

# Robust and Heterogenous Odds Ratio: Estimating Price Sensitivity for Unbought Items

Jean Pauphilet

London Business School, London, UK, jpauphilet@london.edu

*Problem definition:* Mining for heterogeneous responses to an intervention is a crucial step for data-driven operations, for instance to personalize treatment or pricing. We investigate how to estimate price sensitivity from transaction-level data. In causal inference terms, we estimate heterogeneous treatment effects when (a) the response to treatment (here, whether a customer buys a product) is binary, and (b) treatment assignments are partially observed (here, full information is only available for purchased items). *Methodology/Results:* We propose a recursive partitioning procedure to estimate heterogeneous odds ratio, a widely used measure of treatment effect in medicine and social sciences. We integrate an adversarial imputation step to allow for robust estimation even in presence of partially observed treatment assignments. We validate our methodology on synthetic data and apply it to three case studies from political science, medicine, and revenue management. *Managerial Implications:* Our robust heterogeneous odds ratio estimation method is a simple and intuitive tool to quantify heterogeneity in patients or customers and personalize interventions, while lifting a central limitation in many revenue management data.

*Key words:* Prescriptive analytics; Pricing; Customer segmentation; Causal inference; Missing data

---

## 1. Introduction

Pricing and discount strategies are one of the most crucial decisions faced by (r)etailers. Given the amount of information available on every customer and shopping experience, merchants have a growing opportunity to personalize both of them. In particular, granular data at a customer or product level enables quantitative analysis of the impact of pricing and discounts on purchase. Motivated by transaction-level data made available as part of the M&SOM 2020 data-driven challenge (Shen et al. 2020), we propose a new methodology to conduct such an analysis. In doing so, we lift two methodological barriers: First, we propose a new partitioning procedure to estimate heterogeneous odds ratio, a popular metric to measure treatment effect in presence of binary responses. Secondly, we propose an adversarial data imputation method to derive robust findings despite the absence of full

information on unpurchased items, a major limitation in revenue management applications involving transaction data (e.g., purchase panel, receipt, or loyalty data).

### 1.1. A causal inference framework

We adopt terminology and notations from the potential outcome framework (Imbens and Rubin 2015, Chapter 2) and characterize observations by three quantities:  $T \in \{0, 1\}$  is the treatment assignment variable, such that  $T = 1$  (resp.  $T = 0$ ) for the treated (resp. control) individuals.  $Y(T)$  captures the response to treatment  $T$ . Finally,  $\mathbf{X}$  denotes a vector of side information or any observed covariates (e.g., gender or age). In this work, we focus on binary treatments  $T \in \{0, 1\}$  (we will relax this assumption in Section 2.5) and binary response to treatment  $Y \in \{0, 1\}$ . The latter assumption holds in medicine application for instance, where “response to treatment” is defined as the occurrence of adverse events such as death. Binary responses to treatment are also common in pricing applications (purchase decision), economics, and social sciences.

Different measures for the effect of treatment  $T$  on outcome  $Y$  have been proposed. For instance, one can estimate the average response to treatment (resp. control)  $\mathbb{E}[Y(1)]$  (resp.  $\mathbb{E}[Y(0)]$ ), or the average treatment effect defined as  $ATE = \mathbb{E}[Y(1) - Y(0)]$ . Conceptually, these quantities summarize and compare the distributions of the two random variables  $Y(1)$  and  $Y(0)$ . Unfortunately, for each patient, one does not observe  $Y(1)$  and  $Y(0)$  simultaneously, but only the response to the assigned treatment, i.e.,  $Y(T)|T = t$  where  $t \in \{0, 1\}$ . This constitutes a major selection bias and the central problem of causal inference. This selection bias is usually addressed by assuming *unconfoundedness*, i.e., independence of potential outcomes and treatment assignment conditionally on the observed covariates (Rosenbaum and Rubin 1983) and conducting a sensitivity analysis to assess the robustness of the findings to confounding factors (Rosenbaum 2014).

With the widespread availability of patient-level data in the past decade, the interest of the community and industry shifted from average to conditional/personalized treatment effect estimation. Indeed, treatments can now be assigned at a fine-grained scale. Hence, the effect of the treatment on a particular patient category is more relevant than its average on the total population. Mathematically, interest shifted from the population-wide random variables  $Y(0)$  and  $Y(1)$  to their conditional versions,  $Y(0)|\mathbf{X} = \mathbf{x}$  and  $Y(1)|\mathbf{X} = \mathbf{x}$ . Mining for patterns of individual-level differences in treatment effect is also referred to as heterogeneous treatment effect (HTE) estimation. In particular and relevant to our work,

tree-based methods from supervised machine learning have been adapted to the problem of heterogeneous treatment effect using decision trees (Zeileis et al. 2008, Su et al. 2009, Laber and Zhao 2015, Athey and Imbens 2016, Tran and Zheleva 2019, Bertsimas et al. 2019, Lee et al. 2020a) and random forests (Wager and Athey 2018, Athey et al. 2019, Lee et al. 2020b). While estimation of heterogeneous treatment effects can inform personalized treatment, Lakkaraju and Rudin (2017), Kallus (2017) propose a one-step approach to learn personalized policies directly, without an intermediate estimation step. Kallus and Zhou (2018) investigate the question of policy learning in presence of confounding factors.

Recursive partitioning approaches proceed by iteratively dividing the observations into subgroups so as to minimize some splitting criterion. Various splitting criteria have been investigated, ranging from impurity measures (Chan and Loh 2004, Landwehr et al. 2005) to coefficient instability (Zeileis et al. 2008), variance of the estimates (Athey and Imbens 2016), prediction accuracy (Bertsimas et al. 2019, Aouad et al. 2019),  $t$ -statistics (Su et al. 2009), and treatment effectiveness (Bertsimas et al. 2019, Kallus 2017). Despite using different splitting criteria, these approaches unanimously accept the conditional average treatment effect  $CATE = \mathbb{E}[Y(1) - Y(0)|\mathbf{X} = \mathbf{x}]$  as the metric to assess the effect of a treatment. When dealing with binary outcomes  $Y \in \{0, 1\}$ , however, odds ratio are a common and intuitive measure of a treatment effect as well. The odds ratio  $R$  compares the odds of responding when treated,  $\mathbb{P}(Y = 1|T = 1) / \mathbb{P}(Y = 0|T = 1)$ , with the odds of responding when not treated,  $\mathbb{P}(Y = 1|T = 0) / \mathbb{P}(Y = 0|T = 0)$ , i.e.,

$$R := \frac{\mathbb{P}(Y = 1|T = 1)}{\mathbb{P}(Y = 0|T = 1)} \times \frac{\mathbb{P}(Y = 0|T = 0)}{\mathbb{P}(Y = 1|T = 0)}. \quad (1)$$

Defined as such, the odds ratio is simply a measure of association between  $Y$  and  $T$  but can be interpreted as a measure of causal effect under unconfoundedness. If  $R > 1$  (resp.  $R < 1$ ), the treatment increases (resp. decreases) the likelihood of responding positively. On the contrary,  $R = 1$  indicates that treatment and response to treatment are independent. Because it captures a treatment effect relative to the baseline risk, odds ratio are a widely used measure for treatment effect estimation in both randomized and observational studies, especially in meta-analysis that compare results from disparate studies (Engels et al. 2000). The odds ratio relates with (but differs from) the relative risk or risk ratio (see Davies et al. 1998, for a discussion on the relative merits of each metric). Compared to difference in responses, ratio of responses does not only capture the magnitude of the effect but

also its robustness to unmeasured confounding (see Cornfield et al. 1959, for a notorious discussion regarding tobacco and lung cancer). In our opinion, two reasons can explain the lack of interest in estimating heterogeneous odds ratio. First, odds ratios suffer from non-collapsibility, a generalization of Simpson’s paradox, which can hinder their interpretability (Hernán and Robins 2010, Section 4.3)— note, however, that risk ratios are collapsible if there is no confounding. We refer to Greenland et al. (1999) for a discussion on non-collapsibility and its relation to confounding. Second, because ratios are relative measures, some authors expect them to be less heterogeneous than risk differences, while others have suggested that they can be equally heterogeneous but that tests for their heterogeneity has lower power (Poole et al. 2015). This question remains unanswered.

In this paper, we propose a recursive partitioning procedure to evaluate heterogeneous odds ratio. Our procedure relies on Cochran’s  $Q$ -test for heterogeneity and could be applied to a wide variety of treatment effect definitions, including odds ratios, risk ratios and CATE. To the best of our knowledge, we are the first to consider the heterogeneous odds ratio estimation task, which is surprising given the popularity of odds ratio in medicine and social sciences, and the profusion of research on HTE estimation.

## 1.2. Challenges arising from real-world revenue management data

In the realm of the 2020 M&SOM data-driven challenge (Shen et al. 2020), researchers were given access to transaction-level data from the JD.com’s e-tailing activity. This data comprises information about millions of transactions that occurred on the platform over the month of March 2018. For instance, demographic information about the customer, product attributes, pricing, and discount data are available. To estimate the impact of discount (the treatment) on the purchase decision (the response to treatment), our heterogeneous odds ratio estimation method would be a perfect candidate.

Unfortunately, data collection is triggered by a transaction and, by definition, a transaction corresponds to a customer *buying* an item. Consequently, the set of observations only includes customers and products that responded positively, rendering treatment effect estimation impossible. Note that this issue is present in many real-world revenue management settings. To alleviate this problem, we enrich our data set with information about user sessions: for each user, we know the pages this user visited and the products they considered but did not buy. Since these products have not been purchased, we only have partial information about the potential transaction. In particular, we do not know whether these

products were offered to the customer with a discount. In other words, their treatment assignment is unknown. To summarize, the initial data set suffers from strong selection bias since all observations  $(\mathbf{x}_i, t_i, y_i(t_i))$  satisfy  $y_i(t_i) = 1$ . As an answer, we enrich the data set with partial observations of products that have been seen by the user but not bought,  $(\mathbf{x}_i, \text{NA}, 0)$ , hence converting a selection bias problem into a missing data problem.

Missing data is a long-experienced problem in applied statistics. In the context of statistical inference, Rubin (1976) introduced the *Missing At Random* (MAR) assumption and proved that inference was possible under this assumption. However, MAR is not a viable assumption in our case. Indeed, MAR requires that the missingness indicator (here,  $Y$ ) is independent from the missing value (here,  $T$ ), conditioned on the other covariates (here,  $\mathbf{X}$ ). In our case, assuming MAR would negate the existence of an individual treatment effect and is not an acceptable assumption. An alternative is needed. Hence, we develop a robust imputation method integrated within the heterogeneous odds ratio procedure to find robust estimates of the odds ratio in presence of partially observed treatment assignments, without relying on the MAR assumption. In other revenue management applications, out-of-stocks can also be considered as partially observed treatment assignment (Musalem et al. 2010), where the treatment is the assortment offered to the customer. Among others, Vulcano et al. (2012) developed an EM algorithm to properly estimate a demand model under some structural assumptions.

### 1.3. Contributions and structure

Our contributions can be summarized as follows:

- We develop a recursive partitioning procedure to detect heterogeneous effect of treatment. Based on Cochran’s  $Q$ -statistics, we propose a splitting criterion that explicitly aims at maximizing heterogeneity across the subgroups in our partition, while accounting for statistical significance and sample size. Our criterion requires minimal assumptions on the estimator used to quantify the effect of treatment, namely asymptotic normality.
- We implement this procedure in the case of binary treatments and binary responses, using odds ratio as a measure of treatment effect. Our approach, that we refer to as *Heterogeneous Odds Ratio* (HOR) estimation, provides an interpretable partition of the feature space together with asymptotically valid confidence intervals. Compared to approaches based on conditional average treatment effect, HOR estimation readily applies to randomized control trials as well as observational studies, and naturally informs confounding-robust policy improvement. From HOR estimation, we derive personalized treatment

assignment strategies that maximize the odds of responding positively ( $Y = 1$ ). Under this prescriptive lens, we extend our framework to multiple treatments.

- We propose a robust optimization approach integrated in our recursive partitioning scheme to address the issue of partially observed treatment assignment. In practice, this situation emerged for instance when the original data set suffers from a strong selection bias and additional data with unknown treatment assignment is collected. Our approach relies on defining a set of plausible set of treatment assignments and considering adversarially assigned treatments.

- We validate our methodology on synthetic data and apply it to case studies from political science, medicine, and revenue management. In particular, we apply our methodology to evaluate the effect of discount on purchase using JD.com transaction-level data, and identify that the original price and the product type reveal strong heterogeneity in the effectiveness of discounts.

*Structure:* We present the HOR estimation procedure and its implication for policy recommendation in Section 2. Section 3 derives adaptive concentration bounds for our HOR estimates and discusses extensions to random forest estimation. Section 4 addresses the missing treatment assignment problem using random sampling and robust optimization. We validate our approach on synthetic data in Section 5 and then apply it to a collection of examples from social science and medicine (§6), and revenue management (§7).

*Notations:* Nonbold lowercase characters ( $x$ ) denote scalars and bold lowercase characters ( $\mathbf{x}$ ) vectors. Uppercase letters designate random variables, e.g.,  $X$  and  $\mathbf{X}$  respectively denote a random scalar and vector. The symbol  $\perp$  designates independent random variables.

## 2. Policy Learning via Heterogeneous Odds Ratio

In this section, we propose a recursive partitioning procedure for estimating heterogeneous treatment effects in presence of binary response to treatment. We measure the effect of treatment in terms of odds ratio (§2.1) and use Cochran’s  $Q$ -statistics as a splitting criterion (§2.2). The final procedure is described in Section 2.3. We discuss policy implications of odds ratio in Section 2.4 and the extension to multiple treatments in Section 2.5.

### 2.1. Odds ratio: Estimation and confidence intervals

From  $n$  observations of pairs  $(t_i, y_i)$ , construct a  $2 \times 2$  contingency table (Table 1). The odds ratio compares the odds of responding positively when treated,  $N_{1,1}/N_{0,1}$ , with the

**Table 1** Sample  $2 \times 2$  contingency table, where  $N_{t,y} = \sum_i \mathbf{1}(T_i = t)\mathbf{1}(Y_i = y)$ .

	Y = 0	Y = 1
T = 0	$N_{0,0}$	$N_{0,1}$
T = 1	$N_{1,0}$	$N_{1,1}$

odds of responding positively without treatment,  $N_{1,0}/N_{0,0}$ . An estimate for the log-odds ratio  $\log R$  and its variance are given by

$$\log \hat{r} = \log(N_{1,1}) - \log(N_{1,0}) + \log(N_{0,0}) - \log(N_{0,1}), \quad (2)$$

$$\hat{v} = \hat{\sigma}^2(\log \hat{r}) = \frac{1}{N_{1,1}} + \frac{1}{N_{1,0}} + \frac{1}{N_{0,1}} + \frac{1}{N_{0,0}}. \quad (3)$$

By the large-sample normality of  $\log \hat{r}$  (Breslow 1981),  $\log \hat{r} \pm z_{\alpha/2} \hat{\sigma}(\log \hat{r})$ , with  $z_{\alpha/2}$  the  $\alpha/2$  quantile of the standard distribution, is a Wald confidence interval for  $\log R$ . Exponentiating its endpoints provides a confidence interval for  $R$ . This interval was first proposed by Woolf et al. (1955) and performs quite well compared to subsequent proposals (see Lawson 2004, for a comparison). Its robustness for small samples can be improved by adding a  $+1/2$  bias term to each cell value  $N_{t,y}, y \in \{0, 1\}, t \in \{0, 1\}$  (Haldane 1956, Gart and Zweifel 1967, Gart 1966). We use this small-sample correction in our implementation.

## 2.2. Testing for heterogeneity

We now consider  $K$  disjoint subgroups and estimate the sample log-odds ratio  $\log \hat{r}^{(k)}$  and its standard error  $\hat{v}^{(k)}$  on each subgroup separately. Borrowing concepts from hypothesis testing, we define a criterion to measure heterogeneity between subgroups.

Let us illustrate the case with two subgroups,  $K = 2$ . Since the observations are independent and the subgroups are disjoint,  $\log \hat{r}^{(1)}$  and  $\log \hat{r}^{(2)}$  are independent. Hence,  $\log \hat{r}^{(1)} - \log \hat{r}^{(2)}$  is asymptotically normal with mean  $\log r^{(1)} - \log r^{(2)}$  and variance  $\hat{v}^{(1)} + \hat{v}^{(2)}$ . We want to test whether  $r^{(1)} = r^{(2)}$  ( $H_0$ ). Under ( $H_0$ ),  $Z := \frac{(\log \hat{r}^{(1)} - \log \hat{r}^{(2)})^2}{\hat{v}^{(1)} + \hat{v}^{(2)}}$  follows a chi-square distribution with one degree of freedom. In this case,  $Z$  corresponds to McNemar's test statistics. Maximizing heterogeneity between subgroups can be achieved by maximizing the probability of ( $H_0$ ) being rejected, i.e., minimizing

$$\mathbb{P} \left( Z \geq \frac{[\log \hat{r}^{(1)} - \log \hat{r}^{(2)}]^2}{\hat{v}^{(1)} + \hat{v}^{(2)}} \right) = 1 - \chi_1^2 \left( \frac{[\log \hat{r}^{(1)} - \log \hat{r}^{(2)}]^2}{\hat{v}^{(1)} + \hat{v}^{(2)}} \right),$$

where  $\chi_1^2$  denotes the cumulative distribution function of a chi-square random variable with one degree of freedom.

This procedure extends to the case with  $K$  subgroups and is referred to as Cochran's  $Q$ -test (Cochran 1954). We construct  $K$  differences  $\hat{\theta}^{(k)} - \bar{\theta}$ . Here, we denote  $\hat{\theta}^{(k)} := \log \hat{r}^{(k)}$  for concision and define  $\bar{\theta}$  as a weighted average of all  $\hat{\theta}^{(k)}$ 's:

$$\bar{\theta} = \frac{1}{\sum_{k=1}^K \hat{w}^{(k)}} \sum_{k=1}^K \hat{w}^{(k)} \hat{\theta}^{(k)}, \quad \text{with } \hat{w}^{(k)} = 1/\hat{v}^{(k)}.$$

Then, Cochran's  $Q$  statistic is the weighted sum of the squared deviations of the individual effects from their average; that is

$$Q := \sum_{k=1}^K \hat{w}^{(k)} (\hat{\theta}^{(k)} - \bar{\theta})^2. \quad (4)$$

Under mild assumptions and under the null hypothesis that  $\theta^{(1)} = \dots = \theta^{(K)}$ ,  $Q$  asymptotically follows a chi-square distribution with  $K - 1$  degrees of freedom. Hence, minimizing  $1 - \chi_{K-1}^2(Q)$ , or equivalently maximizing  $Q$ , yields greater heterogeneity. We refer to Huedo-Medina et al. (2006), Kulinskaya and Dollinger (2015) for recent discussions on the use of the  $Q$ -test in comparing odds ratio.

In the following section, we will use the  $Q$ -statistics as a measure of heterogeneity and construct a partition of the feature space to maximize it. Related to our work, Su et al. (2009) use the  $t$ -statistics to measure heterogeneity in CATE within subgroups and build their partition. However, Student's  $t$ -test can only be applied to test homogeneity in sample means and cannot be used for comparing (odds) ratios. On this regard, using Cochran's  $Q$ -test is more general and can be used to assess homogeneity of any estimator, provided that the estimator's distribution is asymptotically normal with computable variance.

### 2.3. A recursive partitioning procedure

Let  $\Pi$  be a partition of the covariate space into  $|\Pi|$  disjoint subgroups. Write  $\ell_k$ ,  $k = 1, \dots, |\Pi|$ , an element of  $\Pi$ .  $\ell_k$  will equivalently be referred to as a leaf or a subgroup. By definition, for any  $\mathbf{x}$ , there exists a unique subgroup of  $\Pi$  containing  $\mathbf{x}$ , denoted  $\ell(\mathbf{x}, \Pi)$ .

Consider a dataset of  $n$  observations  $\mathcal{D}_n = \{(\mathbf{x}_i, t_i, y_i), i = 1, \dots, n\}$  and a partition  $\Pi$ . In each subgroup, one can estimate the log-odds ratio (2), its variance (3) and compute the  $Q$ -statistics (4). Let us denote  $Q(\Pi; \mathcal{D}_n)$  its value. The objective is to construct a partition that maximizes the  $Q$ -statistics, i.e., solve

$$\max_{\Pi_{\text{partition}}} Q(\Pi; \mathcal{D}_n).$$



Instead of solving the above maximization problem exactly over the entire set of partitions, we solve it recursively. The benefits of recursive partitioning are fourfold: First, the resulting partition has a natural tree structure, which is highly interpretable and easy to understand for the practitioner. Interpretability of models and explainability of decisions are indeed a growing concern of consumers and regulators (Goodman and Flaxman 2017). In addition, trees constructed in a greedy fashion can be explained sequentially: the first split of the tree explains most of the heterogeneity, followed by the second split, and so on. In other words, depth orders the splits by decreasing importance. This is not the case of globally optimal trees. Secondly, despite a vibrant corpus of works on formulating decision trees as mixed-integer optimization problems and solving them as such to provable optimality (see, e.g., Bertsimas and Dunn 2017, Günlük et al. 2018, Zantedeschi et al. 2020, Aglin et al. 2020, Demirović et al. 2020, Lin et al. 2020), greedy procedures remain the gold-standard for balancing accuracy, interpretability, and scalability. Thirdly, recursive procedures can handle complex non-convex criteria like the  $Q$ -statistics. Finally, while controlling the  $p$ -value of Cochran’s  $Q$ -test ensures that the odds ratio is not constant across all leaves, our recursive procedure will satisfy a stronger property, namely that no leaf can be further divided so as to detect heterogeneity in odds ratio (at a given significance level). From a hypothesis testing perspective, a greedy procedure will control for pairwise heterogeneity between sibling leaves (see Proposition 1 for a formal statement).

We now describe our recursive partitioning procedure. Consider a partition  $\Pi$  and a leaf  $\ell$  of this partition. We investigate whether  $\ell$  can be further divided. A split is defined as a coordinate  $j$  and a threshold  $t$ . The split  $(j, t)$  divides  $\ell$  into two subgroups  $\ell_- = \ell \cap \{\mathbf{x} : x_j < t\}$  and  $\ell_+ = \ell \cap \{\mathbf{x} : x_j \geq t\}$ . There are finitely many splits, so by explicit enumeration, one can find the one that maximizes the  $Q$ -statistics  $Q(\{\ell_-, \ell_+\}, \mathcal{D}_n)$ , denoted  $(j^*, t^*)$ . If the resulting  $p$ -value,  $1 - \chi_1^2(Q(\{\ell_-^*, \ell_+^*\}, \mathcal{D}_n))$ , exceeds a given threshold  $p_{\max}$ , then we do not split  $\ell$ . Otherwise, we remove  $\ell$  from the partition, and add  $\ell_+^*$  and  $\ell_-^*$  instead. The algorithm terminates when no leaf in the current partition can be divided any further. Pseudo-code is given in the online appendix, Algorithm A.1.

The recursive procedure described in Algorithm A.1 generates a partition  $\Pi$  of the feature space that has a natural tree structure. The internal nodes of the tree correspond to the all the sub-groups  $\ell$  considered and split by the algorithm (e.g., its root node is the entire space  $\mathbb{R}^p$ ), while its leaves or terminal nodes are exactly the elements of  $\Pi$ . With a slight

abuse of notation, we will also denote  $\Pi$  the tree associated with the construction of the partition  $\Pi$ . We define *siblings* as two nodes in the tree  $\Pi$  that share the same parent. Partitions produced by our algorithm (Algorithm A.1) satisfy the following property:

**PROPOSITION 1.** *Let  $\Pi$  be the tree resulting from applying Algorithm A.1 on the training data  $\mathcal{D}_n$ . For any siblings  $node_1$  and  $node_2$ , the hypothesis that  $node_1$  and  $node_2$  are homogeneous in terms of odds ratio is rejected by Cochran's Q-test on  $\mathcal{D}_n$ , at a level  $p_{\max}$ .*

#### 2.4. Odds ratio as a treatment recommendation

As mentioned in introduction, comparing the odds ratio to 1 provides an indication on whether the treatment is effective or not. The main advantages of odds ratio as a measure of causal effect are its simplicity, its symmetry with respect to outcome definition, and the fact that it is insensitive to the treatment propensity and baseline effect,  $\mathbb{P}(T = 1)$  and  $\mathbb{P}(Y = 1|T = 0)$  respectively. If we assume that  $Y = 1$  is a positive outcome (e.g., survival or purchase),  $R > 1$  implies that the treatment increases the odds of  $Y$  being positive, hence is beneficial. Odds ratio are prominently used in situations where assigning a reward or value to each outcome is impossible (e.g., in medicine when the outcome is death).

Due to their non-collapsibility, odds ratios usually lack any interpretation either as the change in average odds or the average change in odds, contrary to risk ratios (Cummings 2009). For instance, in revenue management, the risk ratio bounds the allowable discount: Assume  $p_1$  (resp.  $p_0$ ) is the discounted (resp. original) price. Then, from a retailer's perspective, discounting makes sense as long as the expected revenue under discount,  $p_1\mathbb{P}(Y = 1|T = 1)$ , exceeds the baseline revenue  $p_0\mathbb{P}(Y = 1|T = 0)$ , i.e., as long as the risk ratio  $\mathbb{P}(Y = 1|T = 1)/\mathbb{P}(Y = 1|T = 0)$  is greater than  $p_0/p_1$ .

The interpretation of odds ratio is not as straightforward. Define  $\text{odds}(z) := z/(1 - z)$  for any  $z \in [0, 1)$ . The function  $z \mapsto \text{odds}(z)$  is a non-linear, increasing, and convex transformation of  $z$ . From the definition of the odds ratio  $R$  in (1), we have

$$\mathbb{P}(Y = 1|T = 1) = \text{odds}^{-1}(R \times \text{odds}(\mathbb{P}(Y = 1|T = 0))),$$

a non-linear transformation of  $\mathbb{P}(Y = 1|T = 0)$ . Nonetheless, for rare events, i.e.,  $\mathbb{P}(Y = 1) \rightarrow 0$ , since  $\text{odds}(z) \sim_{z \rightarrow 0} z$ , odds ratios are equivalent to risk ratios and share the same interpretation.

Odds ratio also have a natural interpretation in the context of repeated experiments. For instance, consider a situation when a customer comes multiple times to a (virtual) store.

Assume that visits are independent and identical. Then, the number of times a customer will come to the store before making a first purchase follows a geometric distribution over  $\{0, 1, \dots\}$  and is equal to  $\mathbb{P}(Y = 0)/\mathbb{P}(Y = 1)$  on average. In this setting, the odds ratio compares times to purchase with and without treatment. For example, if  $R = 2$ , then treatment (e.g., discount) halves time to purchase. Maximizing the odds ratio is equivalent to minimizing time to purchase. Further assume that the decision-maker incurs a cost  $c$  per visit and receives a reward  $r$  as soon as  $Y = 1$ . For example, if the decision-maker decides to treat,  $c$  can correspond to the marginal cost of treatment. Then, the average reward is positive as long as  $\text{odds}(\mathbb{P}(Y = 1)) > \frac{c}{r - c}$ . In other words, for a fixed reward  $r$ , the odds bound the cost per visit  $c$  that the decision maker can afford.

Note that our splitting criterion only requires asymptotic normality of the estimates of the treatment effect, so Algorithm A.1 can be applied to other definitions of the treatment effect. For instance, it can be applied to the relative risk discussed earlier,  $\mathbb{P}(Y = 1|T = 1)/\mathbb{P}(Y = 1|T = 0)$ , which can be empirically estimated as (notations from Table 1):

$$\frac{N_{1,1}}{N_{1,1} + N_{1,0}} \times \frac{N_{0,0} + N_{0,1}}{N_{0,0}}.$$

As for odds ratio, unconfoundedness is needed to interpret risk ratios as a measure of causal effect and not association only. Furthermore, in observational studies, techniques like inverse propensity score weighting are needed for unbiased estimation. In these settings, the risk ratio between the responses to treatment,  $\mathbb{P}(Y(1) = 1)/\mathbb{P}(Y(0) = 1)$  is equal to the risk ratio between the responses to assigned treatment,  $\mathbb{P}(Y = 1|T = 1)/\mathbb{P}(Y = 1|T = 0)$ , up to a multiplicative factor,  $\mathbb{P}(T = 1)/\mathbb{P}(T = 0)$ . Our recursive partitioning procedure can be used as a non-parametric estimation of these “propensity odds”. Note that, for odds ratio, similar corrections are needed in observational studies. However, using inverse propensity weighting and our partition as a non-parametric model for propensity, the corrective terms would appear both on the numerator and denominator and Equations (2)-(3) would remain valid as is.

## 2.5. Extension to multiple treatments

So far, we restricted our attention to a single treatment. We now consider the situation when one can chose between multiple treatments, i.e.,  $T$  can take values 0 (control),

$1, \dots, m$ . For each treatment  $T = t$ , one can compute the odds ratio associated with this particular treatment:

$$R(t) := \frac{\mathbb{P}(Y = 1|T = t)}{\mathbb{P}(Y = 0|T = t)} \times \frac{\mathbb{P}(Y = 0|T = 0)}{\mathbb{P}(Y = 1|T = 0)}.$$

Consequently, each leaf is characterized by an  $m$ -dimensional vector of odds ratios for the  $m$  possible treatments. A naive way to extend our approach to multiple treatments would be to consider a vector extension of Cochran’s  $Q$ -statistics to test the heterogeneity in the multi-variate response  $(R(t))_{t=1, \dots, m}$  when running Algorithm A.1. Unfortunately, these tests would have little power and might be of little relevance in practice.

Instead, we adopt a more prescriptive and pragmatic approach: we associate each leaf with a scalar value corresponding to the best odds ratio among all treatments, namely  $\max_{t=0, \dots, m} R(t)$ . By doing so, a leaf in Algorithm A.1 will be further divided in two if it identifies two subgroups that have significantly different responses to the best treatment - but what the “best” treatment is might be different for each subgroup. Observe that we consider control ( $t = 0$ ) as a candidate treatment with  $R(0) = 1$ , so the tree can recommend not to treat if no treatment exhibits a positive effect. Kallus (2017) adopt a similar strategy for learning personalized policies for observational studies, when treatment effects are evaluated in terms of conditional average treatment effect.

### 3. Convergence Analysis and Application to Random Forests

In this section, we study the convergence of our recursive partitioning procedure for odds ratio estimation. We prove an adaptive concentration bound that states that, with high probability and simultaneously for any leaf  $\ell$ , the gap between the estimated odds ratio in leaf  $\ell$  and the population-average odds ratio scales as  $\sqrt{\log(n) \log(p)/k}$ . The main result is presented in Section 3.2, after preliminary assumptions and notations (Section 3.1). Section A.2 in the online appendix discusses implications for random forest estimation.

#### 3.1. Notations and assumptions

First, we impose that Algorithm A.1 generates  $(\alpha, k)$ -valid partitions:

**DEFINITION 1.** A partition  $\Pi$  is  $(\alpha, k)$ -valid if it can be generated by a recursive partitioning scheme in which each child node contains at least a fraction  $\alpha \in (0, 0.5]$  of the data points in its parent node and each terminal node contains at least  $k \in \mathbb{N}$  training examples. Given a dataset  $\mathcal{D}_n$ , we denote its set of  $(\alpha, k)$ -valid partitions by  $\mathcal{V}_{\alpha, k}(\mathcal{D}_n)$  (or simply  $\mathcal{V}_{\alpha, k}$  when clear from the context).

This assumption can easily be enforced in practice by considering splits that leave child nodes with a minimum number of training samples.

Our analysis relies on adaptive concentration inequalities for regression trees (Wager and Walther 2015) and require a similar set of assumptions. First, a condition on the dependence of the individual coordinates:

ASSUMPTION 1. (*Weakly dependent features*) *The features  $X \in [0, 1]^p$  are distributed according to a density  $f$  satisfying  $\zeta^{-1} \leq f(x) \leq \zeta$ ,  $\forall x \in [0, 1]$ , for some constant  $\zeta \geq 1$ .*

Second, we require that the minimum leaf size  $k$  grows reasonably fast as  $n, p \rightarrow \infty$ , i.e.,

ASSUMPTION 2. (*Minimum leaf size*) *The minimum leaf-size  $k$  grows with  $n$  at a rate bounded from below by*

$$\lim_{n \rightarrow \infty} \frac{\log(n) \max \{ \log(p), \log \log(n) \}}{k} = 0.$$

For any subgroup  $\ell \subseteq \mathbb{R}^d$ , we consider the  $2 \times 2$  contingency tables (Table 1) for this particular leaf and analyze the convergence of cell counts towards their population-average. In particular, for each cell  $(t, y) \in \{0, 1\}^2$ , we compare

$$\begin{aligned} \hat{\pi}_{t,y}^n(\ell) &= \frac{1}{|\{i : \mathbf{x}_i \in \ell\}|} \sum_{i: \mathbf{x}_i \in \ell} \mathbf{1}(t_i = t, y_i = y), \\ \pi_{t,y}(\ell) &= \mathbb{P}(Y = y, T = t | \mathbf{X} \in \ell), \end{aligned}$$

and we define the corresponding log odds ratio:  $\log(\hat{r}^n(\ell)) = \log(\hat{\pi}_{1,1}^n(\ell)) - \log(\hat{\pi}_{1,0}^n(\ell)) - \log(\hat{\pi}_{0,1}^n(\ell)) + \log(\hat{\pi}_{0,0}^n(\ell))$ ; and similarly for  $\log(r(\ell))$ .

### 3.2. Adaptive concentration bounds

Consider one tree  $\Pi \in \mathcal{V}_{\alpha,k}$ . For any leaf  $\ell \in \Pi$ , Hoeffding's inequality shows that  $\hat{\pi}_{t,y}^n(\ell)$  concentrates around  $\pi_{t,y}(\ell)$  at a rate  $1/\sqrt{k}$ . Then, a standard Chernoff bound technique enables to bound the worst case deviation over all  $\ell \in \Pi$  by

$$\mathbb{P} \left( \sup_{\ell \in \Pi} |\hat{\pi}_{t,y}^n(\ell) - \pi_{t,y}(\ell)| \geq \delta \right) \leq 2 \exp \left( \log \left( \frac{n}{k} \right) - 2k\delta^2 \right),$$

where we bound the total number of leaves in  $\Pi$  by  $n/k$ . In particular, it yields

$$\lim_{n \rightarrow \infty} \mathbb{P} \left( \sup_{\ell \in \Pi} |\hat{\pi}_{t,y}^n(\ell) - \pi_{t,y}(\ell)| \geq \sqrt{\frac{1.1}{2k} \log \left( \frac{n}{k} \right)} \right) = 0. \tag{5}$$

However, this bound is not adaptive in the sense that it is uniform over all leaves  $\ell$  but for a particular partition  $\Pi$ . Our goal is to derive a bound that is also uniform over all potential partitions  $\Pi \in \mathcal{V}_{\alpha,k}$ , to bound the performance of any tree generated by Algorithm A.1.

For the CART algorithm, Breiman et al. (1984) derived a uniform convergence result for regression and classification trees. They first prove the uniform convergence of the empirical leaf distributions (Breiman et al. 1984, Theorem 12.2). Then, under the assumption that the diameter of each leaf shrinks as  $n \rightarrow \infty$  and under some regularity conditions, they conclude on the uniform convergence of the empirical prediction error. Despite its prevalence (Gordon and Olshen 1978, 1980), the vanishing leaf diameter assumption is admittedly strong. Our analysis relies on the recent work of Wager and Walther (2015), who provide an adaptive concentration bound for regression trees. As in Breiman et al. (1984), the key idea is to replace the supremum over all  $\Pi \in \mathcal{V}_{\alpha,k}, \ell \in \Pi$  by a supremum over a finite set of regions of  $\mathbb{R}^d$  whose size moderately increases with the dimension of the problem. Instead of relying on generic  $\varepsilon$ -nets from VC theory (Vapnik and Chervonenkis 1971) like Breiman et al. (1984), Wager and Walther (2015) exhibit a tailored set of approximating rectangles and leverage the fact that partitions are  $(\alpha, k)$ -valid together with Assumption 1. Applied to our context, their result states as follows:

LEMMA 1. (*Wager and Walther 2015, Theorem 1*) *Suppose that we have a sequence of problems with parameters  $(n, p, k)$  satisfying Assumptions 1-2. Then, sample averages of cell counts over all possible valid partitions concentrate around their expectations with high probability, i.e., for all  $(t, y) \in \{0, 1\}^2$ :*

$$\lim_{n,p,k \rightarrow \infty} \mathbb{P} \left( \sup_{\Pi \in \mathcal{V}_{\alpha,k}, \ell \in \Pi} |\hat{\pi}_{t,y}^n(\ell) - \pi_{t,y}(\ell)| \geq \frac{f(n,p,k)}{\sqrt{k}} \right) = 0,$$

$$\text{with } f(n,p,k) := \frac{9}{2} \sqrt{\frac{\log(n/k) (\log(pk) + 3 \log \log(n))}{\log((1-\alpha)^{-1})}}.$$

Notice that under Assumption 2,  $f(n,p,k)/\sqrt{k} \rightarrow 0$ , as  $n, p, k \rightarrow \infty$ . Moreover, in a moderately high-dimensional regime with  $\liminf p/n > 0$ ,  $f(n,p,k)$  scales as  $\sqrt{\log(n) \log(p)}$ . Compared to the bound (5), Lemma 1 is surprisingly strong. The “cost” of adaptively searching over all valid trees in high dimensions only scales logarithmically with  $pk$ .

Based on this lemma, we derive a concentration bound for our odds ratio estimates:

**THEOREM 1.** *Suppose that we have a sequence of problems with parameters  $(n, p, k)$  satisfying Assumptions 1-2. Furthermore, assume that for all  $(t, y) \in \{0, 1\}^2$ , there exists  $\pi_0 > 0$  such that  $\pi_{y,t}(\ell) \geq \pi_0$ , for all  $\Pi \in \mathcal{V}_{\alpha, k}$ ,  $\ell \in \Pi$ . Then,*

$$\lim_{n, p, k \rightarrow \infty} \mathbb{P} \left( \sup_{\Pi \in \mathcal{V}_{\alpha, k}, \ell \in \Pi} |\log(\hat{r}^n(\ell)) - \log(r(\ell))| \geq \frac{8f(n, p, k)}{\pi_0 \sqrt{k}} \right) = 0,$$

$$\text{with } f(n, p, k) := \frac{9}{2} \sqrt{\frac{\log(n/k) (\log(pk) + 3 \log \log(n))}{\log((1 - \alpha)^{-1})}}.$$

*Proof* Denote  $\mathcal{A}_n$  the event

$$\mathcal{A}_n := \left\{ \forall (t, y) \in \{0, 1\}^2, \sup_{\ell \in \mathcal{L}_n} \hat{\pi}_{y,t}^n(\ell) \geq \pi_0/2 \right\}.$$

According to Lemma 1,  $\mathbb{P}(\mathcal{A}_n) \xrightarrow{n \rightarrow \infty} 1$ . Moreover, by concavity of the logarithm, for any  $x, y \geq \pi_0/2$  we have  $|\log(x) - \log(y)| \leq (2/\pi_0)|x - y|$ . Hence, on  $\mathcal{A}_n$ , for any  $\ell \in \mathcal{L}_n$ , we have

$$|\log \hat{r}^n(\ell) - \log r(\ell)| \leq \sum_{(t, y) \in \{0, 1\}^2} |\log(\hat{\pi}_{t, y}^n(\ell)) - \log(\pi_{t, y}(\ell))| \leq \frac{2}{\pi_0} \sum_{(t, y) \in \{0, 1\}^2} |\hat{\pi}_{t, y}^n(\ell) - \pi_{t, y}(\ell)|.$$

Taking the supremum over  $\ell$  on both sides yields, for any  $\delta > 0$ ,

$$\mathbb{P} \left( \sup_{\Pi, \ell} |\log(\hat{r}^n(\ell)) - \log(r(\ell))| \geq \delta \right) \leq \sum_{(t, y) \in \{0, 1\}^2} \mathbb{P} \left( \sup_{\Pi, \ell} |\hat{\pi}_{t, y}^n(\ell) - \pi_{t, y}(\ell)| \geq \frac{\pi_0 \delta}{8} \right) + \mathbb{P}(\mathcal{A}_n^c).$$

Setting  $\delta = 8f(n, p, k)/\pi_0$  concludes the proof.  $\square$

Due to the discrete nature of the splits, decision trees can be notoriously unstable to perturbations in the data. To improve stability, Breiman (2001) originally proposed to introduce randomization in the training process of each tree, train a collection of random trees, and aggregate them in a random forest. Wager and Walther (2015) leveraged their adaptive concentration bounds over trees to derive similar results for random forests. While a thorough analysis of heterogeneous odds ratio estimation via random forest is out of the scope of the present paper, we discuss the different alternatives available to extend our methodology to random forests in Appendix A.2.

#### 4. Adversarial Treatment Assignment Imputation

In this section, we develop a methodology to handle missing treatment assignment, which is a major concern in many pricing and revenue management applications.

#### 4.1. Problem formulation

From a high-level perspective, our HTE estimation problem involves quantities of the form  $\mathbb{E}[\theta(T, Y(T)) | \mathbf{X} = \mathbf{x}]$  for some function  $\theta$  (e.g., cell counts in Table 1) that we estimate empirically within each leaf. However, with transaction-level data, we do not observe the treatment  $T$  whenever  $Y = 0$ . In other words, we cannot estimate the distribution of  $T, Y(T) | \mathbf{X} = \mathbf{x}$ , rendering the estimation of cell counts, odds ratio, and  $Q$ -statistics impossible. As discussed in introduction, the Missing At Random assumption (Rubin 1976) is not suited here for it implies that  $\mathbb{P}(T = 1 | \mathbf{X} = \mathbf{x}) = \mathbb{P}(T = 1 | \mathbf{X} = \mathbf{x}, Y = 1)$  and that the treatment has no effect.

This issue of unobserved treatment assignment is connected with imperfect compliance, i.e., situations when the treatment effectively delivered to each patient,  $D$ , differs from the assigned one,  $T$ . Even in randomized control trials where treatment assignment is random, imperfect compliance can bias estimation of the treatment effect. Efron and Feldman (1991) assume that compliance can be captured by an intensity of treatment delivered (between 0 and 1) and estimate the ATE after making some structural assumptions on how the response depends on treatment intensity. For all-or-nothing compliance, Manski (1990), Balke and Pearl (1997) derive valid upper and lower bounds on the ATE under imperfect compliance. Angrist et al. (1996) propose to use an instrumental variable (e.g., the assigned treatment  $T$ ) to properly estimate the effect of the received treatment  $D$  (see Siddique 2013, for an application). Bargagli-Stoffi et al. (2019), Bargagli-Stoffi and Gnecco (2020) incorporated this IV methodology into a recursive partitioning procedure for HTE estimation. However, these works assume that we observe the imperfect compliance, that is that we observe the assigned and the delivered treatment. On the contrary, in transaction-level data, there is no discrepancy between treatment assigned and delivered, but we do not observe the treatment for any observations that respond negatively ( $Y = 0$ ).

To circumvent this issue, we first make mild distributional assumptions on the distribution of  $T, Y = 0 | \mathbf{X} = \mathbf{x}$ . These assumptions define a set of feasible distributions for  $T, Y = 0 | \mathbf{X} = \mathbf{x}$  and of feasible treatment assignments  $\mathcal{T}$  (Section 4.2). To incorporate uncertainty in treatment assignment into the estimation procedure, we adopt a conservative approach and replace all  $Q$ -statistics by their worst-case value over  $\mathcal{T}$  (Section 4.3). Since finding the worst-case value over  $\mathcal{T}$  is hard, we replace  $\mathcal{T}$  by  $N$  samples drawn uniformly from  $\mathcal{T}$  and provide performance guarantee for this finite sample approximation (Section 4.4).



## 4.2. Ambiguous treatment assignments

In the case of binary treatment, the distribution of  $T, Y = 0 | \mathbf{X} = \mathbf{x}$  is fully parametrized by  $\mathbb{P}(T = 1 | \mathbf{X} = \mathbf{x}, Y = 0)$ . However, in the JD.com data as in many revenue management data, only purchased items are observed. As a result, only the propensity score for purchased products,  $\mathbb{P}(T = 1 | \mathbf{X} = \mathbf{x}, Y = 1)$ , can be estimated. We need to make some structural assumption to relate  $\mathbb{P}(T = 1 | \mathbf{X} = \mathbf{x}, Y = 1)$  and  $\mathbb{P}(T = 1 | \mathbf{X} = \mathbf{x}, Y = 0)$ .

ASSUMPTION 3. *There exists a non-decreasing transformation  $\varphi : [0, 1] \rightarrow [0, 1]$  such that*

$$\mathbb{P}(T = 1 | \mathbf{X} = \mathbf{x}, Y = 0) = \varphi(\mathbb{P}(T = 1 | \mathbf{X} = \mathbf{x}, Y = 1)), \quad \forall \mathbf{x}.$$

Given two contexts  $\mathbf{x}_1, \mathbf{x}_2$  such that  $\mathbb{P}(T = 1 | \mathbf{X} = \mathbf{x}_1, Y = 1) < \mathbb{P}(T = 1 | \mathbf{X} = \mathbf{x}_2, Y = 1)$ , Assumption 3 implies that  $\mathbb{P}(T = 1 | \mathbf{X} = \mathbf{x}_1, Y = 0) \leq \mathbb{P}(T = 1 | \mathbf{X} = \mathbf{x}_2, Y = 0)$ , i.e., the pairwise comparison between  $\mathbf{x}_1$  and  $\mathbf{x}_2$  remains unchanged. Hence, although  $\mathbb{P}(T = 1 | \mathbf{X} = \mathbf{x}, Y = 1)$  does not properly estimate the probability of treatment when  $Y = 0$ , it can be used as a relevant ranking score among items. In short, we assume that discount assignments are different between the two groups but that the drivers of discount are similar. We believe this assumption is fairly general and is broadly applicable beyond revenue management applications.

Assumption 3 alone does not impose any constraint on the resulting odds ratio, which can still take any value between 0 and  $\infty$  (with  $\varphi(z) = 0$  or 1). Problem-specific knowledge can then suggest further restrictions on the transformation  $\varphi$ , such as convexity, bounded Lipschitz constant, or decomposability over a given basis. For example, with our targeted discount application in mind, we can reasonably assume that discount does not negatively impact the probability of purchase, or equivalently:

ASSUMPTION 4. *For any context  $\mathbf{x}$ ,  $\mathbb{P}(T = 1 | \mathbf{X} = \mathbf{x}, Y = 0) \leq \mathbb{P}(T = 1 | \mathbf{X} = \mathbf{x}, Y = 1)$ .*

Assumption 4 is equivalent to the monotonous treatment response assumption (Manski 1997). Recall that, by symmetry, the odds ratio at  $\mathbf{x}$  is equal to  $\text{odds}(z)/\text{odds}(\varphi(z))$ , with  $z = \mathbb{P}(T = 1 | Y = 1, \mathbf{X} = \mathbf{x})$ . Consequently, since  $\text{odds}(\cdot)$  is an increasing function, Assumption 3 implies that the odds ratio is at least equal to 1. The MAR mechanism (which corresponds  $\varphi(z) = z$ ) satisfies Assumptions 3-4 and leads to an odds ratio of precisely 1. Alternatively, the odds ratio can be arbitrarily large by taking  $\varphi(z) \rightarrow 0$ .

Together, Assumptions 3-4 enforces that the quantities  $\mathbb{P}(T = 1|\mathbf{X} = \mathbf{x}_i, Y = 0)$  on our training data belong to the set:

$$\mathcal{T} := \{\mathbf{t} \in [0, 1]^n : t_{\sigma(i)} \leq t_{\sigma(i+1)}, t_i \leq \mathbb{P}(T = 1|\mathbf{X} = \mathbf{x}_i, Y = 1), \forall i\},$$

where  $\sigma$  is a permutation ordering the  $\mathbb{P}(T = 1|\mathbf{X} = \mathbf{x}_i, Y = 1)$ 's in increasing order. Each group of constraints correspond to Assumption 3 and 4 respectively. To construct the set  $\mathcal{T}$  in practice, we fit a machine learning model to predict the propensity score on purchased items,  $\hat{t}(\mathbf{x}) := \hat{\mathbb{P}}(T = 1|\mathbf{X} = \mathbf{x})$ . By applying the function  $\hat{t}$  to unbought items, we estimate  $\hat{t}_i := \mathbb{P}(T = 1|\mathbf{X} = \mathbf{x}_i, Y = 1)$  and construct  $\mathcal{T}$ .

### 4.3. Robust recursive partitioning

A feasible treatment vector is obtained by sampling  $n$  Bernoulli random variables with parameter  $t_i$  for some  $\mathbf{t} \in \mathcal{T}$ . As far as odds ratio estimation is concerned, the quantities  $N_{t,y}$  need to measure the expected number of counts in each cell of the contingency Table 1 so, in the remainder of the section, we assume that treatment assignments can be fuzzy/non-binary for the unbought items and equal to  $t_i$ . Accordingly,  $\mathcal{T}$  should now be interpreted as the set of possible treatment assignments for non-purchased items.

We adapt the HOR estimation procedure (Algorithm A.1) to be robust against uncertainty in treatment assignments. The algorithm is described in pseudo-code in Algorithm 1. We start by computing the treatment assignment vector that minimizes the overall effect of the treatment, i.e., we solve  $\min_{\mathbf{t} \in \mathcal{T}} |1 - \hat{r}(\mathbf{t}, \mathbf{y})|$  where  $\hat{r}(\mathbf{t}, \mathbf{y})$  denotes the empirical odds ratio evaluated on the entire data as in Section 2.1. We then apply our recursive partitioning procedure depth-wise: we first split the root node (depth 0) and then its left and right children (depth 1), and so on. At the end each step, before growing the tree deeper, we update the treatment assignments so as minimize the global heterogeneity  $Q(\Pi, \{(\mathbf{x}_i, t_i, y_i)\})$ . All splits made during this step are then re-assessed: any split that no longer satisfies the splitting condition,  $1 - \chi_1^2(Q(\{\ell_-^*, \ell_+^*\}, \mathcal{D}_n)) \leq p_{\max}$ , is undone. In other words, at each step, we compute the most homogeneous data set according to the  $Q$ -statistics and the stratification from the current partition, and keep only the splits whose validity are robust to this adversarial treatment assignment. We denote this procedure Robust HOR, R-HOR in short.

The main shortcoming of our approach is that, whenever the set of feasible assignment  $\mathcal{T}$  is large, the adversarial can generate treatment assignments that negate the revealed

**Algorithm 1:** Robust recursive partitioning algorithm for HOR (R-HOR)

---

```

Input:  $\{(\mathbf{x}_i, y_i, \hat{t}_i), i = 1, \dots, n\}$ 
Result: Partition  $\Pi$ .
initialization:  $\Pi \leftarrow \{\mathbb{R}^p\}$ ,  $\mathbf{t} \in \arg \min_{\mathcal{T}} |1 - \hat{r}(\mathbf{y}, \mathbf{t})|$ ,  $\mathcal{D}_n \leftarrow \{(\mathbf{x}_i, t_i, y_i), i = 1, \dots, n\}$ ;
for  $d \in \{1, \dots, \text{max\_depth}\}$  do
     $\mathcal{L}_d \leftarrow \{\ell \in \Pi : \text{depth}(\ell) = d\}$  // Note: potentially empty
    for  $\ell \in \mathcal{L}_d$  do
         $(j^*, t^*) \leftarrow \arg \max_{(j,t)} Q(\{\ell_-, \ell_+\}, \mathcal{D}_n)$ ;
        if  $1 - \chi_1^2(Q(\{\ell_-, \ell_+\}, \mathcal{D}_n)) \leq p_{\max}$  and other criterion then
            | split  $\ell$ :  $\Pi \leftarrow (\Pi \setminus \{\ell\}) \cup \{\ell_-, \ell_+\}$ 
        end
    end
     $\mathbf{t} \in \arg \min_{\mathbf{t} \in \mathcal{T}} Q(\Pi, \{(\mathbf{x}_i, t_i, y_i)\})$ ;
    for  $\ell \in \mathcal{L}_d$  do
        if  $1 - \chi_1^2(Q(\{\ell_-, \ell_+\}, \mathcal{D}_n)) > p_{\max}$  then
            | unsplit  $\ell$ :  $\Pi \leftarrow (\Pi \setminus \{\ell_-, \ell_+\}) \cup \{\ell\}$ 
        end
    end
end

```

---

splits. Especially at the first step of the algorithm, when there is only one split to validate. In particular, we previously observed that Assumptions 3-4 are not restrictive and allow the odds ratio to take any value between 1 and  $+\infty$ . We propose two solutions to mitigate this issue. First, we can introduce a reference vector  $\mathbf{t}_0$  and solve instead

$$\min_{\mathbf{t} \in \mathcal{T}} Q(\Pi, \{(\mathbf{x}_i, t_i, y_i)\}) \quad \text{s.t.} \quad \|\mathbf{t} - \mathbf{t}_0\|_1 \leq n\Gamma, \quad (6)$$

where  $\Gamma$  is a parameter controlling the similarity between the treatment assignment  $\mathbf{t}$  and the reference assignment  $\mathbf{t}_0$ . The vector  $\mathbf{t}_0$  can capture prior beliefs on treatment assignments and ensures that the vectors  $\mathbf{t}$  generated at each step of the algorithm remain similar. In our experiments, we take  $\mathbf{t}_0$  equal to the treatment assignment vector obtained at the initialization step. In the robust optimization literature,  $\Gamma$  is referred to as the budget of uncertainty (Bertsimas and Sim 2004) and controls the amount of uncertainty the decision maker wishes to be protected against. We will explore the numerical impact of  $\Gamma$  in Section 7. Second, in our implementation, we replace  $\mathcal{T}$  by a finite-sample approximation. Doing so not only simplifies the minimization problem in (6) but also excludes pathological cases while preserving theoretical guarantees, as described in the following section.

#### 4.4. Finite sample approximation

Algorithm 1 requires finding the worst case treatment assignment  $\mathbf{t} \in \mathcal{T}$ , when worst case is defined in terms of worst odds ratio (initialization) or minimal  $Q$ -statistics. Solving these optimization problems is intractable given the non-convexity of the objectives. Instead, we generate uniformly at random  $N$  potential vectors  $\mathbf{t}^{(j)} \in \mathcal{T}$  and take the worst case over the finite set  $\mathcal{T}^N = \{\mathbf{t}^{(j)}, j = 1, \dots, N\}$ . Naturally, for any criterion  $f$ , we have  $\max_{\mathbf{t} \in \mathcal{T}^N} f(\mathbf{t}) \leq \max_{\mathbf{t} \in \mathcal{T}} f(\mathbf{t})$ , and the larger the  $N$ , the better the approximation, as formalized in the following proposition:

**PROPOSITION 2.** *Denote  $z_N := \max_{\mathbf{t} \in \mathcal{T}^N} f(\mathbf{t})$ . Given a violation probability  $\varepsilon \in (0, 1)$  and a confidence parameter  $\beta \in (0, 1)$ , if  $N \geq \frac{2}{\varepsilon} (1 - \ln \beta)$ , then, with probability at least  $1 - \beta$  (on the samples  $\mathbf{t}^{(j)}$ ), we have*

$$\mathbb{P}_{\mathbf{T}}(f(\mathbf{T}) \leq z_N) \leq \varepsilon,$$

where the probability  $\mathbb{P}_{\mathbf{T}}(\cdot)$  corresponds to the sampling probability over  $\mathcal{T}$

*Proof* We consider a scalar decision variable  $z$  and the robust constraint  $\max_{\mathbf{t} \in \mathcal{T}} f(\mathbf{t}) \leq z$ . Under this lens, our finite sampling approach relates to the scenario approach in robust design optimization and follows immediately from Theorem 1 in Campi et al. (2009).  $\square$

Proposition 2 states that, with probability  $1 - \beta$  over the samples used to construct  $\mathcal{T}^N$ , the inequality  $f(\mathbf{T}) \leq \max_{\mathbf{t} \in \mathcal{T}^N} f(\mathbf{t})$  holds with high probability. Intuitively,  $\beta$  corresponds to the probability of sampling a “bad” set  $\mathcal{T}^N$  and  $\varepsilon$  bounds the probability of  $\mathbf{T}$  falling outside of the convex hull of  $\mathcal{T}^N$ . Note that the distribution of  $\mathbf{T} \in \mathcal{T}$  in Proposition 2 needs to be the same one as the sampling distribution. In our numerical analysis, we sample uniformly from  $\mathcal{T}$ . Accordingly, if pathological cases (e.g.,  $\varphi(z) = 0$  or  $z$ ) are of probability zero, sampling  $N$  scenarios from  $\mathcal{T}$  will exclude these pathological cases with high probability. Regarding the sampling of treatment assignment vectors  $\mathbf{t}^{(j)} \in \mathcal{T}$ , the hit-and-run algorithm (Smith 1984) provides a generic method to sample uniformly from any polyhedron. Given the simplicity of the constraints involved in the definition of  $\mathcal{T}$ , however, we implemented a simple constructive procedure described in Appendix A.3.

## 5. Validation on Synthetic Data

In this section, we assess the validity of our approach on synthetic data. We should acknowledge the fact that whether heterogeneity between patients’ is better captured by the ratio

of responses instead of their difference depends on the problem at hand. As a result, synthetic examples where our procedure outperforms other methods from the HTE literature could be easily constructed, and vice versa. To provide a fair comparison, we focus our analysis on the quality of the  $Q$ -statistics as a splitting criterion. We apply Algorithm A.1 in a setting where responses are continuous and effect of treatment is evaluated in terms of average treatment effect. In this setting, methods from the HTE literature apply and we can compare the methods in their ability to detect heterogeneity.

*Data generation* We adopt a similar methodology as Lee et al. (2020b). For  $j = 1, \dots, p$ ,  $X_j$  is sampled according to a Bernoulli distribution with parameter  $p_j$ , with  $p_j \sim \mathcal{U}(0, 1)$ . We generate the response to treatment and control according to

$$\begin{aligned} Y(0) &= X_1 + \frac{1}{2}X_2 + X_3 + \kappa(X_1, X_2, X_3) + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, 1), \\ Y(1) &= Y(0) + \tau_k(X_1, X_2), \end{aligned}$$

where  $\kappa(x_1, x_2, x_3) = \exp(x_1 - x_2x_3)$  creates non-linear heterogeneity in responses (both to treatment and control) while  $\tau_k(x_1, x_2)$  captures the heterogeneous treatment effect and is one of the following:

- 2-rule case:  $\tau_k(x_1, x_2) = k$  if  $x_1 = 0$  and  $x_2 = 0$ ,  $-k$  if  $x_1 = 1$  and  $x_2 = 1$ , and 0 otherwise.
- 4-rule case:  $\tau_k(x_1, x_2) = k$  if  $x_1 = 0$  and  $x_2 = 1$ ,  $2k$  if  $x_1 = 0$  and  $x_2 = 0$ ,  $-k$  if  $x_1 = 1$  and  $x_2 = 0$ ,  $-2k$  if  $x_1 = 1$  and  $x_2 = 1$ , and 0 otherwise.

The parameter  $k$  controls the magnitude of the treatment effect. Finally, the treatment assignment variable  $T$  is distributed according to the logit model  $\mathbb{P}(T = 1 | \mathbf{X} = \mathbf{x}) = \text{logit}(-1 + x_1 - x_2 + x_3)$ .

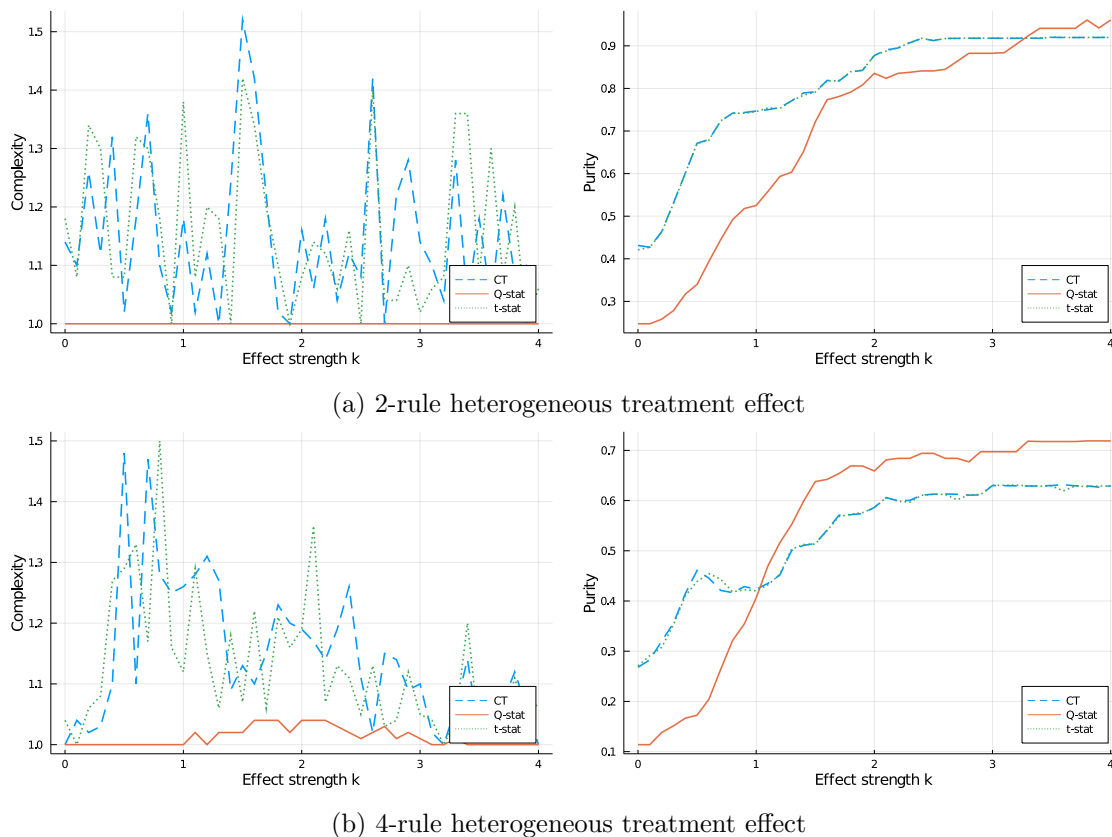
*Our method* Since responses are continuous, we apply our recursive partitioning procedure (Algorithm A.1) to the case where the effect of treatment is measured in terms of conditional average treatment effect (CATE). In-sample estimates of the CATE and its variance are obtained within each leaf and used in the computation of the  $Q$ -statistics for each split. We fix a maximal depth of 10 and  $p_{\max} = 0.1$ .

*Metrics* We sample  $N = 1,000$  observations, for different values of  $k$  ranging between 0 and 4. Consider a partition  $\Pi$  of the feature space. Each rule in the definition of  $\kappa$  (e.g.,  $x_1 = 1$  and  $x_2 = 1$ ) is associated with a set of leaves of  $\Pi$ ,  $\{\ell_r \in \Pi : r \in \mathcal{R}\}$ , such that for any  $r \in \mathcal{R}$ , there is at least one observation  $\mathbf{x}_i$  from the training data such that  $\ell(\mathbf{x}_i, \Pi) = \ell_r$  and  $\mathbf{x}_i$  satisfy the rule. Intuitively,  $\{\ell_r \in \Pi : r \in \mathcal{R}\}$  should be understood

as the smallest subset of leaves needed to cover all the observations satisfying the rule. With this definition, we measure the quality of rule detection in terms of two metrics: *complexity*, i.e., the cardinality of  $\mathcal{R}$ , and *purity*, i.e., the proportion of observations in  $\{\ell_r \in \Pi : r \in \mathcal{R}\}$  that indeed satisfy the rule. We say that  $\Pi$  perfectly detects a rule if it achieves a complexity as low as 1 and a purity of 100%. For each partition, we report complexity and purity, averaged over all rules defining  $\kappa$ .

*Benchmark* We compare our splitting criterion (Q-stat) to alternatives: the  $t$ -statistics of Su et al. (2009) (t-stat), and the causal tree of Athey and Imbens (2016) (CT), as implemented in the R package `causalTree` (Athey and Imbens 2016). Following best practices, we prune the trees obtained in this fashion using a cross-validation criterion that can be different from the splitting one. We consider three cross-validation criteria, denoted `TOT`, `fit` and `CT` in the `causalTree` documentation. When reporting the complexity or the purity of t-stat/CT, we report the value obtained with the cross-validation criterion achieving the best performance, thus leading to optimistic performance of the benchmarks.

*Results* Figure 1a and 1b report the results for the case where  $\kappa$  comprises 2 and 4 rules respectively. From these results, we make the following observations: First, our procedure is a sound procedure to detect heterogeneous treatment effects. It returns partitions with low complexity and a purity comparable with what other methods from the literature achieve. As the strength of the treatment effect  $k$  increases, purity of the leaves increases as well. Compared to the  $t$ -test and CT procedures, our Q-statistics criterion achieves better complexity, i.e., lower and less variable. This improvement is mostly due to fact that our hypothesis testing-based criterion softly penalizes the number of samples in each leaves, leading to relatively shallow trees. On the contrary, the two alternatives grow substantially deeper trees and require a pruning step which is a known source of instability (see Appendix A.4 for results obtained without pruning). Regarding purity, the performance of  $t$ -test and CT are remarkably similar. For low values of  $k$ , they dominate our approach by a substantial margin. This is not surprising since these methods were developed specifically for CATE estimation and leverage the fact that treatment effect can be expressed as a sample average. Our approach is more general, hence, less powerful in this setting, yet remarkably competitive. In particular, the edge of  $t$ -test and CT shrinks as the effect strength or the number of rules increases. Figure A.2 in appendix also suggests that Algorithm A.1 might be more robust to confounding variables. Together, our results suggest



**Figure 1** Complexity (left panel) and purity (right panel) as the effect strength  $k$  increases on synthetic examples with 2-rule (top) and 4-rule (bottom) heterogeneous treatment effect. Metrics are averaged over 20 random datasets, containing  $N = 1,000$  observations each.

that our method is a viable and competitive alternative to the existing methods for HTE, in particular since it can accommodate for various measures of effect, such as odds ratios, unlike CT and  $t$ -test.


## 6. Applications in Social Science and Medicine

In this section, we apply Algorithm A.1 to two data sets from social science and medicine.

### 6.1. Social pressure and voting

Gerber et al. (2008) conducted a large experiment to assess the effect of social pressure on voter turnout, by randomly sending mailings to 200,000 registered voters in the United States. We apply our HOR estimation procedure to mine for response heterogeneity, and compare our optimal treatment recommendations with prior analysis (Zhou et al. 2018).

*Data description* The dataset consists of  $n = 180,002$  voters from the state of Michigan. There are 10 voter characteristics: year of birth, sex, household size, city, and 6 voting indicator variables for the past three general and primary elections (2000, 2002, and 2004).



**Figure 2** Best partition of depth at most 3 obtained for measuring heterogeneity in response to social pressure for the voting dataset.

Gerber et al. (2008) randomly mailed a letter to the households before the 2006 primary election, with one of the four following messages (treatments):

- Civic Duty: “Do your civic duty”.
  - Monitored : “You are being studied”.
  - Self History: The letter contains the household’s past voting records. The letter also indicates that, once the election is over, a follow-up letter on whether the voter has voted will be sent to the household.
  - Neighbors: The letter contains the voting records of the household and their neighbors, and also indicates that “all your neighbors will be able to see your past voting records”.
- A fifth group of households were not mailed anything (control group). The response to treatment is whether a voter has voted to the 2006 primary election (binary).


*Results and insights* We apply our recursive procedure for HOR estimation, in particular its multi-treatment extension developed in Section 2.5. We set  $p_{\max}$  to 0.001 and a maximal depth of 3. The resulting tree is displayed in Figure 2. Consistent with the analysis of Gerber et al. (2008), Zhou et al. (2018), we find that voters display relatively homogeneous responses to mailings and that assigning “Neighbors” to everyone is the best policy.

## 6.2. Application on a medicine example

We also illustrate our approach on medical data from a large clinical trial conducted in the United States starting in 2006.

*Data description* The Prostate, Lung, Colorectal and Ovarian (PLCO) Cancer Screening Trial is a large randomized trial designed and sponsored by the National Cancer Institute to determine the effects of screening on cancer-related mortality. We restrict our attention to colorectal cancer and conduct two analyses, namely evaluate the impact of treatment (i.e., screening) on the incidence of colorectal cancer and the corresponding cancer stage.





**Figure 3** Outcome of the heterogeneous odds ratio estimation procedure on the PLCO data set. Algorithm A.1 is set with a maximal depth of 3 and a  $p_{max}$  value of 0.1.

The trial comprises 154,887 individuals, 50.5% female, aged between 42 and 78 years old at randomization. Out of them, 2,299 contracted and/or died of colorectal cancer. The covariates we use for heterogeneity detection correspond to data available at the beginning of the trial, namely age, gender, prior history of colorectal or any cancer, and eligibility to colorectal analysis based on the Baseline Questionnaire.

*Results and insights* We first consider the entirety of the cohort and evaluate the impact of screening on overall cancer incidence. At a population level, we do observe that intervention leads to lower incidence of colorectal cancer, with odds ratio 0.7984 (95% CI: 0.7387 - 0.8630), with a stronger effect on men than women (see Figure 3a). For the 2,299 patients who contracted a colorectal cancer, we study whether the screening impacted the cancer stage at detection. Indeed, an earlier stage cancer is associated with higher survival (Banister and Broggio 2016). Accordingly, we consider a stage I or II cancer as a “positive” outcome. Screening does have an overall positive impact on cancer stage (OR: 0.98 - 1.36). Yet, it is important to notice that we observe no significant effect on patients aged less than 60 years old (Figure 3b). From a policy perspective, this analysis suggests that enforcing mandatory screening before 60 years old would be inefficient in improving cancer screening and that screening efforts should focus on 60+ year-old individuals.

## 7. Application on a Revenue Management Example

We apply the proposed R-HOR method to data from the e-tailer JD.com (Shen et al. 2020) in order to study the effect of discounts on the purchase probability.

### 7.1. Data description and processing

We aggregated data from different parts of the dataset, which we describe in Appendix A.5.1. We refer to Shen et al. (2020) for a more detailed presentation of the data and summary statistics. The final dataset consists of 566,004 customer-product interactions,

**Table 2** Out-of-sample performance of machine learning algorithms for predicting discount assignment

Model	Hyper-parameters	Out-of-sample <i>AUC</i>
Decision tree	maximal depth, criterion	0.7163
Random forest	maximal depth, criterion, number of trees	0.7295
Nearest neighbors	neighborhood size, norm	0.6914

described with 9 variables concerning the user (number of past purchase, time since first purchase, “Plus” membership indicator, age, gender, education, marital status, purchase power category, city level), 4 variables concerning the product (product type, original price, attribute 1 and 2 missingness indicator), and 2 variables regarding the interaction (day of the week, hour of the day). For each interaction, we know whether the customer eventually purchased the product. If so, we know whether they were offered a discounted price.

### 7.2. Predicting discount assignment

Out of all 566,004 customer-product interactions, 42,184 of them constitute a transaction. 61% of these transactions were offered a discount. We use this data to train a machine learning model to predict the probability of being treated, i.e., discounted, given the 15 observed covariates,  $\mathbb{P}(T = 1 | \mathbf{X} = \mathbf{x}, Y = 1)$ . We randomly split the data into a training (70%) and a test set (30%). We calibrate all models on the training data, using 3-fold cross validation for tuning hyper-parameters. We compare a single decision tree (Breiman et al. 1984), random forest (Breiman 2001) and nearest neighbor classification (Cover and Hart 1967) and report their out-of-sample *AUC* in Table 2. We acknowledge that this comparison is purely illustrative and not exhaustive.

As summarized in Table 2, all machine learning models can predict discount assignment with an *AUC* in the low 0.70s. Out of the three, random forest is the most accurate model and the one we use in the remaining analysis. Among others, machine learning models such as these ones can be used to identify the drivers of discounts in the platform’s current promotion targeting strategy (see, e.g., Table A.1 in the online appendix).

### 7.3. Imputing discount assignments

We now apply the previously-trained random forest model to all observations in our data set corresponding to unbought items, thus obtaining a score  $\hat{t}_i \in [0, 1]$  for each item.

Under the MAR assumption,  $\hat{t}_i$  estimates the probability  $\mathbb{P}(T = 1 | \mathbf{X} = \mathbf{x}_i)$ . Accordingly, for unbought items, we can either consider a single data set where unbought items are softly assigned to treatment, i.e.,  $(\mathbf{x}_i, \hat{t}_i, y_i = 0)$ , or consider  $N_s$  potential treatment assignments

**Table 3** Comparison of the different imputation methods in terms of fraction of discounted items and overall odds ratio. For methods producing multiple imputed data sets, we report average values (and standard deviation)


	MAR	MAR - sampled	MICE	$\mathcal{T}$
Fraction of discounts	62.9%	62.9% (0.1%)	63.7% (0.4%)	34.9% (12.6%)
Odds ratio	0.929	0.929 (0.003)	0.895 (0.015)	4.072 (3.994)

sampled according to  $\hat{\mathbf{t}}$ , i.e.,  $(\mathbf{x}_i, t_i^{(j)}, y_i = 0)$ ,  $j = 1, \dots, N_s$ , with  $t_i^{(j)} \sim \text{Bernoulli}(\hat{t}_i)$ . We will refer to these options as “MAR” and “MAR sampled” respectively. We can also impute the missing treatment assignments using chained equations (Buuren and Groothuis-Oudshoorn 2010), a widely recognized technique for missing data imputation, implemented in the R package `mice`. We perform 5 iterations of the algorithm to generate 5 different imputed data set (default configuration) and refer to this method as “MICE”. Finally, we compare these approaches with the methodology presented in Section 4.4, where we consider  $N_s$  potential treatment assignments  $\hat{\mathbf{t}}^{(j)} \in \mathcal{T}$ . We fix  $N_s = 160$ , which, according to Proposition 2, will lead to a violation probability  $\epsilon$  and a confidence parameter  $\beta$  of 5%.

We compare the four imputation methods in Table 3. First, we observe that “MAR” and “MAR-sampled” lead to similar results, which is justified by the fact that odds ratio can account for fractional treatment assignment. Second, methods relying on the MAR assumption, namely “MAR”, “MAR-sampled” and “MICE”, provide similar results: they replicate the treatment assignment mechanism observed on the purchased items. As a result, the proportion of discounted items is similar to the one observed on the bought items (61.1%) and the overall odds ratio is close to 1. In addition, multiple imputation methods (“MAR-sampled” and “MICE”) produce data sets which tend to be very similar. On the contrary, our less restrictive set of assumptions lead to a greater diversity of scenarios, which is both a blessing and a curse. On the bright side, it allows for cases where discounts can have a (strong) effect. However, scenarios will be generated adversarially within the HOR estimation algorithm so a larger set of treatment assignments also provides the adversarial with more power to revoke our findings.

#### 7.4. Heterogeneous odds ratio estimation

We first consider the vector of treatment assignments  $\mathbf{t}$  leading to the worst odds ratio, i.e., the closest to 1. By doing so, we estimate that offering a discount slightly increases the likelihood of purchase, with an odds ratio of 1.12 (95% CI: 1.10-1.14)—see the contingency table, Table A.2, in the online appendix.




**Figure 4** Outcome of the heterogeneous odds ratio estimation procedure on the JD.com data set. Algorithm A.1 is set with a maximal depth of 1 and a  $p_{max}$  value of 0.01.

This treatment assignment serves as the starting point for our robust heterogeneous assignment. Figure 4a displays the best tree of depth 1 found with this treatment assignment. In particular, we observe that discounts are more effective on cheap products, i.e., product that cost less than 36 (currency unknown). Following the robust estimation procedure (Algorithm 1), this split has to be validated on a range of possible treatment assignment values by solving Problem (6). Figure 4b displays the worst  $p$ -value for Cochran’ $Q$  test applied on this split for different budgets of uncertainty  $\Gamma$ . The higher the  $\Gamma$ , the larger the range of potential treatment assignments  $\mathcal{T} \cap \{\mathbf{t} : \|\mathbf{t} - \mathbf{t}_0\|_1 \leq n\Gamma\}$ , the higher the  $p$ -value can be. In the remaining of the analysis, we fix  $p_{max} = 0.01$  and  $\Gamma = 0.15$ .

Figure 5 displays the best robust partition obtained of depth at most 3. From this partition, we can make the following observations: First, discount has a overall positive impact on purchase probability. Second, product-related variables lead to more heterogeneity in discount effectiveness than customer-related information, suggesting that customers can be relatively homogeneous in their response to discounts. However, their behavior is strongly dependent on the product being discounted. Finally, Two categories of products benefit the most from promotions: cheap products and products from a “1P” supplier.

## 8. Conclusion

Motivated by transaction-level data from a major e-tailer, we propose a recursive partitioning procedure for estimating heterogeneous odds ratio, a widely used measure in medicine, economics, and social sciences. Our procedure can be applied to any definition of the treatment effect for which an asymptotically normal estimator exists. In addition, we develop an adversarial imputation procedure to allow for robust estimation in presence of partially



**Figure 5** Best partition of depth at most 3 obtained for measuring heterogeneity in response to discount on JD.com data.

observed treatment assignment, a central problem in data-driven revenue management. We believe that the integration of problem-specific imputation methods directly into statistical estimation, and their theoretical analysis, constitutes a promising future direction for analytics in operations management.

## Acknowledgments

The authors would like to thank JD.com and the M&SOM 2020 Data Driven Research Challenge committee for offering access to transaction-level data and administering this research challenge, which motivated this work. The authors also thank the National Cancer Institute for access to NCI's data collected by the Prostate, Lung, Colorectal and Ovarian (PLCO) Cancer Screening Trial analyzed in Section 6.2. The statements contained herein are solely those of the authors and do not represent or imply concurrence or endorsement by NCI. Thanks to Ali Aouad, Colin Fogarty, and Nicos Savva for providing early feedback on the work and suggesting relevant references, and to two anonymous referees for their constructive comments.

## References

- Aglin G, Nijssen S, Schaus P (2020) Learning optimal decision trees using caching branch-and-bound search. *AAAI*, 3146–3153.
- Angrist JD, Imbens GW, Rubin DB (1996) Identification of causal effects using instrumental variables. *Journal of the American Statistical Association* 91(434):444–455.
- Aouad A, Elmachtoub AN, Ferreira KJ, McNellis R (2019) Market segmentation trees. *arXiv preprint arXiv:1906.01174* .
- Athey S, Imbens G (2016) Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences* 113(27):7353–7360.
- Athey S, Tibshirani J, Wager S, et al. (2019) Generalized random forests. *The Annals of Statistics* 47(2):1148–1178.
- Balke A, Pearl J (1997) Bounds on treatment effects from studies with imperfect compliance. *Journal of the American Statistical Association* 92(439):1171–1176.

- Bannister N, Broggio J (2016) Cancer survival by stage at diagnosis for england (experimental statistics): adults diagnosed 2012, 2013 and 2014 and followed up to 2015. *Produced in collaboration with Public Health England* .
- Bargagli-Stoffi FJ, De Witte K, Gnecco G (2019) Heterogeneous causal effects with imperfect compliance: a novel bayesian machine learning approach. *arXiv preprint arXiv:1905.12707* .
- Bargagli-Stoffi FJ, Gnecco G (2020) Causal tree with instrumental variable: an extension of the causal tree framework to irregular assignment mechanisms. *International Journal of Data Science and Analytics* 9(3):315–337.
- Bénard C, Biau G, Da Veiga S, Scornet E (2019) Sirius: making random forests interpretable. *arXiv preprint arXiv:1908.06852* .
- 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):164–183.
- Bertsimas D, Sim M (2004) The price of robustness. *Operations Research* 52(1):35–53.
- Breiman L (2001) Random forests. *Machine Learning* 45(1):5–32.
- Breiman L, Friedman J, Stone CJ, Olshen RA (1984) *Classification and regression trees* (CRC press).
- Breslow N (1981) Odds ratio estimators when the data are sparse. *Biometrika* 68(1):73–84.
- Breslow N, Liang K (1982) The variance of the Mantel-Haenszel estimator. *Biometrics* 943–952.
- Buuren Sv, Groothuis-Oudshoorn K (2010) mice: Multivariate imputation by chained equations in r. *Journal of Statistical Software* 1–68.
- Campi MC, Garatti S, Prandini M (2009) The scenario approach for systems and control design. *Annual Reviews in Control* 33(2):149–157.
- Chan KY, Loh WY (2004) Lotus: An algorithm for building accurate and comprehensible logistic regression trees. *Journal of Computational and Graphical Statistics* 13(4):826–852.
- Cochran WG (1954) The combination of estimates from different experiments. *Biometrics* 10(1):101–129.
- Cornfield J, Haenszel W, Hammond EC, Lilienfeld AM, Shimkin MB, Wynder EL (1959) Smoking and lung cancer: recent evidence and a discussion of some questions. *Journal of the National Cancer institute* 22(1):173–203.
- Cover T, Hart P (1967) Nearest neighbor pattern classification. *IEEE transactions on information theory* 13(1):21–27.
- Cummings P (2009) The relative merits of risk ratios and odds ratios. *Archives of Pediatrics & Adolescent Medicine* 163(5):438–445.
- Davies HTO, Crombie IK, Tavakoli M (1998) When can odds ratios mislead? *British Medical Journal* 316(7136):989–991.

- Demirović E, Lukina A, Hebrard E, Chan J, Bailey J, Leckie C, Ramamohanarao K, Stuckey PJ (2020) Murtree: Optimal classification trees via dynamic programming and search. *arXiv preprint arXiv:2007.12652* .
- Efron B, Feldman D (1991) Compliance as an explanatory variable in clinical trials. *Journal of the American Statistical Association* 86(413):9–17.
- Engels EA, Schmid CH, Terrin N, Olkin I, Lau J (2000) Heterogeneity and statistical significance in meta-analysis: an empirical study of 125 meta-analyses. *Statistics in Medicine* 19(13):1707–1728.
- Gart JJ (1966) Alternative analyses of contingency tables. *Journal of the Royal Statistical Society: Series B (Methodological)* 28(1):164–179.
- Gart JJ, Zweifel JR (1967) On the bias of various estimators of the logit and its variance with application to quantal bioassay. *Biometrika* 181–187.
- Gerber AS, Green DP, Larimer CW (2008) Social pressure and voter turnout: Evidence from a large-scale field experiment. *American Political Science Review* 33–48.
- Goodman B, Flaxman S (2017) European union regulations on algorithmic decision-making and a right to explanation. *AI Magazine* 38(3):50–57.
- Gordon L, Olshen RA (1978) Asymptotically efficient solutions to the classification problem. *The Annals of Statistics* 515–533.
- Gordon L, Olshen RA (1980) Consistent nonparametric regression from recursive partitioning schemes. *Journal of Multivariate Analysis* 10(4):611–627.
- Greenland S, Robins JM, Pearl J (1999) Confounding and collapsibility in causal inference. *Statistical Science* 29–46.
- Günlük O, Kalagnanam J, Menickelly M, Scheinberg K (2018) Optimal decision trees for categorical data via integer programming. *arXiv preprint arXiv:1612.03225* .
- Haldane JB (1956) The estimation and significance of the logarithm of a ratio of frequencies. *Annals of Human Genetics* 20(4):309–311.
- Hernán MA, Robins JM (2010) *Causal inference: What if* (Boca Raton: Chapman & Hill/CRC).
- Huedo-Medina TB, Sánchez-Meca J, Marín-Martínez F, Botella J (2006) Assessing heterogeneity in meta-analysis: Q statistic or  $I^2$  index? *Psychological Methods* 11(2):193.
- Imbens GW, Rubin DB (2015) *Causal inference in statistics, social, and biomedical sciences* (Cambridge University Press).
- Kallus N (2017) Recursive partitioning for personalization using observational data. *International Conference on Machine Learning*, 1789–1798 (PMLR).
- Kallus N, Zhou A (2018) Confounding-robust policy improvement. *Advances in Neural Information Processing Systems*, 9269–9279.

- Kulinskaya E, Dollinger MB (2015) An accurate test for homogeneity of odds ratios based on Cochran's Q-statistic. *BMC Medical Research Methodology* 15(1):1–19.
- Laber EB, Zhao YQ (2015) Tree-based methods for individualized treatment regimes. *Biometrika* 102(3):501–514.
- Lakkaraju H, Rudin C (2017) Learning cost-effective and interpretable treatment regimes. *Artificial Intelligence and Statistics*, 166–175.
- Landwehr N, Hall M, Frank E (2005) Logistic model trees. *Machine Learning* 59(1-2):161–205.
- Last M, Maimon O, Minkov E (2002) Improving stability of decision trees. *International Journal of Pattern Recognition and Artificial Intelligence* 16(02):145–159.
- Lawson R (2004) Small sample confidence intervals for the odds ratio. *Communications in Statistics-Simulation and Computation* 33(4):1095–1113.
- Lee HS, Zhang Y, Zame W, Shen C, Lee JW, van der Schaar M (2020a) Robust recursive partitioning for heterogeneous treatment effects with uncertainty quantification. *arXiv preprint arXiv:2006.07917* .
- Lee K, Bargagli-Stoffi FJ, Dominici F (2020b) Causal rule ensemble: Interpretable inference of heterogeneous treatment effects. *arXiv preprint arXiv:2009.09036* .
- Lin J, Zhong C, Hu D, Rudin C, Seltzer M (2020) Generalized and scalable optimal sparse decision trees. *arXiv e-prints arXiv:2006*.
- Manski C (1997) Monotone treatment response. *Econometrica* 65(6):1311–1334, ISSN 0012-9682, URL <http://dx.doi.org/10.2307/2171738>.
- Manski CF (1990) Nonparametric bounds on treatment effects. *The American Economic Review* 80(2):319–323.
- Mantel N, Haenszel W (1959) Statistical aspects of the analysis of data from retrospective studies of disease. *Journal of the National Cancer Institute* 22(4):719–748.
- Meinshausen N, Ridgeway G (2006) Quantile regression forests. *Journal of Machine Learning Research* 7(6).
- Musalem A, Olivares M, Bradlow ET, Terwiesch C, Corsten D (2010) Structural estimation of the effect of out-of-stocks. *Management Science* 56(7):1180–1197.
- Oprescu M, Syrgkanis V, Wu ZS (2019) Orthogonal random forest for causal inference. *International Conference on Machine Learning*, 4932–4941 (PMLR).
- Perakis G, Singhvi D, Skali-Lami O (2019) Xstrees: A tree sampling framework for interpretable tree ensembles. *Available at SSRN 3505431* .
- Poole C, Shrier I, VanderWeele TJ (2015) Is the risk difference really a more heterogeneous measure? *Epidemiology* 26(5):714–718.
- Rosenbaum PR (2014) Sensitivity analysis in observational studies. *Wiley StatsRef: Statistics Reference Online* .



- 
- Rosenbaum PR, Rubin DB (1983) The central role of the propensity score in observational studies for causal effects. *Biometrika* 70(1):41–55.
- Rubin DB (1976) Inference and missing data. *Biometrika* 63(3):581–592.
- Shen M, Tang CS, Wu D, Yuan R, Zhou W (2020) JD.com: Transaction-level data for the 2020 MSOM data driven research challenge. *Manufacturing & Service Operations Management* .
- Siddique Z (2013) Partially identified treatment effects under imperfect compliance: the case of domestic violence. *Journal of the American Statistical Association* 108(502):504–513.
- Smith RL (1984) Efficient monte carlo procedures for generating points uniformly distributed over bounded regions. *Operations Research* 32(6):1296–1308.
- Su X, Tsai CL, Wang H, Nickerson DM, Li B (2009) Subgroup analysis via recursive partitioning. *Journal of Machine Learning Research* 10(2).
- Tran C, Zheleva E (2019) Learning triggers for heterogeneous treatment effects. *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 5183–5190.
- Vapnik VN, Chervonenkis AY (1971) On the uniform convergence of relative frequencies of events to their probabilities. *Measures of Complexity*, 11–30 (Springer).
- Vulcano G, Van Ryzin G, Ratliff R (2012) Estimating primary demand for substitutable products from sales transaction data. *Operations Research* 60(2):313–334.
- 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.
- Wager S, Walther G (2015) Adaptive concentration of regression trees, with application to random forests. *arXiv preprint arXiv:1503.06388* .
- Woolf B, et al. (1955) On estimating the relation between blood group and disease. *Annals of Human Genetics* 19(4):251–253.
- Yuan Y, Shaw MJ (1995) Induction of fuzzy decision trees. *Fuzzy Sets and Systems* 69(2):125–139.
- Zantedeschi V, Kusner MJ, Niculae V (2020) Learning binary trees via sparse relaxation. *arXiv preprint arXiv:2010.04627* .
- Zeileis A, Hothorn T, Hornik K (2008) Model-based recursive partitioning. *Journal of Computational and Graphical Statistics* 17(2):492–514.
- Zhou Z, Athey S, Wager S (2018) Offline multi-action policy learning: Generalization and optimization. *arXiv preprint arXiv:1810.04778* .

## Online Appendix

### Pauphilet, *Robust and Heterogenous Odds Ratio: Estimating Price Sensitivity for Unbought Items*

#### A.1. Recursive procedure for Heterogeneous Odds Ratio estimation

We provide in this section the pseudo-code of the recursive partitioning procedure for heterogeneous odds ratio estimation described in Section 2.3.

The algorithm is initialized with the entire space as the initial partition,  $\Pi = \{\mathbb{R}^p\}$ . At each step, the best split is found by exploring the  $p$  features and all the potential values for  $t$ . For continuous variables, this search can be accelerated by searching threshold values among a predefined list of quantiles as in Aouad et al. (2019). Finally, depth of the leaf or number of training samples in the leaf (total or by treatment regime) can be used as extra criteria to declare a leaf unsplittable, although they are already indirectly captured in the  $Q$ -statistics threshold. To account for multiple hypothesis testing, the threshold value  $p_{\max}$  could be reduced throughout the algorithm.

---

**Algorithm A.1:** Recursive partitioning algorithm for heterogeneous odds ratio estimation

---

**Input:**  $\mathcal{D}_n = \{(\mathbf{x}_i, t_i, y_i), i = 1, \dots, n\}$

**Result:** Partition  $\Pi$ .

initialization:  $\Pi \leftarrow \{\mathbb{R}^p\}$ ;

**for**  $\ell \in \Pi$  **do**

$(j^*, t^*) \leftarrow \arg \max_{(j,t)} Q(\{\ell_-, \ell_+\}, \mathcal{D}_n)$  ;

**if**  $1 - \chi_1^2(Q(\{\ell_-^*, \ell_+^*\}, \mathcal{D}_n)) \leq p_{\max}$  **and other criterion** **then**

        split  $\ell$ :  $\Pi \leftarrow (\Pi \setminus \{\ell\}) \cup \{\ell_-, \ell_+\}$

**end**

**end**

---

#### A.2. Discussion: From trees to forest

Due to the discrete nature of the splits, decision trees can be notoriously unstable to perturbations in the data. To improve stability, Breiman (2001) originally proposed to introduce randomization in the training process of each tree, train a collection of random trees, and aggregate them in a random forest. Wager and Athey (2018) later extended this idea for HTE estimation. Wager and Walther (2015) leveraged their adaptive concentration bounds

over trees to derive similar results for random forests. While a thorough analysis of heterogeneous odds ratio estimation via random forest is out of the scope of the present paper, we discuss in this section the different alternatives available to extend our methodology to random forests.

**Source of randomness:** Following Breiman (2001), randomness can be introduced by calibrating each tree on a bootstrap sample of the training data  $\mathcal{D}_n$ . Randomness can also be introduced in the split selection process, for instance by searching the best split among a random subset of covariates (typically, of size  $\sqrt{p}$ ), or by imposing that at each iteration all variables can be selected with probability at least  $\pi/p$  for some  $\pi \in (0, 1]$ . The later assumption was proposed by Meinshausen and Ridgeway (2006) to ensure that the size of the leaves shrinks as  $n \rightarrow \infty$  and guarantee consistency. Note that stability of decision trees can also be improved by introducing randomness in the nature of the splits directly instead of the training process (see Yuan and Shaw 1995, Last et al. 2002, Bénard et al. 2019, Perakis et al. 2019).

**Aggregation:** Typically, the output of a random forest is obtained by averaging. However, given the non-collapsibility of odds ratios, the outputs of each tree in a forest cannot be simply averaged. Consider an observation  $\mathbf{x}$  and a random forest consisting of  $n_{trees}$  trees. Each tree  $\tau$  places  $\mathbf{x}$  in a leaf  $\ell(\mathbf{x}, \tau)$ . Let us denote  $\pi_{t,y}^{(\tau)}(\mathbf{x})$  the value of  $\hat{\pi}_{t,y}^n(\ell(\mathbf{x}, \tau))$ , and  $N^{(\tau)}(\mathbf{x})$  the number of training samples in that leaf. By generalizing Lemma 1 to random forests as in Wager and Walther (2015), we can show that  $\frac{1}{n_{tree}} \sum_{\tau} \pi_{t,y}^{(\tau)}(\mathbf{x})$  provides consistent estimates of  $\pi_{t,y}(\mathbf{x})$ . Hence, under the assumptions of Theorem 1,

$$\frac{\left( \sum_{\tau} \pi_{1,1}^{(\tau)}(\mathbf{x}) \right) \left( \sum_{\tau} \pi_{0,0}^{(\tau)}(\mathbf{x}) \right)}{\left( \sum_{\tau} \pi_{1,0}^{(\tau)}(\mathbf{x}) \right) \left( \sum_{\tau} \pi_{0,1}^{(\tau)}(\mathbf{x}) \right)}$$

is a consistent estimate of the odds ratio at  $\mathbf{x}$ . Note that this aggregation rule differs from the Mantel-Haenszel estimator (Mantel and Haenszel 1959):

$$\frac{\sum_{\tau} \pi_{1,1}^{(\tau)}(\mathbf{x}) \pi_{0,0}^{(\tau)}(\mathbf{x}) N^{(\tau)}(\mathbf{x})}{\sum_{\tau} \pi_{1,0}^{(\tau)}(\mathbf{x}) \pi_{0,1}^{(\tau)}(\mathbf{x}) N^{(\tau)}(\mathbf{x})}.$$

We refer to Breslow (1981), Breslow and Liang (1982) for a discussion on the estimation of its variance. However, the Mantel-Haenszel estimator is only valid if leaves of different trees can be interpreted as independent strata, which is not the case if they are obtained

via bootstrapped sample of the same training data. Finally, under the assumption that the response follows a logistic model,

$$\log \frac{\mathbb{P}(Y = 1 | \mathbf{X} = x, T = t)}{\mathbb{P}(Y = 0 | \mathbf{X} = x, T = t)} = r(\mathbf{x})t + \boldsymbol{\nu}(\mathbf{x})^\top \mathbf{x},$$

the odds ratio corresponds to the coefficient associated with the treatment variable,  $r(\mathbf{x})$ . Athey et al. (2019) extended the random forest methodology to estimate such non-parametric generalized linear model (here, the functions  $r$  and  $\boldsymbol{\nu}$  are non parametrized): For each  $\mathbf{x}$ ,  $r(\mathbf{x})$  can be estimated by solving a weighted maximum likelihood estimation problem, where each training sample is weighted according to membership in the leaves of the random forest  $\mathbf{x}$  belongs to. Oprescu et al. (2019) later proposed an orthogonalized version of the procedure that is robust to the estimation error in the nuisance component  $\boldsymbol{\nu}(\mathbf{x})$ . Besides positing a logistic model, the main limitation of these approaches, however, is that the odds ratio is not explicitly provided by the random forest but rather as the solution of an optimization problem that depends on the random forest output.

### A.3. Uniform sampling procedure for treatment assignments

Given scores  $\hat{\mathbf{t}} \in [0, 1]^n$ , we randomly sample from  $\mathcal{T}$  following Algorithm A.2.

---

#### Algorithm A.2: Uniform sampling procedure

---

**Input:**  $\hat{\mathbf{t}} \in [0, 1]^n$

**Result:** random sample  $t \in \mathcal{T}$ .

initialization:  $t_i \leftarrow \mathbf{NA}$ ,  $\sigma \leftarrow \text{sortperm}(\hat{\mathbf{t}})$ ;

**while**  $\{i : t_i = \mathbf{NA}\} \neq \emptyset$  **do**

$i \leftarrow \text{random}(\{i : t_i = \mathbf{NA}\})$  ;

$\mathcal{I}_- \leftarrow \{j : \sigma(j) < \sigma(i), t_j \neq \mathbf{NA}\}$ ,  $\mathcal{I}_+ \leftarrow \{j : \sigma(j) > \sigma(i), t_j \neq \mathbf{NA}\}$

$LB \leftarrow \max(t_j, j \in \mathcal{I}_-)$ ,  $UB \leftarrow \min(t_j, j \in \mathcal{I}_+)$  // Note: if  $\mathcal{I}_-$  (resp.  $\mathcal{I}_+$ ) empty, 0 (resp. 1).

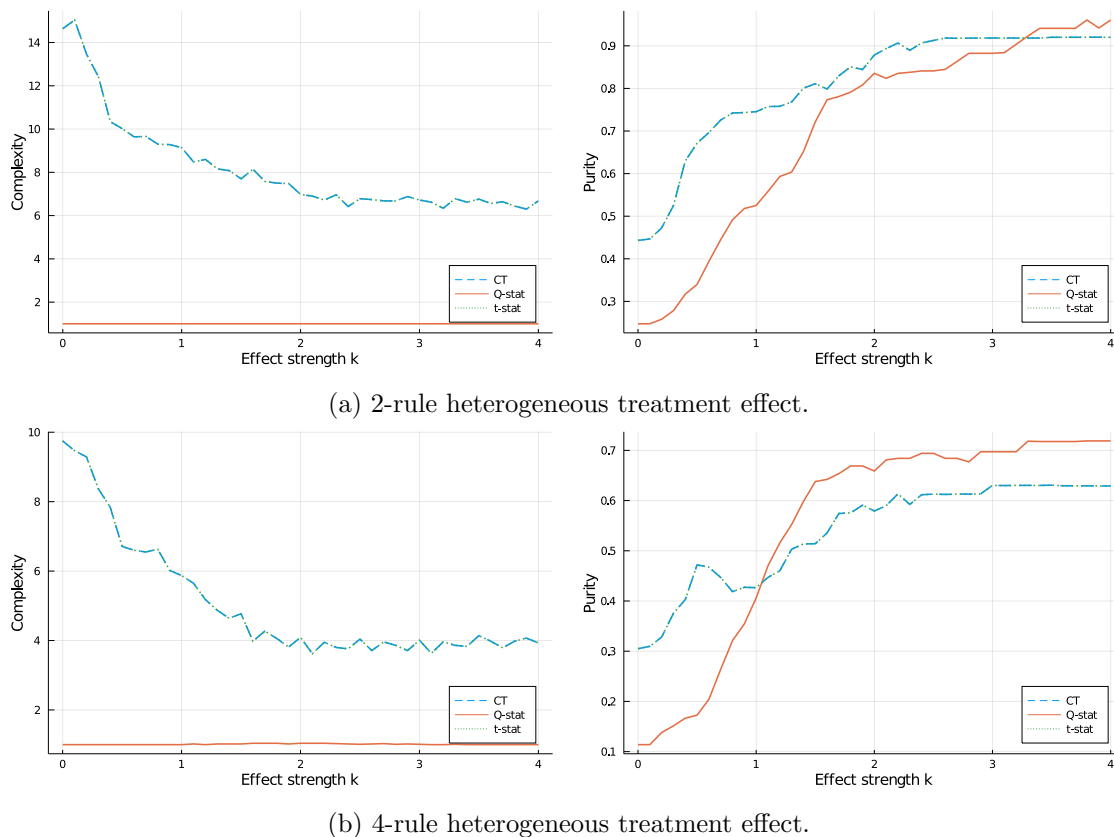
$UB \leftarrow \min(UB, \hat{t}_i)$ ;  $t_i \leftarrow \mathcal{U}([LB, UB])$ ;

**end**

---

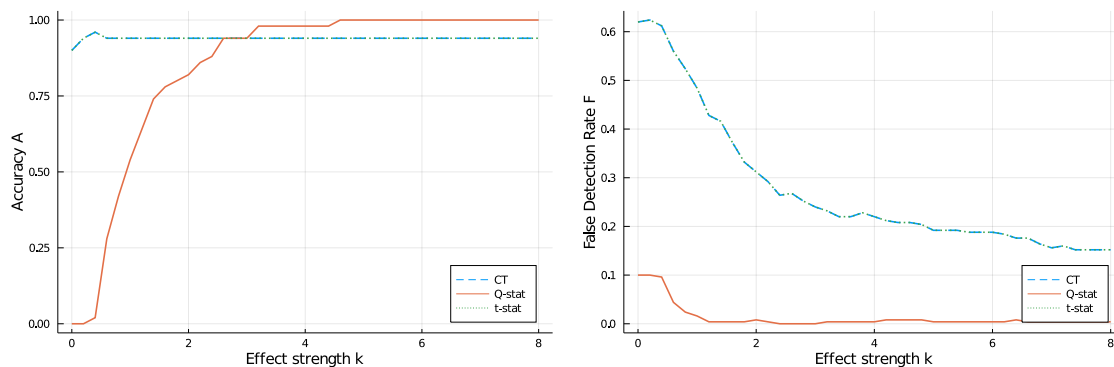
### A.4. Additional numerical experiments on synthetic data

Figures A.1a and A.1b replicate Figures 1a and 1b from Section 5, except that trees obtained by “ $t$ -test” and “CT” are not pruned. These results clearly demonstrate the importance of pruning in reducing the size of the trees, hence achieving lower complexity, without compromise on their quality (purity almost unchanged). Yet, it is also clear from these pictures that pruning introduces instability.



**Figure A.1** Replication of Figure 1 when trees obtained by “*t*-test” and “CT” are not pruned. The left and right panel respectively display the complexity and purity as the effect strength  $k$  increases, on synthetic examples with 2-rule (top) 4-rule (bottom) heterogeneous treatment effect. Metrics are averaged over 20 random datasets, containing  $N = 1,000$  observations each.

To evaluate the robustness of our recursive partitioning procedure to the confounding factors, we conduct a similar set of experiments except that the treatment effect is now defined as  $\tau_k(X_4, X_5)$ . In other words,  $X_1, X_2, X_3$  impact the baseline effect  $Y(0)$  and the propensity score  $\mathbb{P}(T = 1 | \mathbf{X} = \mathbf{x})$ , while  $X_4, X_5$  drive the treatment effect heterogeneity. If  $\mathcal{S}$  denotes the set of features the tree splits on, we evaluate the accuracy  $A = |\mathcal{S} \cap \{4, 5\}|/2$  and false detection rate  $F = |\mathcal{S} \setminus \{4, 5\}|/|\mathcal{S}|$ . With these notations, a good HTE algorithm should split on  $X_4$  and  $X_5$ , i.e.,  $A = 1$ , and  $X_4, X_5$  only, i.e.,  $F = 0$ . Figure A.2 compares the accuracy (left panel) and false detection rate (right panel) of all three methods as the effect strength  $k$  increases, for the 2-rule treatment. Causal trees and *t*-test based trees are pruned. For each metric, the cross-validation criterion for these two methods is chosen so as to optimize this particular metric. Figure A.2 suggests that our recursive partitioning procedure is more robust to confounding than either causal trees or *t*-test based tests.



**Figure A.2** Accuracy (left panel) and False detection rate (right panel) as the effect strength  $k$  increases, to measure the ability of each HTE method to isolate features affecting treatment heterogeneity. Trees obtained by “ $t$ -test” and “CT” are pruned. Metrics are averaged over 20 random synthetic datasets with 2-rule, containing  $N = 1,000$  observations each.

## A.5. Supporting material for the revenue management example

This section provides details on the analysis of the JD.com data (Shen et al. 2020).

### A.5.1. Data processing

We aggregated data from different sources, which we briefly present in this section. We emphasize the data characteristics that are relevant to our analysis but refer to Shen et al. (2020) for a more detailed presentation of the data.

*Clicks data:* This data contains the list of pages/products visited by each user, the channel they used (e.g., desktop, mobile, app), and the corresponding time-stamp. We say that two pages are part of the same session if they are visited by the same user, on the same platform, within 10 minutes. If a user visited the same page multiple times within the same session, we only retain her last visit.

*Users data:* Each user is described by a unique identifier, some statistics on her purchase history (number of past purchases, time of first purchase, loyalty program), and demographics (e.g., gender, education).

*Products data:* We characterize each product by its type and two attributes. Given the large proportion of products with missing attributes, we only use the presence of each attribute as a covariate, not its value (when available). With deanonymized data, product category-specific imputations could be considered as well.

*Orders data:* For each order, we have access to the initial and final price, the difference indicating whether a discount was applied. In our analysis, we do not discriminate between the different types of discounts (direct, quantity, bundle,...). For unbought items, we do

know neither the original price nor whether the price was discounted. For the original price, we impute it using the closest (in time) price for the same product in the training data - the “original price” listed in the data is product-specific and the same for all users. For the discount indicator, we apply our robust estimation strategy from Section 4. We exclude from our analysis products that have never been purchased, products whose price changed more than 62 times (twice a day) over the month of March 2018, and transactions consisting of free products.

All together, the final dataset consists of 566,004 customer-product interactions, described with 9 variables concerning the user (number of past purchase, time since first purchase, “Plus” membership indicator, age, gender, education, marital status, purchase power category, city level), 4 variables concerning the product (product type, original price, attribute 1 and 2 missingness indicator), and 2 variables regarding the interaction (day of the week, hour of the day).

### A.5.2. Predicting discount assignment

A predictive model to predict discount assignment can be used to understand the drivers of discounts (or back-engineer the promotion targeting strategy). Table A.1 reports the most important features for the decision tree and random forest models. In both cases, original price is the most important feature, with a clear edge over the second best one.

**Table A.1** Key parameters driving the application of a promotion on bought items. For each variable, we also report the importance value, according to each model.

Variable rank	Decision tree	Random forest
1	Original price (0.59)	Original price (0.25)
2	Product type (0.08)	Time since 1st purchase (0.15)
3	Time since 1st purchase (0.06)	Hour of the day (0.14)
4	Attribute 2 missing (0.03)	City type (0.06)
5	Hour of the day (0.03)	Number of past purchases (0.05)

### A.5.3. Heterogeneous odds ratio estimation

We first consider the vector of treatment assignments  $\mathbf{t}$  leading to the worst odds ratio, i.e., the closest to 1. By doing so, we obtain the following the  $2 \times 2$  contingency table:

**Table A.2**  $2 \times 2$  contingency table on JD.com’s data with adversarial treatment assignments.

	Y = 0	Y = 1
T = 0	218,069	16,399
T = 1	305,751	25,785