(cf)}(\delta) = 0$. Let the root of the equation be $\widehat{\delta}$. Similar to the count outcome, $\widehat{\delta}$ is a consistent estimator of δ_0, if either the propensity score estimator $\widehat{\pi}_1(z)$ or the outcome predictions $\widehat{\mu}_0^{(-k)}(z;\tau)$ is consistent. Lastly, $D(z)$ can be approximated by $\exp(\widehat{\delta}^\top z)$. Note that the effect of right censoring is corrected by the weight $\widehat{L}_i(R_i)$.

With an estimator $\widehat{D}(z)$ of the CATE, we need to estimate

$$AD(c) = \frac{\mathrm{E}\{\tau - T^{(1)} \wedge \tau | \widehat{D}(Z) \geqslant c\}}{\mathrm{E}\{\tau - T^{(0)} \wedge \tau | \widehat{D}(Z) \geqslant c\}}$$

based on validation data. To this end, We may estimate the $AD(c)$ by constructing doubly robust estimators for $\mathrm{E}(\tau - T^{(r)} \wedge \tau | \widehat{D}(Z) \geqslant c), r = 0, 1$. To this end, we may construct an initial estimator for $\mu_r(z, \tau) = \mathrm{E}(\tau - T^{(r)} \wedge \tau | Z = z)$ using observations in the subgroup $\{z^V \mid \widehat{D}(z^V) \geqslant c\}$. Methods described in Remark 5 can be used to achieve this objective. Denote the resulting estimator by $\widehat{\mu}_{cr}(z; \tau)$. We may estimate $\mathrm{E}(\tau - T^{(r)} \wedge \tau | \widehat{D}(Z) \geqslant c)$ by

$$m_c^{-1} \sum_{\widehat{D}(Z_i^V) \geqslant c} \left[\widehat{\mu}_{cr}(Z_i^V, \tau) + \widehat{L}_{ci}^V(r) \widehat{W}_{ci}^V(r) \left\{ (\tau - T_i^V \wedge \tau) - \widehat{\mu}_{cr}(Z_i^V, \tau) \right\} \right].$$

Then, we propose to estimate $AD(c)$ by the plugin estimator

$$\frac{\widehat{\mu}_1(c; \tau)}{\widehat{\mu}_0(c; \tau)}.$$

3. Numeric Simulation

In this section, we conduct a numerical study to investigate the finite sample performance of the proposed method. In the first set of simulations, we simulate the outcomes $F^{(r)} \mid Z = z$ and $Y^{(r)} \mid F^{(r)} = f, Z = z$ from a uniform distribution over $[0, 0.75]$ and a Poisson distribution $\mathrm{Pois}\{\mu_r(z) f\}$, respectively. With slightly abuse of notations, Z_i (or z_i) stands for the i-th component of the covariate vector Z (or z) rather than the covariate vector of the i-th patient. The 10-dimensional covariate Z 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 distribution of Z_i is standard Gaussian. 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 $R|Z = z$ is generated from a Bernoulli distribution with a probability of $\pi_1(z) = \left\{1 + \exp(z_1 + 0.5 z_2 - 0.5 z_6)\right\}^{-1}$. There are six different settings:

1. The true CATE: $D(z) = \exp(0.75 + 0.125z_1 + 0.05|z_2 + 0.5| - 0.25z_6)$ with

$$\begin{aligned}\mu_1(z) &= \exp\big(\ \ 0.375 + g_1(z_1, z_2, z_6) + 0.30|z_2 + 0.5|\big)\\ \mu_0(z) &= \exp\big(-0.375 + g_2(z_1, z_2, z_6) + 0.25|z_2 + 0.5|\big)\,;\end{aligned}$$

2. The optimal CATE: $D(z) = \exp(0.775 + 0.125z_1 + 0.05z_2 - 0.25z_6)$ with

$$\begin{aligned}\mu_1(z) &= \exp\big(\ \ 0.525 + g_1(z_1, z_2, z_6) + 0.30z_2\big)\\ \mu_0(z) &= \exp\big(-0.250 + g_2(z_1, z_2, z_6) + 0.25z_2\big)\,;\end{aligned}$$

3. The optimal CATE is $\exp(0.025 + 0.125z_1 + 0.05z_2 - 0.25z_6)$ with

$$\begin{aligned}\mu_1(z) &= \exp\left(0.525 + 0.125z_1 + 0.30z_2 + 0.25z_6\right)\\ \mu_0(z) &= \exp\left(0.5 + 0.25z_2 + 0.50z_6\right);\end{aligned}$$

4. The same as setting 1, but the analysis only uses first five covariates, i.e., $Z_1, \cdots, Z_5$;

5. The true CATE: $D(z) = \exp(0.125z_1 + 0.05|z_2 + 0.5| - 0.25z_6)$ with

$$\begin{aligned}\mu_1(z) &= \exp\big(g_1(z_1, z_2, z_6) + 0.30|z_2 + 0.5|\big)\\ \mu_0(z) &= \exp\big(g_2(z_1, z_2, z_6) + 0.25|z_2 + 0.5|\big)\,;\end{aligned}$$

6. The optimal CATE: $D(z) = \exp(0.025 + 0.125z_1 + 0.05z_2 - 0.25z_6)$ with

$$\begin{aligned}\mu_1(z) &= \exp\big(0.15 + g_1(z_1, z_2, z_6) + 0.30z_2\big)\\ \mu_0(z) &= \exp\big(0.125 + g_2(z_1, z_2, z_6) + 0.25z_2\big)\,,\end{aligned}$$

where $g_1(z_1, z_2, z_6) = 0.125z_1 + 0.5|z_1| + 0.5|z_2| + 0.25z_6$ and $g_2(z_1, z_2, z_6) = 0.5|z_1| + 0.5|z_2| + 0.5z_6$. In the first setting, the simple Poisson regressions with a linear predictor are misspecified in both arms and the log-transformed optimal CATE is not a linear function of Z either. The proposed contrast regression is appropriate for the second setting, where the Poisson regression is mis-specified, but the underlying log-transformed CATE is a linear combination of the baseline covariates. The third setting is ideal for the simple regression approach, since

the Poisson regression models are correctly specified in both arms. The fourth setting is the same as the first setting. However, since we omit the last five covariates, Z_6 represents unmeasured confounding in the analysis. In this setting, the propensity score model is also misspecified. The fifth and sixth settings are similar to the first and second settings, respectively, except that the ATEs in the entire cohort are close to the null. In those settings, we expect that the CATE based on $\mathrm{E}(Y^{(1)} - Y^{(0)} \mid Z = z)$ is positively correlated with that based on $D(z) = \mathrm{E}(Y^{(1)} \mid Z = z)/\mathrm{E}(Y^{(0)} \mid Z = z)$.

For each simulated data set with a sample size of $n = 2,000$, we construct the CATE score using six methods.

1. the first method employs the boosting method (the base learner is the regression tree of depth 2) to estimate $\mu_r(z)$ in each arm separately and the CATE score is the ratio of two estimators;

2. the second method fits a naïve Poisson regression model in each arm;

3. the third method uses Bayesian additive regression tree to estimate $\mu_r(z)$ and the CATE score is the difference of two estimators [Lu et al., 2018];

4. the fourth method uses boosting method to predict the modified outcome [Wendling et al., 2018]
$$\frac{Y}{F}\left\{\frac{R}{\widehat{\pi}_1(Z)} - \frac{(1-R)}{\widehat{\pi}_0(Z)}\right\};$$

5. the fifth method is the proposed twin regression with the proposed doubly robust adjustment (the boosting-based estimation of $\mu_r(z)$ serves as the initial predictor);

6. the sixth method is the proposed contrast regression targeting $D(z)$ directly also with the doubly robust adjustment.

The propensity score is always estimated by fitting a standard logistic regression model. We then calculate the true $AD\{\hat{H}^{-1}(q)\}$ and the validation curve, $(1-q, AD\{\widehat{H}^{-1}(q)\})$, based on constructed CATE scores. We have also 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 average

of the resulting validation curves (Figure 1). We also include the validation curve based on the true CATE as the benchmark. Figure 2 shows the distribution of the correlation coefficients between the estimated CATE score and the truth. In addition, for settings 2, 3 and 6, model 7 is correctly specified and $\widehat{\delta}$ from the contrast regression is a consistent estimator of δ_0. We also examine the empirical bias and the coverage level of 95% confidence interval in estimating δ_0 based on 400 replications. The results are summarized in Table 3.

In settings 1, 2, 5 and 6, the proposed methods generates the best CATE score based on the validation curve. The CATE score based on the simple regression (method 2) is not informative and sometimes even yields a wrong direction due to the confounding effect. The CATE score based on nonparametric boosting (method 1) has robust performance but is slightly inferior to the proposed methods, despite the fact that CATE score from the latter is much simpler. The two methods (3 and 4) targeting CATE measured by the difference rather than ratio of $\mu_r(z)$ perform the worst in settings 1 and 2 and slightly better in settings 5 and 6. But in neither case, their performance is satisfactory, implying that the estimation method needs to be tailored for the metric used to measure the treatment effect. In setting 3, the naïve regression performs the best, since the simple regression model is correctly specified. The boosting and two proposed methods are only slightly inferior to the naïve regression. In setting 4, where the propensity score model is mis-specified and there is unmeasured confounding, neither method performs well. The proposed methods still perform relatively better than others. In settings 2, 3 and 6, where $D(z) = \exp(\delta_0^\top z)$, the proposed estimator based on the contrast regression is almost unbiased and the empirical coverage level of the constructed 95% confidence interval is close to the nominal level (Table 3).

In the second set of simulations, we consider the survival outcomes: $T^{(r)} \mid Z = z$ is generated from an exponential distribution with a failure rate $\lambda_r(z)$. The censoring distribution $C^{(r)} \mid Z = z$ is the minimum of a uniform random variable over [0.5, 1] and an exponential random variable with a failure rate of $\exp(0.25 + z_3)$. The treatment assignment $R \mid z$ is a Bernoulli random variable with a probability of $\pi_1(z) = \left\{1 + \exp(-z_1 + 0.5z_2 + 0.5z_6)\right\}^{-1}$. Other than right censoring, the only significant difference from the first set of simulations is that the simple multiplicative model for $\mu_r(z)$ or $D(z)$ is always misspecified. Due to the anticipated similarity to the previous simulations, we only consider the following two

settings:

1. The optimal CATE is approximately $\exp(-0.08 + 0.34z_1 + 0.11z_2 + 0.43z_6)$

$$\lambda_1(z) = \exp\left(-0.35 + 0.125z_1 + 0.3z_2 + 0.5z_6 - 0.5|z_1| + 0.5|z_2|\right)$$
$$\lambda_0(z) = \exp\left(-0.375 + 0.25z_2 + 0.25z_6 - 0.5|z_1| + 0.5|z_2|\right)$$

2. The optimal CATE is approximately $\exp(-0.09 + 0.27z_1 + 0.10z_2 + 0.51z_6)$ with

$$\lambda_1(z) = \exp\left(-0.05 + 0.125z_1 + 0.3z_2 + 0.5z_6\right)$$
$$\lambda_0(z) = \exp\left(-0.075 + 0.25z_2 + 0.25z_6\right)$$

In both settings, log-transformed true CATE is not equal to, but can be approximated by a linear combination of the covariates, which can explain approximately 85% of the variation in log-transformed true CATE. $\tau_0 = 0.75$ in RMTL and the censoring rate is 83% in both settings. In the light of the first set of simulations, we construct the CATE score only using four methods: (1) the first method employs the boosting method coupled with IPW correcting for right censoring to estimate $\mu_r(z, \tau)$ in each arm separately; (2) the second method fits a simple multiplicative regression model $\mu_r(z, \tau) = \exp(\beta_r^\top \widetilde{z})$ in each arm; and the third and fourth methods implement our new proposals. The censoring probability in the IPW is estimated via a Cox regression model, which is misspecified. Again, we calculate the the true validation curve based on constructed CATE scores and 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 average of the resulting validation curves and the distribution of correlation coefficients between the estimated and true CATEs in Figure 3.

In the first setting, the two proposed methods generate the best CATE score based on either the validation curve or the distribution of correlation coefficients. The estimated CATE score based on naïve regression or boosting is not informative. In setting two, where the multiplicative regression model is approximately correct within arm, the naïve regression approach performs the best, as expected. The twin regression performs almost equally well. Both boosting and contrast regression also yield informative CATE scores for subgroup

identification. In general, the simulation results are consistent with those from the first set of simulations.

4. Example

This project includes real-world data from 2791 patients from the NTD MS registry with 1050 patients receiving TERI and 1741 patients receiving DMF. Covariates of interest include age, number of prior treatments, MS duration, prior usage of glatiramer acetate (GA), prior usage of interferon (IFN), number of relapses in the year and in the two years prior to the index therapy, baseline Expanded Disability Status Scale (EDSS), and baseline pyramidal EDSS score. In Table 1, we summarize the distribution of aforementioned covariates by treatment arm. The patients receiving TERI are different from patients receiving DMF in several key patients' characteristics. For example, the patients receiving TERI tend to be older (45 vs 40), with 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.

In the entire cohort, the direct comparison of the normalized relapse rate based on Poisson regression shows that patients receiving DMF have a significantly lower relapse rate than those receiving TERI. The estimated ratio of the relapse rate (TERI vs DMF) is 1.270 (95% confidence interval: 1.121-1.439, $p < 0.001$). After our adjusting the imbalances in baseline covariates using the doubly robust estimation procedure, the estimated relapse rate ratio is 1.299 with a slightly wider confidence interval (1.018, 1.658). The adjusted annual relapse rate is 0.308 for TERI and 0.237 for DMF. 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 used the boosting method for Poisson distribution to construct an initial predictor of the relapse rate, with a depth of 2 for each individual tree and the number of trees selected via 5-fold cross-validation. The proposed CATE score is based on the average of three replicates of 7-fold cross-fitting. The propensity score is constructed based on the standard logistic regression model. The log-transformed CATE score is a linear combination of baseline covariates, with the weights summarized in Table 2, which also includes the estimated standard error for the weights

of covariates in CATE score from the contrast regression, suggesting that GA, number of relapses in the year prior to the therapy, the number of relapses in two years prior to the therapy and baseline EDSS are statistically significant at the 0.05 level in the CATE score. 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 4 shows a scatter plot 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.

We use cross validation to compare and evaluate the performance of the CATE scores more objectively. 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 of cross-validation, the entire data are split into a training set (67%) and a testing set (33%). We estimate the CATE scores based on training set and construct the validation curves in the testing set. After repeating this process 50 times, the empirical median of all 50 validation curves is obtained to measure the performance of the corresponding CATE score. Figure 5 (the upper row) summarizes the results, where we plot the median validation curves in both the training and testing sets to demonstrate the differential performance of four CATE scores. The two proposed CATE scores adjusting for imbalance in baseline covariates appear to have a similarly superior performance in the testing set: both scores suggest a moderate treatment effect heterogeneity. In the aforementioned cross-validation, we also use the estimated CATE score based on the training data 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. If the constructed CATE score is informative, then we expected that the estimated rate ratio in the subgroup of patients with a higher CATE score is also higher. The median of the estimated ratios of the relapse rate (TERI vs DMF) is 1.723 in 50% patients having a CATE score favoring DMF the most and the median of the estimated ratios of the relapse rate is 1.089 in remaining 50% patients based on the twin regression method. These two median ratios become 1.671 and 1.113 based

on the CATE score constructed via the contrast regression method. The distribution of the estimated ratios across different cross-validation replicates are summarized in Figure 6. As a cautionary note, this observed difference in treatment effect may not be adequately stable due to the limited sample size in the validation set (on average, there are only 465 patients in each of the two subgroups). However, the results still exhibit moderate signals of the presence of treatment effect heterogeneity captured by two proposed approaches. A important observation is that the estimated treatment effect heterogeneity doesn't 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. In future studies, it would be beneficial to have an independent cohort for further validation of the constructed CATE score.

In the analysis of the time to the first relapse, we use the ratio of RMTL due to relapse as the metric for the treatment effect, with a truncation time point τ_0 of 4.34 years (52 months). In the naïve comparison without adjusting for baseline covariates, the estimated ratio of mean time lost up to 4.34 years is 1.246 (95% confidence interval: 1.098-1.414, $p < 0.001$), suggesting that time lost to relapse in patients receiving DMF is only 80% of that in patients receiving TERI. After the doubly robust adjustment, the estimated ratio becomes 1.531 (95% confidence interval: 1.313-1.785, $p < 0.001$) with the RMTL of 1.451 and 0.947 years for TERI and DMF, respectively, showing a greater difference between two treatments. We compare CATE scores based on the naïve regression and new proposal. In constructing the proposed CATE scores, we used random forest with 50 trees to generate the initial prediction of the RMTL due to relapse. The log-transformed CATE score is a linear combination of baseline covariates, whose composition is also reported in Table 2. While the contributions of the same covariate to resulting CATE scores can be different, such as the EDSS score at baseline, the resulting CATE scores are correlated. We also performed cross validation to evaluate the performance of the different CATE scores. It appears that the CATE score directly based on boosting overfits the training data and performs similarly as other methods in the testing set. However, the estimated treatment effect heterogeneity is very weak based on any of the four CATE scores. For example, in the cross-validation, the estimated ratio of RMTL in 50% of the patients favoring DMF based on contrast regression

approach is 1.624. The ratio in the rest of the 50% patients is 1.460, only marginally lower. The distributions of two RMTL ratios across replicates substantially overlap. Lastly, we examine the relationship between the CATE scores built for the number of relapses and the time to the first relapse. The two sets of CATE scores are correlated as expected. For example the correlation coefficient between two CATE scores based on the contrast regression is 0.80.

5. Discussion

We 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 CATE. The difference or ratio of the expected response from individual patients may be the most common choice. While we justify that the ATE based ratio or difference in patients with big CATE must also be big, this is not necessarily true for treatment effect measured by OR or hazard ratio. In other words, the ATE and CATE don't always align with each other. When we are interested in treatment effect measured by the OR or hazard ratio, we may not want to identify the high value subgroup according to the CATE on the same metric. Second, the difference in covariates distribution between patients receiving different treatments affects the regression analysis within each arm and the standard methods may induce false treatment-covariate interaction. 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. Lastly, we proposed a set of methods for estimating the CATE measured by the ratio, which may result in very different conclusions in comparison with most current methods, which target on the difference as the treatment effect.

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 of CATE need to be modified with a new arm specific propensity score being constructed to match the

distribution of covariates of the patients in the treatment arm r of the training set to that of the target population. Otherwise, the same CATE score may define a different subgroup of patients in the target population, i.e., $\{z \mid \widehat{D}(z) \geqslant c\}$ may be different from $\{z^V \mid \widehat{D}(z^V) \geqslant c\}$. Consequently, the ATE observed in the high value subgroup $\{z \mid \widehat{D}(z) \geqslant c\}$ may not be reproducible. Furthermore, the validity of the validation results still depends on assumptions made in estimating the causal effect based on observational data. Ideally, the validation data are from an external randomized clinical trial, so that we can estimate the ATE in the identified high value subgroup without systematic biases from even unmeasured confounding.

Qi et al. [2019] discussed the optimal treatment recommendation in the presence of $K > 2$ treatments. The proposed twin regression approach can be used to approximate $\mu_k(z) = \mathrm{E}(Y^{(r)} \mid Z = z), k = 0, 1, \cdots, K$, and select the optimal regiment accordingly. The contrast regression can directly estimate $D_{ij}(z) = \mu_i(z)/\mu_j(z)$ based on the limiting estimating equation:

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

for estimating $D_{ij}(z) = \exp(\delta_{ij}^\top \widetilde{z})$, where $\pi_k(z) = P(R = k \mid Z = z)$ and $w(z, \delta_{ij})$ is a weight function. Appropriate augmentation term based on $\{(I(R = r) - \pi_r(Z), r = 1, \cdots, K\}$ can be introduced for further efficiency improvement. A complication is that resulting estimators don't necessarily have the property that $D_{ij}(z) = D_{il}(z)D_{lj}(z)$, for which further research is warranted.

6. Appendix

6.1. The optimality of the estimating equation

Consider all weighted estimating equations in the form of

$$\begin{aligned} &S_n(\delta; w, \mu) \\ =&n^{-1}\sum_{i=1}^{n}\widetilde{Z}_i w(Z_i)\left(\left[\{1-\pi(Z_i)\}R_iY_i^{(1)} - \pi(Z_i)(1-R_i)Y_i^{(0)}e^{\delta^\top\widetilde{Z}_i}\right] - \mu(Z_i)\{R_i - \pi(Z_i)\}\right). \end{aligned}$$

If $\pi(z) = \pi_1(z) = P(R = 1 \mid Z = z)$ the correct propensity score, the estimating function converges to

$$s(\delta; w, \mu) = \mathrm{E}\left[w(Z)\pi_1(Z)\pi_0(Z)\{\mu_1(Z) - \mu_0(Z)\exp(\delta_0^\top Z)\}\right]$$

and δ_0 is the root of $s(\delta; w, \mu) = 0$. Thus, the root of the estimating equation $S_n(\delta; w, \mu) = 0$, denoted by $\widehat{\delta}_{w,\mu}$, is consistent in estimating δ_0, and $\sqrt{n}(\widehat{\delta}_{w,\mu} - \delta_0)$ converges to a mean zero Gaussian distribution with a variance-covariance matrix of $A_{w,\mu}^{-1} B_{w,\mu} A_{w,\mu}^{-1}$, where

$$A_{w,\mu} = \mathrm{E}\left\{\tilde{Z}\tilde{Z}^\top w(Z)\pi_1(Z)\pi_0(Z)\mu_1(Z)\right\},$$

and

$$B_{w,\mu} = \mathrm{E}\left\{\tilde{Z}\tilde{Z}^\top w(Z)^2 \left(\left[\pi_0(Z_i)R_i Y_i^{(1)} - \pi_1(Z_i)(1 - R_i)Y_i^{(0)} e^{\delta_0^\top \tilde{Z}_i}\right] - \mu(Z_i)\{R_i - \pi_1(Z_i)\}\right)^2\right\}.$$

First, since

$$\mathrm{E}\left(\left[\pi_0(Z_i)R_i Y_i^{(1)} - \pi(Z_i)(1 - R_i)Y_i^{(0)} e^{\delta_0^\top \tilde{Z}_i}\right] \mid Z_i = z\right) = 0$$

and

$$\mathrm{E}\left(R_i - \pi_1(Z_i) \mid Z_i = z\right) = 0,$$

$B_{w,\mu} - B_{w,\bar{\mu}}$ is semi-positive definite, where

$$\begin{aligned}\bar{\mu}(z) =& \frac{\mathrm{E}\left(\left[\pi_0(Z_i)R_i Y_i^{(1)} - \pi_1(Z_i)(1 - R_i)Y_i^{(0)} e^{\delta_0^\top \tilde{Z}_i}\right]\left\{R_i - \pi_1(Z_i)\right\} \mid Z_i = z\right)}{\mathrm{E}\left(\left\{R_i - \pi_1(Z_i)\right\}^2 \mid Z_i = z\right)} \\ =& \frac{\pi_1(z)\pi_0(z)\mu_1(z)}{\pi_1(z)\pi_0(z)} = \mu_1(z).\end{aligned}$$

Secondly, by Cauchy ineqaulity, $A_{w,\mu_1}^{-1}B_{w,\mu_1}A_{w,\mu_1}^{-1} - A_{w_0,\mu_1}^{-1}B_{w_0,\mu_1}A_{w_0,\mu_1}^{-1}$ is also semi-positive definite, where

$$\begin{aligned}
w_0(z) =& \frac{\mathrm{E}\left(\pi_1(Z)\pi_0(Z)\mu_1(Z) \mid Z=z\right)}{\mathrm{E}\left(\left[\pi_0(Z)RY^{(1)} - \pi_1(Z)(1-R)Y^{(0)}e^{\delta_0^\top \tilde{Z}} - \mu_1(Z)\{R-\pi_1(Z)\}\right]^2 \mid Z=z\right)} \\
=& \frac{\pi_1(z)\pi_0(z)\mu_1(z)}{\mathrm{E}\left(\left[\pi_0(z)R(Y^{(1)}-\mu_1(z)\right]^2 + \left[\pi_1(z)(1-R)\left\{Y^{(0)}-\mu_0(z)\right\}e^{\delta_0^\top \tilde{z}}\right]^2 \mid Z=z\right)} \\
=& \frac{\pi_1(z)\pi_0(z)\mu_1(z)}{\pi_0(z)^2\pi_1(z)\mathrm{var}(Y^{(1)} \mid Z=z) + \pi_0(z)\pi_1(z)^2\mathrm{var}(Y^{(0)} \mid Z=z)e^{2\delta_0^\top z}} \\
=& \frac{\mu_1(z)}{\pi_0(z)\mu_1(z) + \pi_1(z)\mu_0(z)e^{2\delta_0^\top z}} = \frac{1}{\pi_0(z) + \pi_1(z)e^{\delta_0^\top z}},
\end{aligned}$$

where we used the fact that $\mathrm{var}(Y^{(r)} \mid Z=z) = \mu_r(z)$. Therefore $A_{w,\mu}^{-1}B_{w,\mu}A_{w,\mu}^{-1} - A_{w_0,\mu_1}^{-1}B_{w_0,\mu_1}A_{w_0,\mu_1}^{-1}$ is always semi-positive definite, suggesting that an estimating function $S_n(\delta, w_0, \mu_1)$ with a limit

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

yields the optimal variance among regular estimators of δ_0.

6.2. 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 $\{\delta_n\}_{n\geqslant 1}$, $\{\Delta_n\}_{n\geqslant 1}$, and $\{\tau_n\}_{n\geqslant 1}$ 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 $n \geqslant 3$ and $P \in \mathcal{P}_n$, (a) $E\{m(G; \delta_0, \mu_0, \pi_1)\} = 0$, and Ω contains a ball of radius $c_1 n^{-1/2}\log n$ centered at δ_0; (b) the map $(\delta, \mu, \pi) \to E\{m(G; \delta, \mu, \pi)\}$ is twice continuously Gateaux-differentiable on $\Omega \times \mathcal{T}$; (c) for all $\delta \in \Omega$, $2\|E\{m(G; \delta, \mu_0, \pi_1)\}\| \geqslant \|J_0(\delta - \delta_0)\| \wedge c_0$, where J_0 is the Jacobian matrix of $\delta \mapsto E[m(G; \delta, \mu_0, \pi_1]$ at δ_0; (d) the score $m(g; \delta, \mu, \pi)$ obeys the Neyman orthogonality condition*

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

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

Assumption 4. *Let* K *be a fixed integer. For all* $n \geqslant 3$ *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* $(\widehat{\mu}_0^{-k}, \widehat{\pi}_1^{-k})_{1 \leqslant k \leqslant K}$ *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)* $\mathcal{F}_{1,(\mu,\pi)} = \{m_j(g; \delta, \mu, \pi) \mid j = 1, \cdots, d+1, \delta_0 \in \Omega\}$ *is suitably measurable and its uniform covering entropy obeys*

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

for $\epsilon_N \in (0, 1]$*, where* $F_{1,(\mu,\pi)}$ *is a measurable envelope for* $\mathcal{F}_{1,(\mu,\pi)}$ *that satisfies* $\|F_{1,(\mu,\pi)}\|_{P,q} \leqslant c_1$*; (c)* $r_n = \sup_{(\mu,\pi)\in\mathcal{T}_n, \delta_0 \in \Omega} \|E\{m(G; \delta, \mu, \pi)\} - E\{m(G; \delta_0, \mu_0, \pi_0)\}\| \leqslant \delta_n \tau_n$*; (d)* $r_n' \log^{1/2}(1/r_n') \leqslant \delta_n$*, where*

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

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

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

(f) all eigenvalues of the matrix $E\big[m(G; \delta_0, \mu_0, \pi_1) m^\top (G; \delta_0, \mu_0, \pi_1)\big]$ *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* $\delta_n \geqslant n^{-1/2+1/q} \log(n)$ *and that* $n^{-1/2} \log(n) \leqslant \tau_n \leqslant \delta_n$ *for all* $n \geqslant 1$ *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*

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

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 $\epsilon_\pi \leqslant \pi(z) \leqslant 1 - \epsilon_\pi$ Q-almost everywhere.

We proved Assumption 3(a) in the main text; see equation (8), by using the assumption in (7): $\mu_1(z) = \mu_0(z) \exp(\delta_0^\top \widetilde{z})$, 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 $0 < \epsilon_\pi \leqslant \pi(z) \leqslant 1 - \epsilon_\pi$,

$$\sup_{z \in \mathcal{Z}, \delta \in \Omega} \{\exp(\delta^\top \widetilde{z})\pi(z) + 1 - \pi(z)\}^{-1} \leqslant C_0 \tag{9}$$

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 $\mathrm{E}\big[m(G; \delta, \mu_0, \pi_1)\big]$ is differentiable with respect to δ with Jacobian

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

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 $1 \leqslant i, j \leqslant d + 1$, 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

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

if $\epsilon \leqslant \lambda_0/2(d+1)$, 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,

$$2\|\mathrm{E}\{m(G;\delta,\mu_0,\pi_1)\}\| \geqslant \|J_0(\delta-\delta_0)\|,$$

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 continuous 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 $(\mu, \pi + rd_\pi) \in \mathcal{T}$,

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

and

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

By Dominated Convergence Theorem, we may exchange the $\lim_{r\to 0}$ with the expectation due to the presence of integrable envelop functions (using Assumption 1). Therefore, the Gateuaux derivative with respect to μ along the direction of d_μ exists and is

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

The Gateuaux derivative with respect to π along the direction of d_π also exists and is

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

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

To examine the orthogonality condition, i.e., 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

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

Then, for

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

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 (9) 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

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

Therefore

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

which verifies Assumption 3(d).

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

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

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

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

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

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

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

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 $|t| \leqslant \sup_{\delta\in\Omega,z\in\mathcal{Z}} |\delta^\top \widetilde{z}|$. Therefore, $m(G;\delta,\mu,\pi)$ is Lipschitz in δ. For all $(\mu,\pi) \in \mathcal{T}$, (9) 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 $\log N(\epsilon_N, \Omega, \|\cdot\|_\infty) \leqslant \widetilde{v} \log(\widetilde{a}/\epsilon_N)$ on the parameter space imply that $\sup_Q \log N(\epsilon_N \|F_{1,(\mu,\pi)}\|_{Q,2}, \mathcal{F}_{1,\mu,\pi}, \|\cdot\|_{Q,2}) \leqslant v \log(a/\epsilon_N)$. Thus Assumption 4(b) is verified.

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

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 $r_n \leqslant \widetilde{C}_1 \|\pi - \pi_1\|_{P,2} \leqslant \delta_n \tau_n$ for adequately large n, using the definition of $\mathcal{T}_n$.

Next,

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

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

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

Thus, $r_n'\log^{1/2}(1/r_n')\leqslant\widetilde{C}_2a_n^{3/8}n^{-1/4}\sqrt{\log(n)}\leqslant\delta_n$ 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

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

where $h_1(z, r) = (1 - e^{r d_f(z)})\pi_0(z)\pi_1(z)\mu_0(z) - r(1 - e^{r d_f(z)})\pi_1(z)\mu_0(z) d_\pi(z) + r^2 e^{r d_f(z)} d_\mu(z) d_\pi(z)$, and $h_2(z, r) = e^{f_0(z) + r d_f(z)}\{\pi_1(z) + r d_\pi(z)\} + \pi_0(z) - r d_\pi(z)$. 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

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

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

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

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

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

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 $\mathrm{var}(Y^{(r)}\mid Z=z)\geqslant\sigma_L>0$.

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

□

The Proof of Theorem 2: First, for $c>0$, we show that $AD(c)\geqslant c$, by showing that $\mathrm{E}(Y^{(1)}\mid D(Z)\geqslant c)\geqslant c\mathrm{E}(Y^{(0)}\mid D(Z)\geqslant c)$. This follows from the following application of the tower property of conditional expectations,

$$\begin{aligned}\mathrm{E}(Y^{(1)}\mid D(Z)\geqslant c)&=\mathrm{E}\{\mathrm{E}(Y^{(1)}\mid Z)\mid D(Z)\geqslant c\} \\ &=\mathrm{E}\{D(Z)\mathrm{E}(Y^{(0)}\mid Z)\mid D(Z)\geqslant c\} \\ &\geqslant\mathrm{E}\{c\mathrm{E}(Y^{(0)}\mid Z)\mid D(Z)\geqslant c\} \\ &=c\mathrm{E}\{\mathrm{E}(Y^{(0)}\mid Z)\mid D(Z)\geqslant c\} \\ &=c\mathrm{E}\{Y^{(0)}\mid D(Z)\geqslant c\}.\end{aligned}$$

Now, let $0 < c' < c$. Expanding $AD(c')$ gives

$$\begin{aligned} AD(c') &= \frac{\mathrm{E}\{Y^{(1)} \mid D(Z) \geqslant c'\}}{\mathrm{E}\{Y^{(0)} \mid D(Z) \geqslant c'\}} \\ &= \frac{\mathrm{E}\{Y^{(1)} \mid D(Z) \geqslant c\}p + \mathrm{E}\{Y^{(1)} \mid D(Z) \in [c', c)\}(1-p)}{\mathrm{E}[Y^{(0)} \mid D(Z) \geqslant c\}p + \mathrm{E}\{Y^{(0)} \mid D(Z) \in [c', c)\}(1-p)}, \end{aligned}$$

where $p = P\{D(Z) \geqslant c \mid D(Z) \geqslant c'\}$. Now, consider the expression in the numerator, which we expand and bound as

$$\begin{aligned} &\mathrm{E}\{Y^{(1)} \mid D(Z) \geqslant c\}p + \mathrm{E}\{Y^{(1)} \mid D(Z) \in [c', c)\}(1-p) \\ &\quad = AD(c)\mathrm{E}\{Y^{(0)} \mid D(Z) \geqslant c\}p + \mathrm{E}\{Y^{(1)} \mid D(Z) \in [c', c)\}(1-p) \\ &\quad = AD(c)\mathrm{E}\{Y^{(0)} \mid D(Z) \geqslant c\}p + \mathrm{E}\{D(Z)\mathrm{E}(Y^{(0)} \mid Z) \mid D(Z) \in [c', c)\}(1-p) \\ &\quad \leqslant AD(c)\mathrm{E}\{Y^{(0)} \mid D(Z) \geqslant c\}p + c\mathrm{E}\{Y^{(0)} \mid D(Z) \in [c', c)\}(1-p) \\ &\quad \leqslant AD(c)\left[\mathrm{E}\{Y^{(0)} \mid D(Z) \geqslant c\}p + \mathrm{E}\{Y^{(0)} \mid D(Z) \in [c', c)\}(1-p)\right], \end{aligned}$$

where in the last inequality we used the fact that $AD(c) \geqslant c$. Plugging this into the expression for $AD(c')$ gives

$$AD(c') \leqslant \frac{AD(c)(\mathrm{E}\{Y^{(0)} \mid D(Z) \geqslant c\}p + \mathrm{E}\{Y^{(0)} \mid D(Z) \in [c', c)\}(1-p))}{\mathrm{E}\{Y^{(0)} \mid D(Z) \geqslant c\}p + \mathrm{E}\{Y^{(0)} \mid D(Z) \in [c', c)\}(1-p)} = AD(c).$$

6.3. Examples of CATE vs ATE measured by OR and HR

OR example:

Consider the case where the response rate of a binary outcome in the control arm is

$$p_i = \begin{cases} 1 - \frac{i}{n+1} & \text{if } \ i \text{ is even}, \\ \frac{i}{n+1} & \text{if } \ i \text{ is odd}, \end{cases}$$

the response rate in the treatment arm is $p_i\theta_i/(1 - p_i + p_i\theta_i)$, $\log(\theta_i) = \log(2) - \log(4)(i - 1)/(n - 1)$, and $n = 100$. Thus, these 100 patients are sorted according to their conditional OR, which is monotone decreasing from 2 to -2. The marginal OR in 10 patients with the highest conditional OR is only 1.14, though the conditional OR of each patient is above 1.76. On the other hand, we can select a different subgroup of 10 patients, in which the marginal OR is as high as 1.66. Figure 7 plots the marginal ORs in a sequence of subgroups based on

(1) the optimal grouping maximizing the marginal ORs and (2) the grouping according to the conditional OR of each patient. The former can be substantially higher than the latter.

HR example:

Consider the following example, where $\lambda_1(t|z)/\lambda_0(t|z)$ is constant, i.e., there is no treatment effect heterogeneity, but $\exp\{(\beta_1 - \beta_0)^\top z\}$ depends on z. Let

$$\lambda_0(t \mid Z) = 2\{Z^{-1}I(t < 1) + Z^{-2}I(1 \leqslant t < 2) + Z^{-3}I(t \geqslant 2)\},$$

$\lambda_1(t \mid Z) = 0.5\lambda_0(t \mid Z)$, and $Z \sim U(0, 2)$, where $I(\cdot)$ is the indicator function. Solving the score equations from the mis-specified Cox model results in $(\beta_1, \beta_0) \approx (-1.85, -1.55)$ in the absence of any censoring, implying nontrivial treatment effect heterogeneity.

REFERENCES

S. Athey and G. Imbens. Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences*, 113(27):7353–7360, 2016.

S. Athey, J. Tibshirani, S. Wager, et al. Generalized random forests. *The Annals of Statistics*, 47(2):1148–1178, 2019.

H. Bang and J. M. Robins. Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61(4):962–973, 2005.

S. Basu, J. B. Sussman, J. Rigdon, L. Steimle, B. T. Denton, and R. A. Hayward. 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(10):e1002410, 2017.

L. Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.

T. Cai, L. Tian, P. H. Wong, and L. Wei. Analysis of randomized comparative clinical trial data for personalized treatment selections. *Biostatistics*, 12(2):270–282, 2010.

S. Chen, L. Tian, T. Cai, and M. Yu. A general statistical framework for subgroup identification and comparative treatment scoring. *Biometrics*, 73(4):1199–1209, 2017.

S. Cheng, J. P. Fine, and L. Wei. Prediction of cumulative incidence function under the proportional hazards model. *Biometrics*, pages 219–228, 1998.

V. Chernozhukov, D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. Newey, and J. Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 31:C1–C68, 2018.

J. C. Foster, J. M. Taylor, and S. J. Ruberg. Subgroup identification from randomized clinical trial data. *Statistics in medicine*, 30(24):2867–2880, 2011.

J. Friedman, T. Hastie, R. Tibshirani, et al. Additive logistic regression: a statistical view of boosting (with discussion and a rejoinder by the authors). *The annals of statistics*, 28 (2):337–407, 2000.

D. P. Green and H. L. Kern. Modeling heterogeneous treatment effects in survey experiments with bayesian additive regression trees. *Public opinion quarterly*, 76(3):491–511, 2012.

G. W. Imbens and D. B. Rubin. *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press, 2015.

J. D. Kang and J. L. Schafer. Demystifying Double Robustness: A Comparison of Alternative Strategies for Estimating a Population Mean from Incomplete Data. *Statistical Science*, 22(4):523–539, 2007.

S. R. Künzel, J. S. Sekhon, P. J. Bickel, and B. Yu. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the National Academy of Sciences*, 116(10):4156–4165, 2019. ISSN 0027-8424. doi: 10.1073/pnas.1804597116. URL https://www.pnas.org/content/116/10/4156.

W.-Y. Loh, X. He, and M. Man. A regression tree approach to identifying subgroups with differential treatment effects. *Statistics in medicine*, 34(11):1818–1833, 2015.

M. Lu, S. Sadiq, D. J. Feaster, and H. Ishwaran. Estimating individual treatment effect in observational data using random forest methods. *Journal of Computational and Graphical Statistics*, 27(1):209–219, 2018.

X. Nie and S. Wager. Quasi-Oracle Estimation of Heterogeneous Treatment Effects. *arXiv:1712.04912 [stat.ML]*, 2019.

S. Powers, J. Qian, K. Jung, A. Schuler, N. H. Shah, T. Hastie, and R. Tibshirani. Some methods for heterogeneous treatment effect estimation in high dimensions. *Statistics in medicine*, 37(11):1767–1787, 2018.

Z. Qi, D. Liu, H. Fu, and Y. Liu. Multi-armed angle-based direct learning for estimating optimal individualized treatment rules with various outcomes. *Journal of the American Statistical Association*, pages 1–33, 2019.

J. M. Robins and A. Rotnitzky. Comment on: Inference for semiparametric models: some questions and an answer. *Statistica Sinica*, 11(4):920–936, 2001.

L. Tian, L. Zhao, and L. Wei. Predicting the restricted mean event time with the subject's baseline covariates in survival analysis. *Biostatistics*, 15(2):222–233, 2013.

L. Tian, A. A. Alizadeh, A. J. Gentles, and R. Tibshirani. A Simple Method for Estimating Interactions Between a Treatment and a Large Number of Covariates. *Journal of the American Statistical Association*, 109(508):1517–1532, 2014.

H. Uno, B. Claggett, L. Tian, E. Inoue, P. Gallo, T. Miyata, D. Schrag, M. Takeuchi, Y. Uyama, L. Zhao, et al. Moving beyond the hazard ratio in quantifying the between-group difference in survival analysis. *Journal of clinical Oncology*, 32(22):2380, 2014.

M. J. Van der Laan and S. Rose. *Targeted learning: causal inference for observational and experimental data*. Springer Science & Business Media, 2011.

S. Wager and S. Athey. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242, 2018.

T. Wendling, K. Jung, A. Callahan, A. Schuler, N. Shah, and B. Gallego. Comparing methods for estimation of heterogeneous treatment effects using observational data from health care databases. *Statistics in medicine*, 37(23):3309–3324, 2018.

Y. Xie, J. E. Brand, and B. Jann. Estimating heterogeneous treatment effects with observational data. *Sociological methodology*, 42(1):314–347, 2012.

B. Zhang, A. A. Tsiatis, M. Davidian, M. Zhang, and E. Laber. Estimating optimal treatment regimes from a classification perspective. *Stat*, 1(1):103–114, 2012.

L. Zhao, L. Tian, T. Cai, B. Claggett, and L.-J. Wei. Effectively Selecting a Target Population for a Future Comparative Study. *Journal of the American Statistical Association*, 108 (502):527–539, 2013.

Y. Zhao, D. Zeng, A. J. Rush, and M. R. Kosorok. Estimating Individualized Treatment Rules Using Outcome Weighted Learning. *Journal of the American Statistical Association*, 107(499):1106–1118, 2012.

Y.-Q. Zhao, D. Zeng, E. B. Laber, R. Song, M. Yuan, and M. R. Kosorok. Doubly robust learning for estimating individualized treatment with censored data. *Biometrika*, 102(1): 151–168, 2014.

X. Zhou, N. Mayer-Hamblett, U. Khan, and M. R. Kosorok. Residual Weighted Learning for Estimating Individualized Treatment Rules. *Journal of the American Statistical Association*, 112(517):169–187, 2017.

Table 1: 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 2: The estimated weights in constructed CATE scores (TERI vs DMF)

	ratio of relapse rate			ratio of RMTL		
	naïve reg.	twin reg.	contrast reg.	naïve reg.	twin reg.	contrast reg.
intercept	0.692	0.476	0.670 (0.711)*	1.061	1.310	1.017
age	0.013	0.013	0.017 (0.013)	0.008	0.011	0.010
# prior treatments	-0.303	0.011	-0.088 (0.195)	-0.255	-0.342	-0.210
MS duration (years)	0.022	0.045	0.028 (0.028)	0.004	0.003	0.007
GA	-0.584	0.517	-0.700 (0.349)	-0.630	-0.801	-0.640
IFN	-0.304	-0.024	-0.185 (0.318)	-0.390	-0.583	-0.410
# relapses (prior year)	-0.258	-0.661	-0.811 (0.271)	-0.217	-0.193	-0.126
# relapses (prior two years)	0.191	0.360	0.444 (0.201)	0.195	0.139	0.130
EDSS	-0.046	-0.247	-0.233 (0.114)	-0.072	-0.060	-0.117
pyramidal EDSS	0.006	0.027	-0.004 (0.160)	0.060	0.030	0.073

* : the estimated standard error of the weight.

Table 3: Empirical bias and coverage level of the 95% confidence interval in estimating δ_0 based on 400 replications under three different settings: $(\delta_0 = (\delta_{00}, \delta_{01}, \delta_{02}, \cdots, \delta_{10})^\top)$

Covariate	Setting 2		Setting 3		Setting 6	
	bias	coverage	bias	coverage	bias	coverage
$\delta_{01} = 0.125$	-0.018	92.5%	-0.030	94.5%	-0.014	94.0%
$\delta_{02} = 0.05$	-0.015	92.5%	-0.014	92.8%	-0.013	90.8%
$\delta_{03} = 0$	0.004	93.8%	0.002	94.2%	-0.002	94.2%
$\delta_{04} = 0$	-0.002	94.0%	0.002	93.8%	0.004	94.8%
$\delta_{05} = 0$	-0.005	94.5%	-0.004	93.0%	-0.007	93.8%
$\delta_{06} = -0.25$	0.013	91.5%	0.008	94.2%	0.016	93.2%
$\delta_{07} = 0$	-0.004	93.0%	0.008	91.8%	0.002	93.8%
$\delta_{08} = 0$	-0.002	94.5%	-0.003	93.0%	0.000	92.8%
$\delta_{09} = 0$	0.005	92.5%	0.000	93.8%	-0.011	93.2%
$\delta_{10} = 0$	-0.006	92.0%	-0.007	93.5%	-0.004	91.8%

Figure 1: The ATE in subgroups of patients identified by different CATE scores including true CATE, boosting (method 1), naïve regression (method 2), Bayesian additive regression tree (method 3), modified outcome boosting (method 4), twin regression (method 5), contrast regression (method 6) in six simulation settings.

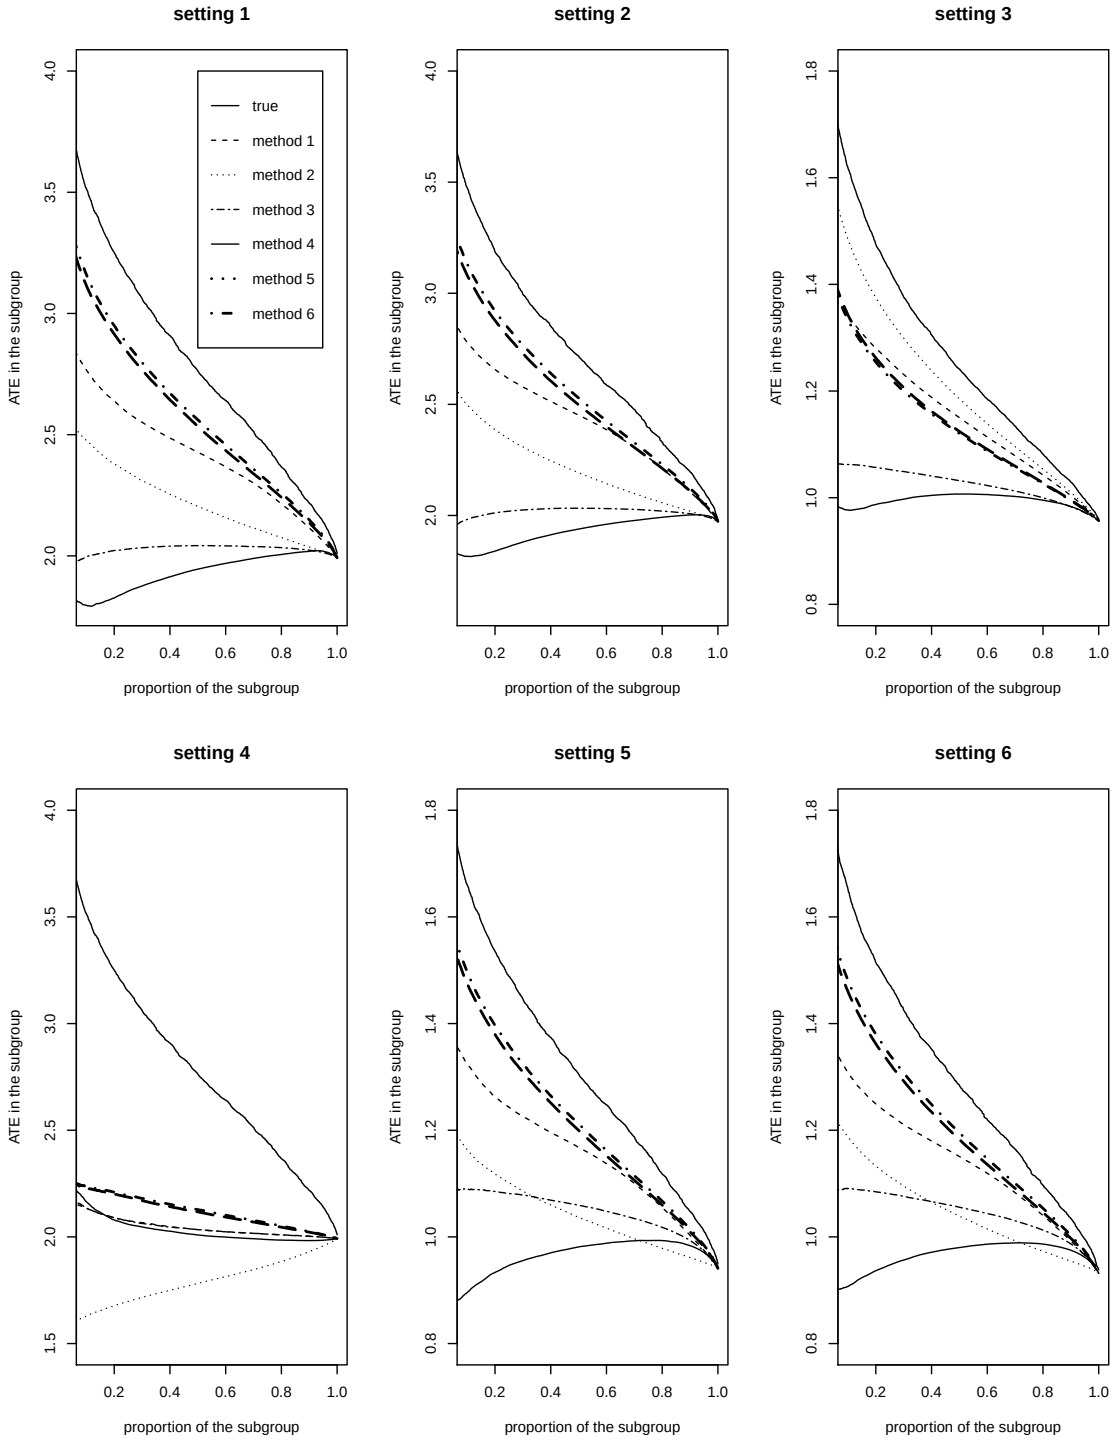

Figure 2: The distribution of correlation coefficients between estimated CATE and true CATE in six simulation settings; there are six methods considered: boosting, naïve regression, Bayesian additive regression tree, modified outcome boosting, twin regression (light gray), and contrast regression (dark gray)

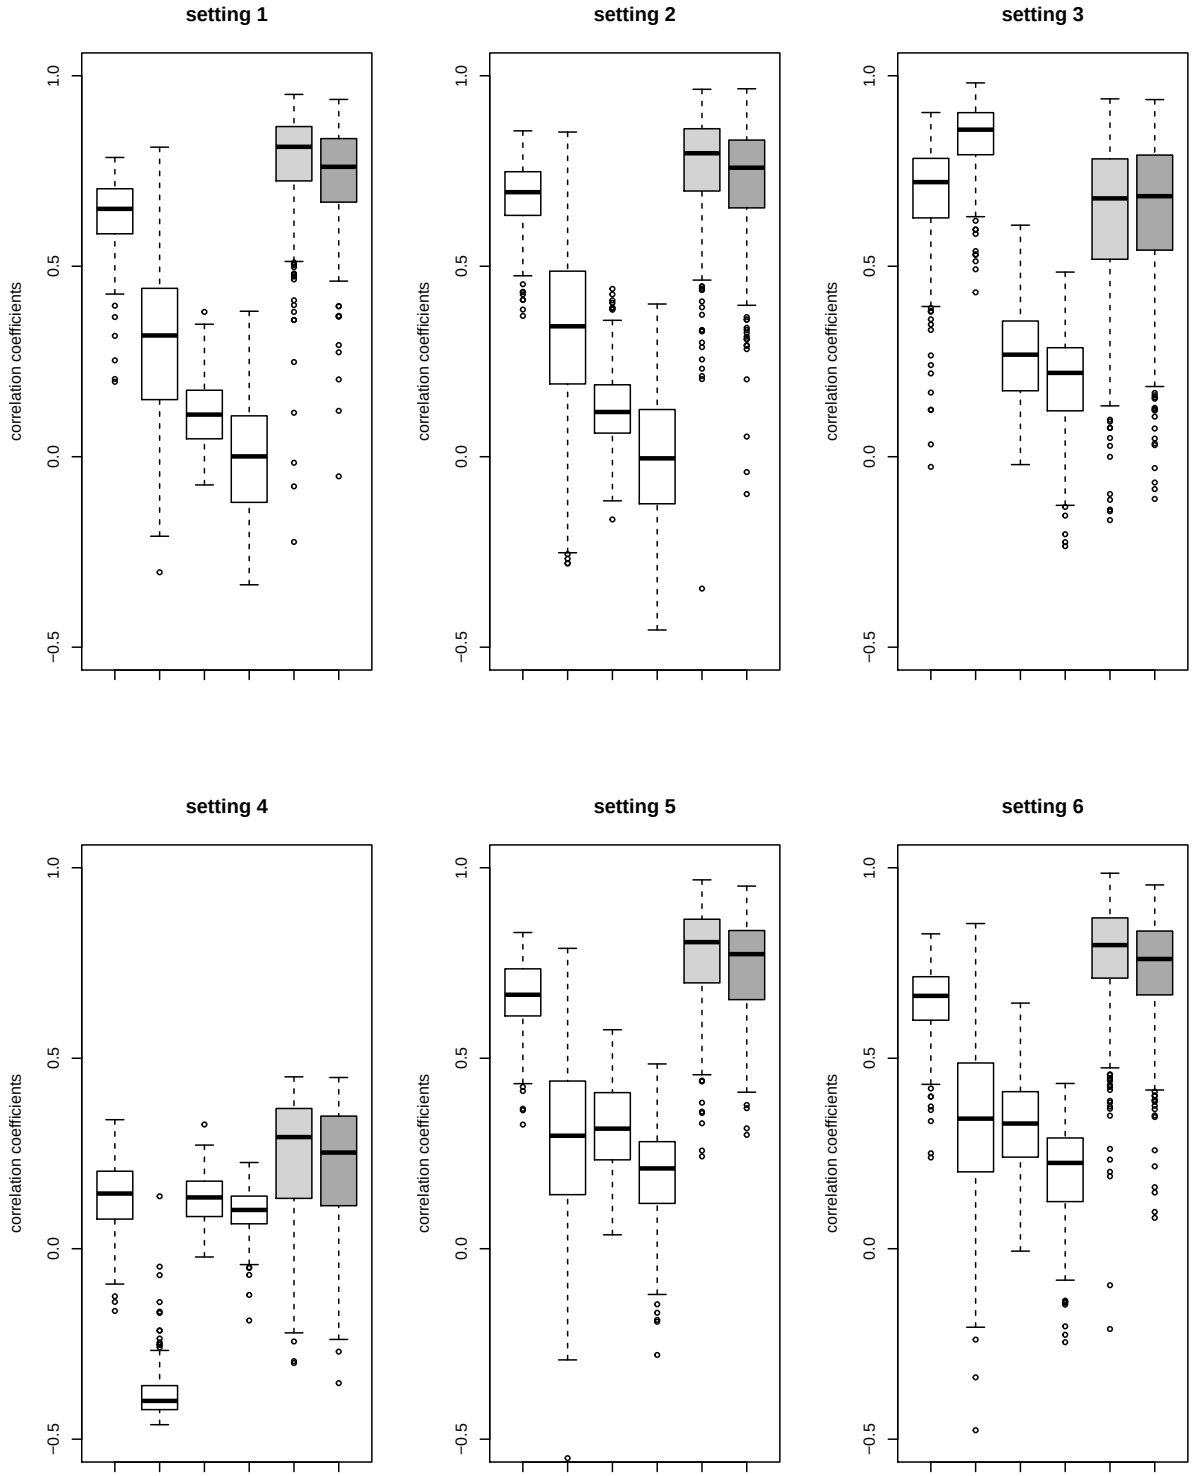

Figure 3: The first row: the ATE in subgroups of patients identified by different CATE scores (boosting, standard regression and the proposed methods) in two simulation settings and the ATE in subgroup of patients sorted according to the true $D(z)$ (solid curve); the second row: the distribution of correlation coefficients between true CATE and estimated CATE scores including that based on twin regression (light gray) and contrast regression (dark gray)

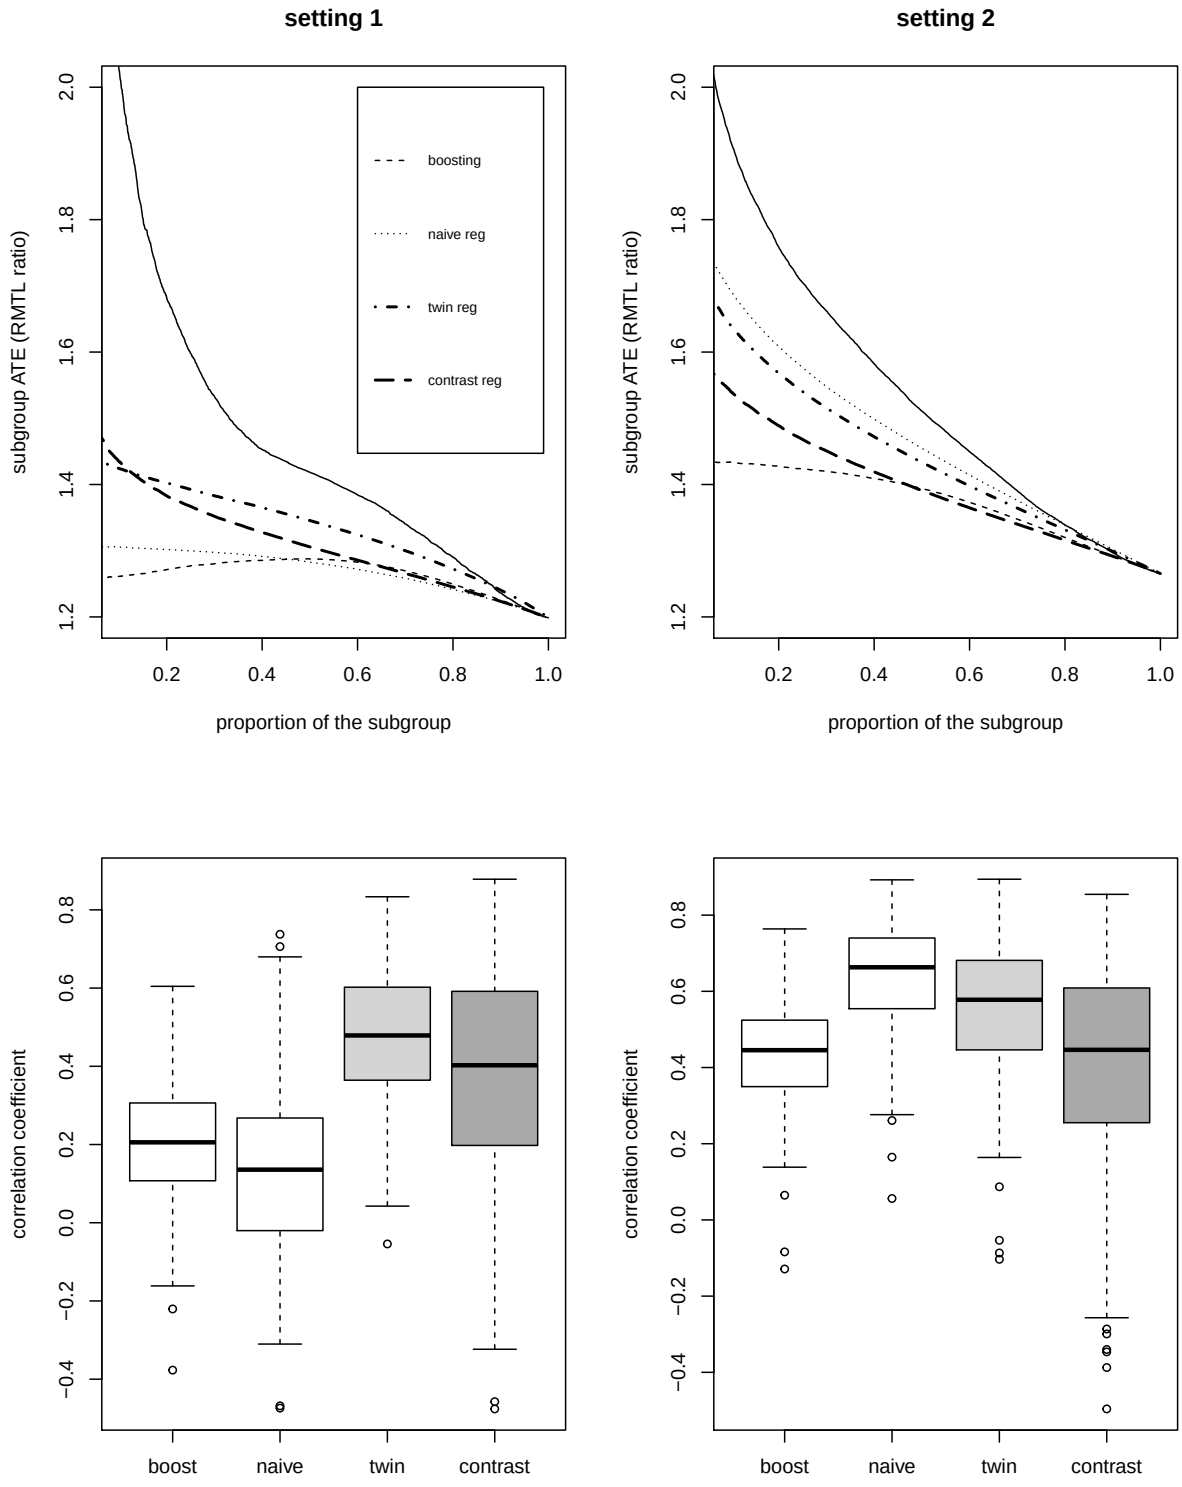

Figure 4: The log-transformed CATE scores based on the standard regression approach and the proposed doubly robust adjustment method: CATE for the ratio of relapse rates

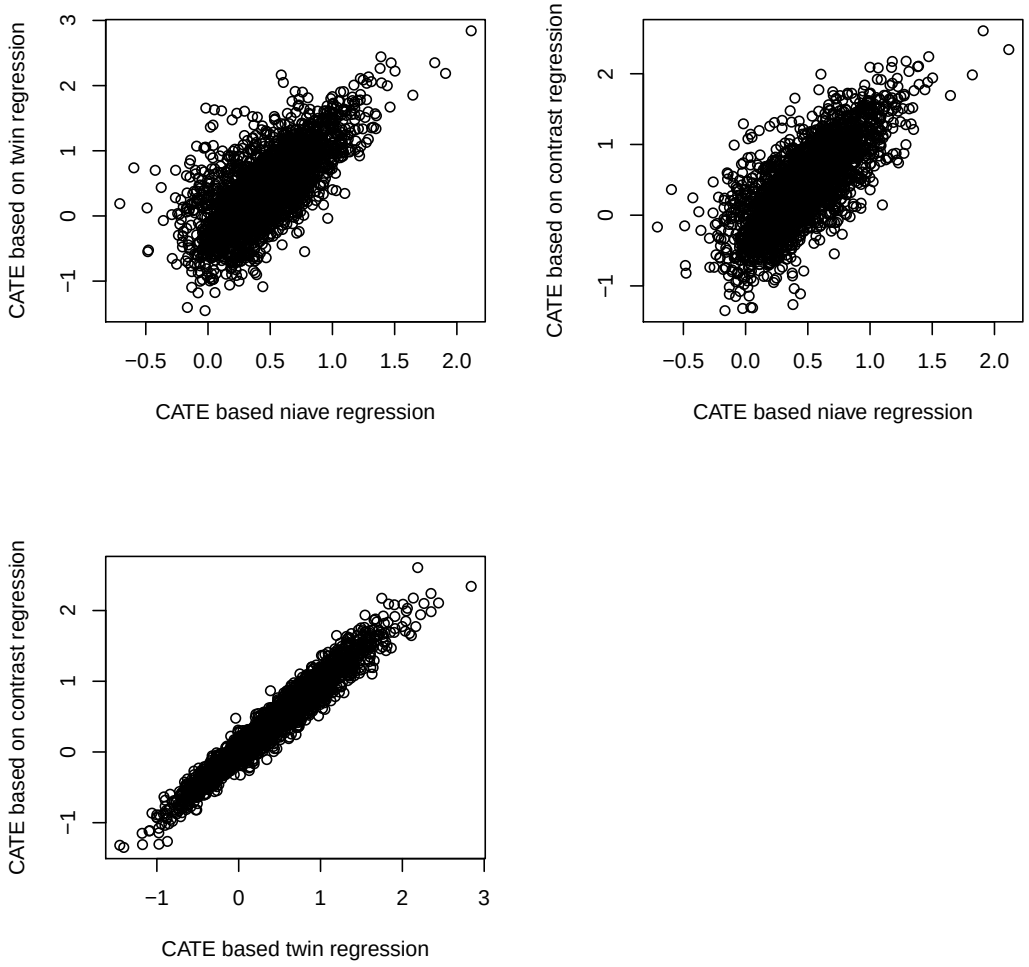

Figure 5: The ATE in subgroups of patients based on the CATE scores constructed in the training set (machine learning, naïve regression and two proposed methods) in the NTD registry: relapse rate ratio of TERI vs DMF for the number of relapses (the first row); RMTL ratio of TERI vs DMF for the time to the first relapse (the second row)

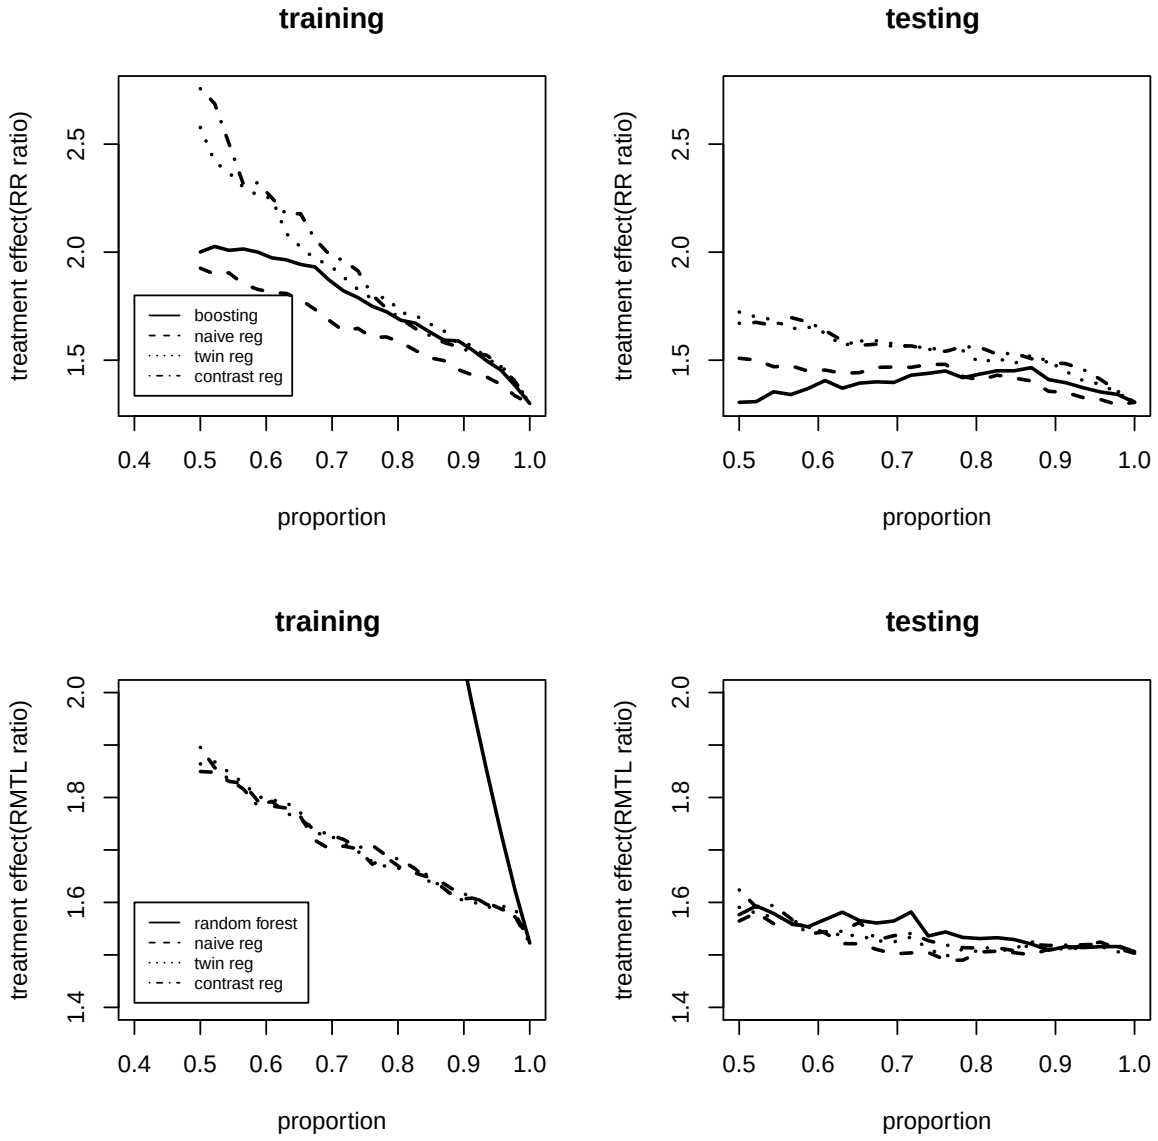

Figure 6: The cross-validated ATE (TERI vs DMF) of subgroups of patients identified by different CATE scores (machine learning, naïve regression and the proposed method) in the NTD registry. relapse rate ratio of TERI vs DMF for the number of relapses (the first row); RMTL ratio of TERI vs DMF for the time to the first relapse (the second row); gray box: the subgroup with higher CATE score; empty box: the subgroup with lower CATE score.

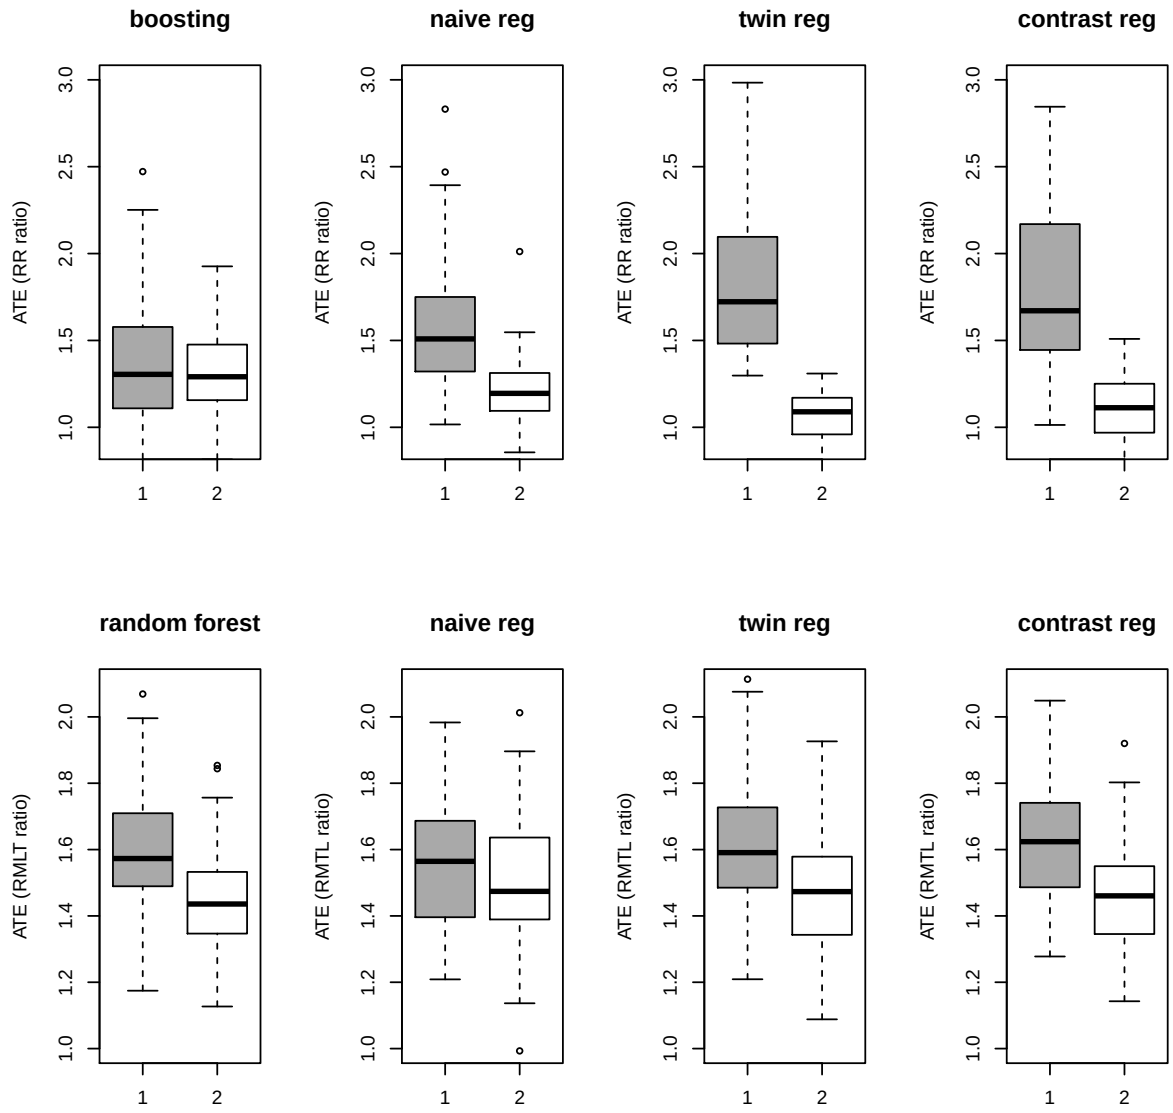

Figure 7: The marginal ORs in nested subgroups according to the size of conditional OR (thick solid curve), marginal ORs in nested subgroups maximizing the marginal ORs consecutively (thick dotted curve), and conditional OR (thin solid curve).