

# Learning Optimal Prescriptive Trees from Observational Data

Nathanael Jo

University of Southern California, Los Angeles, CA 90089, nathanael.jo@gmail.com

Sina Aghaei

University of Southern California, Los Angeles, CA 90089, saghaei@usc.edu

Andrés Gómez

University of Southern California, Los Angeles, CA 90089, gomezand@usc.edu

Phebe Vayanos

University of Southern California, Los Angeles, CA 90089, phebe.vayanos@usc.edu

We consider the problem of learning an optimal prescriptive tree (i.e., an interpretable treatment assignment policy in the form of a binary tree) of moderate depth, from observational data. This problem arises in numerous socially important domains such as public health and personalized medicine, where interpretable and data-driven interventions are sought based on data gathered in deployment – through passive collection of data – rather than from randomized trials. We propose a method for learning optimal prescriptive trees using mixed-integer optimization (MIO) technology. We show that under mild conditions our method is asymptotically exact in the sense that it converges to an optimal out-of-sample treatment assignment policy as the number of historical data samples tends to infinity. Contrary to existing literature, our approach: 1) does not require data to be randomized, 2) does not impose stringent assumptions on the learned trees, and 3) has the ability to model domain specific constraints. Through extensive computational experiments, we demonstrate that our asymptotic guarantees translate to significant performance improvements in finite samples, as well as showcase our uniquely flexible modeling power by incorporating budget and fairness constraints.

*Key words:* prescriptive trees, causal inference, interpretability, observational data, mixed-integer optimization.

---

## 1. Introduction

Prescriptive analytics is concerned with determining the best treatment for an individual based on their personal characteristics. This problem arises in a variety of domains, from optimizing online advertisements for different users (Li et al. 2010) to assigning suitable therapies for patients suffering from a medical condition (Flume et al. 2007). In this paper, we aim to design such personalized policies based on data collected in deployment rather than during randomized trials. In such settings, historical treatments are not assigned at random but rather based on a (potentially unknown) policy (e.g., social workers or doctors triaging clients/patients based on their personal

characteristics or symptoms). Being able to design personalized treatment assignment policies without running randomized trials is extremely important since such trials are often considered unethical, impractical, or may take too long to conduct.

We are particularly motivated by decision- and policy-making in high-stakes domains – for example, to assign scarce housing resources to those experiencing homelessness (Azizi et al. 2018); or to match patients to scarce organs for transplantation (Bertsimas et al. 2013). In such contexts (contrary to, say, online advertising), policies are designed *offline* and need to be (a) highly *interpretable*, so that system stakeholders (e.g., policy-makers and populations on which the system is deployed) can audit and scrutinize the decision process; and (b) *optimal*, so as to optimize outcomes on the population among models in a given class. We take the view that one should opt for simpler models that are themselves interpretable rather than learning black-box models and then explaining their decisions – the latter of which has a burgeoning literature, see e.g., Linardatos et al. (2020). There are, admittedly, trade-offs between designing and learning interpretable models that maximize performance versus justifying a black-box model’s decisions ex-post; we refer the interested reader to Rudin (2019) for a thorough discussion.

With these needs in mind, we focus our attention on the design of *optimal prescriptive trees*. A prescriptive tree takes the form of a binary tree. In each branching node of the tree, a binary test is performed on a feature. Two branches emanate from each branching node, with each branch representing the outcome of the test. If a datapoint passes (resp. fails) the test, it is directed to the left (resp. right) branch. A treatment is assigned to all leaf nodes. Thus, each path from root to leaf represents a treatment assignment rule that assigns the same treatment to all datapoints that reach that leaf. Similar to decision trees in the context of machine learning (Rudin 2019), prescriptive trees are some of the most interpretable models. We also note that we focus on binary trees in line with most works on decision trees (e.g., the C4.5 and CART (Breiman 2017) algorithms). In general, non-binary trees can be modeled using binary trees with more branching decisions, and there is increasingly work being done on non-binary trees (see e.g., Mazumder et al. (2022)).

The goal of learning prescriptive trees is to select branching decisions and treatment assignments such that the expected value of the outcomes of the treatments on the population are maximized. We say that a prescriptive tree is *optimal* (in-sample, given a training set) if there exists a mathematical proof that no other tree yields better expected outcomes in the population used for training the method. Accordingly, we will say that a tree is *optimal out-of-sample* if there exists a proof that no other tree yields better expected outcomes in the entire population.

### 1.1. Problem Statement

We now formalize the problem we study. In our discussion, we use terminology and concepts from Hernán and Robins (2019). The goal is to design a personalized policy  $\pi : \mathcal{X} \rightarrow \mathcal{K}$  in the form of a

prescriptive tree that maps an individual’s characteristics  $x \in \mathcal{X} \subseteq \mathbb{R}^F$  to a treatment from a finite set of candidates indexed in the set  $\mathcal{K}$ . We note that the set  $\mathcal{K}$  may also include a “no treatment” option. Each individual is characterized by their covariates  $X \in \mathbb{R}^F$  and their *potential outcomes*  $Y(k) \in \mathcal{Y} \subseteq \mathbb{R}$  under each treatment  $k \in \mathcal{K}$ . The joint distribution of  $X$  and  $\{Y(k)\}_{k \in \mathcal{K}}$  is unknown. However, we have access to  $I$  i.i.d. historical observations  $\mathcal{D} := \{(X_i, K_i, Y_i)\}_{i \in \mathcal{I}}$  indexed in the set  $\mathcal{I} := \{1, \dots, I\}$ , where  $X_i$  denote the covariates of the  $i$ th observation,  $K_i$  is the treatment assigned to it, and  $Y_i = Y_i(K_i)$  is the *observed outcome*, i.e., the outcome under the treatment received. We use the convention that higher values of  $Y_i$  are more desirable. Critically, as in all observational experiments, we cannot control the historical treatment assignment policy and the outcomes  $Y_i(k)$  for  $k \neq K_i$  remain *unobserved*. Formally, our aim is to learn, from the observational data  $\mathcal{D}$ , a policy  $\pi \in \Pi_d$  that maximizes the quantity

$$Q(\pi) := \mathbb{E}[Y(\pi(X))],$$

where  $\Pi_d$  denotes the set of prescriptive trees of maximum depth  $d$ , and  $\mathbb{E}(\cdot)$  denotes the expectation operator with respect to the joint distribution  $\mathbb{P}$  of  $X, K, Y(1), \dots, Y(k)$ . A major challenge in learning the best policy  $\pi$  is the fact that we cannot observe the *counterfactual* outcomes  $Y_i(k)$ ,  $k \in \mathcal{K}$ ,  $k \neq K_i$  that were not received by datapoint  $i$ . Indeed, the missing data make it difficult to identify the best possible treatment for each datapoint. In fact, without further assumptions on the historical policy, it is impossible to even identify the treatment that is the best *in expectation* for any given datapoint.

## 1.2. Background on Causal Inference

In this section, we provide a brief overview of the tools from causal inference that motivate the assumptions we make in our work and our proposed methods.

**1.2.1. Identifiability, Randomized Experiments, and Observational Studies.** Learning an optimal prescriptive tree from the observed data requires that the expected values of the counterfactual outcomes be *identifiable*, i.e., possible to be expressed as a function of the observed data. A sufficient condition for identifiability is for the data to have been collected during a (*marginally*) *randomized* or *conditionally randomized* experiment, see Hernán and Robins (2019).

In a randomized experiment, all individuals have the same chance of getting any given treatment independent of their characteristics  $x$ . Formally, we have  $K \perp\!\!\!\perp X$ . When treatment assignment is randomized, we say that treatment groups are *exchangeable* because it is irrelevant which group received a particular treatment. Exchangeability means that for all  $k$  and  $k' \in \mathcal{K}$ , those who got treatment  $k$  in the data would have the same outcome distribution under  $k'$  as those who got

treatment  $k'$ . Equivalently, exchangeability means that the counterfactual outcome and the actual treatment are independent, or  $Y(k) \perp\!\!\!\perp K$ , for all values of  $k$ .

In a conditionally randomized experiment, treatment assignment probabilities depend on the characteristics  $x$  of each individual. A conditionally randomized experiment will, in general, not result in exchangeability because treatment groups may differ systemically from one another. However, treatment groups are exchangeable *conditionally on  $X$*  since no information other than  $X$  is used to decide the treatment. Thus, conditional randomization ensures *conditional exchangeability*, i.e.,  $Y(k) \perp\!\!\!\perp K \mid X$ , for all values of  $k$ . We can then infer the outcome distribution under treatment  $k'$  for those with covariates  $x$  that got treatment  $k$  in the data by looking at their counterparts with the same (or similar)  $x$  that received treatment  $k'$ .

As mentioned previously, in high-stakes domains we typically only have access to data from an *observational study*. In such contexts, it is common practice in the literature on causal inference to assume that treatment is assigned at random conditional on  $X$ , although this may be an approximation and not possible to check in practice. The observational study can then be viewed as a conditionally randomized experiment, provided: *a)* interventions are well defined and are recorded as distinct treatment values in the data; *b)* the conditional probability of receiving a treatment depends only on the covariates  $X$ ; and *c)* the probability of receiving any given treatment conditional on  $X$  is positive, i.e.,  $\mathbb{P}(K = k \mid X = x) > 0$  almost surely for all  $k$ . This last condition is termed *positivity* and holds in practice for randomized and conditionally randomized experiments.

We now discuss several methods from the literature for evaluating the performance of a *counterfactual* treatment assignment policy  $\pi$  from (conditionally) randomized experiments or, more generally, from observational studies that satisfy the above conditions. In particular, in the following and throughout the paper, we assume that treatments in the data have been assigned according to a (potentially unknown) *logging policy*  $\mu$ , where  $\mu(k, x) := \mathbb{P}(K = k \mid X = x)$  and  $\mu(k, x) > 0$  for all  $k \in \mathcal{K}$  and  $x \in \mathcal{X}$ .

**1.2.2. Inverse Propensity Weighting.** The performance of a counterfactual treatment assignment policy  $\pi$  can be evaluated using Inverse Probability Weighting (IPW).

IPW was originally proposed to estimate causal quantities such as expected values of counterfactual outcomes, average treatment effects, risk ratios, etc., see Horvitz and Thompson (1952). It relies on reweighting the outcome for each individual  $i$  in the dataset by the inverse of their *propensity score*, given by  $\mu(K_i, X_i)$ . Such reweighting has the effect of creating a *pseudo-population* where all individuals in the data are hypothetically given all treatments (thus simulating a randomized treatment assignment). Doing so allows for the estimation of the distribution of the unobserved counterfactual outcomes for all possible values of  $X$ .

Importantly, IPW can also be used to estimate the performance of a counterfactual policy  $\pi$  by reweighting each individual  $i$  in the data by  $\frac{\mathbb{1}(\pi(X_i)=K_i)}{\mu(K_i, X_i)}$  as illustrated in Figure 1, see also Bottou et al. (2013). This reweighting simulates the performance of the policy  $\pi$  by leveraging the conditional exchangeability property of the data. Using this estimator, the quantity  $Q(\pi)$  is estimated from the data as

$$Q_I^{\text{IPW}}(\pi) := \frac{1}{I} \sum_{i \in \mathcal{I}} \frac{\mathbb{1}(\pi(X_i) = K_i)}{\hat{\mu}(K_i, X_i)} Y_i, \quad (1)$$

where  $\hat{\mu}$  is an estimator of  $\mu$ , which is obtained for instance using machine learning, by fitting a model to  $\{(X_i, K_i)\}_{i \in \mathcal{I}}$ . If  $\mu$  is known or if  $\hat{\mu}$  converges almost surely to  $\mu$ , the IPW estimator of  $Q(\pi)$  is statistically consistent. However, it may suffer from high variance, particularly if some of the propensity scores  $\mu(K, X)$  are small, see Dudík et al. (2011).

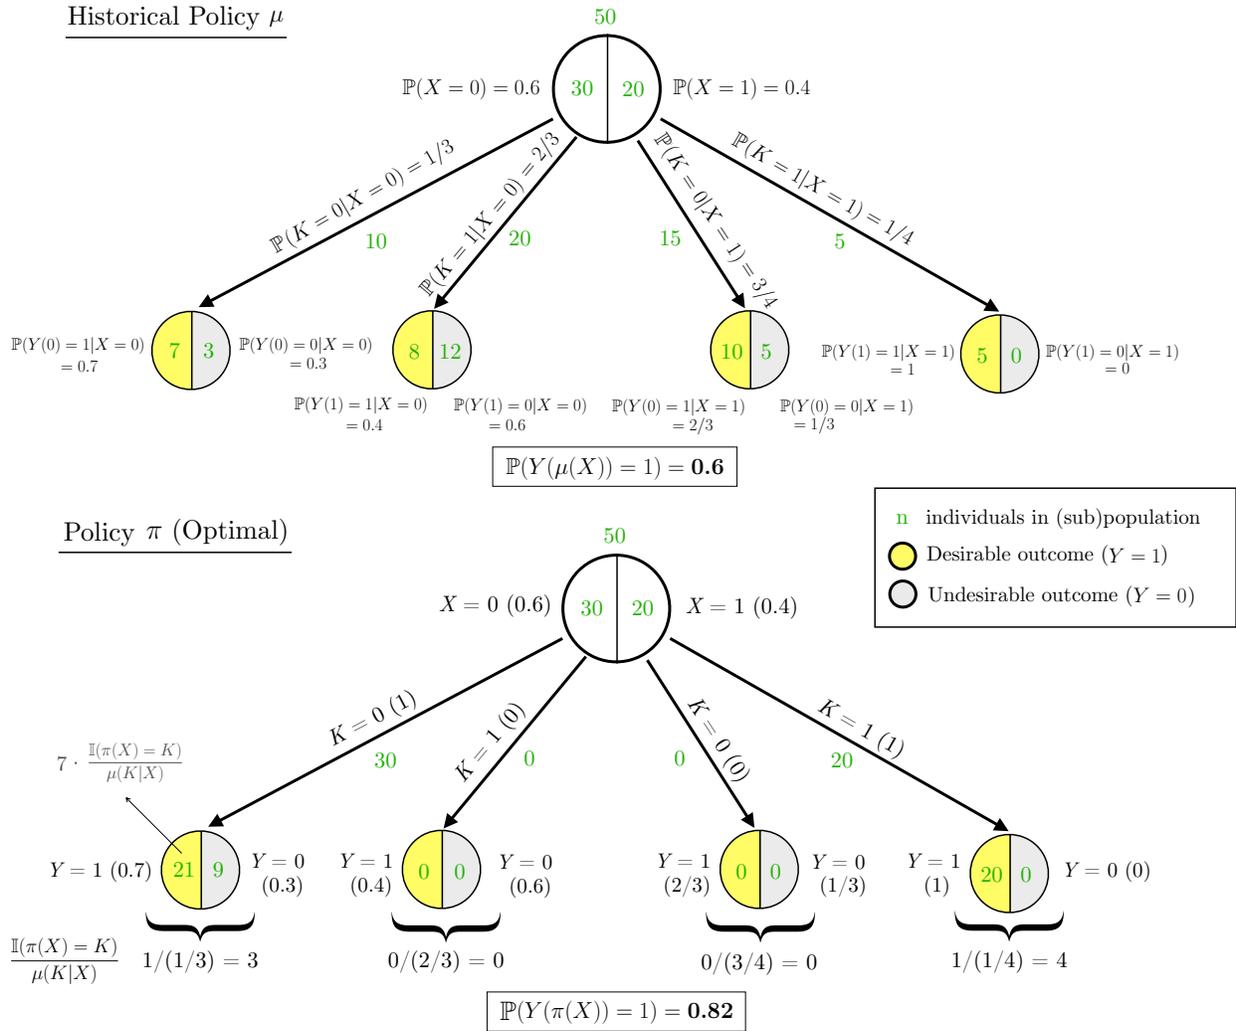
**1.2.3. Direct Method.** An alternative estimator of the performance of the counterfactual policy  $\pi$  is the Direct Method (DM).

DM is based upon the so-called Regress and Compare (R&C) approach, which proceeds in three steps to design a policy that maximizes the expected value of the outcomes (Kallus 2018). First, it partitions the dataset  $\mathcal{D}$  by treatments. Second, for all  $k \in \mathcal{K}$ , it learns a model  $\hat{\nu}_k(X)$  of  $\mathbb{E}(Y|K = k, X)$  using the subpopulation that was assigned treatment  $k$ . Third, given any  $X \in \mathbb{R}^F$ , it estimates  $Y(k)$  as  $\hat{\nu}_k(X)$  and assigns the best treatment according to  $\hat{\nu}$ , that is,  $\pi^{\text{R\&C}}(X) = \arg \max_{k \in \mathcal{K}} \hat{\nu}_k(X)$ . R&C has been used, for example in the online setting by Bastani and Bayati (2020) and Qian and Murphy (2011). Usually, variants of linear regression are used to get estimators  $\hat{\nu}$ , see Bastani and Bayati (2020), Goldenshluger and Zeevi (2013), Li et al. (2010).

The R&C approach, however, does not allow the evaluation of arbitrary policies nor does it enable the design of policies that belong to an arbitrary class. DM adapts the regress and compare framework to resolve these shortcomings. It proposes to estimate  $Q(\pi)$  via

$$Q_I^{\text{DM}}(\pi) := \frac{1}{I} \sum_{i \in \mathcal{I}} \hat{\nu}_{\pi(X_i)}(X_i). \quad (2)$$

Both R&C and DM result in poor performance if  $\hat{\nu}_k(X)$ ,  $k \in \mathcal{K}$ , are biased estimators or make predictions that are far from the true expected outcomes, see Beygelzimer and Langford (2009). As pointed out by Dudík et al. (2011),  $\hat{\nu}_k(X)$  are constructed without information about  $\pi$  and thus may approximate  $\mathbb{E}(Y|K = k, X)$  poorly in areas that are relevant for  $Q(\pi)$ . That said, if  $\hat{\nu}_k$  are statistically consistent for all  $k \in \mathcal{K}$ , the DM estimator of  $Q(\pi)$  will also be consistent.



**Figure 1** Figure illustrating the evaluation of the performance of the counterfactual policy  $\pi$  (bottom tree) from data collected using the logging policy  $\mu$  (top tree) using the IPW estimator—notation conventions based on Hernán and Robins (2019). The historical data presents 50 instances, each with a single binary covariate  $X^1 \in \{0, 1\}$ .  $X^1 = 1$  (resp. 0) indicates that the patient is sick (resp. healthy). There are two treatment options ( $\mathcal{K} := \{0, 1\}$ ), where  $K = 1$  (0) is to (not) treat the patient. The potential outcomes are binary ( $Y(k) \in \{0, 1\}$ ,  $k \in \mathcal{K}$ ). The expected reward under policy  $\mu$  is 0.6, while policy  $\pi$  has a 0.82 probability of getting a positive outcome—in fact, it can be shown that  $\pi$  is optimal.

**1.2.4. Doubly Robust Approach.** Doubly robust (DR) estimation is a family of techniques that combine two estimators. If either one of the two estimators is accurate, then the doubly robust approach is also accurate, which dampens the errors brought by the individual estimators.

In particular, realizing the drawbacks of IPW and DM, Dudík et al. (2011) proposed a doubly robust approach that estimates the counterfactual performance of a policy  $\pi$  as

$$Q_I^{\text{DR}}(\pi) := \frac{1}{I} \sum_{i \in \mathcal{I}} \left( \hat{v}_{\pi(X_i)}(X_i) + (Y_i - \hat{v}_{K_i}(X_i)) \frac{\mathbb{1}(\pi(X_i) = K_i)}{\hat{\mu}(K_i, X_i)} \right). \quad (3)$$

Intuitively, DR uses IPW to estimate the difference between the true outcome and the prediction made by the DM estimator,  $Y_i - \hat{\nu}_{K_i}(X_i)$ , thus removing any bias caused by DM, assuming the IPW estimator is statistically consistent. Provided at least one of  $\hat{\mu}$  or  $\hat{\nu}$  converges almost surely to  $\mu$  or  $\nu$ , respectively, the DR estimator of  $Q(\pi)$  will be asymptotically consistent. Furthermore, in practice,  $Q_I^{\text{DR}}$  usually has a smaller variance compared to  $Q_I^{\text{IPW}}$  but a higher variance than  $Q_I^{\text{DM}}$ , see Dudík et al. (2011).

### 1.3. Related Works

In this section, we position our work in the fields of operations research, causal inference, and machine learning.

**1.3.1. Estimating Heterogeneous Causal Effects.** To estimate heterogeneous causal effects, Athey and Imbens (2016) and Wager and Athey (2018) use recursive partitioning to produce causal trees and causal forests, respectively. Non-tree-based methods can be found in Abrevaya et al. (2015), which use inverse propensity weighting, and Fan et al. (2020), which extends Abrevaya et al. (2015) to account for high-dimensional data. Powers et al. (2018) adapt three methods—random forests, boosting, and MARS (Multivariate Adaptive Regression Splines)—to also estimate treatment effects in high-dimensional data. The overall goal of these approaches is to find good estimates of the treatment effects rather than to design treatment assignment policies. While some of these, such as the method from Athey and Imbens (2016), can be *adapted* to design tree policies, these policies will not be interpretable due to the need to combine the predictive trees associated with many different treatments.

**1.3.2. Mixed-Integer Optimization.** MIO has been gaining traction as a tool to solve challenging learning problems. It has been used for example to tackle sparse regression problems (Wilson and Sahinidis 2017, Atamtürk and Gómez 2019, Bertsimas et al. 2020b, Hazimeh et al. 2020, Xie and Deng 2020, Gómez and Prokopyev 2021), verification of neural networks (Fischetti and Jo 2018, Khalil et al. 2018, Tjandraatmadja et al. 2020), and sparse principal component analysis (Dey et al. 2018, Bertsimas et al. 2020a), among others. More importantly in the context of this paper, MIO methods have been proposed to learn optimal decision trees (Bertsimas and Dunn 2017, Verwer and Zhang 2019, Aghaei et al. 2019, 2020, Elmachtoub et al. 2020, Mišić 2020).

MIO technology has also been used in causal inference. Existing approaches have mostly focused on matching (Bennett et al. 2020, Zubizarreta and Keele 2017), i.e., to estimate causal effects by pairing each treatment instance with a control instance that has similar covariates (Rubin 2006). A related approach to matching that utilizes MIO technology is subset selection to achieve *balance* (i.e., similar distributions) between the covariates in the treatment and control groups (Nikolaev et al. 2013). Finally, Mintz et al. (2017) use MIO to estimate the effects of weight-loss interventions in myopic agents and design optimal policies from the perspective of a healthcare provider.

**1.3.3. Learning Prescriptive Trees.** In a recent work closely related to this paper, Kallus (2017) proposes an MIO formulation to learn optimal prescriptive trees that relies on a critical assumption: that trees are sufficiently deep such that the historical treatment assignment and covariates are independent *at each leaf*. Building upon this approach, Bertsimas et al. (2019) seek to ameliorate it by augmenting the objective function with a term aimed to improve accuracy of the outcome predictions. The paper uses coordinate descent with multiple starts to train the prescriptive trees (as opposed to MIO technology), and empirically verifies that the resulting solutions are indeed high-quality. We provide an in-depth analysis of both methods in Section 2 and compare to them in our experiments.

#### 1.4. Contributions

In this work, we do an in-depth analysis of existing methods to learn prescriptive trees, and propose new MIO formulations for solving such problems. In particular, our key contributions are:

- (a) We analyze the assumption made by the existing methods to learn prescriptive trees discussed in Section 1.3.3. We demonstrate, by means of examples, that when the assumption fails to hold *for all feasible trees*, the trees returned as “optimal” by these methods may be severely suboptimal. They may even branch on noise covariates that are not predictive of outcomes nor treatment assignments. Moreover, our examples show that this assumption is very strict in the sense that trees deep enough to allow branching on all levels of all covariates (including covariates that are independent of outcomes and treatment assignments) may be needed for this assumption to hold at an optimal tree.
- (b) Motivated by the limitations of existing approaches, we propose novel MIO formulations for learning optimal prescriptive trees based on tools from causal inference. Contrary to methods from the literature, our approaches enable the design of shallow and/or deep trees, thus allowing the decision-maker to tune the trade-off between prescription accuracy and interpretability. We demonstrate that, under mild conditions, our methods are asymptotically exact in the sense that they converge to an optimal out-of-sample treatment assignment policy as the number of historical data samples tends to infinity. Based on extensive computational experiments on both synthetic and real data, we demonstrate that our asymptotic guarantees translate to significant out-of-sample performance improvements even in finite samples.

The remainder of this paper is organized as follows. In Section 2, we present the existing prescriptive tree formulations from the literature and study their performance in the context of two simple synthetic examples. We introduce our proposed MIO formulations for learning optimal prescriptive tree in Section 3. Finally, we summarize our computational experiments in Section 4. All proofs and detailed computational results are provided in the Electronic Companion.

## 2. Existing Prescriptive Tree Formulations

This section presents the methods introduced by Kallus (2017) and Bertsimas et al. (2019) – from hereon, we refer to these as K-PT and B-PT, respectively. It further analyzes their advantages and drawbacks using two simple examples.

To describe these two approaches, we work with the notation introduced in Kallus (2017), which views any tree  $\pi$  as producing a partition  $\mathcal{X} = \bigcup_{n \in \mathcal{L}^\pi} \mathcal{X}_n^\pi$  such that  $\mathcal{X}_n^\pi \cap \mathcal{X}_m^\pi = \emptyset$  whenever  $n \neq m$ , where  $\mathcal{L}^\pi$  collects the set of leaf nodes of  $\pi$  and  $\mathcal{X}_n^\pi \subseteq \mathcal{X}$  represents the set of covariate values that are associated with datapoints landing in leaf  $n \in \mathcal{L}^\pi$ . The approaches K-PT and B-PT rely on the following *crucial* assumption about the structure of all feasible prescriptive trees.

**Assumption 1 (Kallus (2017))** *All prescriptive trees  $\pi \in \Pi_d$  produce a partition  $\{\mathcal{X}_n^\pi\}_{n \in \mathcal{L}^\pi}$  that is sufficiently fine such that  $K \perp\!\!\!\perp X \mid X \in \mathcal{X}_n^\pi \quad \forall n \in \mathcal{L}^\pi$ .*

Assumption 1 is implicitly made in Kallus (2017) when the result in Corollary 3 is used to obtain the MIO formulation. This assumption is automatically satisfied in randomized experiments (since  $K \perp\!\!\!\perp X$  unconditionally) or in conditionally randomized experiments provided the partition is “sufficiently fine” – for example, if all feasible trees are such that  $\mathcal{X}_n^\pi$  is a singleton for all  $n \in \mathcal{L}^\pi$ .

### 2.1. K-PT

Given a policy  $\pi$  with partition  $\{\mathcal{X}_n\}_{n \in \mathcal{L}}$  satisfying the “sufficiently fine” condition in Assumption 1, Kallus (2017) proves that the *observable* quantity  $\mathbb{E}(Y|K = k, X \in \mathcal{X}_n)$  is an unbiased estimator of  $\mathbb{E}(Y(k)|X \in \mathcal{X}_n)$ . In other words, to estimate the potential outcomes under treatment  $k$  at leaf  $n$ , it suffices to average outcomes for those that actually received treatment  $k$  at leaf  $n$  under  $\mathbb{P}$ ; this quantity can be estimated from the observed data. This statement is intuitive under Assumption 1, which implies that all individuals with covariates in the set  $\mathcal{X}_n$  have the same chance of getting any given treatment, i.e., they are exchangeable within  $\mathcal{X}_n$ . It follows that, under Assumption 1, the outcomes of a policy  $\pi$  that assigns treatment  $\pi(\mathcal{X}_n) \in \mathcal{K}$  to all individuals that land in leaf  $n$  can be estimated as

$$\sum_{n \in \mathcal{L}} \mathbb{P}(X \in \mathcal{X}_n) \cdot \mathbb{E}(Y|K = \pi(\mathcal{X}_n), X \in \mathcal{X}_n). \quad (4)$$

Accordingly, this quantity can be estimated from data as

$$Q_I^{\text{K-PT}}(\pi) := \sum_{n \in \mathcal{L}} \frac{\sum_{i \in \mathcal{I}} \mathbb{1}[X_i \in \mathcal{X}_n]}{|\mathcal{I}|} \widehat{Y}_n(\pi(\mathcal{X}_n)),$$

where  $\widehat{Y}_n(k)$  is the sample average approximation of  $\mathbb{E}(Y|K = k, X \in \mathcal{X}_n)$ , defined through

$$\widehat{Y}_n(k) := \frac{1}{\sum_{i \in \mathcal{I}} \mathbb{1}[K_i = k, X_i \in \mathcal{X}_n]} \sum_{i \in \mathcal{I}} \mathbb{1}[K_i = k, X_i \in \mathcal{X}_n] Y_i. \quad (5)$$

**Analysis of K-PT.** Kallus (2017) proved that  $Q_I^{\text{K-PT}}(\pi)$  is an unbiased estimator of  $Q(\pi)$  when Assumption 1 holds, which is the case in *marginally* randomized settings for example. A salient advantage of K-PT is that, if Assumption 1 *does* hold, it avoids predicting  $\mu$  and thus avoids the bias caused by potentially inaccurate estimates  $\hat{\mu}$ . Unfortunately, though, in *conditionally* randomized settings, decision trees will in general not satisfy Assumption 1 and the tree obtained by maximizing  $Q_I^{\text{K-PT}}(\pi)$  may be severely suboptimal for the problem of maximizing  $Q(\pi)$ . In fact, as we now demonstrate, K-PT may actively choose trees that violate Assumption 1 even when the complete out-of-sample distribution is available *and* the tree is deep enough to allow branching on all covariates that are predictive of outcomes and treatment assignments.

EXAMPLE 1. Consider the following data distribution. Let the covariate vector  $X = (X^1, X^2) \in \{0, 1\}^2$ , where  $X^1$  and  $X^2$  are independent and Bernoulli distributed with parameter 0.5. We interpret  $X^1$  to denote the severity of a patient’s condition, where  $X^1 = 0$  (resp.  $X^1 = 1$ ) indicates a healthy (resp. sick) patient.  $X^2$  is a noise variable, unrelated to both the outcomes and the treatment assignments. There are two possible treatments, where  $K = 1$  (resp.  $K = 0$ ) indicates that the patient was treated (resp. not treated). Treatments are assigned according to a conditionally randomized experiment, with  $\mathbb{P}(K = 0|X^1 = 0) = 0.9$  and  $\mathbb{P}(K = 0|X^1 = 1) = 0.1$ , i.e., 90% of healthy patients do not receive the treatment, and 90% of sick patients do. Thus, the assumptions of conditional exchangeability and positivity hold in this setting, see Section 1.2.1 for definitions and pointers to the relevant literature. The above numbers imply that  $\mathbb{P}(K = 0) = \mathbb{P}(K = 1) = 0.5$  in the population. Table 1 shows the expected values of the potential outcomes in dependence of the covariate values. For simplicity, we assume that: 1) the potential outcomes in Table 1 have zero variance, meaning there is no uncertainty in the outcomes conditional on  $X$ ; and 2) we have sufficient historical data at our disposal so that we can work with the true values in (4) rather than with their empirical estimates provided in the definition of  $Q_I^{\text{K-PT}}(\pi)$ .

$X^1$	$\mathbb{E}(Y(0))$	$\mathbb{E}(Y(1))$
0	1	0.8
1	0	0.2

**Table 1** Companion table to Example 1. Expectation of potential outcomes in dependence of the covariate values. A value  $X^1 = 0$  (resp.  $X^1 = 1$ ) of the covariate indicates a healthy (resp. sick) patient. Larger values of the outcomes are preferred.

From Table 1, it can be seen that the (unique) optimal policy treats sick patients only. In particular, this optimal policy can be modeled by a prescriptive tree of depth one (i.e., a tree with a single branching node and two leafs) that branches on  $X^1$ , does not treat any of the healthy

patients (left leaf) and treats all sick patients (right leaf). The expected outcome under this policy is given by

$$\mathbb{E}(Y(0)|X^1 = 0)\mathbb{P}(X^1 = 0) + \mathbb{E}(Y(1)|X^1 = 1)\mathbb{P}(X^1 = 1) = 1 \times \frac{1}{2} + 0.2 \times \frac{1}{2} = 0.6.$$

One may similarly calculate the expected outcomes under a policy that branches on the noise feature  $X^2$ . In this case, no matter which treatment is assigned at each leaf, the resulting policies would have a lower expected outcome of 0.5. These candidate policies and the corresponding expected outcomes under each policy are illustrated in the top row of Figure 2.

The optimal policies described above can be obtained by an omniscient decision-maker who knows the true counterfactuals  $\mathbb{E}(Y(k))$ . In practice, however, decision makers can only obtain  $\mathbb{E}(Y|K = k)$  from the observational data. Naturally, incorrect estimates could potentially lead to poor decisions. In particular, we now argue that K-PT indeed produces suboptimal trees due to estimation bias. Consider the tree of depth one that maximizes the objective function of K-PT given by (4). Suppose that we branch on covariate  $X^j$  at the sole branching node of the tree,  $j \in \{1, 2\}$ . Then, at the leaf where  $X^j = x^j \in \{0, 1\}$ , K-PT estimates the counterfactual outcomes  $\mathbb{E}(Y(k)|X^j = x^j)$  as  $\mathbb{E}(Y|K = k, X^j = x^j)$ . However, these estimates can be incorrect. For example,

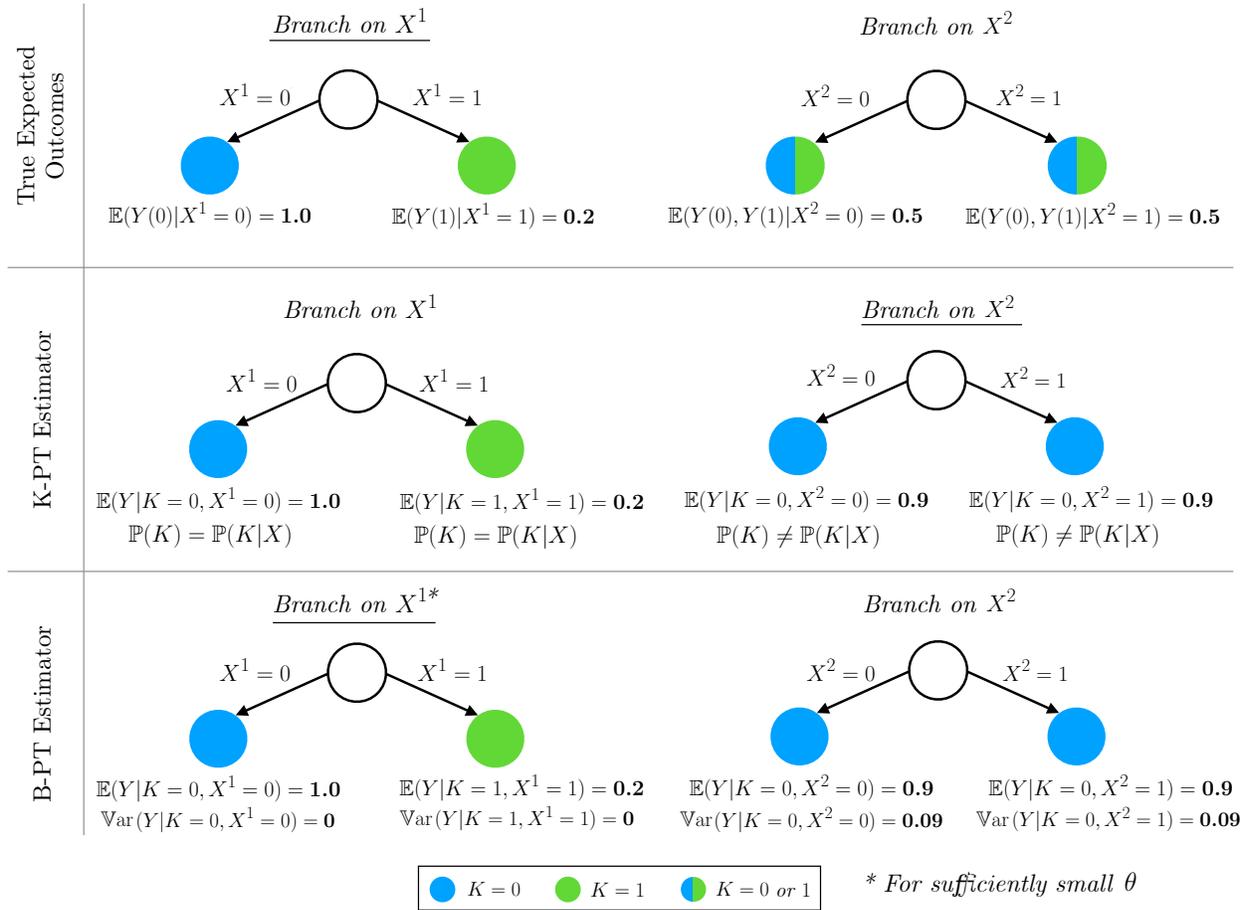
$$\mathbb{E}(Y(0)|X^2 = 0) = \sum_{x^1 \in \{0, 1\}} \mathbb{E}(Y(0)|X^1 = x^1, X^2 = 0)\mathbb{P}(X^1 = x^1|X^2 = 0) = 1 \times \frac{1}{2} + 0 \times \frac{1}{2} = 0.5,$$

whereas

$$\begin{aligned} \mathbb{E}(Y|K = 0, X^2 = 0) &= \sum_{x^1 \in \{0, 1\}} \mathbb{E}(Y|K = 0, X^1 = x^1, X^2 = 0)\mathbb{P}(X^1 = x^1|K = 0, X^2 = 0) \\ &= \sum_{x^1 \in \{0, 1\}} \mathbb{E}(Y|K = 0, X^1 = x^1) \frac{\mathbb{P}(K = 0|X^1 = x^1)\mathbb{P}(X^1 = x^1)}{\mathbb{P}(K = 0)} = 0.9. \end{aligned}$$

The difference between these two quantities is intuitive. Individuals that were not treated ( $K = 0$ ) in the data are more likely to be healthy (with a large associated outcome). The K-PT estimate  $\mathbb{E}(Y|K = 0, X^2 = 0)$  thus *overestimates*  $\mathbb{E}(Y(0)|X^2 = 0) = 0.5$ . Similarly, since the vast majority of individuals that were treated ( $K = 1$ ) in the data are sick (with low outcomes), the K-PT estimate  $\mathbb{E}(Y|K = 1, X^2 = 0) = 0.26$  is an *underestimator* for  $\mathbb{E}(Y(1)|X^2 = 0) = 0.5$ . Repeating this calculation for all combinations of  $K$  and  $X^2$  gives us  $\mathbb{E}(Y|K = 0, X^2 = 0) = \mathbb{E}(Y|K = 0, X^2 = 1) = \mathbb{E}(Y|K = 0)$  and similarly for  $\mathbb{E}(Y|K = 1, X^2 = 0) = \mathbb{E}(Y|K = 1, X^2 = 1) = \mathbb{E}(Y|K = 1)$ . Hence, when splitting on  $X^2$ , the K-PT estimator will choose to not treat anyone ( $K = 0$ ) because the estimator's expected values are  $0.9 > 0.26$ , see the middle right subfigure of Figure 2.

We can conduct a similar analysis for the tree that branches on  $X^1$ . In this case, K-PT estimates  $\mathbb{E}[Y|K = 0, X^1 = 0] = 1$  precisely as stated in Table 1, since the historical policy assigns treatments



**Figure 2** Companion figure to Example 1. Each row shows the estimates of performance of different prescriptive trees (branch on  $X^1$ -left- or  $X^2$ -right-) under different estimators. In each case, only the best treatment assignments at each leaf (as measured by the corresponding estimator) are shown. In this example, the true optimal policy (top left) is to branch on  $X^1$  and treat only those who are sick. Indeed, the expected outcome of the best tree branching on  $X^1$  (top left) is  $1.0 \cdot 0.5 + 0.2 \cdot 0.5 = 0.6$ , which is larger than the expected outcome of the best  $X^2$ -branching tree (top right), which is  $0.5 < 0.6$ . However, the method from Kallus (2017) (middle), which is discussed in Section 2.1, will prefer to branch on  $X^2$  (noise covariate) and not treat any patient (middle right). Indeed, it estimates the outcomes of the best  $X^1$ - (resp.  $X^2$ -) branching tree as 0.6 (resp. 0.9). We note that Assumption 1 is not satisfied at any of the leaves of the tree deemed as best by this estimator. In this case, the method of Bertsimas et al. (2019), see Section 2.2, is able to identify the truly optimal tree (bottom left); it correctly branches on  $X^1$  provided  $\theta$  is sufficiently small. Note that this method penalizes the variances corresponding to all treatments, but for simplicity the figure only depicts variances corresponding to the chosen treatment.

solely based on feature  $X^1$ . Repeating the decisions for all combinations of  $K$  and  $X^1$  gives us a tree where K-PT will choose to treat ( $K=1$ ) for those with  $X^1=1$  and not treat otherwise, see the middle left subfigure of Figure 2. When choosing the optimal branching decision of the two options, K-PT will **incorrectly** choose to branch on  $X^2$  because it estimates a higher expected outcome ( $(0.9 + 0.9)/2 = 0.9$ ) than when branching on  $X^1$  ( $(1.0 + 0.2)/2 = 0.6$ ).

Note that in a depth 1 tree that only branches on  $X^1$ , Assumption 1 is satisfied and the estimated outcome of 0.6 provided by K-PT is correct. However, K-PT prefers the alternative tree which branches on the noise variable and does not satisfy Assumption 1, in this case resulting in an overly optimistic estimation of the outcomes. We conclude that if not *all feasible trees* in the problem satisfy Assumption 1, the method from Kallus (2017) may fail.

## 2.2. B-PT

Bertsimas et al. (2019) also observe that K-PT may fail in certain conditions. They posit that failings can be attributed to poor estimates of the outcomes at the leaves. Thus, they augment the K-PT estimator with a regularization term, which penalizes trees that induce a high variance in the outcomes conditional on each treatment at the leaves, i.e.,

$$\sum_{k \in \mathcal{K}} \sum_{n \in \mathcal{L}} \text{Var}(Y|X \in \mathcal{X}_n, K = k).$$

Specifically, they optimize the quantity

$$Q_I^{\text{B-PT}}(\pi) = \theta (Q_I^{\text{K-PT}}(\pi)) - (1 - \theta) \sum_{i \in \mathcal{I}} \left( Y_i - \widehat{Y}_{\mathcal{X}_{n(i)}}(K_i) \right)^2,$$

where  $n(i) \in \mathcal{L}$  denotes the leaf datapoint  $i$  belongs to, and  $\theta \geq 0$  is a parameter to be tuned.

**Analysis of B-PT.** Recall that Kallus (2017) estimates  $\mathbb{E}(Y(k)|X \in \mathcal{X}_n)$  through  $\mathbb{E}(Y|K = k, X \in \mathcal{X}_n)$  and shows that this estimator is unbiased under Assumption 1. By adding a regularization term that seeks to promote trees with low variance  $\text{Var}(Y|K = k, X \in \mathcal{X}_n)$ , Bertsimas et al. (2019) improves the statistical properties of the K-PT estimator (lower variance). Therefore, if Assumption 1 holds (so that  $\mathbb{E}(Y|K = K_i, X \in \mathcal{X}_{n(i)}) = \mathbb{E}(Y(K_i)|X \in \mathcal{X}_{n(i)})$ ), then B-PT can lead to improved performance over K-PT (for a suitable choice of  $\theta$ ). However, B-PT does not directly address the issues discussed in Example 1. Indeed, when Assumption 1 does not hold, then  $\mathbb{E}(Y|K = K_i, X \in \mathcal{X}_{n(i)})$  and  $\mathbb{E}(Y(K_i)|X \in \mathcal{X}_{n(i)})$  can be drastically different from one another, and the policies chosen by B-PT may still be suboptimal.

In the specific case of Example 1, B-PT can in fact recover the best tree if  $\theta$  is sufficiently small. Nonetheless, as we now show, a small modification of Example 1 results in a setting in which B-PT has *additional* incentives to choose the worst tree, which branches on irrelevant features.

EXAMPLE 1 (CONTINUED). We now revisit Example 1 and use the method from Bertsimas et al. (2019) to identify an optimal tree. Figure 2 (bottom) shows the estimated outcomes and variances associated with each candidate tree. Since the tree that branches on the relevant feature  $X^1$  has smaller variance in this case, it will be chosen by B-PT provided that  $\theta$  is sufficiently small. ■

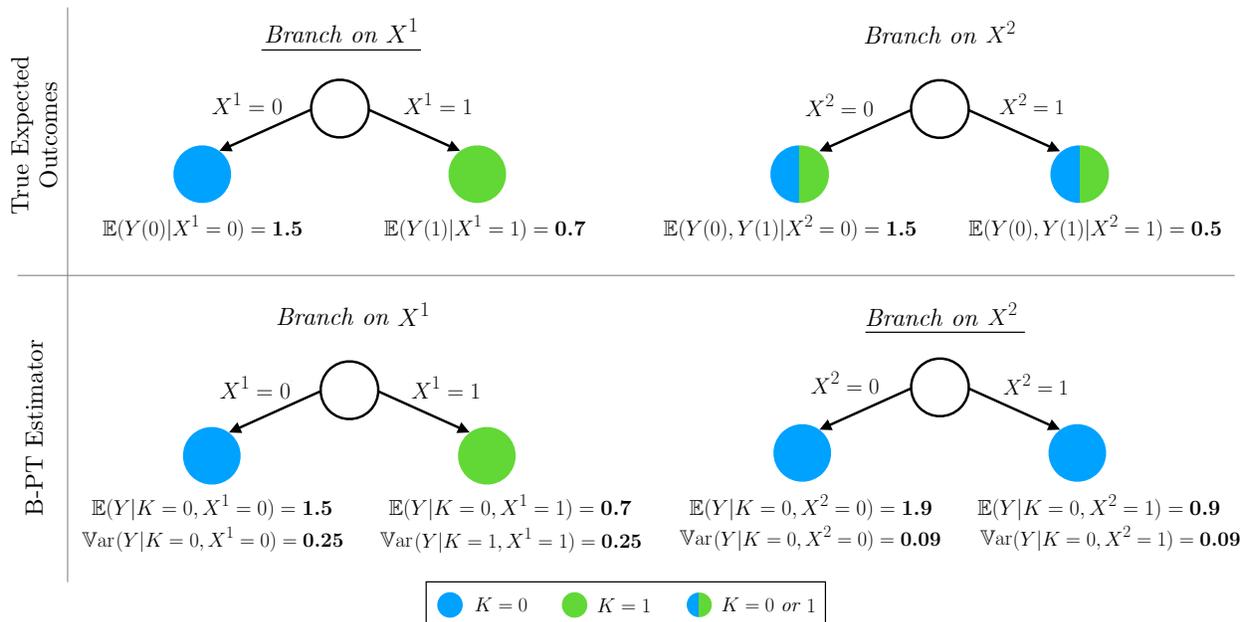
EXAMPLE 2. We now consider a variant of Example 1 where we add to all potential outcomes the quantity  $(1 - X^2)$ , see Table 2. As before,  $\mathbb{P}(K = 0|X^1 = 0) = 0.9$  and  $\mathbb{P}(K = 0|X^1 = 1) = 0.1$ , i.e., 90% of healthy patients do not receive the treatment, and 90% of sick patients do, also implying that  $\mathbb{P}(K = 0) = \mathbb{P}(K = 1) = 0.5$ . Since  $X^2$  now affects the values of the outcomes, we have  $\mathbb{E}(Y|K = 0, X^2 = 1) \neq \mathbb{E}(Y|K = 0, X^2 = 0)$  and similarly for  $K = 1$ . However, this change does *not* affect the value  $\mathbb{E}(Y(1) - Y(0)|X)$  and is thus *irrelevant* in deciding which treatment is preferable. Therefore, under suitable identifiability conditions, there exists an optimal prescriptive tree of depth one, which still splits on  $X^1$ . Figure 3 (top) shows the outcomes, for a depth one tree, under the two possible branching decisions if the counterfactuals are known – the best tree indeed branches on the only predictive feature  $X^1$ . However, when using estimates  $\mathbb{E}[Y|K = k, X = x]$  for the counterfactuals (Figure 3, bottom), the tree that branches on the irrelevant feature  $X^2$  has both a higher estimated outcome and a smaller variance. Therefore, B-PT will always choose to branch on  $X^2$  and not treat anyone in the population, no matter the choice of  $\theta$ . In this case B-PT has even more incentive to select a policy based on the irrelevant feature and both K-PT and B-PT will select a suboptimal tree. ■

$X^1$	$X^2$	$\mathbb{E}(Y(0))$	$\mathbb{E}(Y(1))$
0	0	2	1.8
1	0	1	1.2
0	1	1	0.8
1	1	0	0.2

Table 2 Companion table to Example 2. Expectation of the potential outcomes in dependence of the covariate values. Similar to Example 1 (Continued), we let the potential outcomes in Table 1 have zero variance implying that potential outcomes conditioned on the covariate values are perfectly known.

### 3. Proposed Formulations

In this section, we present our proposed MIO formulations for learning optimal prescriptive trees from observational data. We assume that the covariates are taken from a finite set of integers, i.e.,  $\mathcal{X} \subset \mathbb{Z}^F$ , and let  $\Theta(f)$  capture all levels of each feature  $f \in \mathcal{F}$ . To handle categorical features, one can one-hot encode the variables. To handle continuous variables, one can pass in finite discretizations that are decided a priori. In Section 4.4, we show that discretizing does not lead to a degradation in performance in our experiments; in fact, it may sometimes serve as a form of regularization and produce results that generalize better out-of-sample than passing in continuous variables. We also note that discretization has a rich literature with many proposed solutions, see e.g., Dougherty et al. (1995), Kotsiantis and Kanellopoulos (2006), Rucker et al. (2015) for reviews of methods. Importantly, and in sharp contrast with the existing literature on prescriptive trees, we *do not*



**Figure 3** Companion figure to Example 2. The layout and interpretation of the figure exactly parallel that in Figure 2. The optimal tree (top left) has the objective  $0.5(1.5) + 0.5(0.7) = 1.1$ , which is higher than the  $X^2$ -branching tree (top right),  $0.5(1.5) + 0.5(0.5) = 1.0 < 1.1$ . However, B-PT will incorrectly branch on  $X^2$  no matter the regularization strength  $\theta$ , since the  $X^1$ -branching tree (bottom left) has a lower estimated value and also higher variance.

make Assumption 1. Instead, we directly optimize the estimators (1), (2), or (3) from the causal inference literature. This enables us to design interpretable prescriptive trees that are guaranteed to be optimal out-of-sample as the number of historical samples grows.

As will become clear shortly, our proposed approaches give rise to *weighted* classification trees. We thus adapt the MIO formulations of Aghaei et al. (2020) for learning optimal classification trees from the predictive to the prescriptive setting. We extend our notation from the problem statement in Section 1.1 to describe optimal prescriptive trees. The building block of our formulation is a perfect binary tree of depth  $d$  whose nodes are numbered 1 through  $(2^{d+1} - 1)$  in the order in which they appear in a breadth-first search. Note that we can in general extend our method to learn non-binary trees similar to Menickelly et al. (2016). We let  $\mathcal{B} := \{1, \dots, 2^d - 1\}$  denote the set of branching nodes and  $\mathcal{T} := \{2^d, \dots, 2^{d+1} - 1\}$  collect all terminal nodes (see Figure 4, left). From the perfect binary tree, we build a flow graph as follows. We connect a source node  $s$  to the root node, and connect all nodes other than  $s$  to  $|\mathcal{K}|$  sink nodes (denoted by  $t_k, \forall k \in \mathcal{K}$ ) – see Figure 4, right. All links in the graph are directed from source to sink and have capacity 1.

### 3.1. Formulation Based on Doubly Robust Estimation

Equipped with the flow graph associated with a prescriptive tree, we now formulate an MIO problem to optimize (3) over all trees of maximum depth  $d$ . For every branching node  $n \in \mathcal{B}$ ,

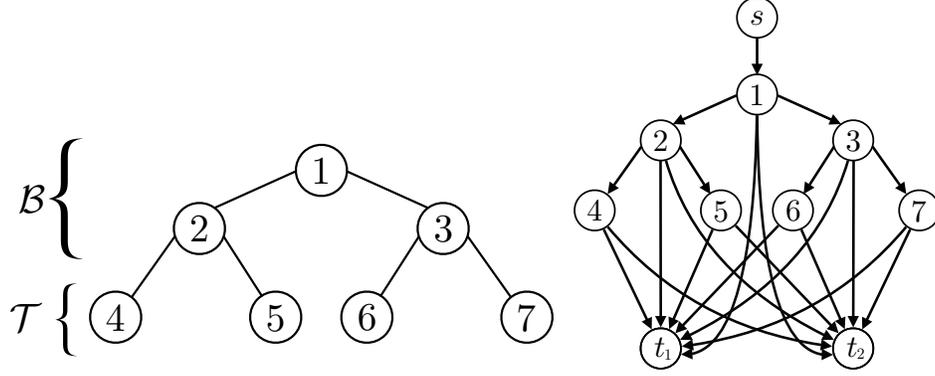


Figure 4 A prescriptive tree with depth 2 (left) and its associated flow graph (right).

feature  $f \in \mathcal{F}$ , and threshold  $\theta \in \Theta(f)$ , we let the binary variable  $b_{nf\theta}$  indicate if feature  $f$  with threshold  $\theta$  is selected for branching at node  $n$ . Accordingly, for every node  $n \in \mathcal{T} \cup \mathcal{B}$ , we let the binary variable  $p_n$  indicate if node  $n$  is a treatment node, in which case a treatment must be assigned to all datapoints that land at that node and no further branching is allowed. For  $k \in \mathcal{K}$ , we let  $w_{nk} \in \{0, 1\}$  equal 1 if and only if (iff) treatment  $k$  is selected at node  $n$ .

In contrast to Aghaei et al. (2020) which decides the flow of each datapoint individually, we propose an *aggregated* variant wherein we decide the flow of all datapoints with same covariates at once. Thus, we let  $z_x^{a(n),n}$  indicate that datapoints with covariates  $x \in \mathcal{X}$  will pass through arc  $(a(n), n)$ . This idea, based on the observation that datapoints with the same covariates follow the same path from root to sink, allows us to often sharply reduce the size of our formulations and speed up computation.

The MIO formulation is as follows:

$$\text{maximize } \sum_{x \in \mathcal{X}} \sum_{k \in \mathcal{K}} \sum_{n \in \mathcal{B} \cup \mathcal{T}} z_x^{n,t_k} \sum_{i: X_i=x} \left[ \hat{\nu}_k(X_i) + \frac{\mathbb{1}[k = K_i](Y_i - \hat{\nu}_{K_i}(X_i))}{\mu(K_i, X_i)} \right] \quad (6a)$$

$$\text{subject to } \sum_{f \in \mathcal{F}} \sum_{\theta \in \Theta(f)} b_{nf\theta} + p_n + \sum_{m \in \mathcal{A}(n)} p_m = 1 \quad \forall n \in \mathcal{B} \quad (6b)$$

$$p_n + \sum_{m \in \mathcal{A}(n)} p_m = 1 \quad \forall n \in \mathcal{T} \quad (6c)$$

$$z_x^{a(n),n} = z_x^{n,\ell(n)} + z_x^{n,r(n)} + \sum_{k \in \mathcal{K}} z_x^{n,t_k} \quad \forall n \in \mathcal{B}, x \in \mathcal{X} \quad (6d)$$

$$z_x^{a(n),n} = \sum_{k \in \mathcal{K}} z_x^{n,t_k} \quad \forall x \in \mathcal{X}, n \in \mathcal{T} \quad (6e)$$

$$z_x^{s,1} = 1 \quad \forall x \in \mathcal{X} \quad (6f)$$

$$z_x^{n,\ell(n)} \leq \sum_{f \in \mathcal{F}} \sum_{\theta \in \Theta(f): x_f \leq \theta} b_{nf\theta} \quad \forall n \in \mathcal{B}, x \in \mathcal{X} \quad (6g)$$

$$z_x^{n,r(n)} \leq \sum_{f \in \mathcal{F}} \sum_{\theta \in \Theta(f): x_f > \theta} b_{nf\theta} \quad \forall n \in \mathcal{B}, x \in \mathcal{X} \quad (6h)$$

$$z_x^{n,t_k} \leq w_{nk} \quad \forall x \in \mathcal{X}, n \in \mathcal{T}, k \in \mathcal{K} \quad (6i)$$

$$\sum_{k \in \mathcal{K}} w_{nk} = p_n \quad \forall n \in \mathcal{B} \cup \mathcal{T} \quad (6j)$$

$$w_{nk} \in \{0, 1\} \quad \forall n \in \mathcal{T}, k \in \mathcal{K} \quad (6k)$$

$$b_{nf\theta} \in \{0, 1\} \quad \forall n \in \mathcal{B}, f \in \mathcal{F}, \theta \in \Theta(f) \quad (6l)$$

$$p_n \in \{0, 1\} \quad \forall n \in \mathcal{B} \cup \mathcal{T} \quad (6m)$$

$$z_x^{a(n),n}, z_x^{n,t_k} \in \{0, 1\} \quad \forall n \in \mathcal{B} \cup \mathcal{T}, x \in \mathcal{X}, k \in \mathcal{K}. \quad (6n)$$

Here, we have where  $\hat{\mu}$  and  $\hat{\nu}_k$  are as defined in Sections 1.2.2 and 1.2.3, respectively.  $\mathcal{A}(n)$  denotes the set of all ancestors of node  $n \in \mathcal{B} \cup \mathcal{T}$ ,  $r(n)$  (resp.  $\ell(n)$ ) is the right (resp. left) descendant of  $n$ . The objective (6a) maximizes the doubly robust objective (3). Constraints (6b) ensure that we branch at all branching nodes where no treatment is assigned, at the node itself or at any of its ancestors. Similarly, constraints (6c) impose that a treatment must be assigned at all terminal nodes, unless a treatment has already been assigned to one of its ancestors. Constraints (6d) and (6e) are flow conservation constraints, whereby any datapoint that flows into a node  $n$  must either exit to its right or left descendant, or flow directly to one of the sink nodes. Constraints (6f) ensure that a flow of at most one can enter the source node for each datapoint. Constraints (6g) and (6h) guarantee that datapoints flow according to the branching decisions specified by variables  $b$ . Constraints (6i) state that datapoints can only flow to the sink associated with their assigned treatment. Finally, constraints (6j) combined with integrality of the  $w_{nk}$  variables ensure that if a treatment is assigned at a node, exactly one of the available treatments must be assigned. Note that we can relax integrality on the  $w$  and  $z$  variables and still reach an integral solution. We also note that in our formulation,  $\mathcal{X}$  could comprise only the features we observe in the data, rather than the entire feature space.

### 3.2. Formulation Based on Inverse Propensity Weighting Estimator

To optimize the inverse propensity weighting objective (1), we can replace (6a) to get

$$\begin{aligned} & \text{maximize} \quad \sum_{x \in \mathcal{X}} \sum_{k \in \mathcal{K}} \sum_{n \in \mathcal{B} \cup \mathcal{T}} z_x^{n,t_k} \sum_{i: X_i=x} \frac{\mathbb{1}[K_i = k] Y_i}{\mu(K_i, X_i)} \\ & \text{subject to} \quad (6b)-(6n). \end{aligned} \quad (7)$$

### 3.3. Formulation Based on Direct Method Estimator

Formulation (6) can similarly be adapted to optimize the direct method objective (2) by simply dropping the second term in the objective (6a), i.e.,

$$\begin{aligned} & \text{maximize} \quad \sum_{x \in \mathcal{X}} \sum_{k \in \mathcal{K}} \sum_{n \in \mathcal{B} \cup \mathcal{T}} z_x^{n,t_k} \sum_{i: X_i=x} \hat{\nu}_k(X_i) \\ & \text{subject to} \quad (6b)-(6n). \end{aligned} \quad (8)$$

### 3.4. Analysis of IPW, DR, and DM

It can be shown that, for the data presented in Examples 1 and 2, if  $\mu$  and/or  $\nu$  are known, a depth one tree trained according to formulations (6), (7), and (8) will be optimal out-of-sample, provided the number of samples in the data is sufficiently large. In this section, we go one step further and demonstrate that, under suitable conditions on the historical data generation process and the estimators, the tree returned by our optimal MIO formulations (6), (7), and (8) will converge almost surely to an optimal tree.

We introduce some notation to help formalize our claim. For any tree based policy  $\pi \in \Pi_d$ , define

$$Q^{\text{IPW}}(\pi) := \mathbb{E} \left[ \frac{\mathbb{I}(K = \pi(X))Y}{\mu(K, X)} \right] \quad \text{and} \quad Q_{N,I}^{\text{IPW}}(\pi) := \frac{1}{I} \sum_{i=1}^I \frac{\mathbb{I}(K_i = \pi(X_i))Y_i}{\hat{\mu}^N(K_i, X_i)},$$

where  $\hat{\mu}^N(K, X)$  is an estimator of  $\mu(K, X)$  trained over an i.i.d. sample of cardinality  $N$  independent from  $\mathcal{D}$ . We denote by  $v^*$  and  $\mathcal{S}^*$ , the optimal value and the set of optimal solutions of  $\max_{\pi \in \Pi_d} Q(\pi)$ , respectively. Similarly, we denote by  $\hat{v}_{N,I}^{\text{IPW}}$  and  $\hat{\mathcal{S}}_{N,I}^{\text{IPW}}$ , the optimal value and the set of optimal solutions of  $\max_{\pi \in \Pi_d} Q_{N,I}^{\text{IPW}}(\pi)$ , respectively. Finally, we use the abbreviation ‘‘w.p.1’’ for ‘‘with probability one’’.

**PROPOSITION 1.** *Suppose that: conditional exchangeability holds, i.e.,  $Y(k) \perp\!\!\!\perp K|X$ , for all values of  $k \in \mathcal{K}$ ; the potential outcomes  $Y(k)$  are bounded, for all values of  $k \in \mathcal{K}$ ; and, both  $\mu(k, x) > 0$  and  $\hat{\mu}^N(k, x) > 0$  for all  $x \in \mathcal{X}$ ,  $k \in \mathcal{K}$ , and  $N$  sufficiently large. If  $\hat{\mu}^N(k, x)$  converges almost surely to  $\mu(k, x)$  for all  $x \in \mathcal{X}$ ,  $k \in \mathcal{K}$ , then  $\hat{v}_{N,I}^{\text{IPW}} \rightarrow v^*$  w.p.1 and  $\hat{\mathcal{S}}_{N,I}^{\text{IPW}} \subseteq \mathcal{S}^*$  w.p.1 for  $I$  and  $N$  large enough.*

Similar to the IPW case, we now show that under suitable conditions, the MIO formulation (6) for the DR estimator returns an optimal out-of-sample tree. We introduce additional notation to help formalize our claim. For any tree based policy  $\pi \in \Pi_d$ , define

$$Q^{\text{DR}}(\pi) := \mathbb{E} \left[ \hat{\nu}_{\pi(X)}(X) + (Y - \hat{\nu}_{\pi(X)}(X)) \frac{\mathbb{I}(K = \pi(X))}{\hat{\mu}(K, X)} \right]$$

$$\text{and} \quad Q_{N,I}^{\text{DR}}(\pi) := \frac{1}{I} \sum_{i=1}^I \left( \hat{\nu}_{\pi(X_i)}^N(X_i) + (Y - \hat{\nu}_{\pi(X_i)}^N(X_i)) \frac{\mathbb{I}(K_i = \pi(X_i))}{\hat{\mu}^N(K_i, X_i)} \right),$$

where  $\hat{\nu}_K^N(\cdot)$  is an estimator of  $\nu_K(\cdot)$ , trained over an i.i.d. sample of cardinality  $N$  independent from  $\mathcal{D}$ . Let  $\hat{\mu}(K, X)$  (resp.  $\hat{\nu}_K(X)$ ) be the limit of  $\hat{\mu}^N(K, X)$  (resp.  $\hat{\nu}_K^N(X)$ ) as  $N$  goes to infinity. We denote by  $\hat{v}_{N,I}^{\text{DR}}$  and  $\hat{\mathcal{S}}_{N,I}^{\text{DR}}$ , the optimal value and the set of optimal solutions of  $\max_{\pi \in \Pi_d} Q_{N,I}^{\text{DR}}(\pi)$ , respectively.

**PROPOSITION 2.** *Suppose that: conditional exchangeability holds, i.e.,  $Y(k) \perp\!\!\!\perp K|X$ , for all values of  $k \in \mathcal{K}$ ; the potential outcomes  $Y(k)$  are bounded, for all values of  $k \in \mathcal{K}$ ; both  $\mu(k, x) > 0$  and*

$\hat{\mu}^N(k, x) > 0$  for all  $x \in \mathcal{X}$ ,  $k \in \mathcal{K}$  and  $N$  sufficiently large; and both  $\hat{\mu}(k, x)$  and  $\hat{\nu}_k(x)$  exist and are bounded for all  $x \in \mathcal{X}$ ,  $k \in \mathcal{K}$ . If either  $\hat{\mu}^N(k, x)$  converges almost surely to  $\mu(k, x)$  or  $\hat{\nu}_k^N(x)$  converges almost surely to  $\nu_k(x)$  for all  $x \in \mathcal{X}$ ,  $k \in \mathcal{K}$ , then  $\hat{\nu}_{N,I}^{DR} \rightarrow v^*$  w.p.1 and  $\hat{\mathcal{S}}_{N,I}^{DR} \subseteq \mathcal{S}^*$  w.p.1 for  $I$  and  $N$  large enough.

At the end, we show that under suitable conditions, the MIO formulation (8) returns an optimal out-of-sample tree. We introduce additional notation to help formalize our claim. For any tree based policy  $\pi \in \Pi_d$ , define

$$Q^{\text{DM}}(\pi) := \mathbb{E} [\hat{\nu}_{\pi(X)}(X)] \quad \text{and} \quad Q_{N,I}^{\text{DM}}(\pi) := \frac{1}{I} \sum_{i=1}^I (\hat{\nu}_{\pi(X_i)}^N(X_i)).$$

We denote by  $\hat{\nu}_{N,I}^{\text{DM}}$  and  $\hat{\mathcal{S}}_{N,I}^{\text{DM}}$  the optimal value and the set of optimal solutions of  $\max_{\pi \in \Pi_d} Q_{N,I}^{\text{DM}}(\pi)$ , respectively.

**PROPOSITION 3.** *Suppose that: conditional exchangeability holds, i.e.,  $Y(k) \perp\!\!\!\perp K|X$ , for all values of  $k \in \mathcal{K}$ ; and  $\hat{\nu}_k(x)$  exists, for all  $k \in \mathcal{K}$ . If  $\hat{\nu}_k^N(x)$  converges almost surely to  $\nu_k(x)$  for all  $x \in \mathcal{X}$ ,  $k \in \mathcal{K}$ , then  $\hat{\nu}_{N,I}^{\text{DM}} \rightarrow v^*$  w.p.1 and  $\hat{\mathcal{S}}_{N,I}^{\text{DM}} \subseteq \mathcal{S}^*$  w.p.1 for  $I$  and  $N$  large enough.*

### 3.5. Extensions

Our formulations in Sections 3.1, 3.2, and 3.3 have strong modeling power. For example, they can be used to learn trees that interpretable and/or fair, to design randomized policies, and to impose budget constraints relevant for under-resourced settings. We discuss these in the present section.

**Intepretability.** In settings where learning interpretable policies is desired, see e.g., Rudin (2019) and Azizi et al. (2018), one may limit the number of branching nodes in the tree to  $M$  by augmenting formulations (6), (7), and (8) with the constraint

$$\sum_{n \in \text{BUT}} p_n \leq M.$$

**Budget Constraints.** In low resource settings such as organ allocation (Bertsimas et al. 2013, Zenios et al. 2000), housing allocation (Azizi et al. 2018), and security (Xu et al. 2018), many treatments/interventions often have limited supply. Tree based policies that satisfy such capacity constraints can be learned by augmenting formulations (6), (7), and (8) with constraints

$$\sum_{n \in \text{BUT}} \sum_{x \in \mathcal{X}} \sum_{i \in \mathcal{I}: X_i = x} z_x^{n,t_k} \leq |\mathcal{I}| C_k \quad \forall k \in \mathcal{K},$$

where  $C_k$  denotes the percentage of instances that can be assigned treatment  $k$ .

**Fairness in Treatment Assignment.** Our formulations can be used to learn trees where treatments are assigned fairly across groups. Concretely, we will say that a policy satisfies treatment assignment parity if the probability of assigning a particular treatment is equal across protected groups. This constraint is similar to notions of fairness in the classification setting, for which there is a rich literature – see, e.g., Dwork et al. (2012), Zemel et al. (2013), Bolukbasi et al. (2016), Hardt et al. (2016), Zafar et al. (2017), Olfat and Aswani (2018), Jo et al. (2022b). We let  $\mathcal{P}$  collect all protected covariates (e.g., LGBTQ status or categories of race). Accordingly, we let  $P_x \in \mathcal{P}$  represent the value of the protected feature(s) of covariate  $x$ . It is worth noting that the features in  $\mathcal{P}$  are usually not included in covariates  $X$  to make sure that the policy does not make decisions based on protected features. To ensure that the learned policy satisfies treatment assignment parity up to a bias  $\delta$ , we may augment (6), (7), and (8) with the constraint

$$\left| \frac{\sum_{n \in \mathcal{B} \cup \mathcal{T}} \sum_{x \in \mathcal{X}: P_x = p} \sum_{i \in \mathcal{I}: X_i = x} z_x^{n, t_k}}{|\{i \in \mathcal{I} : P_i = p\}|} - \frac{\sum_{n \in \mathcal{B} \cup \mathcal{T}} \sum_{x \in \mathcal{X}: P_x = p'} \sum_{i \in \mathcal{I}: X_i = x} z_x^{n, t_k}}{|\{i \in \mathcal{I} : P_i = p'\}|} \right| \leq \delta \quad \forall p, p' \in \mathcal{P} : p \neq p', k \in \mathcal{K}.$$

In a similar fashion, we will say that a policy satisfies *conditional* treatment assignment parity if the probability of assigning a particular treatment is equal across protected groups, conditional on some legitimate feature(s) that affect the outcome. For instance, in prescribing housing resources to people experiencing homelessness, one may require that all individuals with the same vulnerability have the same likelihood of receiving a certain treatment. We let  $\mathcal{A}$  be the set of features indicative of risk, and let  $A_i \in \mathcal{A}$  be the value of the risk feature of datapoint  $i$ . Conditional treatment assignment parity is satisfied up to a bias  $\delta$  and for all features  $a \in \mathcal{A}$  by adding the following constraint to (6), (7), and (8),

$$\left| \frac{\sum_{n \in \mathcal{B} \cup \mathcal{T}} \sum_{x \in \mathcal{X}: P_i = p, A_x = a} \sum_{i \in \mathcal{I}: X_i = x} z_x^{n, t_k}}{|\{i \in \mathcal{I} : P_i = p \cap A_i = a\}|} - \frac{\sum_{n \in \mathcal{B} \cup \mathcal{T}} \sum_{x \in \mathcal{X}: P_x = p', A_x = a} \sum_{i \in \mathcal{I}: X_i = x} z_x^{n, t_k}}{|\{i \in \mathcal{I} : P_i = p' \cap A_i = a\}|} \right| \leq \delta$$

$$\forall p, p' \in \mathcal{P} : p \neq p', k \in \mathcal{K}, a \in \mathcal{A}.$$

**Fairness in Treatment Outcomes.** Our formulations can also be used to ensure fairness in expected outcomes across groups. In general, this requires the average expected outcomes of a protected group  $p \in \mathcal{P}$  to be above some threshold  $\gamma_p \in \mathbb{R}$ :

$$\sum_{n \in \mathcal{B} \cup \mathcal{T}} \sum_{x \in \mathcal{X}: P_x = p} \sum_{k \in \mathcal{K}} \sum_{i: X_i = x} z_x^{n, t_k} \left[ \hat{\nu}_k(X_i) + \frac{\mathbb{1}[k = K_i](Y_i - \hat{\nu}_{K_i}(X_i))}{\mu(K_i, X_i)} \right] \geq \gamma_p \quad \forall p \in \mathcal{P}.$$

For example, Bertsimas et al. (2013) set  $\gamma_p := \sum_{i \in \mathcal{I}: P_i = p} Y_i$ , i.e., the expected outcomes of all protected groups under the learned policy should be greater than or equal to what was observed

in the data. Alternatively, the constraint can be used to impose max-min fairness to protect the outcomes of the groups that are worst off, see Rawls (1974). In this case, one sets  $\gamma_p := \gamma$  for all  $p \in \mathcal{P}$ , where  $\gamma$  is the largest value for which the MIO problem remains feasible.

**Randomized Treatment Assignment Policies.** The addition of fairness and budget constraints may make problems (6), (7), and (8) infeasible. In such cases, it may be desirable to design *randomized* policies where datapoints that land at the same leaf are assigned each treatment with a certain probability (rather than all getting the same treatment). This can be achieved by relaxing integrality on the variables  $w$  and  $z$  in formulations (6), (7), and (8). The variable  $w_{nk}$  can then be interpreted as the probability of assigning treatment  $k$  to datapoints that fall on node  $n$ .

## 4. Experiments

We now evaluate the empirical performance of our proposed formulations in Section 3 in two problem settings: a synthetic setting from the literature and a real setting based on a warfarin dosing dataset. In both cases, we take the viewpoint that data comes from an observational study, which can be viewed as a conditionally randomized experiment. We benchmark against the approaches of Kallus (2017) and Bertsimas et al. (2019) for learning prescriptive trees, see Section 2. As before, we refer to these as K-PT and B-PT, respectively. We also compare our method to four other, non-MIO based, approaches that can similarly learn or be adapted to learn prescriptive policies: causal forests and causal trees (Athey and Imbens 2016), policy tree (Zhou et al. 2023), and regress & compare (see Section 1.2.3). We describe these methods further in Section 4.3. Since Bertsimas et al. (2019) do not propose an MIO formulation, we adapt the method from Kallus (2017) to implement B-PT. For completeness, we provide the MIO formulations that we have implemented for these two approaches in Appendices EC.1 and EC.2, respectively. We evaluate the performance of all approaches as the probability of correct treatment assignment is varied in the historical data.

### 4.1. Dataset Description

**Synthetic Data.** For our experiments on synthetic data, we adapt the data generation process from Athey and Imbens (2016). In this problem, there are two treatment possibilities indexed in the set  $\mathcal{K} = \{0, 1\}$ . The covariate vector has two independent and identically distributed features,  $X = (X^1, X^2)$ , where each  $X^j \sim \mathcal{N}(0, 1)$ . The potential outcome of datapoint  $i$  with covariates  $X_i$  under treatment  $k$  is

$$Y_i(k) = \phi(X_i) + \frac{1}{2}(2k - 1) \cdot \kappa(X_i) + \epsilon_i, \quad (9)$$

where  $\phi(x) := \frac{1}{2}x^1 + x^2$  models the mean effect,  $\kappa(x) := \frac{1}{2}x^1$  models the treatment effect, and  $\epsilon_i \sim \mathcal{N}(0, 0.1)$  is noise added to the outcome that is independent of the covariates.

To study different settings of observational experiments, we vary, in the data, the probability  $p$  of assigning the treatment that is best *in expectation* for each unit. We let  $p \in \{0.1, 0.25, 0.5, 0.75, 0.9\}$  – since there are 2 treatments, 0.5 corresponds to the (marginally) randomized setting, while the other settings correspond to conditionally randomized experiments. For each  $p$ , we randomly generate 5 training and test sets, each with 500 and 10,000 datapoints, respectively. We use the training set to estimate  $\mu$  and  $\nu$ , and to train formulations (6), (7), and (8); we use the test set to evaluate the learned policy. Since this is a synthetic dataset, we have access to counterfactuals in the test set, which we use for evaluation purposes.

**Warfarin Dosing.** For our study on real data, we employ a dataset for personalized warfarin dosing for which counterfactuals are available. Warfarin is the most widely used oral anticoagulant agent, but despite its prevalence, determining one’s optimal warfarin dosage is difficult because it can vary widely depending on demographic variables, clinical factors, and genetics (Consortium 2009). The publicly available dataset that we use was collected by the International Warfarin Pharmacogenetics Consortium and published at the Pharmacogenetics and Pharmacogenomics Knowledge Base, see Consortium (2009). The advantage of using this dataset is that we can model a patient’s true outcomes when given varying doses of warfarin, which allows us to evaluate the performance of arbitrary counterfactual policies. Consortium (2009) published a learned affine function  $f(x) = \beta x + c$  that determines the optimal warfarin dose based on a patient’s age, weight, race, VKORC1 genotype, CYP2c9 genotype, and whether or not the patient is currently taking amiodarone or an enzyme reducer, see equation (EC.3) in Section EC.3. We calculate a patient  $i$ ’s optimal dosage using  $f(X_i) + \epsilon$ , where  $\epsilon \sim \mathcal{N}(0, 0.02)$ . This dosage is then discretized into three groups (i.e.,  $|\mathcal{K}| = 3$ ) using the same convention as Consortium (2009):  $K_i^{\text{opt}} = 0$  ( $\leq 3$  mg/day), 1 (between 3 and 7 mg/day), and 2 ( $\geq 7$  mg/day). The observed outcome is  $Y_i(K_i) = 1$ , if  $K_i = K_i^{\text{opt}}$ ; and  $= 0$  otherwise.

To study different settings of observational experiments, we consider three different treatment assignment mechanisms, as follows. For the marginally randomized setting, we assign each treatment with probability  $1/3$ , and repeat to produce 5 datasets. To simulate data that is based on a more informed policy, we assign treatments based on modified versions of  $f$  obtained by perturbing its coefficients so that patients closer to the boundary of each treatment bucket are at higher risk of receiving an incorrect treatment. Specifically, for each coefficient  $a$  of  $f$ , we generate  $a' \sim U(a - a \cdot r, a + a \cdot r)$ , where  $r$  denotes the range of sampling. We assign patient  $i$  the treatment corresponding to this modified function, discretized into three groups as discussed previously. In testing, we found that  $r \in [0.05, 0.12]$  was the ideal range to simulate a reasonable treatment policy where the probability of correct treatment assignment is between 0.6 and 0.9. We fix 2 values

of  $r$ , 0.06 and 0.11, and for each value, we generate 5 sets of randomly sampled coefficients (from hereon, we refer to this variability as a “realization”). At this point we have 3 experiment designs –each with 5 realizations – making up 15 datasets. Each dataset is then split randomly 5 times into training and test sets of 3,000 and 1,386 instances, respectively, yielding a total of 75 train-test pairs. Similar to the synthetic data, only the training set is used to estimate  $\mu, \nu$ , and to train formulations (6), (7), and (8).

## 4.2. Experimental Setup

Since the proposed methods rely on estimating the propensity scores and/or counterfactuals, we discuss the estimation models we use for both datasets. We also discuss the preprocessing steps required. All approaches are allotted a solve time of 4 hours and utilize 6 Intel Xeon E5-2640 v4, 2.40GHz CPUs, each having 4GB of memory. The MIO-based approaches use the Gurobi<sup>1</sup> solver (version 10.0.0).

**Synthetic Data.** The synthetic data has real covariates but our methods only allow for binary features to split on. We thus discretize each covariate feature into ten buckets corresponding to deciles from a normal distribution. Since the outcome distribution is known by construction (it is linear), we learn the potential outcomes using both linear regression (LR) and lasso regression with  $\alpha = 0.08$  (Lasso), knowing that the latter will be a slightly less accurate predictor. We also use true propensity scores—they are known by construction—as well as two methods to estimate propensity scores: logistic regression (Log) and decision trees (DT). Therefore, there will be six doubly robust methods, each corresponding to a pair of models for IPW and DM.

**Warfarin Dosing.** Most of the features are already binary with the exception of age, height, and weight, so these features were split into 5 buckets, where each bucket contains approximately the same number of datapoints. A variety of other preprocessing and data imputation steps were done based on recommendations from Consortium (2009). A patient’s possible outcomes are binary and thus we use the predicted class probabilities from  $\hat{\nu}$  to train our trees. Recall that in order to predict counterfactual outcomes for a given treatment option, we can only train a predictive model on the datapoints that were given said treatment in the historical policy; therefore, we evaluate a model’s performance on how well it generalizes on the entire population – particularly, we measure mean square loss and AUC. We find that, depending on the data generation process, different models of  $\hat{\nu}$  yield the best performance in the population: for the marginally randomized setting, random forests (RF) with balanced class weights were best; for the non-randomized settings, a combination of weighted RFs and logistic regression (LR) were best to account for class imbalance.

<sup>1</sup> See <https://www.gurobi.com/products/gurobi-optimizer/>

The propensity scores are learned by fitting an ML model to predict historical treatment assignment  $K$  from covariates  $X$  – we use decision trees (DT) because it yields the best performance in the population (lowest mean square loss and highest AUC).

**Policy Evaluation.** We evaluate a learned policy  $\pi$  in terms of its out-of-sample probability of assigning the optimal treatment to a patient (OOSP). In the next section, we compare all methods using this metric.

### 4.3. Other Benchmark Methods from the Literature

**Regress and Compare (R&C).** We employ the regress and compare approach as described in Section 1.2.3, which serves as a highly personalized (but uninterpretable) policy. For the experiments on synthetic data, we use linear regression to match the potential outcome functions (9). For the experiments on warfarin dosing, we use three models: 1) logistic regression (LR) with balanced class weights, 2) random forests (RF) with balanced class weights, and 3) the model (either LR and RF) with custom class weights resulting in the best performance in the population for each dataset, as described in Section 4.2.

**Causal Forests and Causal Trees (CF & CT).** Athey and Imbens (2016) propose a method that adapts the CART and random forest algorithm to learn heterogeneous treatment effects in a population. Each causal tree recursively partitions in a way that maximizes differences in treatment effects across splits (typically, difference between outcomes of the treatment and control groups). A causal forest grows hundreds to thousands of causal trees and aggregates results across the trees. However, CF and CT are concerned with estimation (rather than prescription) and they only consider binary treatment options. To produce a fairer comparison with our method, we propose the following approach, which was adapted from Kallus (2017). We first designate a baseline treatment option  $k_0$  (e.g.,  $k_0 = 0$ ). For all other treatment options  $k \in \mathcal{K} \setminus \{k_0\}$ , we use CF or CT to estimate  $\delta^k(x) := \mathbb{E}[Y|K = k, X = x] - \mathbb{E}[Y|K = k_0, X = x]$  (which we denote by  $\hat{\delta}^k(x)$ ). We also let  $\hat{\delta}^{k_0}(x)$  be 0 for all  $x$ . Finally, we prescribe treatments via  $\arg \min_{k \in \mathcal{K}} \hat{\delta}^k(x)$  to all individuals  $x \in \mathcal{X}$ . We implement both CF and CT using the grf package in R<sup>2</sup>.

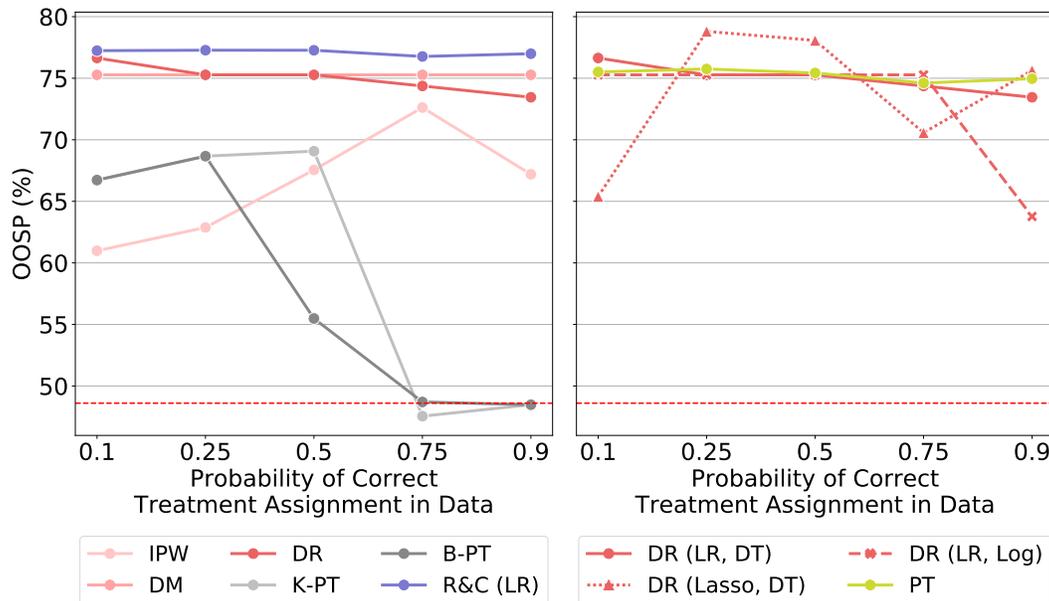
**Policy Tree (PT).** Zhou et al. (2023) propose to learn optimal prescriptive trees using a similar doubly robust objective. One critical difference from our method is that PT learn their trees recursively rather than using MIO. Therefore, while PT can learn trees with continuous data and in less time, it cannot handle additional constraints (see Section 3.5). As such, to further showcase the flexibility of our method, we will run experiments incorporating budget and fairness constraints in Section 4.5. We implement PT via the R package policytree<sup>3</sup> without discretizing the features.

<sup>2</sup> <https://github.com/grf-labs/grf>

<sup>3</sup> <https://github.com/grf-labs/policytree>

#### 4.4. Experimental Results

We now analyze the results for our method in relation to the aforementioned works. We first discuss our experiments on the synthetic data, and then move to a discussion on the warfarin experiments.

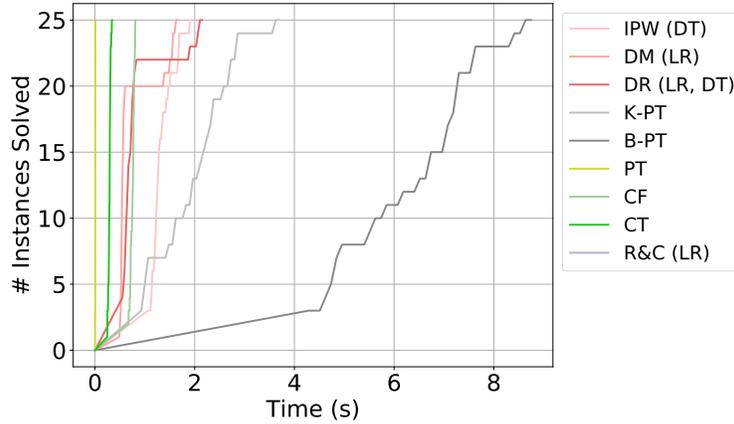


**Figure 5** Results on the synthetic dataset. The left graph compares the out-of-sample probability of correct treatment assignment of the best models for IPW, DM, and DR with a linear regressor regress & compare (R&C) approach, as well as methods described in Kallus (2017) (K-PT) and Bertsimas et al. (2019) (B-PT). The right graph shows the method described in Zhou et al. (2023) (PT), as well as our DR methods where at least one of the combined methods is correct, i.e., linear regression (LR) for estimating  $\nu$  and decision trees (DT) for estimating  $\mu$ . Lasso regression (Lasso) and logistic regression (Log) are both suboptimal models. Both figures are averages from trees of depth  $d=1$  with 2 leaf nodes.

##### 4.4.1. Synthetic Data

*Optimization Performance.* OOSP and solve times for experiments on the synthetic data are summarized in Figures 5 and 6, respectively (more details can be found in Table EC.1). For the MIO-based methods (ours, K-PT, B-PT) and PT, we train trees of depth 1 because one split is sufficient given our data generation process (see equation (9)). All methods solve to all instances in less than 9 seconds. In the following, we omit our analyses on CF and CT for the synthetic data for brevity; refer to our analysis on the warfarin dataset for an in-depth discussion on the advantages and drawbacks for CF/CT.

*Out-of-sample Performance.* Figure 5 shows the average OOSP for each method in dependence of the probability of correct treatment assignment,  $p$ . From the left subfigure, it can be seen that the DR and DM methods have competing performance at around 75% OOSP across  $p$ , while IPW



**Figure 6** Comparison of computational performance for the synthetic data. All methods solved to optimality in less than 9 seconds. We exclude DR (Lasso, DT) and DR (LR, Log) for brevity because the computational times are similar to DR (LR, DT), the best performing model.

performs consistently worse ranging from 60-72% OOSP. DM’s superior performance aligns with our expectations since the linear regressor is an accurate predictor of counterfactual outcomes, and consequently DR benefits from this modeling choice as well.

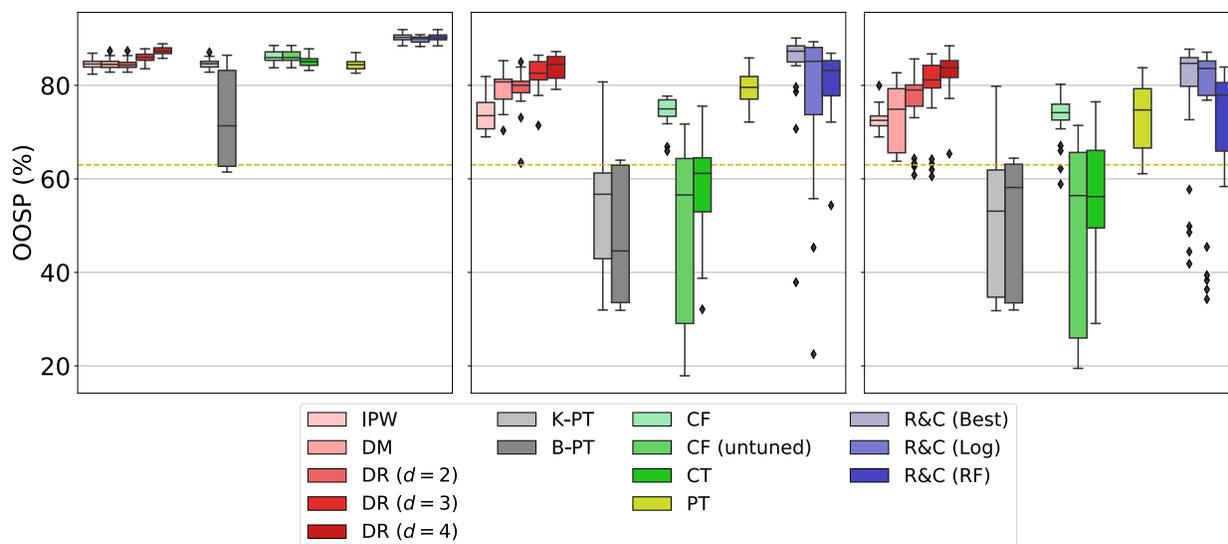
Further, we observe that PT maintains a similar performance to our DR method as expected, since both methods optimize the same objective. Recall that PT learns from the raw, continuous features while DR takes in the discretized versions of these features; the competing performance between these methods indicates that, in this application, there is virtually no loss in performance when learning from discretized data.

We also empirically test DR’s robustness property as discussed in Section 1.2.4. The dotted and dashed lines on the Figure 5 (right) represent instances when one of the two estimators in DR is inaccurate (i.e., uses Lasso to estimate  $\mu$  or Log to estimate  $\nu$ ). In these experiments, DR favors the other (better) estimator as expected, resulting in a relatively consistent performance across  $p$  despite the suboptimal predictors.

These results directly contrast K-PT and B-PT, which are not consistent across  $p$ . Both methods obtain decent performances when  $p = 0.5$  at 70% OOSP, i.e., when the historical policy is randomized – this behavior is expected given their methods rely on Assumption 1. However, both methods drop to around 50% (i.e., similar to random treatment assignment) when  $p \in \{0.75, 0.9\}$ . Their poor performance in these settings is particularly relevant since most observational data is based on an informed policy that presumably treats most patients correctly. The likely explanation for this drop in performance is that a tree of moderate depth is unlikely to partition the dataset sufficiently fine for Assumption 1 to hold.

Finally, we plot R&C (using linear regression) on Figure 5 (left), which corresponds to the best performance of any predictive model (since our counterfactuals are constructed via a linear

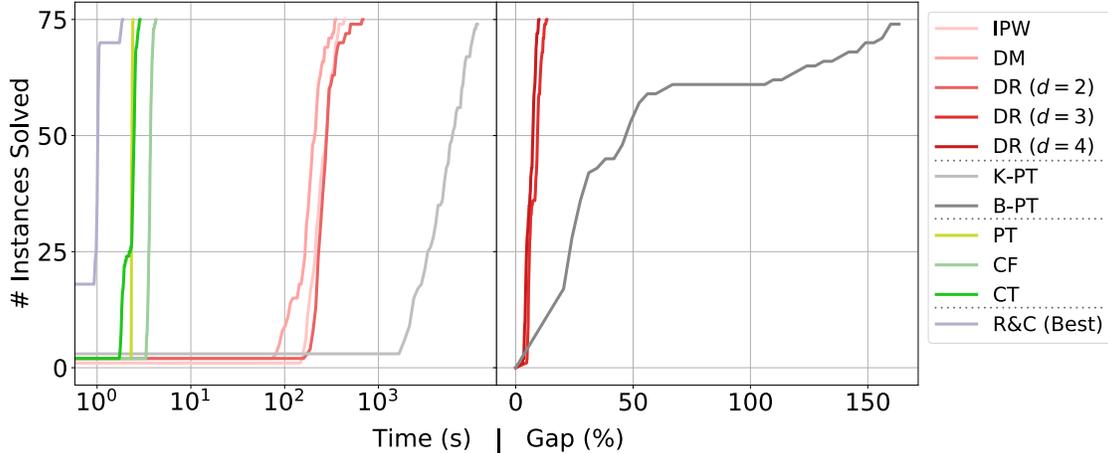
function, see equation (9)). The difference in performance between R&C and our trees quantifies the “price of interpretability”, i.e., the trade-off in performance we observe when opting for a more interpretable model compared to a higher-performing, less interpretable model. Following Jo et al. (2022a), we measure interpretability via a variant of “decision complexity”: the minimum number of parameters required for a model to determine a given datapoint’s treatment assignment. With this definition, our trees of depth 1 have a decision complexity of 3 (1 branching node and 2 leaf nodes), while the R&C model has a complexity of 6: 2 linear models corresponding to the 2 treatment options, each having 2 coefficients and 1 bias term. The price of interpretability is a small gap of on average 2 percentage points (p.p.) in optimal treatment assignment.



**Figure 7** Out-of-sample probability of correct treatment assignment (OOSP) over all 5 realizations on the warfarin dataset. The left graph shows the distribution for the randomized experiments. The middle (resp. right) figure shows the distribution of results for  $r = 0.06$  (resp.  $0.11$ ). Unless mentioned otherwise, all MIO-based methods (IPW, DM, DR, K-PT, B-PT) and PT are averages over trees with depth  $d = 2$  with 4 leaf nodes.

#### 4.4.2. Warfarin Dosing.

*Optimization Performance.* The OOSP and computational times for our experiments on the warfarin data are summarized in Figures 7 and 8, respectively (more details can be found in Table EC.2). Unless otherwise noted, results for all MIO-based methods (ours, K-PT, B-PT) and PT are from trees of depth  $d = 2$ . We chose this depth because 1) splitting on three features yield high enough performance, and 2) K-PT and B-PT did not scale to deeper trees and resulted in large optimality gaps, making for an unfair comparison. In the following, we also display results only for DR ( $d \in \{3, 4\}$ ) in order to analyze the performance improvements on deeper optimal trees.



**Figure 8** Comparison of computational performance for the warfarin data. All instances using the methods DR ( $d \in \{3, 4\}$ ) and B-PT did not solve to optimality within the 4-hour time limit, resulting in optimality gaps. All other methods solved to optimality within the time limit. We exclude CF (untuned) and R&C (Log, RF) for brevity because the computational times are similar to CF and R&C (Best), respectively.

In Figure 8, we see that the non-MIO methods find solutions consistently in less than 10 seconds. In contrast, our methods (for trees of depth 2) solve to optimality in the order of  $10^2$  seconds, and K-PT in  $10^3$  seconds. Our methods for deeper trees ( $d \in \{3, 4\}$ ) do not solve to optimality within the 4-hour time limit, but have small optimality gaps (less than 20%). Meanwhile, B-PT suffers from extremely large optimality gaps of  $> 150\%$  due to the quadratic objective.

*Out-of-sample Performance.* We first compare our methods’ performance across the three experimental designs. In the marginally randomized setting (Figure 7, left), IPW, DM, and DR all achieve similar performance of on average 84.5% OOSP. In the conditionally randomized settings ( $r \in \{0.06, 0.11\}$ , middle and right subfigures), DR performs at a little less than 80% OOSP, while DM and IPW perform marginally worse average at 76% and 73% on average, respectively. This overall decrease in performance is expected given that the conditionally randomized settings are harder to solve, but the fact that DR maintains decent performance compared to DM and IPW indicates that this is generally the strongest method to use, as discussed in Section 1.2.4.

When we increase the depth of our trees (DR,  $d \in \{3, 4\}$ ), performance noticeably increases at on average 86.0% and 87.3% OOSP, respectively, despite all of the instances not solving to optimality within the time limit. Increasing depth, however, has diminishing returns – ultimately, increasing a tree’s depth introduces a trade-off between interpretability and marginal increases in performance that only practitioners can make in their respective domains.

Similar to the synthetic experiments, we observe that K-PT performs well in the marginally randomized setting with competing performance to our methods, while B-PT suffered from optimality gaps. However, in the conditionally randomized settings, both K-PT and B-PT not only

have high variance, but also perform consistently worse than the simple predictor that only predicts  $K = 0$  for the entire population (yellow dashed line). While the simple predictor averages at around 63% OOSP, K-PT and B-PT perform at on average 51% and 48% respectively. We note that all instances of K-PT solved to optimality, indicating that its objective is ill-suited for conditionally randomized experiments, as analyzed in Section 2.1. B-PT, on the other hand, suffered from optimality gaps, but we can assume that its optimal performance might only be slightly better than K-PT since it suffers from similar problems as outlined in Section 2.2.

With regard to the CT and CF, both have similar performances to ours in the marginally randomized setting. CF, being a more complex method, has a slightly higher OOSP at 86.2%. However, in the conditionally randomized settings, both methods have similar performances to K-PT and B-PT when their hyperparameters are untuned. As is the case with many black-box models, it becomes necessary to tune their hyperparameters so as to avoid overfitting and to tailor to the data. Indeed, when we tune CF, we see drastic performance increases at 74% on average. In theory, CF could perform better than our DR method if more hyperparameters were searched since its model complexity is much higher. This process of hyperparameter tuning, however, illustrates a key trade-off between training a black-box causal model like CF – which not only requires tuning but also an analysis on heterogeneity ex-post – versus training an interpretable model (ours) – which requires careful inference of counterfactual predictions and propensity weighting but is in itself transparent. Moreover, we emphasize that CT and CF are designed to estimate heterogeneous treatment effects and not learn prescriptive policies. While we have adapted their method to do the latter, both approaches can only handle two treatment options, and thus when presented with more than two options, they can only make pairwise (local) decisions. In contrast, methods that are designed to prescribe policies like ours have a global view of treatment assignment. In the presence of more than two treatment options, CT – despite it being a simple decision tree – also becomes much less interpretable because it may produce partitions that differ from one treatment to another, and thus when comparing these treatments to build a prescriptive policy, CT ends up yielding a more complicated decision rule.

We also compare our DR method to PT, which again optimizes the same objective except that PT can learn from continuous data. Similar to our observations for the synthetic experiments, PT maintains roughly the same performance as DR – with the exception of the experiments when  $r = 0.11$ . Here, we observe that PT has a higher variance and performs marginally worse at 73% OOSP compared to 76% for DR. We posit that this degradation in performance is a problem of overfitting when PT is given the full suite of continuous features. On the training set, PT performs similarly to DR (see Figure EC.1 in Appendix EC.5), but in Figure 7 its trees fail to generalize

as well out-of-sample. As such, discretizing continuous variables has the benefit of occasionally serving as a form of regularization to avoid overfitting.

Finally, we quantify the “price of interpretability” similar to the synthetic experiments when we compare our method (say, DR  $d = 2$ ) with the R&C (Best) approach. In this case, DR’s decision complexity is 7 (3 branching and 4 leaf nodes) while R&C’s decision complexity is in the order of  $10^5$  (an exact number is hard to calculate given they are averages over a combination of random forests and logistic regression models, but for each random forest we train  $10^2$  trees, each with around  $10^3$  complexity). In the randomized setting and when  $r = 0.06$  (middle subfigure), R&C performs on average 5.75 p.p. better than DR ( $d = 2$ ). That gap shrinks to 0.77 p.p. when  $r = 0.11$  (right subfigure). This price of interpretability is even smaller when we compare DR ( $d = 4$ ) and R&C (Best); R&C performs on average 2.9, 0.0, and 0.0 p.p. better than DR ( $d = 4$ ) on the randomized,  $r = 0.06$ , and  $r = 0.11$  settings, respectively.

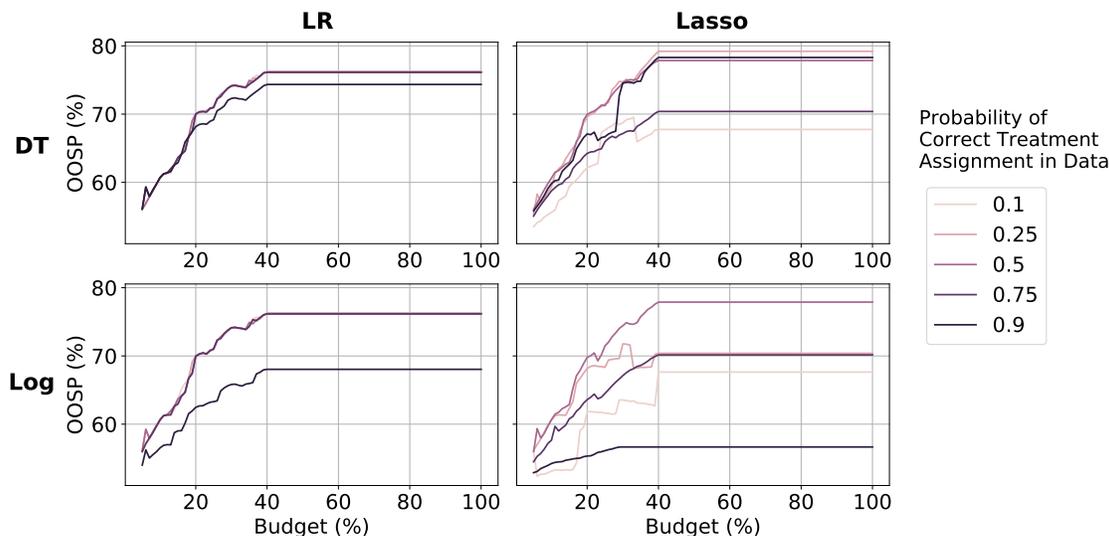
#### 4.5. Experiments with Budget and Fairness Constraints

As discussed in Section 3.5, our method has flexible modeling power; the MIO implementation allows for the addition of various constraints, which other methods like PT and CF cannot incorporate because they are learned recursively or via a heuristic. In this section, we perform additional experiments on our method that consider budget constraints on synthetic data and fairness constraints in the Warfarin dosing setting.

**Budget Constraints on Synthetic Data.** Equation 9 – which models the potential outcomes as functions of covariates  $X = (X^1, X^2)$  – controls treatment assignment to only be affected by  $X^1$ . In particular, the optimal split is easily modeled by a tree of depth 1 where observations with  $X^1 > 0$  (resp.  $X^1 \leq 0$ ) will be assigned treatment  $K = 1$  (resp.  $K = 0$ ), i.e., approximately half of the population receives either treatment. When we impose a constraint such that  $K = 1$  can only be given to  $< 50\%$  of the population, the optimal split will shift towards larger values of  $X^1$ . We model such behavior in this set of experiments. We run the DR method similar to the experiments described in Section 4.2, with the addition of a constraint  $\sum_{n \in \mathcal{B} \cup \mathcal{T}} \sum_{x \in \mathcal{X}} \sum_{i \in \mathcal{I}: X_i = x} z_x^{n, t_1} \leq |\mathcal{I}|C$ , with  $C \in [0.05, 0.4]$  in 0.01 increments.

Figure 9 shows OOSP for the doubly robust method in dependence of  $C$ , using either linear or lasso regression to estimate  $\nu$  and either a decision tree or logistic regression to estimate  $\mu$ . Over all choices of estimators, there is a clear trend whereby tightening the budget on  $K = 1$  decreases performance; our unconstrained performance is as high as 78% but goes down to 57% OOSP when  $C = 0.05$ . This trend is expected as the policy can only assign  $K = 1$  to fewer people when on average around half of the population needs said treatment. The effect of using different estimators is also stark, particularly in settings where the historical policy is not randomized ( $p \neq 0.5$ ). The

experiments using LR and DT (top left of Figure 9) reflect the best-performing estimators. Indeed, similar to the results in Section 4.4.1, we see that using LR/DT leads to consistent performance across all experimental designs. In contrast, when we use a worse model for  $\hat{\nu}$  (Lasso, right column in Figure 9) and/or a worse  $\hat{\mu}$  (Log, bottom row in Figure 9), we observe that performance is degraded when  $p \neq 0.5$ . As is the case with Figure 5 (right), we see that the doubly robust estimator somewhat maintains a high performance when either one of the estimators ( $\hat{\nu}$  or  $\hat{\mu}$ ) is correct.



**Figure 9** Out-of-sample probability of correct treatment assignment (OOSP) (y-axis) for the doubly robust method when the available budget – % of the population that can receive treatment  $K = 1$  – varies (x-axis) and using different estimators for  $\hat{\mu}$  and  $\hat{\nu}$ . Different line colors correspond to different probabilities of correct treatment assignment in the historical policy,  $p$ .

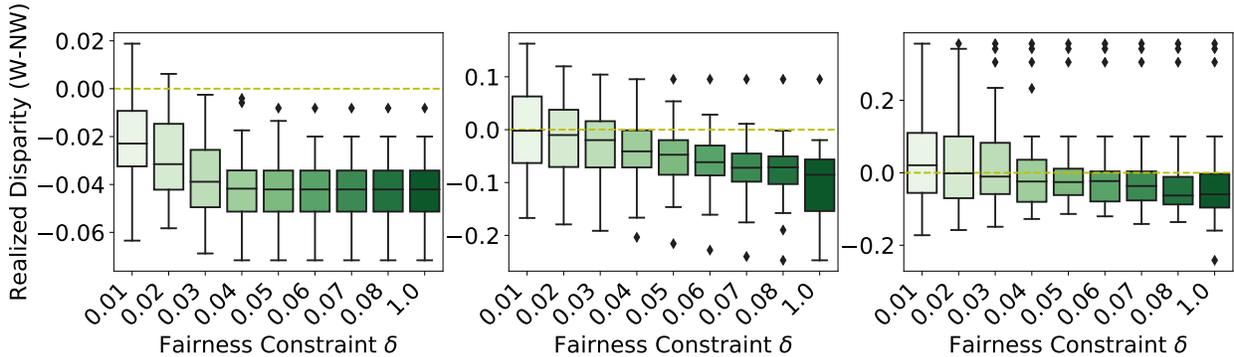
**Fairness Constraints on Warfarin Dosing.** The optimal trees learned in Section 4.4 yield the worst outcomes for White patients: over all trees, the median disparity in realized outcomes between White and non-White patients is 0.04 and 0.08 for the randomized and non-randomized settings, respectively, with no observations achieving full parity or the reverse disparity (being biased toward non-White people). Since our outcomes correspond to correct treatment assignment, this disparity can be interpreted as 4 and 8 percentage points fewer White people who receive their optimal treatment compared to non-White people. To close this gap, we run additional experiments imposing that the expected outcomes (as estimated by the DM objective) between White and non-White people be at most  $\delta$ , i.e.,

$$|M_{\text{white}} - M_{\text{!white}}| \leq \delta,$$

where

$$M_p = \frac{1}{|\{i \in \mathcal{I} : P_i = p\}|} \sum_{n \in \mathcal{B} \cup \mathcal{T}} \sum_{k \in \mathcal{K}} \sum_{x \in \mathcal{X} : P_x = p} z_x^{n,t_k} \sum_{i \in \mathcal{I} : X_i = x} \hat{v}_k(X_i).$$

We note that while the DR objective leads to generally better performance, we found that the DM method estimates outcomes that are closest to the true expected outcomes; we believe the DR method is less-suited to reflect true expected outcomes because the conditionally randomized settings have many observations with extremely low propensity weights, which severely biases the estimates. We let  $\delta \in [0.01, 0.08]$  in 0.01 increments, and display realized disparity between the two groups as  $\delta$  varies in Figure 10 across the three experimental designs. As expected, decreasing  $\delta$  has a direct relationship with a decrease in realized outcome disparity. However, we note that enforcing  $\delta$  does not guarantee that actual disparity is within  $\delta$  because we can only expect the constraint to hold in expectation rather than almost surely in a finite sample. For the randomized setting (left subfigure), enforcing  $\delta = 0.01$  leads to a median disparity of around  $-0.02$  across all seeds. Similarly, for the conditionally randomized settings, enforcing  $\delta = 0.01$  and  $0.02$  leads to a median of full parity for  $r = 0.06$  and  $r = 0.11$ , respectively.



**Figure 10** Disparity of realized outcomes between White (W) and non-White (NW) patients (y-axis) as fairness parameter  $\delta$  varies (x-axis), over the randomized setting (left) and non-randomized settings ( $r \in \{0.06, 0.11\}$  in the middle and right, respectively). The dashed yellow line indicates true parity.

## 5. Summary

We presented MIO formulations to learn treatment assignment policies in the form of prescriptive trees from observational data. We showed that our methods are asymptotically exact, which sets us apart from existing literature on prescriptive trees that 1) require data to originate from marginally randomized experiments; 2) require the learned trees to be very deep to yield correct treatment assignments; or 3) cannot handle budget and/or fairness constraints. Our experiments show that our methods perform consistently better across different experiment designs, in some

cases performing  $3\times$  better than the state-of-the-art. We also showcase our method’s flexible modeling power by running additional experiments incorporating budget and fairness constraints: a feature that no other method in the literature can handle efficiently.

## Acknowledgments

N. Jo acknowledges support from the Epstein Institute at the University of Southern California. P. Vayanos and S. Aghaei are funded in part by the National Science Foundation under CAREER award number 2046230. They are grateful for this support. N. Jo, P. Vayanos, and S. Aghaei gratefully acknowledge support from the Hilton C. Foundation, the Homeless Policy Research Institute, and the Home for Good foundation under the “C.E.S. Triage Tool Research & Refinement” grant. A. Gómez is funded in part by the National Science Foundation under grant 2006762.

## References

- Abrevaya J, Hsu YC, Lieli RP (2015) Estimating conditional average treatment effects. *Journal of Business & Economic Statistics* 33(4):485–505.
- Aghaei S, Azizi MJ, Vayanos P (2019) Learning optimal and fair decision trees for non-discriminative decision-making. *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 1418–1426.
- Aghaei S, Gómez A, Vayanos P (2020) Strong optimal classification trees. Technical report, University of Southern California, URL [http://www.optimization-online.org/DB\\_FILE/2021/01/8220.pdf](http://www.optimization-online.org/DB_FILE/2021/01/8220.pdf).
- Atamtürk A, Gómez A (2019) Rank-one convexification for sparse regression. arXiv:1901.10334.
- Athey S, Imbens G (2016) Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences of the United States of America* 113(27), ISSN 10916490, URL <http://dx.doi.org/10.1073/pnas.1510489113>.
- Azizi MJ, Vayanos P, Wilder B, Rice E, Tambe M (2018) Designing fair, efficient, and interpretable policies for prioritizing homeless youth for housing resources. *International Conference on the Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, 35–51 (Springer).
- Bartle RG (2001) *A modern theory of integration*, volume 32 (American Mathematical Soc.).
- Bastani H, Bayati M (2020) Online decision making with high-dimensional covariates. *Operations Research* 68(1):276–294.
- Bennett M, Vielma JP, Zubizarreta JR (2020) Building representative matched samples with multi-valued treatments in large observational studies. *Journal of Computational and Graphical Statistics* 29(4):744–757.
- Bertsimas D, Cory-Wright R, Pauphilet J (2020a) Solving large-scale sparse pca to certifiable (near) optimality. arXiv:2005.05195.
- Bertsimas D, Dunn J (2017) Optimal classification trees. *Machine Learning* 106(7):1039–1082.

- Bertsimas D, Dunn J, Mundru N (2019) Optimal Prescriptive Trees. *INFORMS Journal on Optimization* 1(2), ISSN 2575-1484, URL <http://dx.doi.org/10.1287/ijoo.2018.0005>.
- Bertsimas D, Farias VF, Trichakis N (2013) Fairness, efficiency, and flexibility in organ allocation for kidney transplantation. *Operations Research* 61(1):73–87.
- Bertsimas D, Kallus N, Weinstein AM, Zhuo YD (2017) Personalized diabetes management using electronic medical records. *Diabetes care* 40(2):210–217.
- Bertsimas D, Van Parys B, et al. (2020b) Sparse high-dimensional regression: Exact scalable algorithms and phase transitions. *The Annals of Statistics* 48(1):300–323.
- Beygelzimer A, Langford J (2009) The offset tree for learning with partial labels. *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, 129–138.
- Bolukbasi T, Chang KW, Zou JY, Saligrama V, Kalai AT (2016) Man is to computer programmer as woman is to homemaker? debiasing word embeddings. *Advances in neural information processing systems* 29.
- Bottou L, Peters J, Quiñero-Candela J, Charles DX, Chikering DM, Portugaly E, Ray D, Simard P, Snelson E (2013) Counterfactual reasoning and learning systems: The example of computational advertising. *Journal of Machine Learning Research* 14(65):3207–3260, URL <http://jmlr.org/papers/v14/bottou13a.html>.
- Breiman L (2017) *Classification and regression trees* (Routledge).
- Consortium IWP (2009) Estimation of the warfarin dose with clinical and pharmacogenetic data. *New England Journal of Medicine* 360(8):753–764.
- Dey SS, Mazumder R, Wang G (2018) A convex integer programming approach for optimal sparse pca. arXiv:1810.09062.
- Dougherty J, Kohavi R, Sahami M (1995) Supervised and unsupervised discretization of continuous features. *Machine learning proceedings 1995*, 194–202 (Elsevier).
- Dudík M, Langford J, Li H (2011) Doubly robust policy evaluation and learning. *Proceedings of the 28th International Conference on Machine Learning, ICML 2011*.
- Dwork C, Hardt M, Pitassi T, Reingold O, Zemel R (2012) Fairness through awareness. *Proceedings of the 3rd innovations in theoretical computer science conference*, 214–226.
- Elmachtoub A, Liang JCN, McNellis R (2020) Decision trees for decision-making under the predict-then-optimize framework. *International Conference on Machine Learning*, 2858–2867 (PMLR).
- Fan Q, Hsu YC, Lieli RP, Zhang Y (2020) Estimation of conditional average treatment effects with high-dimensional data. *Journal of Business & Economic Statistics* 1–15.
- Fischetti M, Jo J (2018) Deep neural networks and mixed integer linear optimization. *Constraints* 23(3):296–309.

- Flume PA, O’Sullivan BP, Robinson KA, Goss CH, Mogayzel Jr PJ, Willey-Courand DB, Bujan J, Finder J, Lester M, Quittell L, et al. (2007) Cystic fibrosis pulmonary guidelines: chronic medications for maintenance of lung health. *American journal of respiratory and critical care medicine* 176(10):957–969.
- Goldenshluger A, Zeevi A (2013) A linear response bandit problem. *Stochastic Systems* 3(1):230–261.
- Gómez A, Prokopyev OA (2021) A mixed-integer fractional optimization approach to best subset selection. *INFORMS Journal on Computing* 176(10):957–969.
- Hardt M, Price E, Srebro N (2016) Equality of opportunity in supervised learning. *Advances in neural information processing systems* 29.
- Hazimeh H, Mazumder R, Saab A (2020) Sparse regression at scale: Branch-and-bound rooted in first-order optimization. arXiv:2004.06152.
- Hernán M, Robins J (2019) *Causal Inference: What If*. Chapman & Hall/CRC (Taylor & Francis), ISBN 9781420076165, URL [https://books.google.com/books?id=\\_KnHIAAACAAJ](https://books.google.com/books?id=_KnHIAAACAAJ).
- Horvitz DG, Thompson DJ (1952) A generalization of sampling without replacement from a finite universe. *Journal of the American statistical Association* 47(260):663–685.
- Jo N, Aghaei S, Benson J, Gómez A, Vayanos P (2022a) Learning optimal fair classification trees. *arXiv preprint arXiv:2201.09932* .
- Jo N, Tang B, Dullerud K, Aghaei S, Rice E, Vayanos P (2022b) Fairness in contextual resource allocation systems: Metrics and incompatibility results. *arXiv preprint arXiv:2212.01725* .
- Kallus N (2017) Recursive partitioning for personalization using observational data. *34th International Conference on Machine Learning, ICML 2017*, volume 4.
- Kallus N (2018) Balanced policy evaluation and learning. *Advances in Neural Information Processing Systems*, volume 2018-December, ISSN 10495258.
- Khalil EB, Gupta A, Dilkina B (2018) Combinatorial attacks on binarized neural networks. arXiv:1810.03538.
- Kotsiantis S, Kanellopoulos D (2006) Discretization techniques: A recent survey. *GESTS International Transactions on Computer Science and Engineering* 32(1):47–58.
- Li L, Chu W, Langford J, Schapire RE (2010) A contextual-bandit approach to personalized news article recommendation. *Proceedings of the 19th international conference on World wide web*, 661–670.
- Linardatos P, Papastefanopoulos V, Kotsiantis S (2020) Explainable ai: A review of machine learning interpretability methods. *Entropy* 23(1):18.
- Lunceford JK, Davidian M (2004) Stratification and weighting via the propensity score in estimation of causal treatment effects: a comparative study. *Statistics in medicine* 23(19):2937–2960.
- Mazumder R, Meng X, Wang H (2022) Quant-bnb: A scalable branch-and-bound method for optimal decision trees with continuous features. *International Conference on Machine Learning*, 15255–15277 (PMLR).

- Menickelly M, Günlük O, Kalagnanam J, Scheinberg K (2016) Optimal generalized decision trees via integer programming. *CoRR*, *abs/1612.03225* .
- Mintz Y, Aswani A, Kaminsky P, Flowers E, Fukuoka Y (2017) Behavioral analytics for myopic agents. *arXiv preprint arXiv:1702.05496* .
- Mišić VV (2020) Optimization of tree ensembles. *Operations Research* 68(5):1605–1624.
- Nikolaev AG, Jacobson SH, Cho WKT, Sauppe JJ, Sewell EC (2013) Balance optimization subset selection (boss): An alternative approach for causal inference with observational data. *Operations Research* 61(2):398–412.
- Olfat M, Aswani A (2018) Spectral algorithms for computing fair support vector machines. *International Conference on Artificial Intelligence and Statistics*, 1933–1942 (PMLR).
- Powers S, Qian J, Jung K, Schuler A, Shah NH, Hastie T, Tibshirani R (2018) Some methods for heterogeneous treatment effect estimation in high dimensions. *Statistics in medicine* 37(11):1767–1787.
- Qian M, Murphy SA (2011) Performance guarantees for individualized treatment rules. *Annals of statistics* 39(2):1180.
- Rawls J (1974) Some reasons for the maximin criterion. *The American Economic Review* 64(2):141–146.
- Rubin DB (2006) *Matched sampling for causal effects* (Cambridge University Press), ISBN 8580000722116.
- Rucker DD, McShane BB, Preacher KJ (2015) A researcher’s guide to regression, discretization, and median splits of continuous variables. *Journal of Consumer Psychology* 25(4):666–678.
- Rudin C (2019) Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence* 1(5):206–215.
- Tjandraatmadja C, Anderson R, Huchette J, Ma W, Patel K, Vielma JP (2020) The convex relaxation barrier, revisited: Tightened single-neuron relaxations for neural network verification. *arXiv:2006.14076*.
- Verwer S, Zhang Y (2019) Learning optimal classification trees using a binary linear program formulation. *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 1625–1632.
- Wager S, Athey S (2018) Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association* 113(523):1228–1242.
- Wilson ZT, Sahinidis NV (2017) The alamo approach to machine learning. *Computers & Chemical Engineering* 106:785–795.
- Xie W, Deng X (2020) Scalable algorithms for the sparse ridge regression. *SIAM Journal on Optimization* 30(4):3359–3386.
- Xu H, Wang K, Vayanos P, Tambe M (2018) Strategic coordination of human patrollers and mobile sensors with signaling for security games. *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.

- Yıldız S, Vielma JP (2013) Incremental and encoding formulations for mixed integer programming. *Operations Research Letters* 41(6):654–658.
- Zafar MB, Valera I, Røgriguez MG, Gummadi KP (2017) Fairness constraints: Mechanisms for fair classification. *Artificial intelligence and statistics*, 962–970 (PMLR).
- Zemel R, Wu Y, Swersky K, Pitassi T, Dwork C (2013) Learning fair representations. *International conference on machine learning*, 325–333 (PMLR).
- Zenios SA, Chertow GM, Wein LM (2000) Dynamic allocation of kidneys to candidates on the transplant waiting list. *Operations Research* 48(4):549–569.
- Zhou Z, Athey S, Wager S (2023) Offline multi-action policy learning: Generalization and optimization. *Operations Research* 71(1):148–183.
- Zubizarreta JR, Keele L (2017) Optimal multilevel matching in clustered observational studies: A case study of the effectiveness of private schools under a large-scale voucher system. *Journal of the American Statistical Association* 112(518):547–560.

## Electronic Companion

### EC.1. K-PT MIO Formulation

In this section, we provide an adapted version of the formulation from Kallus (2017) that we use in our experiments in Section 4.4. As we will soon discuss, we adapt the original formulation for both brevity and so that it better aligns with the notation we use. Extending the notation from the problem statement, we construct a perfect binary tree and number the nodes 1 through  $2^{d+1} - 1$  in the order in which they appear on a breadth-first search. We let  $R_{nm} \in \{1, -1\}$  capture the relationship of node  $n$  with its ancestor(s)  $m \in \mathcal{A}(n)$ , i.e.,  $R_{nm}$  equals 1 iff we use  $m$ 's right branch to reach  $n$ , and  $-1$  otherwise.

We now define the decision variables used in the K-PT formulation. For every branching node  $n \in \mathcal{B}$  and feature  $f \in \mathcal{F}$ , we let the binary variable  $b_{nf}$  indicate if feature  $f$  is selected for branching at node  $n$ . We let the binary variable  $\chi_{in}$  equal 1 iff datapoint  $i$  goes left on branching node  $n$ . Further, we define membership variables  $\lambda_{in} \in \{0, 1\}$ , which takes value 1 iff datapoint  $i$  flows to terminal node  $n \in \mathcal{T}$ . For  $k \in \mathcal{K}$ , we let  $w_{nk} \in \{0, 1\}$  equal 1 iff treatment  $k$  is selected at terminal node  $n$ . Also for each terminal node  $n$ , define  $\rho_n$  to be the average treatment outcome of all datapoints in that node. Let  $\eta_{in}$  be the product of  $\lambda_{in}$  and  $\rho_n$ , which captures the average outcome associated with datapoint  $i$ .

In order to linearize  $\eta$ , Kallus (2017) introduces big- $M$  constraints. To this end, let  $\bar{Y}_i = Y_i - \min_{j \in \mathcal{I}} Y_j$ ,  $\bar{Y}_{max} = \max_i \bar{Y}_i$ , and  $M = \bar{Y}_{max} (\max_{k \in \mathcal{K}} \sum_{i \in \mathcal{I}} \mathbb{I}[K_i = k])$ . The approach now reads:

$$\text{maximize } \sum_{i \in \mathcal{I}} \sum_{n \in \mathcal{T}} \eta_{in} \quad (\text{EC.1a})$$

$$\text{subject to } \lambda_{in} \leq \frac{1 + R_{nm}}{2} - R_{nm} \chi_{im} \quad \forall i \in \mathcal{I}, n \in \mathcal{T}, m \in a(n) \quad (\text{EC.1b})$$

$$\lambda_{in} \geq 1 - \sum_{\substack{m \in a(n) \\ R_{mn}=1}} \chi_{im} + \sum_{\substack{m \in a(n) \\ R_{mn}=-1}} (-1 + \chi_{im}) \quad \forall i \in \mathcal{I}, n \in \mathcal{T} \quad (\text{EC.1c})$$

$$\sum_{f \in \mathcal{F}} b_{nf} = 1 \quad \forall n \in \mathcal{B} \quad (\text{EC.1d})$$

$$\chi_{in} = \sum_{\substack{f \in \mathcal{F}: \\ X_f^i = 0}} b_{nf} \quad \forall i \in \mathcal{I}, n \in \mathcal{B} \quad (\text{EC.1e})$$

$$\eta_{in} \leq \bar{Y}_{max} \lambda_{in}, \eta_{in} \leq \rho_n \quad \forall i \in \mathcal{I}, p \in \mathcal{T} \quad (\text{EC.1f})$$

$$\eta_{in} \geq \rho_n - \bar{Y}_{max} (1 - \lambda_{in}) \quad \forall i \in \mathcal{I}, p \in \mathcal{T} \quad (\text{EC.1g})$$

$$\sum_{i: K_i = k} (\eta_{in} - \lambda_{in} \bar{Y}_i) \leq M(1 - w_{nk}) \quad \forall p \in \mathcal{T}, k \in \mathcal{K} \quad (\text{EC.1h})$$

$$\sum_{i: K_i = k} (\eta_{in} - \lambda_{in} \bar{Y}_i) \geq M(w_{nk} - 1) \quad \forall p \in \mathcal{T}, k \in \mathcal{K} \quad (\text{EC.1i})$$

$$\sum_{k \in \mathcal{K}} w_{nk} = 1 \quad \forall n \in \mathcal{T} \quad (\text{EC.1j})$$

$$w_{nk} \in \{0, 1\} \quad \forall n \in \mathcal{T}, k \in \mathcal{K} \quad (\text{EC.1k})$$

$$b_{nf} \in \{0, 1\} \quad \forall n \in \mathcal{B}, f \in \mathcal{F} \quad (\text{EC.1l})$$

$$\lambda_{in} \in \{0, 1\} \quad \forall i \in \mathcal{I}, n \in \mathcal{T} \quad (\text{EC.1m})$$

$$\chi_{in} \in \{0, 1\} \quad \forall n \in \mathcal{B}, i \in \mathcal{I} \quad (\text{EC.1n})$$

$$\eta_{in} \in \mathbb{R}_+ \quad \forall n \in \mathcal{T}, i \in \mathcal{I} \quad (\text{EC.1o})$$

$$\rho_n \in \mathbb{R}_+ \quad \forall n \in \mathcal{T}. \quad (\text{EC.1p})$$

The objective function (EC.1a) sums the associated averages over all datapoints and terminal nodes. Constraints (EC.1b) and (EC.1c) mode the flow of datapoints from the root down to the terminal nodes. Constraints (EC.1d) state that, at each branching node, the tree must branch on exactly one feature, whereas (EC.1e) defines  $\chi$  using the associated branching decisions. Constraints (EC.1f) and (EC.1g) both use big- $M$  constraints to linearize  $\eta$ . Constraints (EC.1h) and (EC.1i) ensure that the predicted outcome for a given treatment at each leaf of the tree is indeed the average observed outcome. Lastly, constraint (EC.1j) ensures that exactly one treatment is assigned at every terminal node.

Note that we have here converted the formulation from Kallus (2017) to a maximization problem to account for the different representation of outcome  $Y$  (higher values preferred). Further, we make the following changes from the formulation presented in Kallus (2017): *(i)* constraint (EC.1c) is slightly different, since the original constraint does not result in a correct behavior; *(ii)* for simplicity, we removed constraints inspired from Yıldız and Vielma (2013) from the formulation and replaced it with constraint (EC.1d), which serves the same function; *(iii)* we removed constraint  $\sum_{i:K_i=k} \lambda_{in} \geq N_{\min}$  from the formulation to allow for an equivalent comparison with our proposed methods, and in tuning the parameter for our experiments, we did not find statistical improvements.

## EC.2. B-PT MIO Formulation

Bertsimas et al. (2019) briefly discuss a coordinate descent algorithm to learn an optimal decision tree with the loss function form

$$\min_T \text{error}(T, D) + \alpha \cdot \text{complexity}(T),$$

where  $T$  is the tree being optimized,  $D$  is the training data, the “error” function measures how well the tree fits training data, and the second term is a penalization term that controls the trade-off

between the quality of the tree and the tree's complexity/size. The algorithm iterates through different branching decisions and treatment assignment options until no possible improvements are found (i.e., the tree is a local minimum). This process is repeated for different randomly generated starting trees, and “the lowest objective function is selected as the final solution”. The documentation for this method is limited and there is no existing codebase allowing us to replicate this algorithm. Hence, in this section, we describe the MIO implementation that we use in our experiments for B-PT, which is adapted from Kallus (2017) to optimize the objective function proposed by Bertsimas et al. (2019).

Building from Kallus (2017), Bertsimas et al. (2019) added a regularization term that penalizes trees whose leaves have high variance. Let  $\beta_{nk}$  be the average outcome over all datapoints that were assigned treatment  $k$  in terminal node  $n$ . Let  $g_i$  be the *empirical* average of the outcomes at the terminal node where datapoint  $i$  lands over all datapoints that were assigned the same treatment as  $i$  (i.e.  $g_i = \beta_{nK_i}$ ). Finally, let  $\theta$  control the regularization strength. The formulation becomes:

$$\text{maximize } \theta \sum_{i \in \mathcal{I}} \sum_{n \in \mathcal{T}} \eta_{in} - (1 - \theta) \sum_{i \in \mathcal{I}} (Y_i - g_i)^2 \quad (\text{EC.2a})$$

$$\text{subject to } \lambda_{in} \leq \frac{1 + R_{nm}}{2} - R_{nm} \chi_{im} \quad \forall i \in \mathcal{I}, n \in \mathcal{T}, m \in a(n) \quad (\text{EC.2b})$$

$$\lambda_{in} \geq 1 - \sum_{\substack{m \in a(n) \\ R_{mn}=1}} \chi_{im} + \sum_{\substack{m \in a(n) \\ R_{mn}=-1}} (-1 + \chi_{im}) \quad \forall i \in \mathcal{I}, n \in \mathcal{T} \quad (\text{EC.2c})$$

$$\sum_{f \in \mathcal{F}} b_{nf} = 1 \quad \forall n \in \mathcal{B} \quad (\text{EC.2d})$$

$$\chi_{in} = \sum_{\substack{f \in \mathcal{F}: \\ X_f^i = 0}} b_{nf} \quad \forall i \in \mathcal{I}, n \in \mathcal{B} \quad (\text{EC.2e})$$

$$\eta_{in} \leq \bar{Y}_{max} \lambda_{in}, \eta_{in} \leq \rho_n \quad \forall i \in \mathcal{I}, p \in \mathcal{T} \quad (\text{EC.2f})$$

$$\eta_{in} \geq \rho_n - \bar{Y}_{max} (1 - \lambda_{in}) \quad \forall i \in \mathcal{I}, p \in \mathcal{T} \quad (\text{EC.2g})$$

$$\sum_{i: K_i=k} (\eta_{in} - \lambda_{in} \bar{Y}_i) \leq M(1 - w_{nk}) \quad \forall p \in \mathcal{T}, k \in \mathcal{K} \quad (\text{EC.2h})$$

$$\sum_{i: K_i=k} (\eta_{in} - \lambda_{in} \bar{Y}_i) \geq M(w_{nk} - 1) \quad \forall p \in \mathcal{T}, k \in \mathcal{K} \quad (\text{EC.2i})$$

$$\sum_{k \in \mathcal{K}} w_{nk} = 1 \quad \forall n \in \mathcal{T} \quad (\text{EC.2j})$$

$$g_i - \beta_{nK_i} \leq M(1 - \lambda_{in}) \quad \forall i \in \mathcal{I}, n \in \mathcal{T} \quad (\text{EC.2k})$$

$$g_i - \beta_{nK_i} \geq M(\lambda_{in} - 1) \quad \forall i \in \mathcal{I}, n \in \mathcal{T} \quad (\text{EC.2l})$$

$$w_{nk} \in \{0, 1\} \quad \forall n \in \mathcal{T}, k \in \mathcal{K} \quad (\text{EC.2m})$$

$$b_{nf} \in \{0, 1\} \quad \forall n \in \mathcal{B}, f \in \mathcal{F} \quad (\text{EC.2n})$$

$$\lambda_{in} \in \{0, 1\} \quad \forall i \in \mathcal{I}, n \in \mathcal{T} \quad (\text{EC.2o})$$

$$\chi_{in} \in \{0, 1\} \quad \forall n \in \mathcal{B}, i \in \mathcal{I} \quad (\text{EC.2p})$$

$$\eta_{in} \in \mathbb{R}_+ \quad \forall n \in \mathcal{T}, i \in \mathcal{I} \quad (\text{EC.2q})$$

$$\rho_n \in \mathbb{R}_+ \quad \forall n \in \mathcal{T} \quad (\text{EC.2r})$$

$$g_i \in \mathbb{R}_+ \quad \forall n \in \mathcal{T}. \quad (\text{EC.2s})$$

The objective function (EC.2a) now penalizes high variance. All constraints remain the same with the exception of (EC.2k) and (EC.2l), which force  $g_i$  to be  $\beta_{nK_i}$  when datapoint  $i$  lands in terminal node  $n$ . There are no additional constraints that define  $\beta_{nk}$  because the optimal solution for  $\min_{\beta_{nk}} \sum_{i=1}^m (a_i - \beta_{nk})^2$  is  $\beta_{nk}^* = \sum_{i=1}^m \frac{a_i}{m}$ , where  $\{a_1, \dots, a_m\}$  are datapoints in terminal node  $n$  that were assigned treatment  $k$  in the data.

### EC.3. Warfarin Dosage

We present an equation from Consortium (2009) that we use in our experiments in Section 4.4 to determine the optimal warfarin dose for a patient. In particular, it allows us to generate patient counterfactuals and evaluate the learned policies for our experiments. Note that VKORC1 and CYP2C9 denote genotypes. If  $W$  is the optimal weekly warfarin dosage, then the equation becomes:

$$\begin{aligned} \sqrt{W} = & 5.6044 - 0.2614 \times \text{Age in decades} + 0.0087 \times \text{Height in cm} + 0.0128 \times \text{Weight in kg} \\ & - 0.8677 \times \text{VKORC1 A/G} - 1.6974 \times \text{VKORC1 A/A} - 0.4854 \times \text{VKORC1 genotype unknown} \\ & - 0.5211 \times \text{CYP2C9*1/*2} - 0.9357 \times \text{CYP2C9*1/*3} - 1.0616 \times \text{CYP2C9*2/*2} \\ & - 1.9206 \times \text{CYP2C9*2/*3} - 2.3312 \times \text{CYP2C9*3/*3} - 0.2188 \times \text{CYP2C9 genotype unknown} \\ & - 0.1092 \times \text{Asian race} - 0.2760 \times \text{Black or African American} - 0.1032 \times \text{Missing or Mixed race} \\ & + 1.1816 \times \text{Enzyme inducer status} - 0.5503 \times \text{Amiodarone status} \end{aligned} \quad (\text{EC.3})$$

## EC.4. Experiment Results (Raw)

In this section, we provide tables that contain the raw results for the experiments summarized in Section 4.4.

**Table EC.1** Companion table to the experiments on synthetic data in Section 4.4.1. The first and second number correspond to the average and standard deviation of the optimality gap, solving time, out-of-sample (OOS) regret, and OOS probability of correct treatment assignment (OOSP) across 5 random samples.

Depth	Method	Model	Gap	Solving Time (s)	OOS Regret	OOSP (%)
1	IPW	DT	0.00 ± 0.00	1.34 ± 0.22	123.18 ± 95.63	66.25 ± 15.84
1	IPW	Log	0.00 ± 0.00	1.31 ± 0.31	138.12 ± 114.65	63.89 ± 18.00
1	DM	LR	0.00 ± 0.00	0.76 ± 0.40	64.22 ± 38.06	75.28 ± 9.87
1	DM	Lasso	0.00 ± 0.00	0.57 ± 0.09	200.68 ± 72.26	53.44 ± 11.00
1	DR	DT, LR	0.00 ± 0.00	0.86 ± 0.46	65.76 ± 40.68	75.01 ± 10.13
1	DR	DT, Lasso	0.00 ± 0.00	2.08 ± 0.82	77.04 ± 70.23	73.70 ± 12.83
1	DR	Log, LR	0.00 ± 0.00	0.95 ± 0.56	78.78 ± 55.88	72.98 ± 11.56
1	DR	Log, Lasso	0.00 ± 0.00	1.79 ± 1.01	111.47 ± 68.44	67.16 ± 11.88
1	K-PT	-	0.00 ± 0.00	1.97 ± 0.76	161.83 ± 103.62	60.10 ± 17.46
1	B-PT	-	0.00 ± 0.00	6.28 ± 1.36	179.04 ± 106.95	57.62 ± 17.63
1	PT	DT, LR	0.00 ± 0.00	0.01 ± 0.00	62.01 ± 37.00	75.25 ± 10.29
-	CF	-	0.00 ± 0.00	0.75 ± 0.04	76.03 ± 57.20	73.13 ± 12.05
-	CT	-	0.00 ± 0.00	0.29 ± 0.02	92.20 ± 74.63	71.53 ± 13.64
-	R&C	LR	0.00 ± 0.00	0.23 ± 0.00	51.81 ± 32.87	77.11 ± 10.09

**Table EC.2** Companion table to the experiments on synthetic data in Section 4.4.2. The first and second number correspond to the average and standard deviation of the optimality gap, solving time, out-of-sample (OOS) regret, and OOS probability of correct treatment assignment (OOSP) across 25 random samples.

Depth	Method	Model	Gap	Solving Time (s)	OOS Regret	OOSP (%)
2	IPW	DT	0.00 ± 0.00	255.45 ± 65.13	317.39 ± 82.90	77.10 ± 5.98
2	DM	RF/Log	0.00 ± 0.00	192.03 ± 64.13	288.79 ± 88.46	79.16 ± 6.38
2	DR	DT, RF/Log	0.00 ± 0.00	284.58 ± 83.83	279.32 ± 87.00	79.85 ± 6.28
2	K-PT	-	0.00 ± 0.03	5435.72 ± 2792.71	522.12 ± 264.43	62.33 ± 19.08
2	B-PT	-	52.30 ± 43.78	14333.53 ± 674.39	601.39 ± 239.26	56.61 ± 17.26
2	PT	DT, RF/Log	0.00 ± 0.00	2.35 ± 0.03	257.73 ± 76.39	81.40 ± 5.51
3	IPW	DT	13.56 ± 4.15	14405.89 ± 0.50	246.35 ± 48.07	82.23 ± 3.47
3	DM	RF/Log	2.10 ± 1.05	14406.33 ± 0.55	252.92 ± 69.22	81.75 ± 4.99
3	DR	DT, RF/Log	8.08 ± 2.38	14406.65 ± 0.70	243.00 ± 78.18	82.47 ± 5.64
4	IPW	DT	9.51 ± 2.91	14412.07 ± 1.17	221.00 ± 39.47	84.05 ± 2.85
4	DM	RF/Log	1.48 ± 0.73	14412.46 ± 1.05	243.23 ± 72.31	82.45 ± 5.22
4	DR	DT, RF/Log	6.26 ± 1.90	14413.10 ± 1.57	210.95 ± 48.01	84.78 ± 3.46
-	CF	-	0.00 ± 0.00	3.67 ± 0.19	304.47 ± 94.22	78.03 ± 6.80
-	CF (untuned)	-	0.00 ± 0.00	3.75 ± 0.33	542.52 ± 334.07	60.86 ± 24.10
-	CT	-	0.00 ± 0.00	2.31 ± 0.32	459.35 ± 223.37	66.86 ± 16.12
-	R&C	Best	0.00 ± 0.00	0.95 ± 0.19	225.89 ± 163.98	83.70 ± 11.83
-	R&C	Log	0.00 ± 0.00	0.12 ± 0.01	261.47 ± 198.86	81.14 ± 14.35
-	R&C	RF	0.00 ± 0.00	0.83 ± 0.03	253.59 ± 125.24	81.70 ± 9.04

**Table EC.3** Mean and standard deviation statistics over 5 random samples in the synthetic experiments, with a constraint that treatment  $K = 1$  is available to at most  $C\%$  in the population.

Depth	Method	Model	Budget Range ( $C$ )	Gap	Solving Time (s)	OOS Regret	OOSP (%)
1	DR	LR, DT	0.05-0.09	0.00 ± 0.00	0.16 ± 0.03	158.32 ± 49.39	58.35 ± 9.03
1	DR	LR, DT	0.10-0.14	0.00 ± 0.00	0.22 ± 0.13	132.80 ± 43.61	61.54 ± 9.04
1	DR	LR, DT	0.15-0.19	0.00 ± 0.00	0.23 ± 0.10	113.91 ± 44.06	65.44 ± 9.48
1	DR	LR, DT	0.20-0.24	0.00 ± 0.00	0.20 ± 0.11	85.58 ± 32.88	70.01 ± 8.02
1	DR	LR, DT	0.25-0.29	0.00 ± 0.00	0.23 ± 0.13	77.34 ± 32.68	72.25 ± 8.12
1	DR	LR, DT	0.30-0.34	0.00 ± 0.00	0.23 ± 0.13	69.97 ± 34.71	73.71 ± 8.56
1	DR	LR, DT	0.35-0.40	0.00 ± 0.00	0.23 ± 0.14	66.32 ± 36.18	75.04 ± 8.86
1	DR	LR, Log	0.05-0.09	0.00 ± 0.00	0.19 ± 0.11	163.39 ± 46.47	57.67 ± 8.51
1	DR	LR, Log	0.10-0.14	0.00 ± 0.00	0.22 ± 0.15	139.03 ± 41.97	60.64 ± 8.44
1	DR	LR, Log	0.15-0.19	0.00 ± 0.00	0.23 ± 0.13	119.83 ± 45.13	64.53 ± 9.18
1	DR	LR, Log	0.20-0.24	0.00 ± 0.00	0.22 ± 0.16	93.20 ± 40.82	68.85 ± 8.22
1	DR	LR, Log	0.25-0.29	0.00 ± 0.00	0.24 ± 0.16	85.11 ± 42.53	70.99 ± 8.61
1	DR	LR, Log	0.30-0.34	0.00 ± 0.00	0.24 ± 0.16	77.97 ± 44.93	72.40 ± 9.07
1	DR	LR, Log	0.35-0.40	0.00 ± 0.00	0.26 ± 0.18	74.07 ± 46.99	73.80 ± 9.65
1	DR	Lasso, DT	0.05-0.09	0.00 ± 0.00	0.29 ± 0.13	172.73 ± 59.95	56.89 ± 10.02
1	DR	Lasso, DT	0.10-0.14	0.00 ± 0.00	0.33 ± 0.17	150.23 ± 63.71	60.15 ± 10.97
1	DR	Lasso, DT	0.15-0.19	0.00 ± 0.00	0.37 ± 0.19	129.98 ± 70.03	63.92 ± 12.14
1	DR	Lasso, DT	0.20-0.24	0.00 ± 0.00	0.46 ± 0.24	110.86 ± 71.28	67.15 ± 12.24
1	DR	Lasso, DT	0.25-0.29	0.00 ± 0.00	0.47 ± 0.27	98.02 ± 64.04	69.78 ± 11.28
1	DR	Lasso, DT	0.30-0.34	0.00 ± 0.00	0.46 ± 0.30	84.12 ± 60.11	72.13 ± 10.78
1	DR	Lasso, DT	0.35-0.40	0.00 ± 0.00	0.41 ± 0.23	79.02 ± 64.46	73.62 ± 11.83
1	DR	Lasso, Log	0.05-0.09	0.00 ± 0.00	0.27 ± 0.14	180.83 ± 65.74	55.80 ± 10.47
1	DR	Lasso, Log	0.10-0.14	0.00 ± 0.00	0.29 ± 0.16	164.75 ± 74.07	57.94 ± 11.64
1	DR	Lasso, Log	0.15-0.19	0.00 ± 0.00	0.34 ± 0.18	148.86 ± 75.23	60.86 ± 12.30
1	DR	Lasso, Log	0.20-0.24	0.00 ± 0.00	0.34 ± 0.20	130.36 ± 73.55	63.94 ± 12.10
1	DR	Lasso, Log	0.25-0.29	0.00 ± 0.00	0.40 ± 0.25	124.98 ± 73.15	65.19 ± 12.03
1	DR	Lasso, Log	0.30-0.34	0.00 ± 0.00	0.36 ± 0.23	117.06 ± 76.89	66.58 ± 12.74
1	DR	Lasso, Log	0.35-0.40	0.00 ± 0.00	0.36 ± 0.22	114.78 ± 76.44	67.13 ± 12.96

**Table EC.4** Mean and standard deviation statistics over 25 random samples in the warfarin experiments, with a constraint that the difference between estimated outcomes between White and non-White patients is at most  $\delta$ .

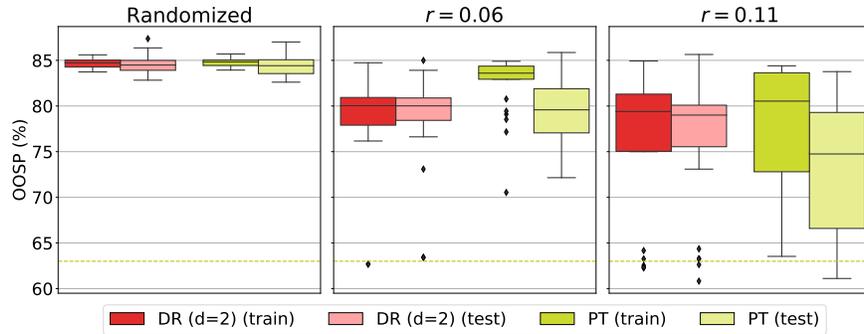
Depth	Method	Model	Fairness ( $\delta$ )	Gap	Solving Time (s)	Est. Outcome Disparity	Actual Disparity	OOS Regret	OOSP (%)
2	DM	RF/Log	0.01	0.00 ± 0.00	316.31 ± 112.04	-0.01 ± 0.01	0.01 ± 0.10	326.02 ± 106.02	0.76 ± 0.08
2	DM	RF/Log	0.02	0.00 ± 0.00	284.67 ± 125.32	-0.02 ± 0.01	-0.00 ± 0.10	321.20 ± 107.86	0.77 ± 0.08
2	DM	RF/Log	0.03	0.00 ± 0.00	290.46 ± 130.17	-0.02 ± 0.02	-0.01 ± 0.10	312.77 ± 105.23	0.77 ± 0.08
2	DM	RF/Log	0.04	0.00 ± 0.00	280.67 ± 116.49	-0.03 ± 0.02	-0.02 ± 0.09	303.54 ± 99.35	0.78 ± 0.07
2	DM	RF/Log	0.05	0.00 ± 0.00	271.45 ± 123.14	-0.03 ± 0.03	-0.03 ± 0.09	297.82 ± 96.61	0.79 ± 0.07
2	DM	RF/Log	0.06	0.00 ± 0.00	268.73 ± 151.18	-0.03 ± 0.03	-0.03 ± 0.09	294.58 ± 93.16	0.79 ± 0.07
2	DM	RF/Log	0.07	0.00 ± 0.00	241.31 ± 111.07	-0.04 ± 0.04	-0.04 ± 0.09	290.84 ± 92.17	0.79 ± 0.07
2	DM	RF/Log	0.08	0.00 ± 0.00	234.19 ± 105.39	-0.04 ± 0.04	-0.04 ± 0.09	290.78 ± 92.97	0.79 ± 0.07

## EC.5. Policytree Overfitting

In Figure EC.1, we illustrate that Policytree (PT) – despite optimizing the same objective as our doubly robust method (DR) – leads to worse out-of-sample performance because it learns from continuous features (as opposed to our method, which takes in discrete features). While PT’s performance on the training set is more often than not better than DR, it does not generalize well on the testing set; in contrast, our method maintains similar performance between the training and testing sets.

## EC.6. Scale of Proposed MIO

Table EC.5 shows the growth of the number of constraints, continuous variables, and binary variables of formulations (6), (7), and (8). This growth is a function of the characteristics of the



**Figure EC.1** A comparison of the doubly robust method's performance to the formulation proposed by Zhou et al. (2023) (PT), on the training and testing sets. Both methods are over trees of depth  $d=2$ .

observational data (number of treatment options, size of covariates, etc.) and tree depth. It also assumes that we relax the integrality requirement on  $z$  and  $w$ .

**Table EC.5** Summary of growth of the number of constraints and decision variables in formulations (6), (7), and (8).

# Constraints	# Continuous vars.	# Binary vars.
$\mathcal{O}(2^d  \mathcal{K}   \mathcal{X} )$	$\mathcal{O}(2^{d+1}  \mathcal{K}   \mathcal{X} )$	$\mathcal{O}(2^d  \mathcal{F}  \max_f \Theta(f))$

## EC.7. Proofs

*Proof of Proposition 1.* We introduce additional notation to help formalize our claim. For any tree based policy  $\pi \in \Pi_d$ , define

$$Q_N^{\text{IPW}}(\pi) := \mathbb{E} \left[ \frac{\mathbb{I}(K = \pi(X))Y}{\hat{\mu}^N(K, X)} \right].$$

To simplify the analysis, without loss of generality, we assume that the potential outcomes have been normalized, ensuring that  $|Y(k)| \leq 1$  holds for all  $k \in \mathcal{K}$ . We note that if all policies  $\pi \in \Pi_d$  are optimal, the statements in the proposition follow immediately, and thus henceforth focus on the case where  $\Pi_d \setminus \mathcal{S}^*$  is not empty.

We begin by showing that  $\hat{v}_{N,I}^{\text{IPW}} \rightarrow v^*$  w.p.1. Fix  $\epsilon > 0$  and a policy  $\pi \in \Pi_d$ . Since  $\mu$  and  $\hat{\mu}^N$  are bounded away from 0,  $\exists m_0 > 0$  and  $N_0 \in \mathbb{Z}_+$  such that

$$\mu(K, X) > m_0 \text{ and } \hat{\mu}^N(K, X) > m_0 \quad \forall K \in \mathcal{K}, X \in \mathcal{X}, \text{ and } N \geq N_0. \quad (\text{EC.4})$$

Define the function  $g : (m_0, \infty) \rightarrow \mathbb{R}$  through  $g(x) := \frac{1}{x}$ . Since  $g$  is uniformly continuous on the interval  $(m_0, \infty)$ , equation (EC.4) implies that there exists  $\delta_1 > 0$  such that

$$|\hat{\mu}^N(K, X) - \mu(K, X)| < \delta_1 \quad \implies \quad \left| \frac{1}{\hat{\mu}^N(K, X)} - \frac{1}{\mu(K, X)} \right| < \frac{\epsilon}{2} \quad \forall K \in \mathcal{K}, X \in \mathcal{X}. \quad (\text{EC.5})$$

Moreover, since  $\hat{\mu}^N$  converges almost surely to  $\mu$  and the fact that the support set of  $\mu$  is finite, it follows that there exists  $N^\pi \geq N_0$  such that

$$\max_{X,K} |\hat{\mu}^N(K, X) - \mu(K, X)| < \delta_1 \quad \text{w.p.1} \quad \forall N \geq N^\pi. \quad (\text{EC.6})$$

Therefore, it follows from equations (EC.5) and (EC.6) that

$$\max_{X,K} \left| \frac{1}{\hat{\mu}^N(K, X)} - \frac{1}{\mu(K, X)} \right| < \frac{\epsilon}{2} \quad \text{w.p.1} \quad \forall N \geq N^\pi.$$

We then have

$$\begin{aligned} \sup_{X,K,Y} \left| \frac{\mathbb{I}(K = \pi(X))Y}{\hat{\mu}^N(K, X)} - \frac{\mathbb{I}(K = \pi(X))Y}{\mu(K, X)} \right| &\leq \max_{X,K} \left| \frac{\mathbb{I}(K = \pi(X))}{\hat{\mu}^N(K, X)} - \frac{\mathbb{I}(K = \pi(X))}{\mu(K, X)} \right| \\ &< \frac{\epsilon}{2} \quad \text{w.p.1} \quad \forall N \geq N^\pi, \end{aligned}$$

where the supremum and maximum above are taken over  $X \in \mathcal{X}$ ,  $Y \in \mathcal{Y}$ ,  $K \in \mathcal{K}$ . The first inequality above follows since  $|Y| \leq 1$  and implies that

$$\begin{aligned} &\mathbb{E} \left( \frac{\mathbb{I}(K = \pi(X))Y}{\mu(K, X)} - \frac{\mathbb{I}(K = \pi(X))Y}{\hat{\mu}^N(K, X)} \right) \\ &= \sum_{x,k,y} \left( \frac{\mathbb{I}(k = \pi(x))y}{\mu(k, x)} - \frac{\mathbb{I}(k = \pi(x))y}{\hat{\mu}^N(k, x)} \right) \mathbb{P}(X = x, K = k, Y = y) \\ &< \frac{\epsilon}{2} \quad \text{w.p.1} \quad \forall N \geq N^\pi. \end{aligned} \quad (\text{EC.7})$$

Note that in the equation above, for the sake of clarity and ease of notation, we employ summation instead of integration. By definition of  $Q^{\text{IPW}}$  and  $Q_N^{\text{IPW}}$  and from our conditional exchangeability assumption, see Hernán and Robins (2019), it then follows that

$$|Q^{\text{IPW}}(\pi) - Q_N^{\text{IPW}}(\pi)| = |Q(\pi) - Q_N^{\text{IPW}}(\pi)| < \frac{\epsilon}{2} \quad \text{w.p.1} \quad \forall N \geq N^\pi. \quad (\text{EC.8})$$

At the same time, by the strong law of large numbers, there exists  $I^\pi$  such that

$$|Q_{N,I}^{\text{IPW}}(\pi) - Q_N^{\text{IPW}}(\pi)| < \frac{\epsilon}{2} \quad \text{w.p.1} \quad \forall I \geq I^\pi. \quad (\text{EC.9})$$

Then, equations (EC.8) and (EC.9) imply that

$$|Q_{N,I}^{\text{IPW}}(\pi) - Q(\pi)| \leq \epsilon \quad \text{w.p.1} \quad \forall N \geq N^\pi \text{ and } I \geq I^\pi. \quad (\text{EC.10})$$

Since the set  $\Pi_d$  of decision trees of depth at most  $d$  is finite, and the union of a finite number of sets each of measure zero also has measure zero, the pointwise convergence given in equation (EC.10) also results in uniform convergence, meaning that, there exist  $N_1 := \max_{\pi \in \Pi_d} N^\pi$  and  $I_0 := \max_{\pi \in \Pi_d} I^\pi$  such that

$$\max_{\pi \in \Pi_d} |Q_{N,I}^{\text{IPW}}(\pi) - Q(\pi)| \leq \epsilon \quad \text{w.p.1} \quad \forall I \geq I_0 \text{ and } N \geq N_1.$$

Since  $|\hat{v}_{N,I}^{\text{IPW}} - v^*| \leq \epsilon$  and the choice of  $\epsilon$  was arbitrary, then w.p.1  $\hat{v}_{N,I}^{\text{IPW}} \rightarrow v^*$  for all  $I \geq I_0$  and  $N \geq N_1$ . This completes the first part of the proof.

We now show that  $\hat{\mathcal{S}}_{N,I}^{\text{IPW}} \subseteq \mathcal{S}^*$  w.p.1 for  $I$  and  $N$  large enough. Consider the difference in objective value between the best and second best policies, given by

$$\rho := \max_{\pi \in \Pi_d \setminus \mathcal{S}^*} Q(\pi) - v^*.$$

Since for any  $\pi \in \Pi_d \setminus \mathcal{S}^*$  it holds that  $Q(\pi) < v^*$  and since the set  $\Pi_d$  is finite, it follows that  $\rho < 0$ . Let  $I$  and  $N$  be large enough such that  $\epsilon_{N,I} < -\frac{\rho}{2}$ . Then, w.p.1,  $\hat{v}_{N,I}^{\text{IPW}} > v^* + \frac{\rho}{2}$ , and for any  $\pi \in \Pi_d \setminus \mathcal{S}^*$  it holds that  $Q_{N,I}^{\text{IPW}}(\pi) < v^* + \frac{\rho}{2}$ . It follows that if  $\pi \in \Pi_d \setminus \mathcal{S}^*$ , then  $Q_{N,I}^{\text{IPW}} < \hat{v}_{N,I}^{\text{IPW}}$  and hence  $\pi$  does not belong to set  $\hat{\mathcal{S}}_{N,I}^{\text{IPW}}$ . As a result, for large enough  $I$  and  $N$ , w.p.1 the inclusion  $\hat{\mathcal{S}}_{N,I}^{\text{IPW}} \subseteq \mathcal{S}^*$  holds, which completes the proof.  $\square$

We make use of the following auxiliary lemma in the proof of Proposition 2.

**LEMMA EC.1 (Bartle (2001)).** *Suppose  $f_n : D \rightarrow \mathbb{R}$  and  $g_n : D \rightarrow \mathbb{R}$  are sequences of functions which converge uniformly to  $f, g : D \rightarrow \mathbb{R}$ , respectively. Then, if  $f$  and  $g$  are bounded, i.e.,  $\exists B > 0$  such that  $|f(x)| < B$  and  $|g(x)| < B$ ,  $f_n g_n$  converges uniformly to  $fg$ .*

*Proof of Proposition 2.* We introduce additional notation to help formalize our claim. For any tree based policy  $\pi \in \Pi_d$ , define

$$Q_N^{\text{DR}}(\pi) := \mathbb{E} \left[ \hat{v}_{\pi(X)}^N(X) + (Y - \hat{v}_{\pi(X)}^N(X)) \frac{\mathbb{I}(K = \pi(X))}{\hat{\mu}^N(K, X)} \right].$$

Since  $\hat{\mu}(K, X)$  and  $\hat{v}_K(X)$  are bounded, there exists  $B > 0$  such that  $|\hat{\mu}(K, X)| < B$  and  $|\hat{v}_K(X)| < B \quad \forall X \in \mathcal{X}$  and  $K \in \mathcal{K}$ . Fix  $\epsilon > 0$  and a policy  $\pi \in \Pi_d$ . From the definition of  $\hat{\mu}$ , uniform continuity of the inverse function (bounded away from 0), and the assumption that  $\mathcal{X}$  is finite and  $\mathcal{Y}$  is bounded,  $\exists N^\pi \in \mathbb{Z}_+$  such that

$$\begin{aligned} \sup_{X, K, Y} \left| \frac{\mathbb{I}(K = \pi(X))Y}{\hat{\mu}^N(X, K)} - \frac{\mathbb{I}(K = \pi(X))Y}{\hat{\mu}(X, K)} \right| &< \max_{X, K} \left| \frac{\mathbb{I}(K = \pi(X))}{\hat{\mu}^N(X, K)} - \frac{\mathbb{I}(K = \pi(X))}{\hat{\mu}(K, X)} \right| \\ &< \frac{\epsilon}{4(B+1)} \quad \forall N \geq N^\pi. \end{aligned} \quad (\text{EC.11})$$

Also, by definition of  $\hat{v}$ ,  $\exists N_0 \in \mathbb{Z}_+$  such that

$$\max_{X, K} |\hat{v}_K^N(X) - \hat{v}_K(X)| < \frac{\epsilon}{4(B+1)} \quad \forall N \geq N_0. \quad (\text{EC.12})$$

Thus, using (EC.11), (EC.12), and the assumption that  $\hat{v}$  and  $\hat{\mu}(K, X)$  are bounded, by Lemma EC.1,  $\exists N_1^\pi \geq N^\pi, N_0$  such that

$$\max_{X, K} \left| \frac{\mathbb{I}(K = \pi(X))\hat{v}_K^N(X)}{\hat{\mu}^N(K, X)} - \frac{\mathbb{I}(K = \pi(X))\hat{v}_K(X)}{\hat{\mu}(K, X)} \right| < \frac{2\epsilon B}{4(B+1)} \quad \forall N \geq N_1^\pi. \quad (\text{EC.13})$$

Moreover, from (EC.11), (EC.12), and (EC.13), we can conclude that for  $N \geq N_1^\pi$ ,

$$\sup_{X,K,Y} \left| \left\{ \hat{v}_{\pi(X)}^N(X) + (Y - \hat{v}_{\pi(X)}^N(X)) \frac{\mathbb{I}(K = \pi(X))}{\hat{\mu}^N(K, X)} \right\} - \left\{ \hat{v}_{\pi(X)}(X) + (Y - \hat{v}_{\pi(X)}(X)) \frac{\mathbb{I}(K = \pi(X))}{\hat{\mu}(K, X)} \right\} \right| < \frac{\epsilon}{2}.$$

Similar to the proof of Proposition 1, for a given policy  $\pi \in \Pi_d$ , for  $N \geq N_1^\pi$ ,

$$|Q^{\text{DR}}(\pi) - Q_N^{\text{DR}}(\pi)| < \frac{\epsilon}{2}.$$

Finally, we know that if w.p.1, either  $\hat{\mu}(k, x) = \mu(k, x)$  or  $\hat{\nu}_k(x) = \nu_k(x)$ , for all  $x \in \mathcal{X}$  and  $k \in \mathcal{K}$ , then  $Q(\pi) = Q^{\text{DR}}(\pi)$ , see for example Lunceford and Davidian (2004). Since the assumption in the previous clause is satisfied w.p.1 (as a result of the almost sure convergence of either  $\hat{\mu}^N$  or  $\hat{\nu}_k^N$ ), we conclude that,

$$|Q(\pi) - Q_N^{\text{DR}}(\pi)| < \frac{\epsilon}{2} \quad \text{w.p.1} \quad \forall N \geq N_1^\pi.$$

The rest of the proof is omitted as it can be derived by following the same logic as in the proof of Proposition 1. □

*Proof of Proposition 3.* We omit the proof due to its similarity to the proof of Proposition 2. □