

Group Average Treatment Effects for Observational Studies*

Daniel Jacob[†]

This Draft: March 27, 2020

Abstract

The paper proposes an estimator to make inference of heterogeneous treatment effects sorted by impact groups (GATES) for non-randomised experiments. The groups can be understood as a broader aggregation of the conditional average treatment effect (CATE) where the number of groups is set in advance. In economics, this approach is similar to pre-analysis plans. Observational studies are standard in policy evaluation from labour markets, educational surveys and other empirical studies. To control for a potential selection-bias, we implement a doubly-robust estimator in the first stage. We use machine learning methods to learn the conditional mean functions as well as the propensity score. The group average treatment effect is then estimated via a linear projection model. The linear model is easy to interpret, provides p-values and confidence intervals, and limits the danger of finding spurious heterogeneity due to small subgroups in the CATE. To control for confounding in the linear model, we use Neyman-orthogonal moments to partial out the effect that covariates have on both, the treatment assignment and the outcome. The result is a best linear predictor for effect heterogeneity based on impact groups. We find that our proposed method has lower absolute errors as well as smaller bias than the benchmark doubly-robust estimator. We further introduce a bagging type averaging for the CATE function for each observation to avoid biases through sample splitting. The advantage of the proposed method is a robust linear estimation of heterogeneous group treatment effects in observational studies.

JEL classification: C01, C14, C31, C63

Keywords: *causal inference, machine learning, simulation study, conditional average treatment effect, multiple splitting, sorted group ATE (GATES), doubly-robust estimator*

*Financial support from the Deutsche Forschungsgemeinschaft via the IRTG 1792 “High Dimensional Non Stationary Time Series”, Humboldt-Universität zu Berlin, is gratefully acknowledged. All correspondence may be addressed to the author by e-mail at daniel.jacob@hu-berlin.de.

[†]School of Business and Economics, Humboldt-Universität zu Berlin, Spandauer Straße 1, 10178 Berlin, Germany.

1 Introduction

When evaluating the causal effect of some policy, marketing action or other treatment indicator, it might not be sufficient to report the average treatment effect (ATE). The estimation of heterogeneous effects, e.g. the conditional (on covariates) average treatment effect (CATE), provides further insight into causal mechanisms and helps researchers and practitioners to actively adjust the treatment assignment towards an efficient allocation. The more information in terms of characteristics is provided, the better can heterogeneity be observed. If we have little deterministic information, it might be that heterogeneity effects are overlooked. The trade-off here is that the more covariates datasets have, the more complex they get. We might not know the structural form of our functions of interest or we simply have more covariates than observations. Even if the number of raw covariates is smaller than the number of observations, including quadratic or even higher order interaction increases the number of covariates and can easily exceed the sample size. This is why parametric models are often insufficient when applied on high-dimensional, non-linear datasets [Chernozhukov, Chetverikov, Demirer, Duflo, Hansen, Newey, and Robins \(2018a\)](#).

Under the assumption of approximate sparsity, meaning that out of our rich covariate space only a few are dependent on our variables of interest (the outcome and the treatment assignment), recent methods for treatment effect estimation use machine learning models that have shown to be superior in high-dimensional prediction problems [Hastie, Tibshirani, and Friedman \(2009\)](#). The idea is to learn nuisance functions and regularize the parameter space while making as little assumptions as possible. This is especially helpful when the data does not come from randomised experiments where treatment is randomly assigned to the individuals. In observational studies, self-selection into treatment can arise and introduce a bias that has to be corrected for (i.e. self-selection bias) ([Heckman, Ichimura, Smith, and Todd, 1998](#)). For the ATE one would use the nuisance parameter to orthogonalize the effect that covariates have on both, the treatment assignment and the outcome variable. See [Chernozhukov et al. \(2018a\)](#) for a recent approach, which they call double machine learning.

Methods to estimate the CATE are, among others, the generalized random forest, which builds on the idea of controlling for observed confounders through a tree structure and then estimates the CATE within each final leaf ([Athey, Wager, and Tibshirani, 2019](#)). Another approach is causal boosting, which uses the idea of causal trees as a weak learner ([Powers, Qian, Jung, Schuler, Shah, Hastie, and Tibshirani, 2018](#)). What the aforementioned methods lack, however, is that they are built on tree algorithms and therefore do not allow a flexible estimation of heterogeneous treatment effects in terms of the underlying algorithm (e.g. LASSO or Neural Networks). A recent method called R-learner does provide such flexibility and shows competitive performance in the estimation of the CATE ([Nie and Wager, 2017](#)). Other models, known as meta-learners, decompose the modelling procedure into sub-regression functions, which can be solved using any supervised learning method. This can, for example, be done by a two-model approach (TMA) where a response function (conditional mean) on the

treated and another one on the non-treated observations is estimated. The difference between the two functions can thus be interpreted as the CATE (Künzel, Sekhon, Bickel, and Yu, 2019). This approach falls into the category of indirect estimation since its goal is to minimize the squared error loss based on the outcome variable. A more efficient approach is to directly estimate the treatment effect and regularize based on the treatment effect itself rather than the observed outcome. This leads to a reduced variance of the prediction (Hitsch and Misra, 2018). In order to do this, we first have to estimate a proxy function of the treatment effect. The literature refers to this approach as the treatment effect projection (Hitsch and Misra, 2018) or the modified outcome method (Knaus, Lechner, and Strittmatter, 2018). The idea is to use inverse probability weighting or the doubly-robust estimator as proposed by (Robins and Rotnitzky, 1995). Using the estimates from the two-model approach in combination with inverse probability weighting (IPW) decreases the variance of the estimator and controls for observed confounding (see e.g. Lunceford and Davidian (2004)). Additional orthogonalization using the two conditional mean functions produced by the TMA further decreases the bias of the parameter of interest (Lee, Okui, and Whang, 2017). The doubly-robust estimator can be used in high-dimensional settings to estimate a reduced dimensional conditional average treatment effect function. Functional limit theory can be derived for the case where the nuisance functions are trained via machine learning methods, which are then applied on the doubly-robust estimator (Fan, Hsu, Lieli, and Zhang, 2019; Zimmert and Lechner, 2019). Recent papers study and evaluate different models that are designed for the estimation of heterogeneous treatment effects (see e.g. Knaus (2020); Knaus et al. (2018); Künzel et al. (2019); Powers et al. (2018)).

The difficulty, however, is that machine learning methods are often a black box that is not easy to interpret. This fact hinders the information on drivers for effect heterogeneity. In this paper, we, therefore, build on the ideas of Chernozhukov, Demirer, Duflo, and Fernandez-Val (2018b), who concentrate on estimating group average treatment effects (GATE) in randomised experiments. The groups are built on the distribution from the CATE (e.g. quantiles to get five groups). A semi-parametric model is then used to identify the best linear predictor for the group treatment effect, providing standard errors and confidence intervals. The heterogeneity between the groups can further be interpreted through covariates, which shed some light on the question of what characteristics determine the differences between groups. In this paper, we extend the approach to estimating the GATE parameter towards the use in observational studies and also towards the possibility to estimate a best linear CATE based on group heterogeneity. In a paper that is very close to ours the authors use a linear model with the aforementioned orthogonalization step to estimate group treatment effects based on observed covariates (Park and Kang, 2019). They define the groups by conditioning on observed covariates, for example, for a binary variable they estimate the average treatment effect for the two levels of the variable.

The advantage of the method proposed here is a robust estimation of group heterogeneous treatment effects that is comparable with other models thus keeping its flexibility in the choice of machine learning methods and at the same time its ability to interpret the results. The latter is especially useful in all areas of empirical economics

like policy or labour markets interventions. It also has the advantage to control for potential self-selection bias. The idea of going beyond the average, but not as deep as to estimate conditional average treatment effects for many covariates, is first considered in [Chernozhukov, Fernández-Val, and Luo \(2018c\)](#). They provide standard errors and confidence bands for the estimated sorted group effects and related classification analysis and provide confidence sets for the most and least affected groups. While they only use parametric estimators, a non-parametric attempt to estimate group average treatment effects and also provide insights from the heterogeneity in terms of observed covariates is proposed by [Fan et al. \(2019\)](#) and [Zimmert and Lechner \(2019\)](#). They use a two-step estimator of which the second step consists of a kernel estimator.

Our contribution is to use the doubly-robust estimator and thus keep the flexibility in using any kind of machine learning method to learn the nuisance parameter in the first step. In a second step, we use Neyman-orthogonal scores to cancel out the effect that covariates can have on both, the outcome and the treatment selection. Using the residuals, we set up a linear model to learn the group average treatment effects. When grouping the heterogeneous treatment effect, we find that we can get more precise estimates in terms of mean absolute error (MAE) and Bias with our proposed method. Especially the linearization increases the accuracy of our estimates. We also show that the CATE function, which we estimate via a doubly-robust estimator, should be weighted over multiple iterations. This type of bagging further decreases the MAE of the CATE function. We split our paper in three main parts. First, we state the methodology for randomized experiments and second, the extensions for observational studies. Third, we employ a extensive simulation study that include selection bias, high-dimensionality and non-linearity in the data generating process.

2 Generic Machine Learning for Group ATE

2.1 Potential Outcome Assumptions

Throughout this paper, we make use of the potential outcome theorem ([Rosenbaum and Rubin, 1983](#)) and state four necessary assumptions. The first assumption is the ignorability of treatment, conditional on observed covariates (X), from the two potential outcomes. It is also known as unconfoundedness or simply conditional independence:

$$(Y_i^1, Y_i^0) \perp\!\!\!\perp D_i | X_i. \quad (1)$$

With Y^1 denoting the potential outcome under treatment and Y^0 if not being treated. D is the treatment assignment variable.

The second assumption, the Stable Unit Treatment Value Assumption (SUTVA), guarantees that the potential outcome of an individual is unaffected by changes in the treatment assignment of others. This assumption might be violated if individuals can interact with each other (peer and social effects).

The third assumption, called overlap, guarantees that for all $x \in \text{supp}(X)$, the probability of being in the treatment group (i.e. the propensity score, $e_0(X)$), is bounded away from 0 and 1:

$$\begin{aligned} 0 < P(D = 1|X = x) < 1. \\ e_0(X) = P(D = 1|X = x). \end{aligned} \tag{2}$$

We control for the common support by estimating the propensity score and balance the treatment and control group based on the distribution. We hence exclude all observations that have a propensity score lower 0.02 or higher than 0.95. The fundamental problem of causal inference is that we only observe one of the two potential outcomes at the same time. The counterfactual for a nontreated (treated) person, namely, what would have happened if this person were (not) treated, is always missing. We can represent this statement through a switching regression where the observed outcome (Y_i) depends on the two potential outcomes and the treatment assignment:

$$Y_i = Y_i^0 + D(Y_i^1 - Y_i^0). \tag{3}$$

We further assume that, for the estimation of standard errors, the following moments exist: $E[|Y^j|^q] < \infty$ for $q \geq 4$ and $j = 0, 1$.

2.2 Randomized Control Trial

To provide valid estimation and inference for a causal interpretation of parameters, [Chernozhukov et al. \(2018b\)](#) focus on features of the CATE. One of the main features is the **Sorted Group Average Treatment Effect (GATES)**. The idea is to find groups of observations depending on the estimated treatment effect heterogeneity. Their proposed method relies on a two-model approach in the first step. Here, two response functions are trained separately for the treated and non-treated observations. This approach will be biased if the data sample is from an observational study. In randomized control trials, the difference between the two functions provides an estimate of the treatment effect for every observation. We refer to $S(X)$ as our proxy-predictor:

$$\begin{aligned} \tau(X) &= E[Y|D = 1, X] - E[Y|D = 0, X], \\ S(X) &= g_1(X) - g_0(X). \end{aligned} \tag{4}$$

Here $g_D(X) = E(Y|D, X)$ is the regression model of the outcome variable on X separately for $D \in \{0, 1\}$. The two functions can be estimated with a broad range of supervised machine learning methods. The target parameters are

$$E[\tau(X)|G_k] \quad G_k := \{S(X) \in I_k\}, \quad k = 1, \dots, K, \tag{5}$$

where G is an indicator of a group membership with $I_k = [\ell_{k-1}, \ell_k)$ and ℓ_k is the k/K -quantile of $\{S_i\}_{i \in M}$. Subscript M denotes that these are all out-of-sample predictions. We will for readability not always refer to the sets but always make use of sample splitting when using machine learning methods. The groups are ex-post defined by the predicted score function in the first stage. Figure 2.1 shows an illustration of how the GATE parameter are defined.

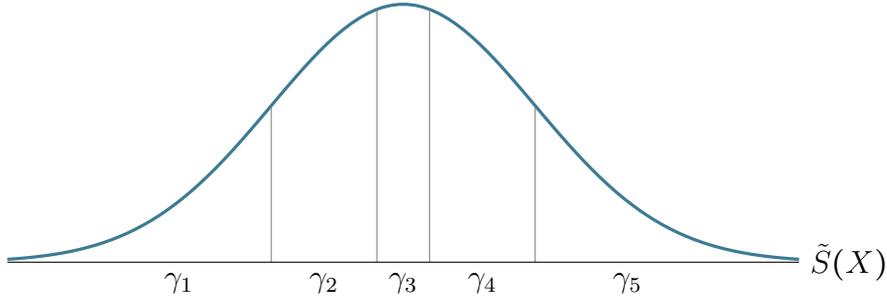


Figure 2.1: Illustration of a treatment effect function and how we define the GATE.

If the treatment effect for the groups are consistent, it asymptotically holds that

$$\mathbb{E}[\tau(X)|G_1] \leq \mathbb{E}[\tau(X)|G_2] \leq \dots \leq \mathbb{E}[\tau(X)|G_k], \quad (6)$$

which is the monotonicity restriction. Furthermore, it can be tested whether there is a homogeneous effect if $\mathbb{E}[\tau(X)|G_k]$ would be equal for all k groups. The weighted linear projection equation to recover the GATES parameter is:

$$YH = \beta^\top A_1 H + \sum_{k=1}^K \gamma_k \cdot \mathbf{I}(S(X) \in I_k) + \nu, \quad (7)$$

with $A_1 = (1, B(X))$ and $B(X) = \mathbb{E}[Y|D=0, X]$ being the baseline function without treatment. See pseudo-code of Algorithm 1, which describes the implementation of this method. The weights H represent the Horvitz-Thompson transformation ([Horvitz and Thompson, 1952](#)):

$$H = H(D, X) = \frac{D - e(X)}{e(X)(1 - e(X))}. \quad (8)$$

This estimator, which is applied to account for different proportions of observations within strata in a target population, is equivalent to the simple inverse probability weighting estimator. These estimators, however, might exhibit a high variance if the identification (the precision) of the propensity scores is lacking ([Lunceford and Davidian, 2004](#)).

The main identification result is that the projection coefficients γ_k can be represented in the following way:

$$\gamma = (\gamma)_{k=1}^K = (E[\tau(X)|G_k])_{k=1}^K. \quad (9)$$

Algorithm 1: GATES

```

1 for  $b=1$  to  $B$  do
2   Split Data in  $k = 2$  samples:  $I^a$  and  $M$  with  $I^a \cup M$ 
3   Train  $Y_i^0 = g_0(X_i, D = 0) + U_{0i}$ , with  $i \in I^a$ 
4   Train  $Y_i^1 = g_1(X_i, D = 1) + U_{1i}$ , with  $i \in I^a$ 
5     Predict  $\hat{Y}_i^0 = \hat{g}_0(X_i)$ , with  $i \in M$ 
6     Predict  $\hat{Y}_i^1 = \hat{g}_1(X_i)$ , with  $i \in M$ 
7     Calculate  $S_b(X|i) = \hat{Y}_i^1 - \hat{Y}_i^0$ 
8   Train  $D_i = e_0(X_i) + V_i$ , with  $i \in I^a$ 
9     Predict  $\hat{D}_i = \hat{e}(X_i)$ , with  $i \in M$ 
10    Calculate  $\hat{V}_i = D_i - \hat{e}(X_i)$ , with  $i \in M$ 
11  Estimate GATES parameters  $(\gamma)$  with weighted OLS using  $M$  (see
    equation 7)
12 end
13 Average  $\gamma$  over  $B$  iterations:  $\tilde{\gamma} = \text{median}\{\gamma\}$ 

```

The algorithm is based on [Chernozhukov et al. \(2018b\)](#).

There are two potential sources of uncertainty. One is estimation uncertainty regarding our parameter of interest, keeping sample splitting fixed. The second source is exactly due to the sample splitting. To account for this, the p-values, as well as the confidence intervals, need to be adjusted. ([Chernozhukov et al., 2018b](#)) show, that sample-splitting-adjusted p-values can have the following form.

$$\mathbb{P}(p_A \leq \alpha/2 | \text{Data}) \geq 1/2, \quad (10)$$

with p_A being the realized p-value given the auxiliary sample and α is the significance level. Given that we use medians to average the parameters γ_k over B bootstrap repetitions it holds that for at least 50% of the random data splits out of B , the p-value is at largest $\alpha/2$. Small values provide evidence that the group parameter is different from zero.

2.3 Extensions for Observational Studies

To use the best linear predictor for group heterogeneity in observational studies, we change and extend the TMA in the first step and the linear model in the second step. First, we replace the two-model approach by a doubly-robust estimator. This means

we not only weight by the inverse of the propensity score but also orthogonalize the outcome variable by subtracting the conditional mean. We also make use of sample splitting as a form of cross-fitting. The auxiliary sample is applied to estimate the score function via the doubly-robust estimator and the main sample to predict the final score function, which is used in the linear step. In this way, we limit the danger of overfitting. The resulting function is a more robust version of the CATE for each individual as well as for the GATE function. The two steps are described in more detail in the following.

The function in equation 12 is calculated using the training data (the I^a sample). In a second step, a new supervised model is trained on the transformed outcome using I^a while predictions are made on the test set M to get an unbiased estimate (see equation 12). Algorithm 2 describes this process.

$$\hat{S}_i = \hat{g}_1(X_i) - \hat{g}_0(X_i) + \frac{D_i(Y_i - \hat{g}_1(X_i))}{\hat{e}(X_i)} - \frac{(1 - D_i)(Y_i - \hat{g}_0(X_i))}{(1 - \hat{e}(X_i))} \quad (11)$$

$$\hat{S}_i = l(X_i) + \omega \quad (12)$$

In equation 11, $\hat{g}_1(X_i) - \hat{g}_0(X_i)$ is equivalent to the score-function from the two-model approach. Simulation evidence from [Knaus et al. \(2018\)](#); [Powers et al. \(2018\)](#) suggests that estimators based on \hat{S}_i might be more stable because of the doubly-robust property and that the performance is competitive for the estimation of heterogeneous treatment effects in observational studies. The doubly-robust property states that the estimator is consistent and unbiased if only one of the models, the regression or the propensity score, is correctly specified ([Robins, Rotnitzky, and Zhao, 1994](#); [Robins and Rotnitzky, 1995](#)). [Lunceford and Davidian \(2004\)](#); [Williamson, Forbes, and White \(2014\)](#); [Belloni, Chernozhukov, and Hansen \(2014\)](#) study the theoretical properties of the doubly-robust estimator and highlight implications for practice. One of the findings is that the variance can be decreased when using the doubly-robust estimator instead of a simple inverse probability estimator ([Lunceford and Davidian, 2004](#)). [Chernozhukov and Semenova \(2018\)](#) show that equation 11 is conditionally locally robust to the estimation error of the nuisance parameter.

Next we state some asymptotic results to recover the CATE. From equation 4 it follows that

$$\tau(X) = E\{E[Y|D = 1, X] - E[Y|D = 0, X]|X = x_i\} \quad (13)$$

Let $\eta(X) := (e(X), g_1(X_i), g_0(X_i))$ be the true high dimensional nuisance parameters. Following [Fan et al. \(2019\)](#) we can define

$$\begin{aligned} \psi(D, Y, X, \eta(X)) = & g_1(X_i) - g_0(X_i) \\ & + \frac{D_i(Y_i - g_1(X_i))}{e_0(X_i)} - \frac{(1 - D_i)(Y_i - g_0(X_i))}{(1 - e_0(X_i))} \end{aligned} \quad (14)$$

Theorem 2.1

(i) under Assumptions 1,2,3,4

$$\begin{aligned} \mathbb{E} \left[g_1(X_i) + \frac{D_i(Y_i - g_1(X_i))}{e_0(X_i)} \middle| X = x_i \right] &= \mathbb{E}[Y^1 | X = x_i], \\ \mathbb{E} \left[g_0(X_i) + \frac{(1 - D_i)(Y_i - g_0(X_i))}{1 - e_0(X_i)} \middle| X = x_i \right] &= \mathbb{E}[Y^0 | X = x_i] \end{aligned}$$

(ii) $\mathbb{E}[\psi(D, Y, X, \eta(X)) - \tau(X) | X = x_i] = 0$ given (i).

This moment condition satisfies the Neyman-orthogonality condition. Neyman-orthogonality is a key component in ensuring that the CATE estimators are robust to the regularization bias inherent for the nuisance functions, which are learned via machine learning models.

Next, we set up a linear model for the estimation of the low-dimensional parameters of interest γ_k . Suppose that the data generating process follows a partial linear model that has the following form:

$$Y = \tau(X)D + \mu_0(X) + U, \quad \mathbb{E}[U | X, D] = 0, \quad (15)$$

$$D = e_0(X) + V, \quad \mathbb{E}[V | X] = 0, \quad (16)$$

$$\mathbb{E}[\tau(X)] = \frac{1}{K} \sum_{k=1}^K \gamma_k(X). \quad (17)$$

We define γ as before as $\gamma = (\gamma)_{k=1}^K = (\mathbb{E}[\tau(X) | G_k])_{k=1}^K$. The outcome variable Y depends not only on the treatment effect parameter but also on observed covariates through the function $\mu_0(X)$. The second equation displays the setting in observational studies, namely that the treatment assignment also depends on observed covariates through the function $e_0(X)$. In randomized control trials, it is sufficient to directly learn the regression function $\hat{\tau}(X)D$. We do not need to include the function $\mu_0(X)$ since in RCT the distribution of the covariates are assumed to be the same for both, the control and the treatment group. The linear model in section 2.2 does exactly this but also takes into account that the treatment assignment might be different for strata in the covariate space.

In observational settings, confounding through covariates leads to a bias that we have to account for. We wish to partial out the effect from X on D as well as the effect from X on Y . Following [Chernozhukov et al. \(2018a\)](#), we again make use of Neyman-orthogonal moments, which leads to the orthogonalized regressors $\hat{U} = Y - \hat{\mu}(X)$ and $\hat{V} = D - \hat{e}(X)$. The terms \hat{U} and \hat{V} are the residuals that we use in the linear projection function in Equation 19. The propensity-score estimates ($\hat{e}(X)$) can be derived by using the main sample on the already estimated propensity-score function, which is used in the doubly-robust step. The function $\hat{\mu}(X)$ is estimated using any machine

learning model on the auxiliary sample. For the estimation of the average treatment effect $\hat{\tau}$, the residualized regression function has the following form:

$$\hat{\tau} = \left(\frac{1}{N} \sum_{i \in M} \hat{V}_i \hat{V}_i \right)^{-1} \frac{1}{N} \sum_{i \in M} \hat{V}_i (Y_i - \hat{\mu}(X_i)). \quad (18)$$

In our case, we are not specifically interested in the average treatment effect but the average effect given the quantiles of the CATE function. This leads to the linear projection equation to estimate the group average treatment effect using the main sample of observations:

$$(Y - \mu_0(X)) = \sum_{k=1}^K \gamma_k \cdot (D - e_0(X)) \cdot \mathbf{I}(S(X) \in I_k) + \nu. \quad (19)$$

Let $\tilde{S}(X) = \hat{l}(X)$ and rewrite the empirical analog for the k -th specific group as:

$$\hat{\gamma}_k = \left(\frac{1}{N} \sum_{i \in M} \hat{V}_i \hat{V}_i \cdot \mathbf{I}(\tilde{S}(X_i) \in I_k) \right)^{-1} \frac{1}{N} \sum_{i \in M} \hat{V}_i (Y_i - \hat{\mu}(X_i)) \cdot \mathbf{I}(\tilde{S}(X_i) \in I_k). \quad (20)$$

Through orthogonalization, we can overcome the regularization bias, present when naively employing ML models to estimate the parameter γ in this setting. Note that we can use different ML algorithms for each function. A trade-off to regularization bias is overfitting, which can be taken care of through sample splitting. However, through sample splitting, we can only use $N - n$ observations, given that n observations are in the training data. To account for the uncertainty through sample splitting, we run multiple repetitions via bootstrapping and split the sample randomly in each bootstrap. We average the GATES parameter by taking the median over B repetitions.

Naturally, we can do the same sort of bagging for the score-function. Given that we split our sample into two equal parts (one for training and one for estimation), we get estimates of the CATE for half of the observations in each iteration. If we repeat this procedure often enough, we get estimates for every observation and decrease the uncertainty bias due to sample splitting. We then take medians for each observation to get the final CATE estimate ($\bar{S}_i(X)$):

$$\bar{S}_i(X) = \text{median}\{\hat{S}_b(X)\}. \quad (21)$$

In our simulation study we show exactly how bagging for CATE can influence the results.

Algorithm 2 shows the steps to identify the group treatment effect and the steps to average the CATE function over all repetitions.

Algorithm 2: Extended GATES

```
1 for  $b=1$  to  $B$  do
2   Split Data in  $k = 2$  samples:  $I^a$  and  $M$  with  $I^a \cup M$ 
3   Train  $Y_i^0 = g_0(X_i, D = 0) + U_{0i}$ , with  $i \in I^a$ 
4   Train  $Y_i^1 = g_1(X_i, D = 1) + U_{1i}$ , with  $i \in I^a$ 
5   Train  $D_i = e_0(X_i) + V_i$ , with  $i \in I^a$ 
6   Train  $Y_i = \mu(X_i) + U_i$ , with  $i \in I^a$ 
7     Predict  $\hat{Y}_i^0 = \hat{g}_0(X_i)$ , with  $i \in I^a$ 
8     Predict  $\hat{Y}_i^1 = \hat{g}_1(X_i)$ , with  $i \in I^a$ 
9     Predict  $\hat{D}_i = \hat{e}(X_i)$ , with  $i \in I^a$ 
10    Predict  $\hat{Y}_i = \hat{\mu}(X_i)$ , with  $i \in M$ 
11    Calculate doubly-robust estimator (see equation 11)
12    Train  $\hat{S}_i = l_0(X_i) + W$  with  $i \in I^a$ 
13    Predict  $\tilde{S}_i = \hat{l}(X_i)$  with  $i \in M$ 
14      Calculate  $\hat{V}_i = D_i - \hat{e}(X_i)$ , with  $i \in M$ 
15      Calculate  $\hat{U}_i = Y_i - \hat{\mu}(X_i)$ , with  $i \in M$ 
16      Store  $S_b^*(X)|i = \tilde{S}_i(X)|i$ 
17    Estimate GATES parameters ( $\gamma$ ) with OLS using  $M$  (see equation 19)
18 end
19 Average  $\gamma$  over  $B$  iterations:  $\tilde{\gamma} = \text{median}\{\gamma\}$ 
20 Calculate Density for every  $i$ :  $\hat{S}(X)|i$  given  $S_b^*(X)|i$  over  $b$ 
21 Calculate Final score-function ( $\bar{S}_i(X)$ ) given density of medians for  $i = 1$  to  $N$ 
```

The number of groups can, of course, be increased to e.g. 10. In empirical settings, it would depend on the sample size. If we want to have at least 30 observations within a group, we could have $\frac{N}{30 \times \Lambda}$ groups, with Λ -splits or folds of the dataset in the first stage used for training and testing. Here we consider only two-folds. However, there is no general relationship between the number of folds in cross-fitting and the precision of the estimator (see Chernozhukov et al. (2018a) for an example with different folds).

2.4 Alternative Grouping: The Baseline Effect

In this paper, we focus on group average treatment effects that are based on quantiles from the CATE function. In the linear model we, therefore, condition on these quantiles to estimate the GATE parameters. The condition on which the sample is split into groups can, of course, be changed. Of particular interest could be the GATE based on different levels of the baseline outcome (Y^0). This has some similarity with the traditional “Quantile Treatment Effect”, which is defined as $\delta(\lambda) = F_{Y^1}^{-1}(\lambda) - F_{Y^0}^{-1}(\lambda)$ $0 < \lambda < 1$, where $F_{Y^D}^{-1}$ is the unconditional quantile function and λ a specific quantile. Here, however, we only focus on Y^0 .

Let us think of Y as being a binary variable, indicating if a customer bought something online or if a person participated in a program (e.g. unemployment training).

We might believe that the treatment effect differs among different baseline probabilities. In such cases, we want to find group treatment effects based on quantiles of Y^0 . As an estimator for Y^0 , we can use the doubly-robust estimator and only concentrate on the baseline potential outcome:

$$\mathbb{E} \left[g_0(X_i) + \frac{(1 - D_i)(Y_i - g_0(X_i))}{1 - e_0(X_i)} \middle| X = x_i \right] = \mathbb{E}[Y^0 | X = x_i] \quad (22)$$

The linear model would then have the following form:

$$(Y - \mu_0(X)) = \sum_{k=1}^K \gamma_k \cdot (D - e_0(X)) \cdot \mathbf{I}(Y^0 \in I_k) + \nu. \quad (23)$$

3 Simulation Study

To evaluate the proposed extensions i) doubly-robust first stage, ii) orthogonal semi-parametric second stage, and iii) bagging of CATE, we use simulated data where the true treatment effects are known. In the following, we describe the data generating process (DGP) and show the variations that we consider. In all simulations we consider five groups denoted by γ . For example, γ_1 is the mean from the lower $\frac{100}{K}\%$ quantile from the CATE function while γ_K is the mean of the upper $\frac{100}{K}\%$ quantile. Our method, which we call DO GATES (we refer to the name DO GATES as for Double Orthogonal GATES) estimates γ_k quantiles from a linear model. As a benchmark model, we use the mean-effect from each of the K -quantiles from the final treatment effect function ($\bar{S}(X)$) estimated via the doubly-robust estimator. Below, we show an example of the distribution from the estimated CATE function. A natural question is how the mean for the K -quantile groups, directly build from the CATE function, behaves without the second linear orthogonalization. To answer this, we compare the estimates from the benchmark and our approach with the mean values of $K = 5$ quantiles from the true treatment effect function. For the non-parametric estimation of functions, we use a random forest method (*R-package: ranger*) as the corresponding machine learning algorithm due to its fast performance.

In this simulation, we are interested in finite sample results and estimate our group treatment effects for $N = 2000$ and again for $N = 500$ observations.

3.1 Data Generating Process

The basic model used in this simulation study is a partially linear regression model based on [Robinson \(1988\)](#):

$$Y = \tau_0(X)D + \mu_0(X) + U, \quad \mathbb{E}[U|X, D] = 0, \quad (24)$$

$$D = e_0(X) + V, \quad \mathbb{E}[V|X] = 0, \quad (25)$$

$$\tau(X) = t_0(X) + W \quad \mathbb{E}[W|X] = 0, \quad (26)$$

with Y being a continuous outcome variable. $\tau_0(X)$ is the true treatment effect or population uplift, while D is the treatment status. The vector $X = (X_1, \dots, X_p)$ consists of p different features, covariates or confounder. U , V and W are unobserved covariates, which follow a random normal distribution $= N(0, 1)$.

Equation 25 is the propensity score. In the case of completely random treatment assignment, the propensity score $e_0(X_i) = c$ for all units ($i = 1, \dots, N$). The scalar c can take any value within the interval $(0, 1)$. In the simulation we consider $c = 0.5$ (balanced) and $c = 0.8$ (imbalanced).

The function $\mu_0(X)$ takes the following form:

$$\mu_0(X) = X_{p/2} + X_{p/10} + X_{p/4} \times X_{p/10}. \quad (27)$$

The vector $b = \frac{1}{l}$ with $l \in \{1, 2, \dots, p\}$ represents weights for every covariate.

In the simulation, we focus on different functions of the **treatment assignment**. We use a CDF to create probabilities, which are then used in a Binomial function to create a binary treatment variable. The dependence of covariates within the normal distribution function is represented as $a(X)$, for which we use a variety of functions, namely random assignment with balanced and imbalanced groups, a linear dependence, interaction terms, and non-linear dependence.

$$e_0(X) = \Phi\left(\frac{a(X) - \mu(a(X))}{\sigma(a(X))}\right), \quad (28)$$

$$D \stackrel{ind.}{\sim} \text{Bernoulli}(e_0(X)) \quad \text{such that} \quad D \in \{0; 1\}. \quad (29)$$

random assignment: $e_0(X) = c$ with $c \in (0, 1)$,

linear: $a(X) = X_2 + X_{p/2} + X_{p/4} - X_8$,

interaction: $a(X) = X \times b + X_{p/2} + X_2 + X_{p/4} \times X_8$,

non-linear: $a(X) = X \times b + \sin(X_{p/2}) + X_2 + \cos(X_{p/4} \times X_8)$.

The **treatment effect** is heterogeneous and generated with a linear or a non-linear dependency of X :

$$\text{linear: } \tau_u(X) = X_1 + X_2 > 0 + W \quad \text{with} \quad W \sim N(0, 0.5), \quad (30)$$

$$\text{non-linear: } \tau_u(X) = \sin(X \times b) + X_{5+p/2}. \quad (31)$$

We standardise the treatment effect within the interval $[0.1, 1]$:

$$\tau_0(X) = \frac{\tau_u(X) - \min(\tau_u(X))}{\max(\tau_u(X)) - \min(\tau_u(X))} (1 - 0.1) + 0.1. \quad (32)$$

Our aim is to perform a Monte Carlo simulation, in which we fix the distribution and dependence of the covariates (correlation). This is necessary to also fix the treatment

effect function for each setting. All other functions can have small deviations within each repetition due to the error term and the random generation for the binomial distribution (for the treatment assignment). The idea is to gather samples to approximate the distribution of the true treatment effect. Figure A.1 in the Appendix shows the true treatment effect function for the linear and non-linear case. They are quite normally distributed without heavy tails.

3.2 Simulation Results

In Table 3.1 and 3.2, we show the different scenarios, specifically for the treatment assignment mechanism. We also consider a simulation where both, the $\hat{e}(X)$ function as well as $\hat{g}_D(X_i)$ (and as a consequence also $\hat{\mu}(X)$) is misspecified. We model misspecification by introducing an unobserved confounder in the DGP (namely X_2), which we exclude in the observed dataset. In scenarios *A* to *F*, the data generation process for the treatment effect depends linear on the covariates, whereas in scenarios *G* to *L* the dependency is non-linear (see equations 30 and 31). In Table 3.1 we consider $N = 2000$ observations. In Table 3.2, we set $N = 500$. All other processes are the same. We repeat every DGP 100 times and estimate the mean absolute error (MAE) over all groups and repetitions. Another performance measurement is the squared bias as a comparison between the expected function and the true function. We aggregate our measurement as follows:

$$MAE_{kj} = |\hat{\gamma}_k - \gamma_k|, \quad (33)$$

$$MAE_k = \frac{1}{S} \sum_{j=1}^S (MAE_{kj}), \quad (34)$$

$$MAE = \frac{1}{K} \sum_{k=1}^K (MAE_k). \quad (35)$$

$$Bias^2(\hat{\gamma}_k) = \left(\underbrace{\mathbb{E}[\hat{\gamma}_k]}_{\frac{1}{S} \sum_{j=1}^S \hat{\gamma}_{kj}} - \gamma_k \right)^2, \quad (36)$$

$$Bias^2 = \frac{1}{K} \sum_{k=1}^K (Bias^2(\hat{\gamma}_k)). \quad (37)$$

In Figure 3.1 and 3.2 we plot the results for each group of the simulation study for $N = 2000$ and 500 observations, respectively. We find that especially the lower quantiles show a smaller absolute error of the quantile treatment effect. If the treatment assignment mechanism is randomized, our model performs competitively or slightly worse than the doubly-robust estimator, as can be seen from scenario A, B, G and H. In all observational settings, we find a smaller mean absolute error from our DO GATES method compared to the benchmark model. We find that the highest error from the true treatment effect arises when both models are misspecified as in scenario

Table 3.1: Settings and Monte Carlo averages for $N = 2000$.

Scenarios	A/G	B/H	C/I	D/J	E/K	F/L
N	2000	2000	2000	2000	2000	2000
\mathbb{R}^p	20	20	20	20	20	20
$P(D = 1)$	0.5	0.2 (imbalanced)	linear	interaction	non-linear	linear
Misspecification	No	No	No	No	No	Yes
MAE CATE	0.06/0.05	0.15/0.14	0.62/0.61	0.66/0.67	0.61/0.61	0.91/0.92
MAE DO GATES	0.09/0.10	0.08/0.10	0.33/0.32	0.32/0.36	0.32/0.31	0.75/0.76
$Bias^2$ CATE	0.00/0.00	0.03/0.03	0.47/0.44	0.49/0.50	0.46/0.44	0.99/1.00
$Bias^2$ DO GATES	0.01/0.00	0.00/0.00	0.15/0.15	0.17/0.18	0.15/0.13	0.74/0.75

Notes: Mean absolute error (MAE) of the average treatment effect (estimated from the quantile estimates) over $S = 100$ Monte Carlo simulations. Setting G to L consists of a non-linear treatment effect function.

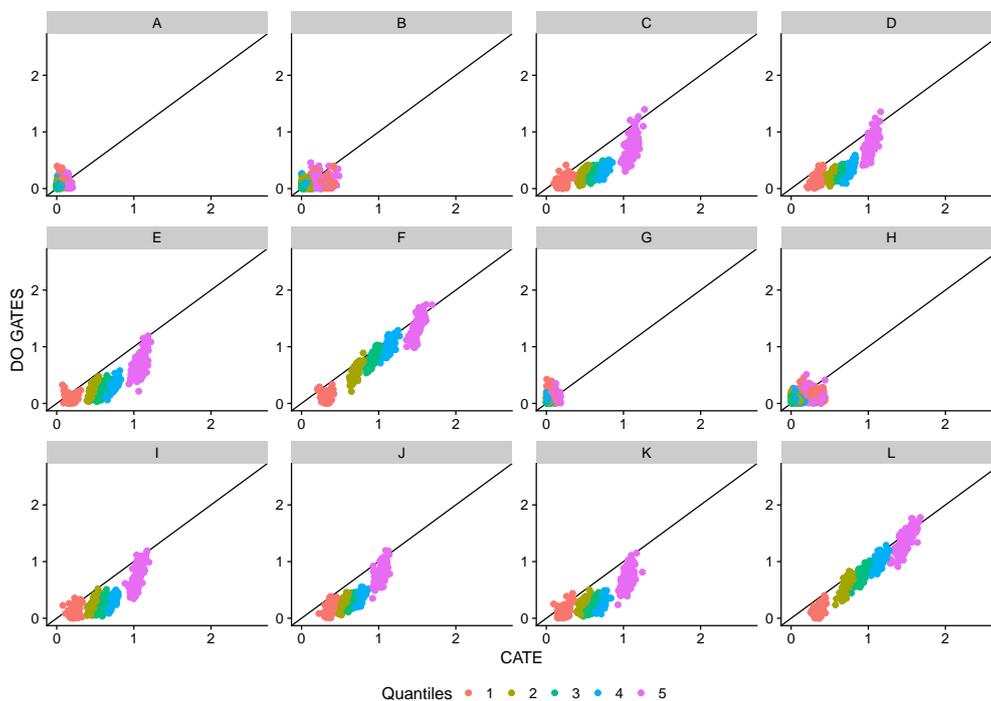


Figure 3.1: Group estimate from the CATE function and DO GATES. Groups are coloured by quantiles (1 = least affected, 5 = most affected). Axes show the absolute error between estimates and true quantile treatment effects. The 45-degree line indicates the equality of both methods. The number of observations equals 2000.

F and L. There is, however, no significant difference in the MAE if we change the dependency of the treatment effect from linear to non-linear.

The bias behaves like the MAE in the sense that only for the balanced randomization the two models show equal bias. In all other settings, we find a smaller bias for our method and again the highest differences when there is selection-bias, which depends linear and or with interactions on the treatment. Table B.1 and B.2 in the Appendix

show the MAE for each group and DGP over all repetitions. We also notice that the deviation between the two methods is highest for the lowest group and decreases towards the upper quantiles. When we set $N = 500$, we find that the MAE and the BIAS increase in both methods but the overall structure stays the same as for $N = 2000$.

Table 3.2: Settings and Monte Carlo averages for $N = 500$.

Scenarios	A/G	B/H	C/I	D/J	E/K	F/L
N	500	500	500	500	500	500
\mathbb{R}^p	20	20	20	20	20	20
$P(D = 1)$	0.5	0.2 (imbalanced)	linear	interaction	non-linear	linear
Misspecification	No	No	No	No	No	Yes
MAE CATE	0.12/0.14	0.16/0.20	1.13/1.10	1.22/1.21	1.11/1.10	1.44/1.41
MAE DO GATES	0.19/0.16	0.21/ 0.19	0.64/0.63	0.71/0.69	0.62/0.64	1.13/1.10
$Bias^2$ CATE	0.02/0.02	0.03/0.05	1.47/1.38	1.60/1.54	1.43/1.40	2.34/2.29
$Bias^2$ DO GATES	0.03/0.02	0.02/0.00	0.67/0.65	0.60/0.59	0.63/0.66	1.67/1.69

Notes: Mean absolute error (MAE) of the average treatment effect (estimated from the quantile estimates) over $S = 100$ Monte Carlo simulations. Setting G to L consist of a non-linear treatment effect function.

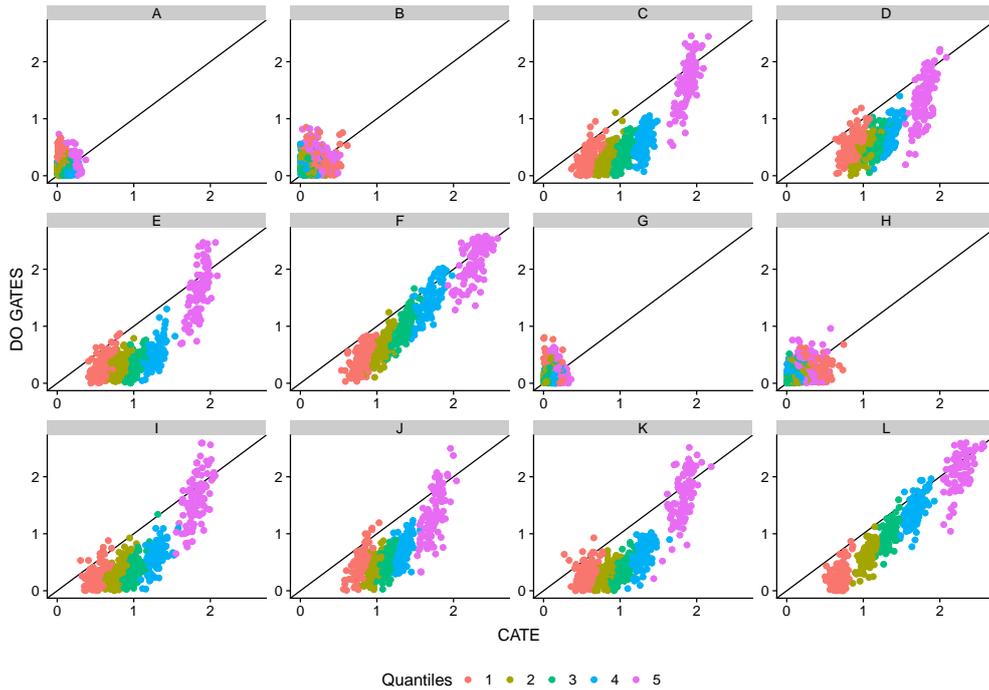


Figure 3.2: Group estimate from the CATE function and DO GATES. Groups are coloured by quantiles (1 = least affected, 5 = most affected). Axes show the absolute error between estimates and true quantile treatment effects. The 45-degree line indicates the equality of both methods. The number of observations equals 500.

3.3 Model Averaging of CATE

Figure 3.3 shows the estimates from the conditional average treatment effect over $B = 100$ bootstrap iterations (we split the sample once in equal parts). For every iteration, we either get an estimate for a specific observation i if this observation is in the test-sample or we don't which we label with a "NA" estimate. After 100 iterations we have at least 39 estimates for each observation. If the specific sample we use for training would always be equal in terms of distributions, we would assume that every iteration estimates the same value for a specific observation. Here we show that this is not the case since the sample splitting does matter, especially in finite sample settings.

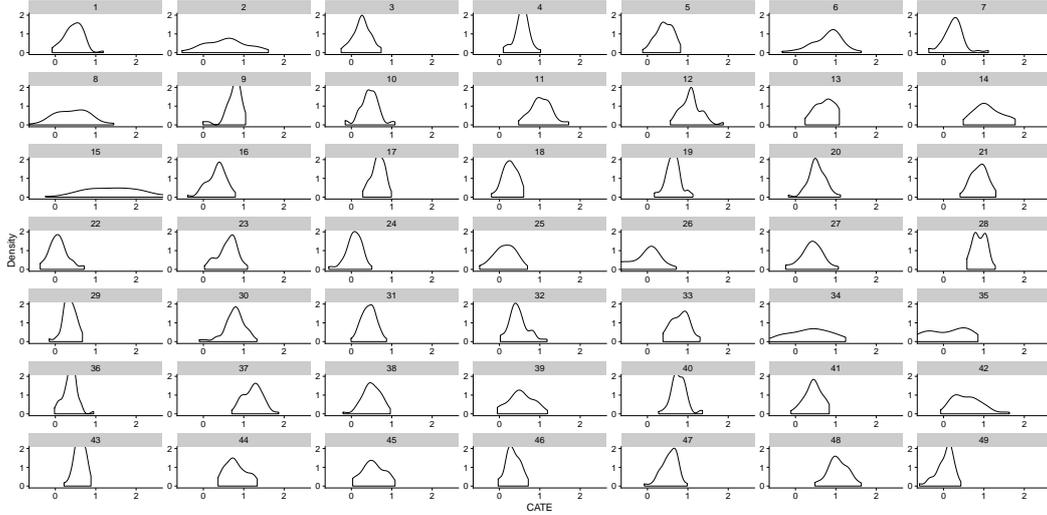


Figure 3.3: Distribution of score function ($\hat{S}_b(X)$) for each of 49 randomly selected individuals. Density is estimated over 100 iterations. Due to random sample splitting, the number of estimates for each individual is between 39 and 59 for 100 iterations.

The simulated data, in this case, has the following properties. $N = 1000$, $X = \mathbb{R}^{20}$, $e_0(X) = 0.5$ and $\tau_u(X) = X_1 + X_2 > 0 + W$ with $W \sim N(0, 0.5)$. We standardize τ_u to $\tau(X) \in [0.1, 1]$. The densities for 49 selected observations show that even in randomised experiments, the point estimates differ due to the sample splitting in the first step. We propose to average the estimates by taking the median over each iteration. This leads to a more stable conditional treatment effect function:

$$\bar{S}_i(X) = \text{median}\{\hat{S}_b(X)\} \quad (38)$$

We show in Figure 3.4 how the absolute error (AE) for each group depends on the number of iterations. The upper plot shows the AE for the CATE function and the lower plot for the DO GATES method. The AE for the CATE function stabilizes after around 50 iterations while for the DO GATES we can decrease the AE further as we

take the median over more iterations. Interestingly, it seems that the variance in AE for the DO GATES method does not mainly depend on the CATE function as we can see after taking the median of 50 iterations. The CATE function is quite stable but the GATES still show some variance in the AE. It could be that this is caused in the orthogonalization step where we train the two conditional mean functions. We also try to average with the arithmetic mean and find no significant difference between the two approaches. However, it could be that in some scenarios, like a very small sample size, single estimates are far away from the true estimate and should be treated as outliers. Therefore, we suggest to use the median for averaging.

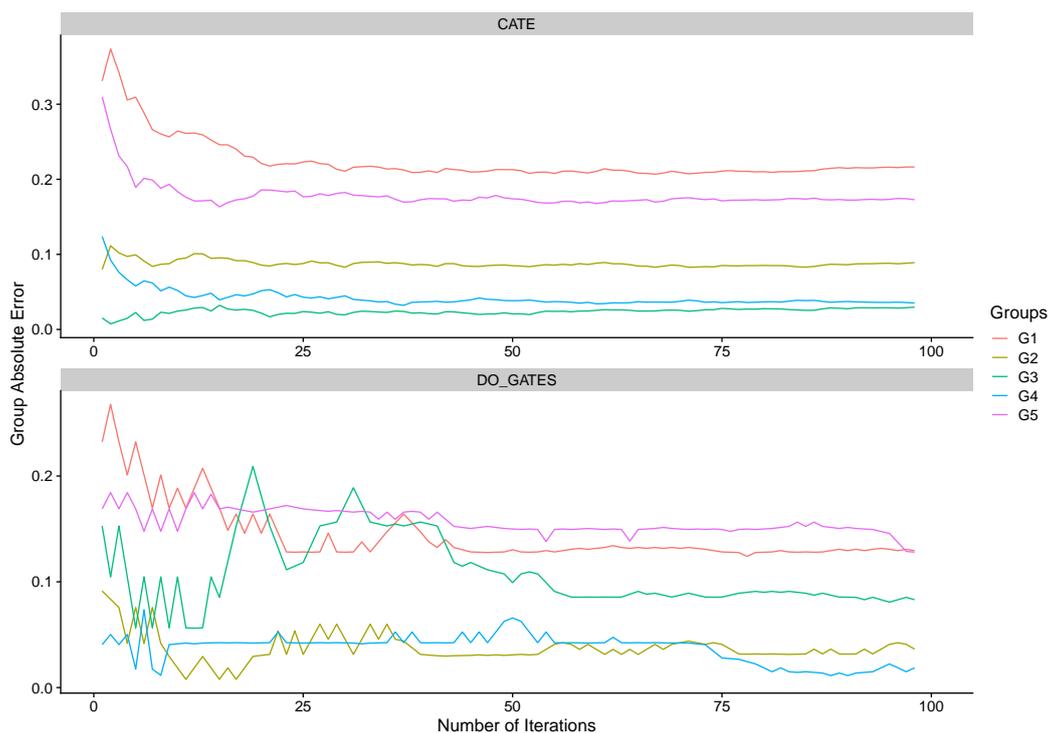


Figure 3.4: Upper plot shows the absolute error for the CATE estimator given B iterations. Lower plot shows the same for the DO GATES method.

4 Conclusion

In this paper, we propose a method to estimate group average treatment effects using machine learning methods and linear estimation for non-randomized control trials. Since flexibility in terms of the model choice, as well as interpretability of the results, is our main interest, we extend the idea of the GATES approach by using a doubly-robust estimator in the first, the non-parametric, step. In the second step, the linear projection function, we use Neyman-orthogonal moments to overcome the regularization bias due to the dependency of the covariates on the outcome and the treatment assignment parameter. This ensures to control for self-selection into treatment which is a realistic challenge in observational studies. Our main interest is to find heterogeneity in terms of quantiles from the treatment effect function and estimate them via a best linear projection. We also propose to weight the CATE function by taking the median over B iterations instead of single estimates. Since we are interested in describing the groups in terms of all observations, random sample splitting in half of the sample only gives us estimates from half of the observations. After approximately $B = 10$ iterations we have at least one estimate for each observation. Therefore, we suggest to use at least 50 or preferably 100 iterations in order to cover the whole sample.

We find that our proposed model has a lower MAE and a lower Bias for the majority of the groups we consider and among different simulation scenarios, compared to the benchmark model. In randomized control trials, both methods perform competitive. As soon as we include selection-bias, the DO GATES method performs better due to its second orthogonalization step. We note that a natural extension is to estimate the ATE by taking the mean over the group estimates. However, this estimator is not the scope of the paper and there are direct approaches to estimate the ATE.

References

- Susan Athey, Stefan Wager, and Julie Tibshirani. Generalized random forests. *The Annals of Statistics*, 47(2):1148–1178, 2019.
- Alexandre Belloni, Victor Chernozhukov, and Christian Hansen. Inference on treatment effects after selection among high-dimensional controls. *The Review of Economic Studies*, 81(2):608–650, 2014.
- Victor Chernozhukov and Vira Semenova. Simultaneous inference for best linear predictor of the conditional average treatment effect and other structural functions. (CWP40/18), July 2018.
- Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 2018a.
- Victor Chernozhukov, Mert Demirer, Esther Duflo, and Ivan Fernandez-Val. Generic machine learning inference on heterogeneous treatment effects in randomized experiments. ArXiv preprint, <https://arxiv.org/abs/1712.04802>, 2018b.
- Victor Chernozhukov, Iván Fernández-Val, and Ye Luo. The sorted effects method: Discovering heterogeneous effects beyond their averages. *Econometrica*, 86(6):1911–1938, 2018c.
- Qingliang Fan, Yu-Chin Hsu, Robert P Lieli, and Yichong Zhang. Estimation of conditional average treatment effects with high-dimensional data. ArXiv preprint, <https://arxiv.org/abs/1908.02399>, 2019.
- Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The elements of statistical learning: data mining, inference, and prediction*, Springer Series in Statistics. Springer New York, 2009.
- James Heckman, Hidehiko Ichimura, Jeffrey Smith, and Petra Todd. Characterizing selection bias using experimental data. Technical report, National bureau of economic research, 1998.
- Günter J Hitsch and Sanjog Misra. Heterogeneous treatment effects and optimal targeting policy evaluation. Available at SSRN 3111957, 2018.
- Daniel G Horvitz and Donovan J Thompson. A generalization of sampling without replacement from a finite universe. *Journal of the American Statistical Association*, 47(260):663–685, 1952.
- Michael C Knaus. Double machine learning based program evaluation under unconfoundedness. ArXiv preprint, <https://arxiv.org/abs/2003.03191>, 2020.

- Michael C Knaus, Michael Lechner, and Anthony Strittmatter. Machine learning estimation of heterogeneous causal effects: Empirical monte carlo evidence. *CEPR Discussion Paper No. DP13402*, 2018.
- Sören R Künzel, Jasjeet S Sekhon, Peter J Bickel, and Bin Yu. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the National Academy of Sciences*, 116(10):4156–4165, 2019.
- Sokbae Lee, Ryo Okui, and Yoon-Jae Whang. Doubly robust uniform confidence band for the conditional average treatment effect function. *Journal of Applied Econometrics*, 32(7):1207–1225, 2017.
- Jared K Lunceford and Marie Davidian. Stratification and weighting via the propensity score in estimation of causal treatment effects: A comparative study. *Statistics in Medicine*, 23(19):2937–2960, 2004.
- Xinkun Nie and Stefan Wager. Quasi-oracle estimation of heterogeneous treatment effects. ArXiv preprint, <https://arxiv.org/abs/1712.04912>, 2017.
- Chan Park and Hyunseung Kang. A groupwise approach for inferring heterogeneous treatment effects in causal inference. ArXiv preprint, <https://arxiv.org/abs/1908.04427>, 2019.
- Scott Powers, Junyang Qian, Kenneth Jung, Alejandro Schuler, Nigam H Shah, Trevor Hastie, and Robert Tibshirani. Some methods for heterogeneous treatment effect estimation in high dimensions. *Statistics in Medicine*, 37(11):1767–1787, 2018.
- James M Robins and Andrea Rotnitzky. Semiparametric efficiency in multivariate regression models with missing data. *Journal of the American Statistical Association*, 90(429):122–129, 1995.
- James M Robins, Andrea Rotnitzky, and Lue Ping Zhao. Estimation of regression coefficients when some regressors are not always observed. *Journal of the American Statistical Association*, 89(427):846–866, 1994.
- Peter M Robinson. Root-n-consistent semiparametric regression. *Econometrica: Journal of the Econometric Society*, pages 931–954, 1988.
- Paul R Rosenbaum and Donald B Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 1983.
- Elizabeth J Williamson, Andrew Forbes, and Ian R White. Variance reduction in randomised trials by inverse probability weighting using the propensity score. *Statistics in Medicine*, 33(5):721–737, 2014.
- Michael Zimmert and Michael Lechner. Nonparametric estimation of causal heterogeneity under high-dimensional confounding. ArXiv preprint, <https://arxiv.org/abs/1908.08779>, 2019.

A Figures

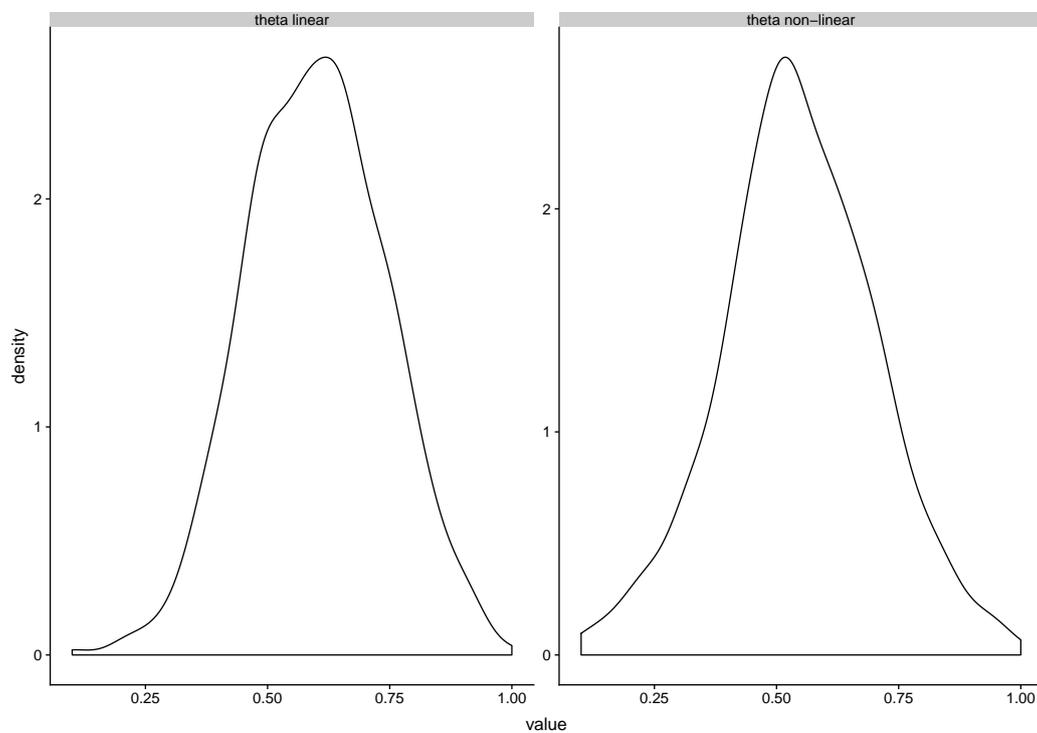


Figure A.1: Density for the true and fixed heterogeneous treatment effect function. Left plot shows the linear treatment effect dependency and right plot the non-linear dependency.

B Tables

Table B.1: Mean absolute error by setting and groups for $N = 2000$.

SETTING	G	DO GATES	CATE	SETTING	G	DO GATES	CATE
A	1	0.18	0.07	G	1	0.18	0.07
A	2	0.07	0.03	G	2	0.08	0.02
A	3	0.04	0.02	G	3	0.05	0.02
A	4	0.07	0.03	G	4	0.07	0.04
A	5	0.12	0.13	G	5	0.14	0.12
B	1	0.12	0.27	H	1	0.15	0.25
B	2	0.07	0.07	H	2	0.08	0.07
B	3	0.06	0.03	H	3	0.06	0.04
B	4	0.07	0.11	H	4	0.08	0.09
B	5	0.12	0.29	H	5	0.14	0.27
C	1	0.11	0.21	I	1	0.12	0.23
C	2	0.25	0.49	I	2	0.23	0.48
C	3	0.25	0.61	I	3	0.23	0.59
C	4	0.30	0.74	I	4	0.29	0.71
C	5	0.73	1.10	I	5	0.73	1.05
D	1	0.16	0.34	J	1	0.16	0.37
D	2	0.23	0.53	J	2	0.25	0.55
D	3	0.27	0.64	J	3	0.30	0.66
D	4	0.35	0.75	J	4	0.35	0.75
D	5	0.78	1.03	J	5	0.77	1.03
E	1	0.09	0.18	K	1	0.12	0.23
E	2	0.25	0.48	K	2	0.24	0.48
E	3	0.26	0.60	K	3	0.23	0.59
E	4	0.30	0.72	K	4	0.27	0.70
E	5	0.73	1.08	K	5	0.70	1.05
F	1	0.10	0.32	L	1	0.11	0.36
F	2	0.54	0.70	L	2	0.56	0.72
F	3	0.78	0.92	L	3	0.78	0.93
F	4	1.00	1.14	L	4	0.99	1.13
F	5	1.36	1.50	L	5	1.36	1.49

Notes: Settings as in Table 3.1. Number of observations: $N = 2000$. Averages over 100 repetitions.

Table B.2: Mean absolute error by setting and groups for N = 500.

SETTING	G	DO GATES	CATE	SETTING	G	DO GATES	CATE
A	1	0.34	0.08	G	1	0.27	0.14
A	2	0.16	0.05	G	2	0.16	0.06
A	3	0.12	0.09	G	3	0.12	0.10
A	4	0.11	0.14	G	4	0.09	0.15
A	5	0.21	0.23	G	5	0.19	0.23
B	1	0.29	0.29	H	1	0.25	0.39
B	2	0.18	0.10	H	2	0.15	0.13
B	3	0.13	0.07	H	3	0.14	0.07
B	4	0.17	0.10	H	4	0.16	0.13
B	5	0.30	0.25	H	5	0.28	0.29
C	1	0.31	0.60	I	1	0.27	0.52
C	2	0.31	0.83	I	2	0.31	0.80
C	3	0.37	1.04	I	3	0.37	1.01
C	4	0.60	1.31	I	4	0.59	1.29
C	5	1.64	1.88	I	5	1.63	1.82
D	1	0.55	0.85	J	1	0.47	0.79
D	2	0.44	1.00	J	2	0.43	0.98
D	3	0.53	1.14	J	3	0.50	1.13
D	4	0.73	1.35	J	4	0.73	1.34
D	5	1.30	1.78	J	5	1.35	1.74
E	1	0.31	0.58	K	1	0.29	0.55
E	2	0.31	0.81	K	2	0.29	0.82
E	3	0.33	1.02	K	3	0.37	1.03
E	4	0.53	1.29	K	4	0.62	1.30
E	5	1.62	1.86	K	5	1.64	1.82
F	1	0.38	0.79	L	1	0.23	0.69
F	2	0.65	1.09	L	2	0.62	1.06
F	3	0.97	1.35	L	3	0.99	1.33
F	4	1.42	1.68	L	4	1.43	1.68
F	5	2.21	2.30	L	5	2.25	2.31

Notes: Settings as in Table 3.2. Number of observations: N = 500. Averages over 100 repetitions.

C Technical Appendix

C.1 Correlated Covariates

We show how we generate the covariates and their correlation. We assume correlation through a uniform distribution of the covariance matrix, which is then transformed to a correlation matrix. Correlated characteristics are more common in real datasets and helps to investigate the performance of ML algorithms, especially the regularization bias, in a more realistic manner. Figure C.1 shows the correlation matrix for 10 randomly selected covariates to give an example of the correlation.

Algorithm 3: Generation of Covariates

- 1 **Generate** random positive definite covariance matrix Σ based on a uniform distribution over the space $k \times k$ of the correlation matrix.
 - 2 **Scale covariance matrix** This equals the correlation matrix and can be seen as the covariance matrix of the standardized random variables $\Sigma = \frac{X}{\sigma(X)}$.
 - 3 **Generate** random normal distributed variables $X_{N \times k}$ with mean = 0 and variance = Σ .
-

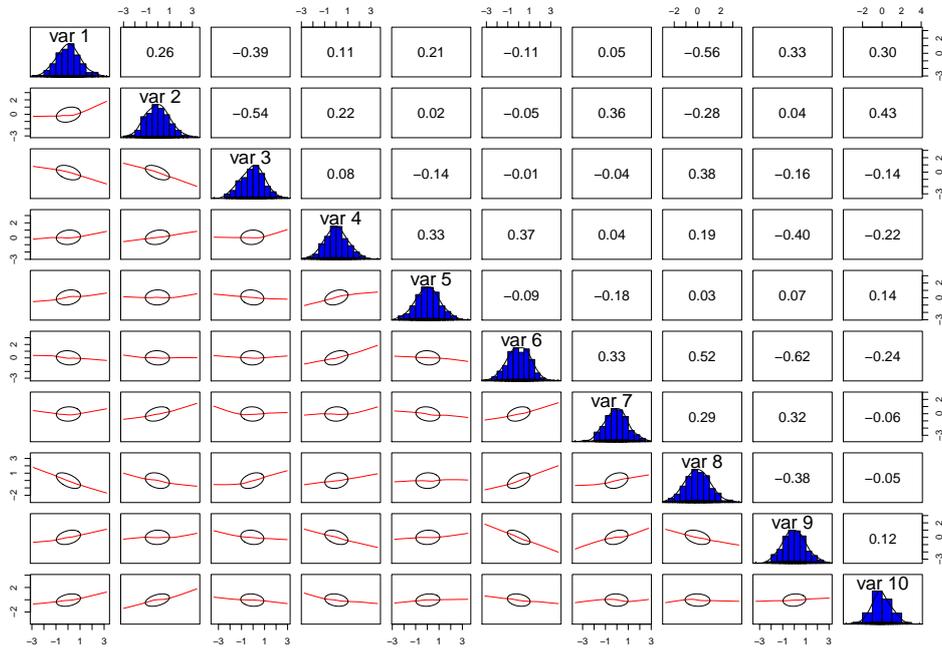


Figure C.1: Correlation Matrix of Covariates. Correlation metric is Bravais-Pearson.