Causal machine learning for personalized discount targeting with technographic trace data

Alex P. Miller
Marshall School of Business
University of Southern California
alex.miller@marshall.usc.edu

Kartik Hosanagar
The Wharton School
University of Pennsylvania
karitkh@wharton.upenn.edu

Abstract

We investigate the potential of using technographic data—user characteristics revealed in the metadata of standard web communication protocols—for the targeting of promotional discounts to online shoppers. While researchers have historically emphasized the importance of purchase history and behavioral data for promotional targeting, such data is typically unavailable about most visitors to online storefronts. Digital trace data, on the other hand, has the potential to be used in targeted promotional campaigns, even for individuals with whom a firm may have no prior relationship. The value of this data, however, has yet to be empirically investigated in this context. To do so, we develop a novel framework that identifies the optimal targeting policy for a flexible set of discount campaigns. We demonstrate how to use machine learning methods and data from an online experiment to estimate this policy and apply our techniques to data from two promotional campaigns at separate online retailers. Using counterfactual policy evaluation, we find (1) our proposed methodology outperforms both a non-targeted baseline and standard benchmark techniques in targeted marketing, and (2) that technographic trace data can be used as an effective means of price discrimination in online retail. We estimate the increases in profit possible using these methods to be on the order of 3%-6%—equal to thousands of dollars of incremental value over the course of the studied promotional campaigns. An explanatory analysis reveals that device-specific variables such as screen size, operating system, and web browser (as opposed to geographic or behavioral variables) prove most valuable for targeting purposes in this context. By empirically quantifying the value of this data for price discrimination, this project adds valuable insight to the growing discussion about the use of personalization technologies for price-related interventions on the web, with implications for e-commerce managers, consumer advocates, and policymakers.

Notes: An earlier draft of this paper was the winner of a "Best Short Paper in Track" award at the *International Conference on Information Systems* 2020. We thank participants for their feedback at the 2020 *Symposium on Statistical Challenges in Electronic Commerce Research* and the 2020 *Conference on Information Systems & Technology*. Research support for this work is funded in part by grants from the Mack Institute and the Baker Retailing Center.

1 Introduction

The use of technology to deliver personalized customer experiences has been a key strategy in retail management for decades (Peppers and Rogers, 1993). As shopping has increasingly moved online, algorithmic product recommendation, service customization, and targeted promotional messaging have all become staples of the retail experience (Ansari and Mela, 2003, Dias et al., 2008). But in addition to the value information technology has for improving customers' shopping experiences, there is a growing appreciation of the power of technology as a means of price discrimination (Wallheimer, 2018). As such, algorithmically-enabled price discrimination, or "personalized pricing", is an increasingly important topic of study in legal, social, and economic disciplines. Our objective in this project is to study the potential value of price discrimination in the e-commerce environment, with a particular focus on evaluating the prospect of using technographic trace data to target promotional discounts to online shoppers.

Though there has been increased interest in the subject in recent years, price discrimination of various sorts has been an essential practice throughout the history of retail. Promotional discounts in particular have been used by firms for decades as a form of passive price discrimination, whereby consumers effectively self-segment by deciding to use (or not use) a discount coupon. (Bawa and Shoemaker, 1987, Blattberg et al., 1995, Bolton, 1989, Narasimhan, 1984, Reibstein and Traver, 1982). But with the advent of customer relationship management (CRM) techniques and software, it became possible to measure and exploit heterogeneity in customer preferences (Drew et al., 2001, Shapiro and Varian, 1998). As a result, firms in many industries started to use CRM data to *target* their discounts at the segment or individual customer level (Bawa and Shoemaker, 1989, Grewal et al., 2011). Within this vein, several empirical studies have investigated the value of using CRM data as a means of price discrimination through the strategic targeting of price discounts. (Johnson et al., 2013, Musalem et al., 2008, Rossi et al., 1996, Zhang and Krishnamurthi, 2004). A common finding within this body of work is the importance of using purchase history data to segment consumers; among the studies that include other variables (e.g., customer demographics) within a targeting framework, the authors find these features to be significantly less valuable than customer relationship data (Gupta and Chintagunta, 1994, Khan et al., 2009).[1]

[1]The importance of behavioral history profiles for targeting has been demonstarted in other settings as well,

While assuming the presence of purchase history data makes sense in many historical contexts, the online retail environment presents a novel set of both opportunities and challenges for customer targeting. Though traditional, CRM-based methods of targeting are still valuable for email and mobile application promotions (Dubé et al., 2017, Ghose et al., 2019, Luo et al., 2014, Sahni, 2015), they are much less useful for targeting interventions of active web browsers. E-commerce firms have the unique ability to deliver algorithmically personalized experiences in real-time to anonymous website visitors for whom they may have no relational history data.[2] For any user merely browsing a firm's online storefront—even if they are not logged in or are otherwise anonymous to the firm—modern e-commerce websites can exploit the presence of digital trace data provided by nearly all modern web browsers to target that person with a discount offer or other form of promotional messaging. In Table 1, we have listed various forms of technographic data generated by web browsers and the information this data can reveal about an individual.

In some sense, the transition to online-shopping can seen as both a curse and a blessing from the targeter's perspective: on one hand, many (if not the vast majority of) online shoppers will not be logged in when browsing an online storefront and, thus, traditional mechanisms for promotional discount targeting will be irrelevant; on the other hand, the presence of digital trace data provides a means of distinguishing users that visit a firm's website, whether or not they have any prior relationship with the firm. Prior work has demonstrated the value of adaptive web personalization using digital trace data in non-promotional contexts (Kobsa et al., 2001, Padmanabhan et al., 2001, Yi et al., 2009, Zhang, 2003), but existing literature contains no research about the value of this data for discount targeting in online retail. And while there is descriptive research that has shown firms engaging in technographic price discrimination (Hamermesh, 2013, Hannak et al., 2014), there is a lack of empirical research on this subject in terms of methodological implementation and quantifiable economic value.

This project makes several contributions to the growing literature on personalized pricing and targeted marketing. First, we derive the theoretically optimal targeting strategy for campaigns involving promotional price discounts and describe a technique for using ma-

such as web and mobile advertising (Rafieian and Yoganarasimhan, 2020, Trusov et al., 2016).

[2]While it has been possible for brick-and-mortar firms to reach prospective customers through the purchase of market intelligence and demographic data on households (Simester et al., 2020), these techniques can be expensive and make little sense for digital-first e-commerce enterprises, for which their addressable market has no geographic constraint.

Table 1: Common trace data provided by modern web browsers

Data Source	Inferrable information
TCP header	
IP address	Approximate location, internet service provider
HTTP header	
`User-Agent`	Operating system, operating system version, browser, browser version, device type
`Referer`	How the user arrived to the website; search query (if provided)
`Cookie`	Whether or not the user has visited website before
Client-side tracking script	
Built-in browser JavaScript objects	Device features, e.g., screen size, color depth, language, timezone, graphics hardware, plugins installed, fonts installed, canvas hash

chine learning to estimate this strategy from experimental data. Our model is specifically designed to accommodate many different forms of price promotions that are common in online retail, making it a flexible framework for estimating and measuring the returns from optimal targeting policies in a variety of campaigns. As a novel theoretical contribution, we find that in contrast to prior literature on targeted marketing in the absence of discounts, the optimal discount targeting strategy depends on segmenting customers based on a calibrated trade-off between their individual baseline purchase rate and heterogeneity in their response to the discount. Next, our methodology demonstrates how to use counterfactual policy evaluation to assess the economic value of a targeted discount strategy using data from a randomized experiment; thus firms can use this technique to accurately estimate the profitability of a campaign prior to implementation. Lastly, we conclude by applying our proposed approach to real-world data from A/B tests at two separate firms and find our method results in significantly higher profits than both non-targeted strategies and existing benchmarks in targeted marketing. This research is the first to quantify the empirical value of technographic data for discount targeting, finding that it can result in gains that are competitive with those observed in prior research based on traditional CRM data.

2 Background & Related Literature

Given the growing interest and concern about the use of digital personalization technologies for price discrimination, it is useful to contextualize our project within these broader topics. We briefly review the economic literature on price discrimination, and then review more recent developments in public policy and consumer advocacy around the potential and practice of personalized pricing. In the subsequent section, we discuss how our work relates to prior methodological research on targeted marketing.

2.1 Price discrimination, welfare, & digital privacy. The traditional lens through which to evaluate the overall effect of price discrimination—in both academic and public policy settings—has been to measure its impact on Marshallian welfare, based on utilitarian notions of producer and consumer surplus. Within this framework, economic models of monopoly markets have consistently found that price discrimination decreases consumer surplus while increasing firm profits and total welfare. (Katz, 1987, Varian, 1989). However, as various researchers have incorporated more complex dynamics into models of personalized pricing—including the effects of competition (Choudhary et al., 2005), quality differentiation (Ghose and Huang, 2009), price-comparison technologies (Chen and Sudhir, 2004), and strategic consumer disclosure of data (Ali et al., 2020, Chen et al., 2020)—the implications for firm profits and consumer surplus are ambiguous, with various models predicting increases and decreases in both quantities. Two empirical studies have specifically investigated the welfare effects of personalized pricing on the internet, with both supporting the monopoly theory of price discrmination, finding that total welfare increases at the expense of consumer surplus (Dubé and Misra, 2019, Shiller, 2014).[3]

Of course, there are other lenses through which to view the practice of personalized pricing besides its implications for utilitarian economic welfare. Various forms of price discrimination remain controversial, as consumers generally view the practice of charging different consumers different prices to be unfair, especially when such decisions are based on personal characteristics (Englmaier et al., 2012, Huang et al., 2005, Kahneman et al., 1986). Further, targeted price discrimination has the potential to result in unintended racial or class discrimination (Ayres and Siegelman, 1995, Larson et al., 2015, Miller and Hosanagar, 2019). There is also a growing concern among consumer advocates and policymakers around the use of big data and digital technologies for personalization more generally (Alreck and Settle, 2007, O'Neil, 2016, The Council of Economic Advisers, 2015,

[3]Though information technology has increased the ability of firms to price discriminate, it is also relevant to consider the ways consumers have benefited more generally from the transition to e-commerce. For example, a macroeconomic analysis of the European market suggests that e-commerce has had substantial benefits for household welfare (Cardona et al., 2015). The effects of shopbots have served to increase competition and keep prices low (Tang et al., 2010); similar effects likely exist for coupon and deal aggregators. We highlight this to note that even purely economic analysis of how novel technologies affect societal welfare can be a very complex and nuanced exercise.

Wagner and Eidenmuller, 2019).

As a result, there has been a vigorous policy debate in recent years with landmark privacy-focused regulations being passed in several major markets—including the EU's General Data Protection Regulation (GDPR) and California's Consumer Privacy Act (CCPA). These polices have generally focused on regulating the use of "personal data" and cross-session, cross-site cookie tracking. One consequence of these laws has been to force many firms to delete CRM data gathered without affirmative consent, with some companies losing up to 70% of their records (Hall, 2018). As such, the value of this data and traditional CRM practices—which, as prior research has shown, have been critical for customer targeting historically—face a future with increased compliance costs, regulatory scrutiny, and market uncertainty.

At the same time, the extent to which technographic trace data falls under the purview of these privacy laws is ambiguous. IP addresses and geolocation have been singled out by some authorities as protected personal data (Meyer, 2018, Reid, 2017), but the collection of other forms of trace data—such as a device's operating system, browser version, and referring domain—do not appear to be subject to the same strict regulations (unless these data are tied to a user's personal profile within a CRM). In any case, given that this data can be anonymously logged and that it requires no historical user profiles or third-party tracking software, it almost certainly faces lower regulatory barriers than many forms of CRM and behavioral profiling data that have emerged in recent years.[4]

In light of these circumstances, research on the value of this form of first-party trace data for targeted marketing and price discrimination may be of interest to several stakeholders, including managers, consumer advocates, and policymakers. While we cannot definitively address total welfare effects of personalized pricing generally, this project is able to provide some empirical insights that are relevant to the broader discussion on personalization technologies and the value of technographic data for price discrimination. In particular, we attempt to quantify the how profitable this data is for the purposes of offering targeted discount offers to online shoppers.

[4]Attempts to use the unique combination of all a user's technographic data for device fingerprinting—by virtue of attempting to identify or link a user's information across sources—is almost certainly prohibited without affirmative consent under EU privacy laws (Laperdrix et al., 2020).

2.2 Statistical methods for targeted marketing. While we focused much of the prior discussion on the history of *discount* targeting and personalized pricing, this project is also related to the literature on targeted marketing more generally and the use of machine learning in marketing applications. Several innovations in statistics and economics have highlighted the potential of machine learning for estimating individual-level, heterogeneous treatment effects using experimental data and high-dimensional covariate information (Chernozhukov et al., 2018, Künzel et al., 2019, McFowland III et al., 2018, Taddy et al., 2016, Wager and Athey, 2017). In marketing in particular, multiple recent papers have demonstrated the value of these innovations for targeting purposes (Gutierrez and Gérardy, 2017, Hitsch and Misra, 2018, Yoganarasimhan et al., 2020).

We also highlight the long history of research on using statistical techniques to develop targeted outreach strategies for direct mail and customer retention campaigns (Cui et al., 2006, Gensch, 1984, Kim et al., 2005). Research in these areas has emphasized the merits of targeting customers least likely to take a desired action (e.g., customers who may be at the lowest risk of contract renewal) (Ascarza and Hardie, 2013, Neslin et al., 2006) or targeting all customers that will respond positively to a given marketing campaign—a staple technique of so-called "uplift" modeling (Ascarza, 2018, Lo, 2002, Radcliffe, 2007). An important insight, highlighted by the recent work of Lemmens and Gupta (2020) in the context managing customer churn, is the value of directly using individual-level profitability as a criterion for targeting. Another recent paper by Yang et al. (2020) proposes a novel technique for targeting customers based on their long-term profitability outcomes.

Many of the interventions studied in the research cited in the previous paragraph do involve promotional discounts, but none of these papers factor in the effects of discounting on the immediate profitability of a promotional campaign. Prior work on targeted promotional pricing has accounted for these effects, but these studies only offer empirical evidence for the value of personalization in B2B settings (Dubé and Misra, 2019) or B2C settings in which relational purchase data is available (Johnson et al., 2013, Rossi et al., 1996). In this project, we are able to study the effects of personalized price discrimination in a novel empirical context, and also to more closely relate the literature on targeted marketing for pricing and non-pricing interventions. In particular, we demonstrate that, in the case of discount interventions, the traditional marketing strategies described above—based on either baseline response rates or heterogeneous response to intervention—are insufficient for estimating the optimal targeting policy. Our model shows (and empirical findings confirm) that an optimal discount targeting strategy depends on a calibrated trade-off between both of these quantities.

Further, as a practical matter, a key benefit of our model is that it can be used to target campaigns with various discount and cost structures. This is important for the modern online retail environment, in which frequent and varied discount campaigns are a key component of both consumer expectations and digital marketing strategy (RetailMeNot, 2018). And given that consumers are known to exhibit heterogeneous response to different types of discounts (Ahmad and Callow, 2018, Broeder and Derksen, 2018, Cao et al., 2018, Chen et al., 2012, Shampanier et al., 2007), there is considerable value in the ability to use a generic, parsimonious, and effective targeting framework in an *ad hoc* fashion across various marketing campaigns.

The remainder of the paper is organized as follows: in Section 3 we describe a theoretical model of discount targeting and derive the optimal policy; then, in Section 4 we describe a framework for estimating this policy using a combination of experimental data and methods from machine learning; Section 5 describes the results of an empirical investigation on the value of our proposed approach. We conclude by discussing the implications of this work for digital marketing managers and policymakers.

3 Decision-theoretic model for optimal discount targeting

3.1 Problem Set-up. We consider an e-commerce firm that observes a continuous stream of users to their online storefront. When a user (indexed by i) arrives, the firm observes $X_i \in \mathcal{X}$, a vector of customer characteristics, and must decide on a treatment $T_i \in \{0, 1\}$ to which the user will be assigned. Without loss of generality, we think of $T_i = 0$ as the control treatment in which no discounts are offered; in the treatment $T_i = 1$, users are offered a promotional discount.

We allow the exact nature of this discount to vary flexibly in our model. While this adds some complexity, this feature is motivated by the fact that, in the real world, firms make promotional offers of varying types — and the profit a firm earns takes a different form depending on the discount type. ; we parameterize the structure of a campaign's discount by allowing it to take the form of either a percentage price discount in the amount of $d \times 100\%$ for some $d \in [0, 1]$ (e.g., "20% off"), or a level discount amount in the amount of $k \geq 0$ dollars (e.g., "\$20 off"). In the e-commerce setting, one can think of each treatment manipulation as a banner at the top of the retailer's website advertising the associated discount.

For each user, the firm observes the amount of revenue spent by the customer at the end of their session, indicated by the variable $R_i \in [0, \infty)$. Note in this model the revenue variable R_i is equal to the nominal price of any goods purchased *before* discounts are considered. We also define $C_i := \mathbf{1}\{R_i > 0\}$ as a shorthand variable to indicate the binary outcome of whether a user's session ends with a purchase. Lastly, we assume the firm may have some non-ignorable marginal cost c for each purchase on their site. Using this notation, we can express the firm's profit for users in a the control condition ($T_i = 0$) as:

$$\pi_i = R_i - c$$

For users in the treatment condition in which they are offered a discount ($T_i = 1$), the firm's observed profit will be given as:

$$\pi_i = (1-d)R_i - (c+k)C_i$$

Profits in this case are calculated as nominal revenue R_i times $(1-d)$ when a percentage discount is offered and, when a conversion is observed (i.e., $C_i = 1$), marginal costs c and the level discount amount k are deducted.[5]

Lastly, to simplify our notation in subsequent derivations, define the *conditional response functions* of a given targeting campaign using the following notation:

$$\mu_t^R(x) := \mathbf{E}\left[R_i \mid T_i = t, X_i = x\right]$$
$$\mu_t^C(x) := \mathbf{E}\left[C_i \mid T_i = t, X_i = x\right] = \Pr\left[C_i = 1 \mid T_i = t, X_i = x\right]$$

These represent the conditional expected values of revenues (R_i) and conversion rates (C_i), respectively, of a user with observed covariate $X_i = x$ under treatment assignment $T_i = t$. Further, we define the *conditional average treatment effect* functions for both revenue and conversion as:

$$\tau^R(x) := \mathbf{E}[R_i \mid T_i = 1, X_i = x] - \mathbf{E}[R_i \mid T_i = 0, X_i = x] = \mu_1^R(x) - \mu_0^R(x)$$
$$\tau^C(x) := \mathbf{E}[C_i \mid T_i = 1, X_i = x] - \mathbf{E}[C_i \mid T_i = 0, X_i = x] = \mu_1^C(x) - \mu_0^C(x)$$

[5] As a simple extension to the offline retail environment, we note that in cases where the firm may incur outreach costs, such as in direct mail, the profit function can be modified to allow for an outreach-dependent cost q, which will need to be accounted for in subsequent derviations: $\pi_i = C_i(R_i(1-d) - c - k) - q$. In this case, the subsequent optimization problem may be constrained by a budget, requiring a linear programming solution; e.g., see Imai and Strauss (2011). Because we focus specifically on digital targeting in this project, we assume $q = 0$ and therefore any budget considerations are ignorable.

3.2 Optimal targeting policy. Given the setup above, we are now poised to investigate the nature of the firm's profit-maximizing targeting policy. The central problem we seek to solve is whether or not the firm should assign a user (specified by some covariate value X_i) to a discount treatment ($T_i = 1$) or a control treatment ($T_i = 0$).

To continue, we formalize the firm's decision problem by introducing notation for the firm's *decision function* or *policy* $\delta : \mathcal{X} \longrightarrow \{0, 1\}$, which maps the space of user-covariates into the space of treatments. That is, for an arbitrary user i, δ is the policy the firm uses to assign treatment so that $T_i = \delta(X_i)$.[6] We will consider decision functions of a particular form, denoted by δ_f, that depend on a thresholding function f for assigning treatments:

$$\delta_f(x) = \begin{cases} 1 & \text{if } f(x) > 0 \\ 0 & \text{if } f(x) \leq 0 \end{cases}$$

We will refer to f as the firm's *targeting function*.

In this notation, the problem of finding the profit-maximizing targeting policy can be expressed in the following form:

$$\max_f \Pi(f) = \mathbf{E}[\pi_i \mid T_i = \delta_f(X_i)] \tag{1}$$

In the following proposition, we derive the exact form of the optimal targeting function in terms of the various quantities described above.

Proposition 1. *Given the discount and cost parameters (d, k, c) associated with a targeting campaign, the firm's optimal score function, i.e, the argument of the maximum of the optimization problem in (1), is given by:*

$$f^*(x) = \tau^R(x) - c\tau^C(x) - \left[d\mu_1^R(x) + k\mu_1^C(x)\right] \tag{2}$$

$$= (1-d)\tau^R(x) - (c+k)\tau^C(x) - \left[d\mu_0^R(x) + k\mu_0^C(x)\right] \tag{3}$$

Proof. By definition, $T_i^* = 1$ if and only if $\mathbf{E}[\pi_i \mid T = 1, X = x] > \mathbf{E}[\pi_i \mid T = 0, X = x]$.

Substituting in the definitions of conditional response and condtional average treatment

[6]To clarify notation, we use T_i to indicate the treatment a user is assigned in the abstract, random variable sense and use $\delta(X_i)$ to indicate the treatment chosen for a user by a given decision rule δ.

effect functions, we have:

$$
\begin{aligned}
\mathbf{E}[\pi \mid T = 1, X = x] &> \mathbf{E}[\pi \mid T = 0, X = x] \\
\mathbf{E}[R(1-d) - cC - kC) \mid T = 1, X = x] &> \mathbf{E}[R - cC \mid T = 0, X = x]) \\
\mu_1^R(x)(1-d) - c\mu_1^C(x) - k\mu_1^C(x) &> \mu_0^R(x) - c\mu_0^C(x) \\
\mu_1^R(x) - \mu_0^R(x) - c(\mu_1^C(x) - \mu_0^C(x)) &> d\mu_1^R(x) + k\mu_1^C(x) \\
\tau^R(x) - c\tau^C(x) &> d\mu_1^R(x) + k\mu_1^C(x) \\
\tau^R(x) - c\tau^C(x) &> d(\mu_0^R(x) + \tau^R(x)) + k(\mu_0^C(x) + \tau^C(x)) \\
(1-d)\tau^R(x) - (c+k)\tau^C(x) &> d\mu_0^R(x) + k\mu_0^C(x)
\end{aligned}
$$

If we let

$$f^*(x) = (1-d)\tau^R(x) - (c+k)\tau^C(x) - [d\mu_0^R(x) + k\mu_0^C(x)]$$

then $f^*(x) > 0$ if and only if $\mathbf{E}[\pi \mid T = 1, X = x] > \mathbf{E}[\pi \mid T = 0, X = x]$.

■

3.3 Comments on optimal policy. Before continuing, we make several remarks about the result from Proposition 1. First, as we will demonstrate in the following section, the decision criteria in Eq. (2) can be readily evaluated by using experimental data to estimate the response and treatment effect functions; Eq. (3) is a completely equivalent expression of this criteria. Note that the left side of Eq. (2) can be described as the expected gain in profit by offering a discount and the right side (in brackets) represents the costs of offering a discount to a user. To see this, note that $\mu_1^R(x)$ represents the total revenue the firm expects from a user by targeting them with a discount; in the case of a percentage discount, $d\mu_1^R(x)$ represents the money the firm loses as a direct consequence of the discount on the user's total spend; in the case of a level discount, $k\mu_1^C(x)$ is the relevant quantity. But $\tau^R(x) - c\tau^C(x)$ represents the *incremental gain* in revenue the firm can expect by targeting that user with a discount treatment. If this gain exceeds the costs, it is profitable to offer that user a discount promotion.

Next, to relate this targeting strategy to prior work on targeted marketing in non-discount settings, we focus on the form of the optimal targeting function in Eq. (3), and consider the simplest discount structure where a firm offers a percentage discount with no marginal costs. In this scenario, the optimal targeting criteria can be written as $(1-d)\tau^R(x) > d\mu_0^R(x)$. Expressed in this form, it is apparent that the optimal targeting condition—even in the simplest of cases—requires individual-level estimates of *both* treatment heterogeneity and baseline response rates. We highlight this as a contrast to existing research on targeted marketing in the context of retention and direct mail campaigns, which has traditionally focused on only one of these quantities (Ascarza and Hardie, 2013, Neslin et al., 2006, Radcliffe, 2007). Ascarza (2018) does compare targeting campaigns based on both quantities, but concludes for the purposes of their research context that it is more beneficial to focus on treatment effects than baseline responses as a targeting criteria. However, when a marketing campaign offers a promotional price incentive, and a firm is interested in maximizing the short-run profits of the campaign in question, the theoretically optimal targeting policy depends on a calibrated trade-off between *both* baseline response rates and responses to treatment, in a way that depends on the exact parameters of the firm's campaign.

Lastly, we reiterate that the policy we derive in Proposition 1 is designed to accommodate multiple different discount campaign scenarios. It is very unlikely that all exogenous parameters in the model will be non-zero in a real-world application, but our generic framework allows us to derive the form of the optimal policy across a variety of campaigns. To highlight how this framework can be used in realistic scenarios, Table 2 compiles a list of cost structures common in several marketing campaigns and applies Proposition 1 to derive the mathematical form of the optimal policy.

Table 2: Realistic discount campaigns with corresponding optimal targeting policies

Description	Relevant cost parameters	Optimal targeting function
Percentage discount	d	$(1-d)\tau^Y(x) - d\mu_0^Y(x)$
Free flat-rate shipping or dollar-off discount	k	$\tau^Y(x) - k\tau^C(x) - k\mu_0^C(x)$
Percentage discount and free shipping	d, k	$(1-d)\tau^Y(x) - k\tau^C(x) - d\mu_0^Y(x) - k\mu_0^C(x)$
Percentage discount with universal free shipping	d, c	$(1-d)\tau^Y(x) - c\tau^C(x) - d\mu_0^Y(x)$

4 Experimentation & estimation framework

While we have established the relevant theoretical foundations for the optimal targeting of discounts in the online retail environment, we have yet to describe how firms can implement this strategy in a feasible way. In this section, we outline a framework that allows for efficient estimation of optimal discount targeting strategies using data from randomized experiments. We elaborate on some of the practical details of this methodology below, but here we lay out our targeting framework at a high level. It consists of three primary phases:

1. *Experimentation:* The firm will choose discount structure (d,k) and run an A/B test in which a randomized subset of users are assigned to the discount treatment condition. In the process, they will gather data on targetable customer features X_i, individual revenue and conversion responses, R_i and C_i, and treatment assignments T_i.

2. *Estimation:* Using the realized experimental data $\mathcal{D} = \{r_i, c_i, x_i, t_i\}$ gathered in the first phase, the firm can use machine learning techniques to estimate the conditional response and treatment effect functions ($\hat{\mu}_0^R, \hat{\mu}_0^C, \hat{\tau}^R$ and $\hat{\tau}^C$); estimation of these functions is discussed in detail below. Factoring in their relevant revenue and cost parameters, the firm can use these functions to estimate a targeting policy based on the optimal criteria derived in Eq. 2.

3. *Targeting:* For customers that arrive to their website moving forward, the firm observes their covariate x, evaluates the targeting criterion using the estimated quantities, $(1-d)\hat{\tau}^R(x) - (c+k)\hat{\tau}^C(x) - [d\hat{\mu}_0^R(x) + k\hat{\mu}_0^C(x)] > 0$, and offers a discount if the criterion is met.

4.1 Estimation of customer-level responses and treatment effects. For our theoretical model to be useful in practice, it must be the case that we are able to estimate conditional response and treatment effect functions with sufficiently high-fidelity at the individual customer level. At a conceptual level, the accuracy of our predictions will depend on two main factors: our prediction algorithm and the data provided to it. We focus on explaining our algorithm and estimation technique in this section and use Section 5 to study whether our approach is profitable with the technographic data commonly available in e-commerce environments.

Conditional response function estimation. We first describe our process for estimating the conditional response functions $\mu_t^R(x)$ and $\mu_t^C(x)$. Recall these functions are supposed to map a set of observable customer characteristics x to their expected response, conditional on treatment assignment. A key challenge in deriving an accurate estimate of expected revenues for each individual customer is an abundance of zero-revenue observations in most e-commerce environments. For example, if we assume that a firm has an overall conversion rate of 3%, then 97% of the observations in any dataset will have a revenue value of zero. To deal with this challenge, we implement a two-stage hurdle model for predicting customer revenue levels (Tu and Liu, 2014). Such models are inspired by the following identity, by which a positive count variable (such as revenue) can be decomposed into the product of two distinct quantities:

$$\mathbf{E}[R_i \mid x_i] = \Pr[R_i > 0 \mid x_i]\mathbf{E}[R_i \mid R_i > 0, x_i]$$

The first quantity, $\Pr[R_i > 0 \mid x_i]$, is the probability that a customer buys anything in a session; the second, $\mathbf{E}[R_i \mid R_i > 0, x_i]$, is the expected revenue observed, conditional on a conversion occurring. In our proposed technique, we estimate both of these quantities by first fitting a probabilistic classifier to predict whether or not a customer will convert; note this quantity is precisely the conditional response function for conversion μ_t^C. Then, a second continuous-valued predictor is fit to estimate revenues, but only on the subset of users that made a purchase. Once these models are fit, we obtain an unconditional estimate of expected revenue (μ_t^R) by merely multiplying the predictions of the two sub-models described above.[7]

While we only need estimates of conditional response functions under control (μ_0^R, μ_0^C) to evaluate the targeting criterion in Eq. (3), we will also need to estimate response functions under treatment (μ_1^R, μ_1^C) in subsequent calculations of heterogeneous treatment effects. As such, we estimate the conditional response functions for both control and treatment groups by training the aforementioned hurdle model separately on each subset of treatment and control data; this process is described in the pseudocode for "EstimateResponseFunctions" provided in Algorithm 1. In summary, by performing this approach to both control and treatment conditions, we are able to derive estimated predictors of conversions and revenues for both groups; we use hat notation to refer to these estimated quantities as $\hat{\mu}_0^C, \hat{\mu}_0^R, \hat{\mu}_1^C$, and $\hat{\mu}_1^R$.

As for the choice of learning algorithms used for the classification and regression tasks, our technique uses gradient-boosted decision trees (GBDT).[8] Gradient boosting models are

[7]By training our conversion model and revenue-conditional-on-conversion models separately, our technique treats these as independent quantities. While this approach would be problematic if our goal was inference on a set of regression parameters (Heckman, 1979), this is not our objective in this project. Rather we wish to maximize the predictive accuracy of our targeting algorithms; the evidence of whether or not this assumption of independence is appropriate is an empirical question and can be assessed by observing our model's performance in Section 5.

[8]While simpler methods, such as penalized logistic and linear regression, are also natural candidates to use,

frequently the best-performing algorithm in popular machine learning competitions and have shown strong performance in a large variety of contexts that rely on tabular data (Martinez and Gray, 2019, Nielsen, 2016). In addition to exhibiting high levels of accuracy in many different environments, GBDT models can be easily adapted for both classification and regression tasks by choosing suitable loss functions. As such, we are able to minimize the complexity of our estimation procedure by using similar optimization techniques and vocabulary across classification and regression tasks.

Conditional average treatment effect estimation. Estimating conditional average treatment effects (CATE) at the individual level is a notoriously difficult problem but one that has been facilitated in recent years by several advances in machine learning and causal inference. Many new techniques in this space have focused specifically on the problem of inference around CATE, which is often achieved using some form of repeated sample-splitting (Chernozhukov et al., 2018, Wager and Athey, 2017). In our case, because inference is not as important as prediction performance, we opt for a simpler technique that is commonly referred to as the "T-learner" for heterogeneous treatment effects. Because we assume treatment assignments will have been exogenously randomized in our experimental data, we can recover asymptotically unbiased estimates of individual-level treatment effects by merely predicting a user's response for both treatment and control conditions and taking the difference (Künzel et al., 2019).[9] In concrete terms, for the case of predicting treatment effects on customer conversion, given the estimates of $\hat{\mu}_0^C$ and $\hat{\mu}_1^C$ described above, we calculate the expected treatment effect for a user with covariate x as $\hat{\tau}^C(x) := \hat{\mu}_1^C(x) - \hat{\mu}_0^C(x)$. Estimating treatment effects on revenue is done similarly: $\hat{\tau}^R(x) := \hat{\mu}_1^R(x) - \hat{\mu}_0^R(x)$. Pseudocode for this process is given in the "EstimateCATE" function in Algorithm 1. Once the response and treatment effect functions have been estimated, we can now evaluate the decision criteria in Eq. (3) to determine whether a new user that arrives at the website should be offered a targeted discount. Building on the estimation processes introduced above, the "EstimateTargetingFunction" code provided in Algorithm 1 describes how the optimal tar-

they consistently demonstrated worse performance in our empirical findings. Further, with modern machine learning software, there is little penalty for using more complex models, so long as they are appropriately tuned. In our project, we use the LightGBM library as the core of our gradient boosting techniques (Ke et al., 2017).

[9] We have implemented versions of our framework using both the methods of Wager and Athey (2017) and Chernozhukov et al. (2018); even with considerable tuning, both were found to perform consistently worse than the T-learning procedure described here. This is because the sample-splitting procedures at the heart of these techniques reduce our effective sample size.

geting function for an arbitrary discount campaign can be derived from raw experimental data.

Hyperparameter optimization. All machine learning algorithms have a number of hyperparameters that must be exogenously specified prior to training that significantly affect model performance. In the case of gradient boosting, important hyperparameters include the number of boosting rounds, the learning rate, and decision tree termination criteria. (See Table 3 for the full list of parameters we use in our procedure and the range of values in our search space.) For standard supervised learning, such parameters are often set by analysts using cross-validation to estimate prediction accuracy (e.g., ROC AUC in the case of classification; squared loss or log-likelihood in the case of regression). While such techniques are suitable for straightforward prediction tasks, they break down in the case of heterogeneous treatment effect estimation. This is because, in most experimental data, a single individual is only exposed to one treatment condition. This makes it is impossible to observe the "true" value of their counterfactual response to treatment, precluding the possibility of minimizing a loss function between predicted and observed data. While several methods exist for getting around this limitation, all techniques in existing literature are focused singularly on maximizing the accuracy of treatment effect estimators (Nie and Wager, 2017, Rolling and Yang, 2014, Schuler et al., 2018).

However, our ultimate economic objective in this project is not to minimize the statistical error of our treatment effect estimators but to maximize firm profits by estimating a targeting policy. We determined earlier that the theoretically optimal targeting policy depends on estimating both response functions and treatment effect functions. As such, there is no guarantee that selecting the best model for either task individually will result in the most profitable outcomes for the task as a whole. In light of these challenges, we employ a tailor-made hyperparameter optimization technique that is designed to maximize expected profits directly. This approach is adaptive to campaigns with varying cost and discount structures and, as such, can be used as an off-the-shelf model selection technique in many different scenarios.

To motivate our technique, note the task of estimating the expected profits from a given targeting policy maps directly on to the problem of off-policy evaluation in the literature on reinforcement learning (Sutton and Barto, 2018). If a user was assigned one condition in our experimental data and a targeting policy would have assigned them to the opposite condition, we cannot observe their counterfactual response and thus must impute this value in a reliable way to estimate the profitability of the proposed policy. However, there are multiple methods for inferring the value of a counterfactual policy from offline records. One approach, known as the *direct method* relies on simply imputing the value of a counterfactual policy on the whole sample by using the subset of observations for which the proposed intervention matches the intervention assigned in the offline data (Li et al., 2012). Another solution is *inverse probability weighting* (IPW), which is known to provide an unbiased estimation of off-policy rewards (Horvitz and Thompson, 1952). While both techniques have been used in the literature on targeted marketing (Lemmens and Gupta, 2020, Yoganarasimhan et al., 2020), they are known to offer poor performance when the counterfactual policy differs substantially from the policy used for data collection.

An innovation to improve on these methods is *doubly robust* (DR) off-policy evaluation, which requires either that we have a good model of the relationship between covariates and outcomes *or* a good model of the data collection policy (Dudík et al., 2014). In our case, where we are using data from random experiments—for which we know exactly the data collection policy—we do not expect the results of these methods to differ significantly, though the DR method can still offer efficiency gains since it has been shown to exhibit both lower bias and lower variance than the direct or IPW methods (Jiang and Li, 2016). Note that the DR approach we employ here has also been recently applied successfully in the domain of targeted marketing by Yang et al. (2020).[10]

Before we can express the formula for the DR estimator, we define a predictive model of how firm profits depend on both treatment assignment and user covariates; let $V(t,x)$ represent such a model. This can be estimated using standard supervised learning techniques from our experimental data by adding the observed treatment assignment t_i as a predictor to the observed features x_i and using this data to predict the observed profit π_i in each user's session. Adapting the technique of Dudík et al. (2014) to our context, the DR estimator takes the form:

$$\hat{\pi}_f^{DR} = \frac{1}{N}\sum_{i=1}^{N}\left(V\left(\delta_f(x_i);x_i\right) + \sum_{t=0,1}\frac{\mathbf{1}\left\{\delta_f(x_i)=t\right\}}{\Pr[t_i=t]}\left(\pi_i - V(t_i;x_i)\right)\right)$$

With the ability to estimate out-of-sample profits for counterfactual targeting policies

[10]Also note that this technique requires significantly fewer assumptions than model-based, simulation approaches that have been used in past literature, which assume their parametric models accurately capture customer behavior (Khan et al., 2009). So long as we know the treatment assignment policy that was used for the observations in our dataset, these techniques are able to provide an unbiased counterfactual assessment of any proposed targeting policy.

using the formula above, we can maximize this quantity as our objective function within a cross-validation framework for parameter tuning. We use an adaptive grid search over the space of hyperparameters and select the combination of parameters that maximizes cross-validated expected profits (Bergstra et al., 2013); refer to Table 3 for details on our parameter search space. Pseudocode for this entire procedure is described in the appendix in Algorithm 2.

Table 3: Hyperparameters used in gradient boosting decision tree models

Parameter name	Distribution used in search space	Quantization interval
Number of boosting rounds	LogUniform(50, 500)	25
Learning rate	LogUniform(0.01, 0.2)	
Maximum number of leaves	Uniform(30, 150)	10
Minimum samples in leaf	Uniform(20, 500)	20
L1 regularization	Uniform(0, 1)	
L2 regularization	Uniform(0, 1)	

Note: The two numbers in each distribution represent the lower and upper limits of the search space used. The presence of a quantization interval indicates the distribution was only sampled at numbers divisible by the interval. The same hyperparameter space is used for both classification and regression tasks.

5 Empirical applications with A/B test data

Up to this point, we have derived the theoretically optimal personalized discount policy and described how a firm might use this result, in combination with experimental data, to optimize a marketing campaign of targeted discounts. However, we have yet to show that our findings have value in real-world settings where there are many reasons our theoretically-optimal strategy might fail. For example, the common sample sizes used in A/B tests and the limited number of features that are observable within technographic trace data can make it difficult to estimate the individual-level response and treatment effect functions required for optimal targeting. If the estimates of $\hat{\mu_0}(x)$ and $\hat{\tau}(x)$ are too noisy, it may be more profitable to fall back on simpler targeting rules that don't require such fine-grained distinction between customers on multiple dimensions. As such, it is important to study our proposed strategy in an empirical setting with practical limitations common in real-world e-commerce environments. We use the remainder of this paper to address this topic.

5.1 Empirical context and dataset. To empirically evaluate our method, we use experimental data from two separate US-based e-commerce firms. The data from Experiment 1 comes from a retailer of women's beauty products; the data from Experiment 2 comes from a novelty apparel company. Summary data from each dataset are provided in Table 4. In each experiment, a discount treatment was randomly assigned to a subset of the website's visitors; for such users, a discount was advertised on the homepage and with a persistent banner across the header of each firm's website. The nature of the discount differed between firms, allowing us to study whether our generalized discounting framework adds value under different discount and cost structures.

Table 4: Summary characteristics for experimental data

Variable	Experiment 1	Experiment 2
Number of sessions	87,675	59,353
Conversion rate	1.4%	4.9%
Average revenue per user	\$0.79	\$2.15
Average revenue per conversion	\$54.09	\$44.82
Proportion randomized to treatment	50%	90%
Average effect size on conversion	+0.08%	+1.0%
Time period	Q1 2014	Q1 2015
Length of experiment	17 days	96 days

For each user in both experiments, we observe their conversion responses, treatment assignments, and a set of technographic characteristics that are commonly accessible to most web servers. As our data has been provided by a collaborating experimentation platform, we did not have explicit control over the variables collected in the experimentation process. As such, we do not have access to all the variables shown in Table 1, but rather a subset that includes many of the most common attributes regularly collected by standard web analytics software. This includes a user's operating system, web browser, screen dimensions, referral source, and—when the user arrived through a search engine that appends this data to the referring URL—search query information. The firm also observed each user's IP address which they map to a user's approximate GPS coordinates using a geolocation service; we also use this data to match each user to a Nielsen designated market area (DMA). The only variables in our dataset that require a client-side tracking script to collect are those related to screen size; all others can be directly inferred from standard metadata provided by internet communication protocols. In principle, a marketer with direct access to web log data would be able to observe the raw trace data with slightly higher fidelity, meaning our analysis should be considered as a lower bound on the informational content contained in these features.[11]

[11]For example, our data does not include the raw `User-Agent` header for each user's session; this can be used to extract information about exact version of operation system and web browser being used by the client.

We provide summary statistics of the features available in our dataset in Table 5. For numeric variables, we report the sample mean μ and standard deviation σ. For categorical variables, we instead report the number of categories K. We also report the Gini coefficient, G, based on the count data for observations across categories within a given variable. Though historically used for evaluating inequality in macro-economic data, the Gini index is a useful summary statistic for characterizing how skew the distribution of counts is within a categorical variable. In this case, a higher Gini coefficient corresponds to a categorical distribution that has more of its observations clumped within a small number categories; lower Gini coefficients indicate there is a more equal spread of observations across categories.

As can be seen in Table 5, many of these features are categorical; even after removing categories with only one observation, our raw data matrix is very high-dimensional. Given the success of single value decomposition (SVD) as a dimensionality-reduction technique in other supervised learning tasks with high-dimensional data, we employ SVD to preprocess our covariate matrix in this application (Sarwar et al., 2000, Wall et al., 2003). In particular, we approximate the categorical features in our data with a truncated SVD of rank 10 (Hansen, 1987). In both experiments, this process reduces the dimension of our data matrices to 15 features.

Table 5: Variables available in dataset for promotional targeting

Variable	Type	Experiment 1		Experiment 2	
Device					
Operating system	Categorical	$K = 22$	$G = 0.81$	$K = 18$	$G = 0.78$
Browser	Categorical	$K = 12$	$G = 0.72$	$K = 15$	$G = 0.75$
Screen height	Numeric	$\mu = 1277.2$	$\sigma = 346.9$	$\mu = 791.4$	$\sigma = 234.2$
Screen width	Numeric	$\mu = 889.7$	$\sigma = 144.6$	$\mu = 871.7$	$\sigma = 553.7$
Behavioral					
HTTP referrer	Categorical	$K = 817$	$G = 0.97$	$K = 327$	$G = 0.98$
Search term	Categorical	$K = 3{,}189$	$G = 0.96$	$K = 2,883$	$G = 0.95$
Existing cookie	Binary	$\mu = 0.32$	—	$\mu = 0.19$	—
Geographic					
DMA	Categorical	$K = 210$	$G = 0.75$	$K = 211$	$G = 0.70$
Latitude	Numeric	$\mu = 37.8$	$\sigma = 4.8$	$\mu = 37.6$	$\sigma = 5.0$
Longitude	Numeric	$\mu = \text{-}95.2$	$\sigma = 19.2$	$\mu = -91.7$	$\sigma = 16.9$

Notes: For numeric variables, we report the sample mean μ and standard deviation σ. For categorical variables, we report the number of categories K and the Gini coefficient G, based on the count data for observations in each category.

5.2 Derivation of optimal policies. We now demonstrate how the differing nature of each firm's discount campaign can be accommodated by our model. In Experiment 1, the discount was for free shipping, with no assumed marginal costs incurred by the firm. Because we do not observe the firm's shipping costs, we assume a flat cost of $7.50 to the firm when a user makes a purchase in the treatment condition.[12] Under this assumption, the free shipping promotion acts as a level discount in the amount of $k = 7.5$; plugging this into the optimal targeting function in Eq. (3), assuming marginal costs $c = 0$ and setting $d = 0$, yields the following expression for the firm's optimal targeting criteria for this campaign:

$$\hat{f}(x) = \hat{\tau}^R(x) - 7.5\hat{\tau}^C(x) - 7.5\hat{\mu}_0^C(x)$$

In Experiment 2, the discount was 20% off the retail price of a user's entire order; at the same time, the firm was offering free shipping on all orders, whether or not the user was in the promotional discount condition. To derive the optimal targeting criteria in this scenario, the level discount parameter will be $k = 0$, and the percentage discount parameter will be $d = 0.20$. If we assume the firm faces a flat shipping cost of $7.50 per order, the fact that they offer free shipping to customers in both control and treatment conditions implies they face an effective marginal cost of $c = 7.5$ on every order. Plugging these values into Eq. (3) yields the targeting criteria:

$$\hat{f}(x) = 0.80\hat{\tau}^R(x) - 7.5\hat{\tau}^C(x) - 0.20\hat{\mu}_0^R(x)$$

5.3 Alternative targeting policies. Before moving on to our empirical findings regarding the performance of our proposed targeting policy, it will be instructive to identify other reasonable policies a firm might use in its place. As a starting place, it makes sense to consider a *non-targeted* (or uniform) policy. The profit-maximizing choice for such a policy will be identical to the optimal policy found in Section 3 but with the individual level estimates of τ and μ replaced by their average values. Using an overline to denote sample means (e.g., $\overline{\tau}^R = \mathbf{E}_x[\tau^R(x)]$), the targeting function for this policy can be written as:

$$f(x) = (1-d)\overline{\tau}^R - (c+k)\overline{\tau}^C - [d\overline{\mu}_0^R + k\overline{\mu}_0^C]$$

This approach treats all users the same (i.e., either assigns all users to the control condition or all users to the treatment condition), but does so in a profit-maximizing way that

[12]This value is based on standard shipping costs observed on the web and is consistent with shipping costs according to e-commerce merchants we have spoken with.

factors in the cost and discount parameters. Additionally, we consider a well-known policy that has been mentioned many times in the literature on targeted marketing and uplift modeling (Lo, 2002, Rzepakowski and Jaroszewicz, 2012), which is to target all customers with positive treatment effect on conversions:

$$f(x) = \tau^C(x)$$

This strategy, which we will refer to as the "uplift" approach, will serve as a useful baseline for considering the value of our decision-theoretic approach relative to existing benchmarks used for customer targeting.

5.4 Evaluation & Empirical Results. Recall that in Section 4.1, we described how to estimate the expected profits of a counterfactual targeting policy from experimental data. We use this same approach for evaluating the profitability of our proposed technique, but with an additional, nested level of cross-validation to honestly assess our method's out-of-sample performance. In particular, we used 100 iterations of Monte Carlo cross-validation, in which for each iteration, two-thirds of our data are used for both optimizing hyperparameters and training the models needed to estimate a policy's target function $\hat{f}$; the remaining one-third of the data is used to estimate the policy's expected profits on out-of-sample data using the doubly-robust technique. Our primary outcome metric is the value of expected profits $\mathbf{E}[\hat{\pi}_f]$ given by the DR estimator, averaged across all 100 iterations. Histograms of the distributions of the profits observed across these iterations are plotted in Figure 1.

To facilitate a comparison between our technique and the uniform approach, we also compute the within-fold "lift" of our targeting policy over the non-targeted baseline. If $\hat{\pi}_0$ are the profits observed from the uniform policy, the level lift of a policy f, $\hat{\Delta}_f$, is given as $\hat{\Delta}_f := \hat{\pi}_f - \hat{\pi}_0$. We also report the lift in percentage terms by comparing the average gains to the average baseline profit values: $\hat{\Delta}_f^{\%} := (\hat{\pi}_f - \hat{\pi}_0)/\hat{\pi}_0 \times 100$. Results of these metrics averaged across Monte Carlo iterations for each of the aforementioned targeting policies are summarized in Table 6. For relevant comparisons, we report p-values using a standard t-test, calculated against the null hypothesis that the mean lift for each policy is exactly zero.

We remark on several aspects of our empirical results. First, we observe that the uplift approach, which offers discounts to all users with a positive estimated treatment effect

Figure 1: Estimated profit per user by policy, estimated over 100 cross-validation splits

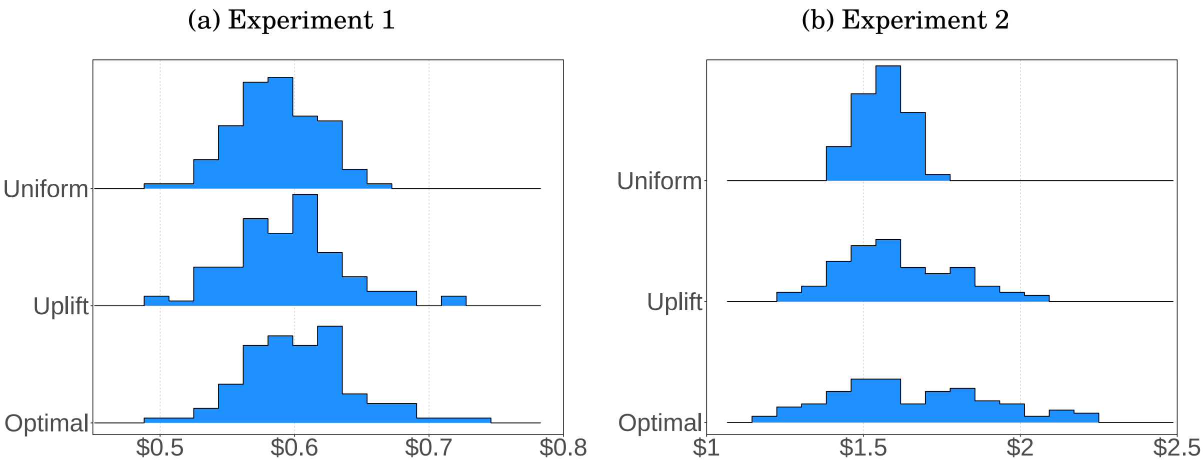

Table 6: Empirical results for proposed targeting policies

Experiment 1

Policy name	% receiving discount	Profits $\mathbf{E}[\hat{\pi}_f]$	Lift $\mathbf{E}[\hat{\Delta}_f]$	% Lift $\mathbf{E}[\hat{\Delta}_f^{\%}]$
Uniform	0.0	0.586 (0.003)	—	—
Uplift	57%	0.594 (0.004)	+0.007 (0.003)**	+1.39%
Optimal	52%	0.604 (0.004)	+0.018 (0.003)***	+3.16%

Experiment 2

Policy name	% receiving discount	$\mathbf{E}[\hat{\pi}_f]$	Lift $\mathbf{E}[\hat{\Delta}_f]$	% Lift $\mathbf{E}[\hat{\Delta}_f^{\%}]$
Uniform	0.0	1.56 (0.008)	—	—
Uplift	82%	1.60 (0.018)	+0.046 (0.018)**	+3.17%
Optimal	77%	1.65 (0.025)	+0.092 (0.025)***	+6.04%

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $^{+}p < 0.1$. Standard error of the mean for each value is reported in parentheses.

independent of the discount rate, is a profitable targeting strategy in both experiments, resulting in a profit gain of +1.39% in Experiment 1 and +3.17% lift in Experiment 2. This need not be a universal characteristic, especially for campaigns with larger discounts for which this strategy may actually decrease profits. That being said, in this case, uplift remains a reasonable policy for targeting discount interventions.

Turning to the performance of our proposed targeting policy, we see an even greater increase in profitability. For Experiment 1, we estimate the use of the optimal policy for discount targeting will result in a profit gain of +3.16%. Though not reported in the results table, we find that this increase is significantly larger than the gain observed from the uplift policy ($+0.011, t = 9.45, p < 0.001$). If the firm were to implement this policy over a time period equivalent to the length of the experiment we observed (i.e., 17 days with 87,000+ total sessions), we estimate they would earn an additional $\$1,578$ over a strategy that uniformly implements the most profitable treatment arm. While modest, our targeting approach clearly has the potential to increase firm profits by non-trivial amounts over relevant time scales. In Experiment 2, we see the predicted gains of our approach are even larger. The optimal targeting policy is estimated to result in a +6.04% increase over the uniform baseline, which is near twice the gains observed from the uplift strategy. (The direct comparison in level gains between the uplift and optimal policies is also statistically significant at the 5% level; $+0.046, t = 2.21, p = 0.014$.) In dollar terms, we estimate that our policy, if applied to the 59,000+ users in this experiment, is estimated to be $\$5,443$ more profitable than a uniform policy; this is an expected gain of nearly $0.10 ($0.092) for every person that visits the website.

Overall, these empirical results indicate that technographic data can be profitably exploited for personalized price discrimination. Further, the estimated value of the data in this context is not far removed from the value of purchase history data that has been calculated in prior research on targeted price promotions. Rossi et al. (1996) estimated that, with access to purchase history data for the purposes of targeting grocery store coupons delivered by direct mail, a firm could earn an additional $0.15 per customer in their database. Similarly, Khan et al. (2009) estimated that transactional history data can increase firm profits by 7.8% in a targeted discount campaign delivered by email to customers of a brick-and-mortar drug store chain. In these cases, however, the firm was required to establish some sort of prior relationship with each targeted customer to obtain their physical or electronic mailing address. In the context studied in this project, we demonstrated that firms can use digital trace data on any customer that visits their website—independent of any prior purchasing behavior—to profitably target them with promotional incentives.

5.5 Quantifying feature importance. To better understand exactly which variables play the largest role in determining targeting outcomes in our application, we use *post-hoc* model explanation techniques from the literature on interpretable machine learning. While these techniques are typically used in the context of standard supervised learning, they can easily be adapted to better understand which features drive targeting decisions in our context. To do so, we can use the treatment assignments implied by our targeting policy as the main dependent variable to be explained and assess which predictor variables have the largest effect on whether or not customers are targeted with a promotion. This analysis can inform managers in deciding which type of data to retain for targeting purposes; it can also add empirical insight to the discussion on public policy and consumer privacy about what types of data in our context are most consequential for online personalization.

To quantify the relative contribution of different features for promotional targeting in our context, we will adapt a measure of global variable importance first introduced by Breiman (2001). The intuition behind this approach is to randomly permute the ordering of data points in each covariate column and measure over many repetitions how this noise reduces model accuracy. While often used as a feature importance measure specifically for techniques based on bagging such as random forests—for which these measures can be estimated with relative computational efficiency using out-of-bag data—there is no reason permutation-based importance measures cannot be used for arbitrary black-box models. Indeed, the approach we use here has been previously introduced as "model reliance" by Fisher et al. (2019), who applied it as a model-agnostic measure of feature importance.

To describe our importance measure concretely: for each fold in our cross-validation procedure, we fit our targeting algorithm, denoted $\delta_{\hat{f}}$, using data from the fold's training data. We then use this algorithm to predict the optimal treatment assignment for each observation in the fold's test data. Those treatment assignments are then used to estimate the counterfactual profits reported in the previous section. We then take the test data X and, for each variable j in our dataset, we randomly permute the rows of our data in the j-th column; denote the permuted test data $X^{(j)}$.[13] After each permutation, we calculate the accuracy with which the targeting algorithm based on the permuted test data $X^{(j)}$ can predict the treatment assignments of the unpermuted test data X. For each fold and for each feature j, we perform 100 permutations and calculate the importance of feature j by taking 1 minus the observed accuracy over each permutation and averaging across both

[13]There are some technical difficulties associated with this approach that are worth mentioning. Our raw data takes the form of weblogs associated with user website visits which have been formatted into a data frame containing the columns described in Table 4. However, before actually training any supervised learning algorithms on our data, this data frame must be processed in several different steps. First, we encode the categorical data using binary dummy variables for each category within a given variable, and then (as described in Section 5), we use a truncated SVD procedure to project the categorical variables into a 10-dimensional continuous vector space. Only after this pre-processing do we train the GBDT algorithms and estimate the targeting function (as described in Algorithm 1 of the appendix). When we say that we "permute" a feature in our dataset, this means we actually permute the user-level feature in the original data frame. To measure the effect of this permutation occurring very early in the ML pipeline, our targeting algorithm has to be engineered so that it can take a raw observation in the data frame, apply all the pre-processing steps, feed the resulting output into the GBDT algorithm to predict the user's response function, and finally derive each user's CATE and optimal treatment assignment. Fortunately, modern machine learning tooling—especially the `pipeline` functionality provided by the open-source `scikit-learn` project—has progressed to the point that makes a task of this complexity relatively painless (Pedregosa et al., 2011).

permutations and cross-validation folds (this technique is also sometimes referred to as *mean decrease in accuracy*).

A common measure of accuracy used for permutation importance scores is the area under the receiver operating characteristic curve (AUC); however, since we are dealing with purely binary outcomes (whether or not a user is targeted with a promotion), we are unable to use the standard AUC metric, which depends on a continuous, thresholded predictor function. Nonetheless, we can use an analogous measure for binary outcomes, which can be considered as "balanced accuracy"; this metric is similar to a standard "accuracy" score, except it gives equal weight to false positives and false negatives so that a completely uninformative predictor will always have a balanced accuracy of 0.5 (Mosley, 2013, Powers, 2020, Youden, 1950). We use this characteristic of the balanced accuracy score to normalize our importance measure so that a feature that fully determines targeting outcomes has a score of one, and a feature that has no effect on targeting outcomes has a score of zero.[14]

Moving on to the empirical findings from this procedure, we will review feature importance results for Experiment 1; the findings of this procedure are qualitatively similar to Experiment 2 and are omitted here. In Figure 2, we have plotted the importance values for each column in our dataset. (Variance of average importance values across folds is on the order of 0.001, and thus we have omitted standard error bars.) Looking at the importance values measured in this analysis, it appears the single most important feature for determining whether a user is targeted vs. not targeted with promotion is "screen width". When considered independently, a user's screen height is the second most important feature, followed by a user's designated market area (DMA). Note that no single variable has an importance score anywhere near 1; this indicates that our model relies on several of these variables simultaneously, and the performance of our targeting algorithm is not driven by any single user characteristic.

[14]The exact formulae we use are provided below, where—because we average importance across 100 cross-validation folds—expectation can be considered as being taken over both the randomness of the permutations of column j and the sampling variation due to cross-validation:

$$\mathbb{I}(j) = 2 \times \left(1 - \mathbf{E}\left[\text{BalancedAccuracy}\left(\delta_{\hat{f}}\left(X^{(j)}\right), \delta_{\hat{f}}(X)\right)\right]\right)$$

where

$$\text{BalancedAccuracy}\left(\delta_{\hat{f}}\left(X^{(j)}\right), \delta_{\hat{f}}(X)\right) = \frac{1}{2}\left(\frac{TP}{TP+FN} + \frac{TN}{TN+FP}\right)$$

In this formula, TP, TN, FP, and FN are true positives, true negatives, false positives, and false negatives respectively, which can be derived from the confusion matrix comparing the optimal treatment assignment for each user in the observed test data, $\delta_{\hat{f}}(X)$, to the treatment assignments derived from the permuted test data, $\delta_{\hat{f}}\left(X^{(j)}\right)$.

Figure 2: Permutation importance of individual features for targeting

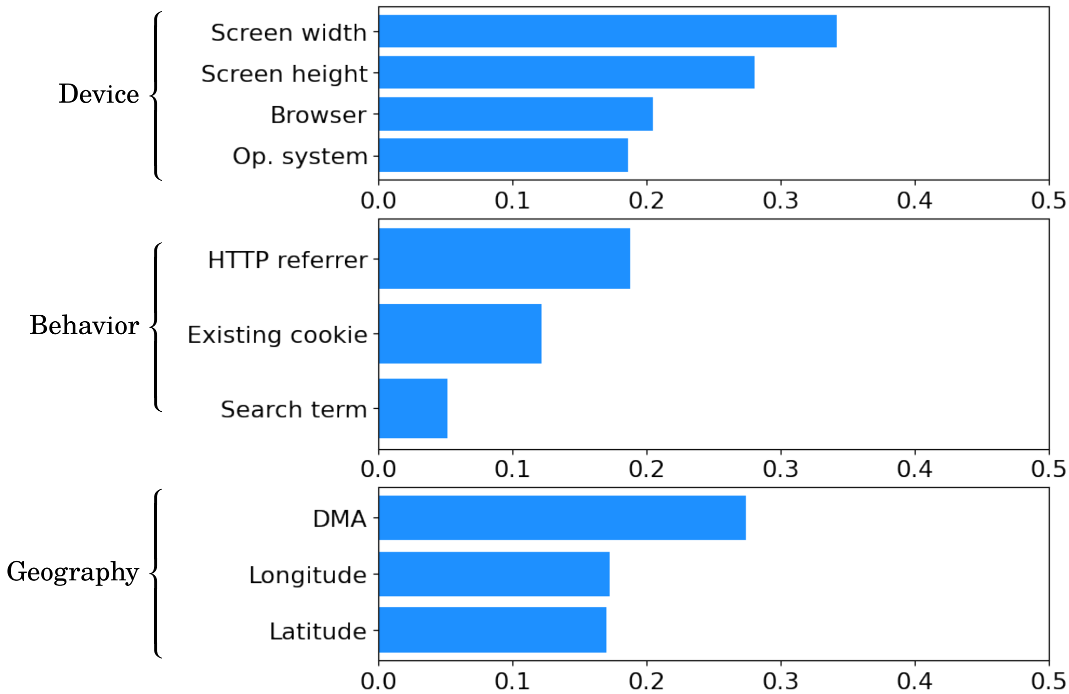

One common criticism of permutation-based feature importance measures is that, by independently permuting individual columns, we break the correlation structure between features, giving an inaccurate assessment of any individual feature's importance for prediction (Hooker and Mentch, 2019). This is particularly relevant in our context since, for example, the screen width and operating system variables are likely to be very correlated: users with small screen widths are likely to be using mobile operating systems (iOS, Android) as opposed to desktop operating systems (Windows, macOS). Permuting these columns independently, as done above, requires our model to make predictions in areas of the feature space that are unrealistic and not likely indicative of our algorithm's performance on real-world data.

One way we can minimize this concern is by also permuting sets of features together, and comparing the marginal benefit in prediction accuracy attributable to various groups of features. To avoid a combinatorial explosion and to give us a more interpretable understanding of which factors affect model performance, we employ a *grouped* variable importance measure, that permutes multiple columns in tandem (Gregorutti et al., 2015). In particular, for each of the three types of data in our feature set—device-related, behavioral, and geographic (corresponding to the three subplots in Firgure 2 and also listed in Table

4)—we can permute all the columns within each group *together*. To do this, we employ the same permutation procedure described above for the individual columns, except in this case we will permute subsets of multiple columns simultaneously. For example, to assess the importance of *device*-related columns (screen height, screen width, browser, and operating system), we take our dataset X and permute the rows in all four of these columns to obtain a permuted dataset $X^{(device)}$. Importantly, we do not permute each column independently, but rather generate one row-wise permutation that applies across all columns in a given group. This preserves the correlation structure within each column group, which addresses a criticism of column-wide permutation procedures that depend on model performance in regions of the feature space that are unrealistic.

In our grouped variable importance analysis reported in Figure 3, we report the importance scores computed for each of the three groups of variables separately, but also the importance scores observed when permuting multiple groups at the same time. For example, the importance score of "Device + Behavior" is computed by permuting all seven of the device-related and behavioral columns in our dataset simultaneously. Combining the groups of variables in this way allows us to better understand the effect of each type of variable for driving model performance.

Figure 3: Permutation importance of grouped features for targeting

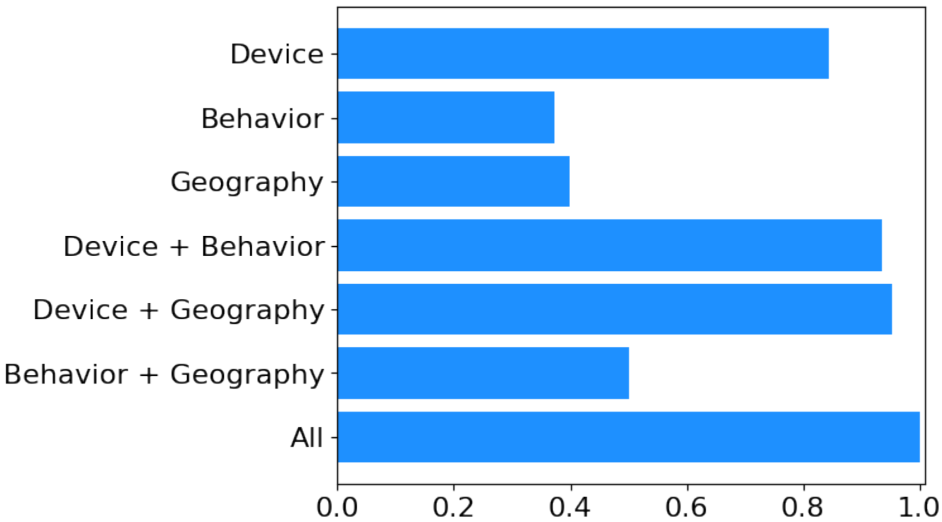

Examining the importance scores, we can see several patterns worth remarking on. First, it appears the most important group of variables in our context are those related to the *device* that users use to access the retailer's website. With an importance score of greater than 0.80, this set of variables is clearly driving most of the targeting algorithm's decisions. Importantly, however, note that the combination of all device-related variables explains much more of the model's decisions than any of the variables independently reported in Figure 2, where the single most important variable (screen width) has an importance score near 0.35. Again, this highlights how—even within each group—there is some signal being added by separate variables included in our model. Considering the other types of data, it appears the behavioral and geographical variables have importance scores near 0.38 and 0.40 respectively. On their own, these types of variables explain some, but clearly not all the variation in targeting decisions. Even when combined (in the "Behavior + Geography" entry), the grouped importance score of these variables only reaches 0.49. The importance scores reported for the "Device + Behavior" and "Device + Geography" entries are quite similar, reaching 0.92 and 0.96, respectively. Thus, even though *most* of the targeting algorithm's decisions are driven by device-related variables, taking away any subset of variables has a measurable effect on the importance score, suggesting that each variable is contributing some non-zero amount to model performance.

These findings are useful for better understanding the black-box nature of the highly non-parametric machine learning algorithm we used for targeting in this application. Additionally, this analysis highlights that not all types of data are created equal, especially when it comes to price discrimination online. This can be important for e-commerce managers, who can use these findings to prioritize the collection or preservation of certain types of data for targeting purposes (in this case, it appears device-related data is most valuable). Our findings may also be of interest to policymakers, regulators, and developers of internet protocols when considering what types of data can and should be shared by default online. As of now, all the variables used in the project were derived from web log data that is automatically transmitted by HTTP protocols and automatically stored by many web servers. Given ongoing discussions in the governing bodies of the European Union and United States about user privacy and algorithmic regulation, these findings may be of interest for demonstrating how different types of data can be used for online personalization. This project may inform the choices that regulators and developers make when designing the next generation of legal statutes and communication protocols. For example, the logging of users' IP addresses is commonly singled out as problematic practice from a privacy perspective; however, our findings suggest that IP addresses (from which we derived our geographical variables) are not the most important type data for promotional targeting purposes. Instead, data that may seem more innocuous—such as screen width and device type, that may reveal little about a person's individual location or identity—may actually be the most important for enabling price discrimination online (at least among variables contained in technographic trace data). We believe this highlights an important disconnect between commonly understood definitions of "privacy" (which often revolve around a person's identity) and the many other factors that lawmakers might be wise to consider when regulating the use of data in online settings.

6 Conclusion

In this project, we set out to study the potential for price discrimination in online retail settings, based solely on user-level technographic trace data. To this end, we developed a framework for using field experiments and machine learning to optimize e-commerce customer targeting strategies in the presence of discount and cost parameters. We found that, whenever discounts are part of a promotional intervention, the optimal targeting strategy depends on individual-level estimates of both consumers' baseline response rates and treatment effects in response to the promotional price. Further, we laid out a procedure for using machine learning to estimate these parameters from experimental data and applied our techniques to real-world data. Using counterfactual policy evaluation techniques on A/B test data from two separate firms, we found that our proposed targeting strategy significantly outperforms both non-targeted baselines and industry standard techniques for customer targeting. Given that our targeting policy is based on non-parametric machine learning methods, we adapted techniques from the interpretable ML literature to better understand what is driving the performance of our targeting algorithm. An analysis based on feature permutation revealed that, while the behavioral and geographic variables inferred from technographic data have a role to play, the most significant predictors of whether or not customers are targeted with a promotion are related to the type and size of the device customers use to access retailers' websites.

This paper adds important insight to the growing discussions around personalized pricing, with implications for managers, consumer advocates, and policymakers. For managers, our empirical results can be used as a ballpark estimate of the profit gains possible using this technique. Further, firms can use the methods described in this project to estimate the profitability of targeting various discount campaigns by merely running pilot experiments and applying our estimation procedure. We caution that, prior to implementation, managers would be wise to consider the potential reputational risks associated with targeted price discrimination. Even though promotional targeting is far from new, awareness of and consumer attitudes about the practice may be changing rapidly. As one proactive measure, prior research has shown that consumers are more willing to purchase from firms that are transparent in their privacy policies (Tsai et al., 2011), indicating that openly messaging how online services are personalized can mitigate this downside risk.

For consumer advocates and policymakers, this research highlights the economic value present in trace forms of technographic data transmitted by standard web protocols, and demonstrates that such data can be exploited for algorithmic price discrimination. While prior work has documented the use of this data for customer targeting, this paper is the first to establish the profitability and measure the value of technographic data in a standard e-commerce setting. We highlight that this analysis is based on data from websites that are not large tech companies or national brands, but rather small to medium-sized businesses. Our findings suggest the practice of using technographic data for promotional targeting—including for price-related interventions—may become more commonplace as firms of this type develop the technical expertise required to do so. Our research also highlights how even forms of information not obviously subject to statutes that apply to "personal data" may be used for online personalization. As such, we hope this work contributes to the broader discussion about promotional targeting on the web and serves as a valuable empirical analysis of the potential of personalization technologies in online retail.

APPENDIX

Algorithm 1: Pseudocode for deriving optimal targeting policy from experimental data

Input: $\mathcal{D} = \{r_i, c_i, x_i, t_i\}$, dataset of revenue, conversion, covariates, and treatment assignment variables observed in experiment;
ξ, set of hyperparameters used for GBDT models;
Output: $\hat{f}$, estimated targeting function associated with optimal policy

function EstimateResponseFunctions$(\mathcal{D}, \xi)$
  for *treatment condition* $t \in \{0, 1\}$ **do**
    $\hat{\mu}_t^C \leftarrow$ train GBDT classifier using data for which $t_i = t$, using hyperparameters ξ
    $\hat{\mu}_t^{R>0} \leftarrow$ train GBDT regressor using data for which $t_i = t$ and $r_i > 0$, using hyperparameters ξ
    $\hat{\mu}_t^R \leftarrow \hat{\mu}_t^C \times \hat{\mu}_t^{R>0}$, hurdle model for unconditional revenue distribution
  return estimated predictors $\hat{\mu}_0^C, \hat{\mu}_0^R, \hat{\mu}_1^C, \hat{\mu}_1^R$

function EstimateCATE$(\hat{\mu}_0^C, \hat{\mu}_0^R, \hat{\mu}_1^C, \hat{\mu}_1^R)$
  $\hat{\tau}^C \leftarrow \hat{\mu}_1^C - \hat{\mu}_0^C$
  $\hat{\tau}^R \leftarrow \hat{\mu}_1^R - \hat{\mu}_0^R$
  return estimated CATE functions $\hat{\tau}^C, \hat{\tau}^R$

function EstimateTargetingFunction$(\mathcal{D}, \xi)$
  $\hat{\mu}_0^C, \hat{\mu}_0^R, \hat{\mu}_1^C, \hat{\mu}_1^R \leftarrow$ EstimateResponseFunctions$(\mathcal{D}, \xi)$
  $\hat{\tau}^C, \hat{\tau}^R \leftarrow$ EstimateCATE$(\hat{\mu}_0^C, \hat{\mu}_0^R, \hat{\mu}_1^C, \hat{\mu}_1^R)$
  $\hat{f} \leftarrow (1-d)\hat{\tau}^R - (c+k)\hat{\tau}^C - [d\hat{\mu}_0^R + k\hat{\mu}_0^C]$, plug estimated response and treatment effect function into theoretically optimal decision formula
  return estimated targeting function $\hat{f}$

Algorithm 2: Hyperparameter optimization pseudocode

Input: $\mathcal{D} = \{r_i, c_i, x_i, t_i\}$, dataset of revenue, conversion, covariates, and treatment assignment variables observed in experiment;
$\text{Grid}(\xi)$, grid values over distribution of hyperparameters (see Table 1);
I, number of grid values to search (set to 100 in our applications);
K, number of folds used for cross-validation (set to 10 in our applications);
Output: ξ^*, optimal set of hyperparameters to use for discount policy
Notation:
d, k, c, discount and marginal cost parameters;
$\mathcal{D}_k$, subset of data in fold k;
$\mathcal{D}_{-k}$, subset of data excluding fold k;

function EstimateHyperparameters$(\mathcal{D}, \text{Grid}(\xi), I, K)$
  for *iteration* $i \in [I]$ **do**
    $\xi_i \leftarrow$ random sample of parameters from Grid(ξ)
    $\{\mathcal{D}_k\} \leftarrow$ randomly divide $\mathcal{D}$ into K-folds   stratified by treatment condition t_i
    for *fold* $k \in [K]$ **do**
      $\hat{f} \leftarrow$ EstimateTargetingFunction($\mathcal{D}_{-k}, \xi$)
      $\hat{\pi}_k \leftarrow \frac{1}{N}\sum_{i=1}^{N}\left(V\left(\delta_f(x_i); x_i\right) + \sum_{t=0,1}\frac{\mathbf{1}\{\delta_f(x_i)=t\}}{\Pr[t_i=t]}(\pi_i - V(t_i; x_i))\right)$
    $\overline{\pi}(\xi_i) \leftarrow \frac{1}{K}\sum_k \hat{\pi}_k$ average estimated counterfactual profits across folds
  $\xi^* \leftarrow \arg\max_{\xi_i} \overline{\pi}(\xi_i)$ returns hyperparameters with largest observed profits

References

Ahmad, S. N. and Callow, M. (2018). "Free Shipping" or "Dollar Off"? The Moderating Effects of List Price and E-Shopping Experience On Consumer Preference For Online Discount. *"International Journal of Electronic Commerce Studies"*, 9(1):55–70.

Ali, S. N., Lewis, G., and Vasserman, S. (2020). Voluntary disclosure and personalized pricing. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 537–538.

Alreck, P. L. and Settle, R. B. (2007). Consumer reactions to online behavioural tracking and targeting. *Journal of Database Marketing & Customer Strategy Management*, 15(1):11–23.

Ansari, A. and Mela, C. F. (2003). E-customization. *Journal of Marketing Research*, 40(2):131–145.

Ascarza, E. (2018). Retention futility: Targeting high-risk customers might be ineffective. *Journal of Marketing Research*, 55(1):80–98.

Ascarza, E. and Hardie, B. G. (2013). A joint model of usage and churn in contractual settings. *Marketing Science*, 32(4):570–590.

Ayres, I. and Siegelman, P. (1995). Race and gender discrimination in bargaining for a new car. *American Economic Review*, pages 304–321.

Bawa, K. and Shoemaker, R. W. (1987). The coupon-prone consumer: some findings based on purchase behavior across product classes. *Journal of Marketing*, 51(4):99–110.

Bawa, K. and Shoemaker, R. W. (1989). Analyzing incremental sales from a direct mail coupon promotion. *Journal of Marketing*, 53(3):66–78.

Bergstra, J., Yamins, D., and Cox, D. (2013). Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures. In *International conference on machine learning*, pages 115–123. PMLR.

Blattberg, R. C., Briesch, R., and Fox, E. J. (1995). How promotions work. *Marketing Science*, 14(3_supplement):G122–G132.

Bolton, R. N. (1989). The relationship between market characteristics and promotional price elasticities. *Marketing Science*, 8(2):153–169.

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

Broeder, P. and Derksen, R. (2018). Exclusivity in online targeted promotions: cross-cultural preferences of consumers. *International Journal of Business and Emerging Markets*, 10(4):396–408.

Cao, Z., Hui, K.-L., and Xu, H. (2018). When discounts hurt sales: The case of daily-deal markets. *Information Systems Research*, 29(3):567–591.

Cardona, M., Duch-Brown, N., Francois, J., Martens, B., and Yang, F. (2015). The macro-economic impact of e-commerce in the EU digital single market. Technical report, Institute for Prospective Technological Studies Digital Economy Working Paper.

Chen, H., Marmorstein, H., Tsiros, M., and Rao, A. R. (2012). When more is less: The impact of base value neglect on consumer preferences for bonus packs over price discounts. *Journal of Marketing*, 76(4):64–77.

Chen, Y. and Sudhir, K. (2004). When shopbots meet emails: Implications for price competition on the Internet. *Quantitative Marketing and Economics*, 2(3):233–255.

Chen, Z., Choe, C., and Matsushima, N. (2020). Competitive Personalized Pricing. *Management Science*.

Chernozhukov, V., Demirer, M., Duflo, E., and Fernandez-Val, I. (2018). Generic machine learning inference on heterogenous treatment effects in randomized experiments. Technical report, National Bureau of Economic Research.

Choudhary, V., Ghose, A., Mukhopadhyay, T., and Rajan, U. (2005). Personalized pricing and quality differentiation. *Management Science*, 51(7):1120–1130.

Cui, G., Wong, M. L., and Lui, H.-K. (2006). Machine learning for direct marketing response models: Bayesian networks with evolutionary programming. *Management Science*, 52(4):597–612.

Dias, M. B., Locher, D., Li, M., El-Deredy, W., and Lisboa, P. J. (2008). The value of personalised recommender systems to e-business: a case study. In *Proceedings of the 2008 ACM conference on Recommender systems*, pages 291–294.

Drew, J. H., Mani, D., Betz, A. L., and Datta, P. (2001). Targeting customers with statistical and data-mining techniques. *Journal of Service Research*, 3(3):205–219.

Dubé, J.-P., Fang, Z., Fong, N., and Luo, X. (2017). Competitive price targeting with smartphone coupons. *Marketing Science*, 36(6):944–975.

Dubé, J.-P. and Misra, S. (2019). Personalized Pricing and Customer Welfare. *Working Paper*.

Dudík, M., Erhan, D., Langford, J., Li, L., et al. (2014). Doubly robust policy evaluation and optimization. *Statistical Science*, 29(4):485–511.

Englmaier, F., Gratz, L., and Reisinger, M. (2012). Price discrimination and fairness concerns. *Available at SSRN 1999757*.

Fisher, A., Rudin, C., and Dominici, F. (2019). All Models are Wrong, but Many are Useful: Learning a Variable's Importance by Studying an Entire Class of Prediction Models Simultaneously. *Journal of Machine Learning Research*, 20(177):1–81.

Gensch, D. H. (1984). Targeting the switchable industrial customer. *Marketing Science*, 3(1):41–54.

Ghose, A. and Huang, K.-W. (2009). Personalized pricing and quality customization. *Journal of Economics & Management Strategy*, 18(4):1095–1135.

Ghose, A., Li, B., and Liu, S. (2019). Mobile targeting using customer trajectory patterns. *Management Science*, 65(11):5027–5049.

Gregorutti, B., Michel, B., and Saint-Pierre, P. (2015). Grouped variable importance with random forests and application to multiple functional data analysis. *Computational Statistics & Data Analysis*, 90:15–35.

Grewal, D., Ailawadi, K. L., Gauri, D., Hall, K., Kopalle, P., and Robertson, J. R. (2011). Innovations in retail pricing and promotions. *Journal of Retailing*, 87:S43–S52.

Gupta, S. and Chintagunta, P. K. (1994). On using demographic variables to determine segment membership in logit mixture models. *Journal of Marketing Research*, 31(1):128–136.

Gutierrez, P. and Gérardy, J.-Y. (2017). Causal Inference and Uplift Modelling: A Review of the Literature. In Hardgrove, C., Dorard, L., Thompson, K., and Douetteau, F., editors, *Proceedings of The 3rd International Conference on Predictive Applications and APIs*, volume 67 of *Proceedings of Machine Learning Research*, pages 1–13, Microsoft NERD, Boston, USA. PMLR.

Hall, D. (2018). How to Legally Recover CRM Contacts After GDPR. `https://people.ai/blog/flushed-your-crm-contacts-down-the-gdpr-drain/`.

Hamermesh, D. S. (2013). Six decades of top economics publishing: Who and how? *Journal of Economic Literature*, 51(1):162–172.

Hannak, A., Soeller, G., Lazer, D., Mislove, A., and Wilson, C. (2014). Measuring price discrimination and steering on e-commerce web sites. In *Proceedings of the 2014 Conference on Internet Measurement Conference*, pages 305–318. ACM.

Hansen, P. C. (1987). The truncatedsvd as a method for regularization. *BIT Numerical Mathematics*, 27(4):534–553.

Heckman, J. J. (1979). Sample selection bias as a specification error. *Econometrica: Journal of the Econometric Society*, pages 153–161.

Hitsch, G. J. and Misra, S. (2018). Heterogeneous treatment effects and optimal targeting policy evaluation. *Available at SSRN 3111957*.

Hooker, G. and Mentch, L. (2019). Please stop permuting features: An explanation and alternatives. *arXiv preprint arXiv:1905.03151*.

Horvitz, D. G. and Thompson, D. J. (1952). A generalization of sampling without replacement from a finite universe. *Journal of the American Statistical Association*, 47(260):663–685.

Huang, J.-H., Chang, C.-T., and Chen, C. Y.-H. (2005). Perceived fairness of pricing on the Internet. *Journal of Economic Psychology*, 26(3):343–361.

Imai, K. and Strauss, A. (2011). Estimation of heterogeneous treatment effects from randomized experiments, with application to the optimal planning of the get-out-the-vote campaign. *Political Analysis*, 19(1):1–19.

Jiang, N. and Li, L. (2016). Doubly robust off-policy value evaluation for reinforcement learning. In *International Conference on Machine Learning*, pages 652–661. PMLR.

Johnson, J., Tellis, G. J., and Ip, E. H. (2013). To whom, when, and how much to discount? A constrained optimization of customized temporal discounts. *Journal of Retailing*, 89(4):361–373.

Kahneman, D., Knetsch, J. L., and Thaler, R. (1986). Fairness as a constraint on profit seeking: Entitlements in the market. *American Economic Review*, pages 728–741.

Katz, M. L. (1987). The welfare effects of third-degree price discrimination in intermediate good markets. *American Economic Review*, pages 154–167.

Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., and Liu, T.-Y. (2017). LightGBM: A highly efficient gradient boosting decision tree. In *Advances in Neural Information Processing Systems*, pages 3146–3154.

Khan, R., Lewis, M., and Singh, V. (2009). Dynamic customer management and the value of one-to-one marketing. *Marketing Science*, 28(6):1063–1079.

Kim, Y., Street, W. N., Russell, G. J., and Menczer, F. (2005). Customer targeting: A neural network approach guided by genetic algorithms. *Management Science*, 51(2):264–276.

Kobsa, A., Koenemann, J., and Pohl, W. (2001). Personalized hypermedia presentation techniques for improving online customer relationships. *Knowledge Engineering Review*, 16(2):111–156.

Künzel, S. R., Sekhon, J. S., Bickel, P. J., and Yu, B. (2019). Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the National Academy of Sciences*, 116(10):4156–4165.

Laperdrix, P., Bielova, N., Baudry, B., and Avoine, G. (2020). Browser fingerprinting: A survey. *ACM Transactions on the Web (TWEB)*, 14(2):1–33.

Larson, J., Mattu, S., and Angwin, J. (2015). Unintended Consequences of Geographic Targeting. *ProPublica*.

Lemmens, A. and Gupta, S. (2020). Managing churn to maximize profits. *Marketing Science, Forthcoming*.

Li, L., Chu, W., Langford, J., Moon, T., and Wang, X. (2012). An unbiased offline evaluation of contextual bandit algorithms with generalized linear models. In *Proceedings of the Workshop on On-line Trading of Exploration and Exploitation 2*, pages 19–36.

Lo, V. S. (2002). The true lift model: a novel data mining approach to response modeling in database marketing. *ACM SIGKDD Explorations Newsletter*, 4(2):78–86.

Luo, X., Andrews, M., Fang, Z., and Phang, C. W. (2014). Mobile targeting. *Management Science*, 60(7):1738–1756.

Martinez, W. and Gray, J. B. (2019). On the Insufficiency of the Large Margins Theory in Explaining the Performance of Ensemble Methods. *arXiv preprint arXiv:1906.04063*.

McFowland III, E., Somanchi, S., and Neill, D. B. (2018). Efficient discovery of heterogeneous treatment effects in randomized experiments via anomalous pattern detection. *arXiv preprint arXiv:1803.09159*.

Meyer, D. (2018). What the GDPR will mean for companies tracking location. *International Association of Privacy Professionals (iapp)*.

Miller, A. and Hosanagar, K. (2019). How Targeted Ads and Dynamic Pricing Can Perpetuate Bias. *Harvard Business Review. https://hbr. org/2019/11/how-targeted-ads-anddynamic-pricing-can-perpetuate-bias*.

Mosley, L. (2013). *A balanced approach to the multi-class imbalance problem*. PhD thesis, Iowa State University.

Musalem, A., Bradlow, E. T., and Raju, J. S. (2008). Who's got the coupon? Estimating consumer preferences and coupon usage from aggregate information. *Journal of Marketing Research*, 45(6):715–730.

Narasimhan, C. (1984). A price discrimination theory of coupons. *Marketing Science*, 3(2):128–147.

Neslin, S. A., Gupta, S., Kamakura, W., Lu, J., and Mason, C. H. (2006). Defection detection: Measuring and understanding the predictive accuracy of customer churn models. *Journal of Marketing Research*, 43(2):204–211.

Nie, X. and Wager, S. (2017). Learning objectives for treatment effect estimation. *arXiv preprint arXiv:1712.04912*.

Nielsen, D. (2016). Tree boosting with XGBoost: Why does XGBoost win "every" machine learning competition? Master's thesis, NTNU.

O'Neil, C. (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy*. Broadway Books.

Padmanabhan, B., Zheng, Z., and Kimbrough, S. O. (2001). Personalization from incomplete data: what you don't know can hurt. In *Proceedings of the seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 154–163.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. (2011). Scikit-learn: Machine learning in Python. *the Journal of machine Learning research*, 12:2825–2830.

Peppers, D. and Rogers, M. (1993). *The One to One Future: Building Relationships One Customer at a Time*. Doubleday Business.

Powers, D. M. (2020). Evaluation: from precision, recall and F-measure to ROC, informedness, markedness and correlation. *arXiv preprint arXiv:2010.16061*.

Radcliffe, N. (2007). Using control groups to target on predicted lift: Building and assessing uplift model. *Direct Marketing Analytics Journal*, pages 14–21.

Rafieian, O. and Yoganarasimhan, H. (2020). Targeting and privacy in mobile advertising. *Available at SSRN 3163806*.

Reibstein, D. J. and Traver, P. A. (1982). Factors affecting coupon redemption rates. *Journal of*

Marketing, 46(4):102–113.
Reid, A. (2017). The European court of justice case of Breyer. *Journal of Information Rights, Policy and Practice*, 1(2).
The Council of Economic Advisers (2015). Big Data and Differential Pricing. `https://www.whitehouse.gov/sites/default/files/docs/Big_Data_Report_Nonembargo_v2.pdf`. Executive Office of the President of the United States.
RetailMeNot (2018). RetailMeNot Survey: Deals and Promotional Offers Drive Incremental Purchases Online, Especially Among Millennial Buyers. `https://prn.to/3q08AA7`.
Rolling, C. A. and Yang, Y. (2014). Model selection for estimating treatment effects. *Journal of the Royal Statistical Society: Series B: Statistical Methodology*, pages 749–769.
Rossi, P. E., McCulloch, R. E., and Allenby, G. M. (1996). The value of purchase history data in target marketing. *Marketing Science*, 15(4):321–340.
Rzepakowski, P. and Jaroszewicz, S. (2012). Uplift modeling in direct marketing. *Journal of Telecommunications and Information Technology*, pages 43–50.
Sahni, N. (2015). Effect of temporal spacing between advertising exposures: evidence from online field experiments. *Quantitative Marketing and Economics*, 13(3):203–247.
Sarwar, B., Karypis, G., Konstan, J., and Riedl, J. (2000). Application of dimensionality reduction in recommender system-a case study. Technical report, Minnesota Univ Minneapolis Dept of Computer Science.
Schuler, A., Baiocchi, M., Tibshirani, R., and Shah, N. (2018). A comparison of methods for model selection when estimating individual treatment effects. *arXiv preprint arXiv:1804.05146*.
Shampanier, K., Mazar, N., and Ariely, D. (2007). Zero as a special price: The true value of free products. *Marketing Science*, 26(6):742–757.
Shapiro, C. and Varian, H. (1998). *Information rules: a strategic guide to the network economy*. Harvard Business Press.
Shiller, B. R. (2014). First-degree price discrimination using big data. `https://www.brandeis.edu/economics/RePEc/brd/doc/Brandeis_WP58R2.pdf`.
Simester, D., Timoshenko, A., and Zoumpoulis, S. I. (2020). Targeting prospective customers: Robustness of machine-learning methods to typical data challenges. *Management Science*, 66(6):2495–2522.
Sutton, R. S. and Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press.
Taddy, M., Gardner, M., Chen, L., and Draper, D. (2016). A nonparametric bayesian analysis of heterogenous treatment effects in digital experimentation. *Journal of Business & Economic Statistics*, 34(4):661–672.
Tang, Z., Smith, M. D., and Montgomery, A. (2010). The impact of shopbot use on prices and price dispersion: Evidence from online book retailing. *International Journal of Industrial Organization*, 28(6):579–590.
Trusov, M., Ma, L., and Jamal, Z. (2016). Crumbs of the cookie: User profiling in customer-base analysis and behavioral targeting. *Marketing Science*, 35(3):405–426.
Tsai, J. Y., Egelman, S., Cranor, L., and Acquisti, A. (2011). The effect of online privacy information on purchasing behavior: An experimental study. *Information Systems Research*, 22(2):254–268.
Tu, W. and Liu, H. (2014). Zero-inflated data. *Wiley StatsRef: Statistics Reference Online*, pages 1–7.
Varian, H. R. (1989). Price Discrimination. *Handbook of Industrial Organization*, 1:597–654.
Wager, S. and Athey, S. (2017). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, (just-accepted).
Wagner, G. and Eidenmuller, H. (2019). Down by Algorithms: Siphoning Rents, Exploiting Biases, and Shaping Preferences: Regulating the Dark Side of Personalized Transactions. *U. Chi. L. Rev.*, 86:581.
Wall, M. E., Rechtsteiner, A., and Rocha, L. M. (2003). Singular value decomposition and principal component analysis. In *A practical approach to microarray data analysis*, pages 91–109. Springer.
Wallheimer, B. (2018). Are you ready for personalized pricing? *Chicago Booth Review*.
Yang, J., Eckles, D., Dhillon, P., and Aral, S. (2020). Targeting for long-term outcomes. *arXiv preprint arXiv:2010.15835*.
Yi, X., Raghavan, H., and Leggetter, C. (2009). Discovering users' specific geo intention in web search. In *Proceedings of the 18th International Conference on World Wide Web*, pages 481–490.
Yoganarasimhan, H., Barzegary, E., and Pani, A. (2020). Design and Evaluation of Personalized Free Trials.

Youden, W. J. (1950). Index for rating diagnostic tests. *Cancer*, 3(1):32–35.
Zhang, D. (2003). Delivery of personalized and adaptive content to mobile devices: a framework and enabling technology. *Communications of the Association for Information Systems*, 12(1):13.
Zhang, J. and Krishnamurthi, L. (2004). Customizing promotions in online stores. *Marketing Science*, 23(4):561–578.