

Smith, Adam N.; Seiler, Stephan; Aggarwal, Ishant

Working Paper

Optimal Price Targeting

CESifo Working Paper, No. 9439

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Smith, Adam N.; Seiler, Stephan; Aggarwal, Ishant (2021) : Optimal Price Targeting, CESifo Working Paper, No. 9439, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/248984>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Optimal Price Targeting

Adam N. Smith, Stephan Seiler, Ishant Aggarwal

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

Optimal Price Targeting

Abstract

We examine the profitability of personalized pricing policies that are derived using different specifications of demand in a typical retail setting with consumer-level panel data. We generate pricing policies from a variety of models, including Bayesian hierarchical choice models, regularized regressions, and classification trees using different sets of data inputs. To compare pricing policies, we employ an inverse probability weighted estimator of profits that explicitly takes into account non-random price variation and the panel nature of the data. We find that the performance of machine learning models is highly varied, ranging from a 21% loss to a 17% gain relative to a blanket couponing strategy, and a standard Bayesian hierarchical logit model achieves a 17.5% gain. Across all models purchase histories lead to large improvements in profits, but demographic information only has a small impact. We show that out-of-sample hit probabilities, a standard measure of model performance, are uncorrelated with our profit estimator and provide poor guidance towards model selection.

JEL-Codes: C110, C330, C450, C520, D120, L110, L810.

Keywords: targeting, personalization, heterogeneity, choice models, machine learning.

Adam N. Smith
University College London / United Kingdom
a.smith@ucl.ac.uk

Stephan Seiler
Imperial College London / United Kingdom
stephan.a.seiler@gmail.com

Ishant Aggarwal
Lloyds Banking Group / London / United Kingdom
iagg1946@outlook.com

This draft: April 29, 2021

The authors are listed in reverse alphabetical order and contributed equally. We thank Rob Donnelly, Elea Feit, Jim Griffin, Bruce Hardie, Günter Hitsch, Unnati Narang, and Max Pachali for helpful comments, as well as seminar participants at CREST, the Frankfurt School of Finance and Management, Penn State, UCL, UC Riverside, University of Illinois at Urbana-Champaign, University of Rochester, and the European Quant Marketing Seminar. We also thank IRI for making the data available. All estimates and analysis in this paper based on data provided by IRI are by the authors and not by IRI. All errors are our own.

Electronic copy available at: <https://ssrn.com/abstract=3822459>

1 Introduction

The personalization of marketing instruments has long been a focus of marketing research and practice. The standard approach to deriving personalized marketing policies such as targeted prices or coupons consists of specifying a model of demand and then deriving optimal actions at the individual customer level that maximize expected profits. In practice, researchers have to choose from a set of models of demand that rely on different functional form and parametric assumptions, and also decide which customer characteristics are allowed to influence the shape of the demand curve. Recent advances in machine learning allow researchers to specify more flexible models by relaxing functional form assumptions and by allowing for high-dimensional vectors of customer characteristics to enter as covariates. In this paper, we explore how model flexibility and different data inputs impact the possible profit gains from designing targeted pricing policies.

We make two primary contributions. First, we develop a framework for comparing the profitability of price targeting strategies derived from different models of demand. We employ an inverse probability weighted estimator of profits that is independent of any demand-side modeling assumptions and can be used to predict and compare profits across any set of targeting policies. This approach allows us to compare models directly on the decision-relevant loss function (i.e., profits) rather than purely statistical measures of fit. Inverse probability weighted profit estimators were first applied in the marketing literature by Hitsch and Misra (2018) and Yoganarasimhan et al. (2020) who used cross-sectional data from field experiments (with random treatment assignment) to measure the gains from targeted marketing. We extend this approach to a setting where choices are observed for a panel of consumers with non-random price variation as is typically available from household-level data sets in retail markets. Our second contribution is to trace out targeting strategies derived from a variety of economic and machine learning models and then use the proposed evaluation framework to compare the performance of these models in terms of profits. For each model, we measure the value of different customer characteristics such as demographic and purchase history variables.

In our empirical application we use household panel data from the mayonnaise product category and solve for the optimal price targeting strategy for the leading brand in the category.¹ We fit a battery of models to the data including Bayesian hierarchical choice models, regularized regressions, neural networks, K -nearest neighbor (KNN) classifiers, and random forests. We find that the profitability of the induced targeting policies varies dramatically across different model specifications. The most profitable price targeting strategy leads to a 17.5% improvement in profits relative to a blanket (i.e., non-targeted) coupon, while others perform very poorly and generate pricing strategies that result in profits that are lower than the blanket coupon. The Bayesian hierarchical logit model (with only a limited set of demographics entering the upper-level model of

¹More specifically, we consider a one-off targeted couponing campaign initiated by the manufacturer. We do not observe coupons in our data (but only untargeted shelf-price discounts) and therefore treat a coupon as equivalent to a price discount that is accompanied by feature and display advertising. We discuss the relationship between (counterfactual) targeted price discounts via coupons and the price variation in our data in more detail in Section 2.3.

preference heterogeneity) generates the most profitable pricing policy, closely followed by regularized regression and the neural network based on an extended set of demographics and purchase history variables. Model refinements such as a feature selection step and a correction for class imbalance have relatively little impact on our results.

With regard to the value of different data inputs, we find that adding information on consumers' purchase histories in the machine learning models greatly improves performance, whereas demographic variables only have a small impact on profits. This pattern holds consistently across all types of models and reaffirms a key result in Rossi et al. (1996). Namely, demographic variables have limited informational content for measuring preference heterogeneity and this remains true even after including more demographic variables and allowing them to enter demand more flexibly. On the other hand, information on consumers' purchase behavior – either encoded in a unit-level likelihood within a Bayesian model or as covariates in a machine learning model – tends to be more valuable for designing personalized pricing policies.

Finally, we contrast our profit-based measure of model performance with a more traditional approach of ranking models by out-of-sample fit. The latter approach assesses model performance purely in terms of predictive power, does not require us to solve for the optimal pricing policy, and does not evaluate model performance in economic terms. We find that an out-of-sample hit probability metric ranks models very differently than our profit-based approach. More specifically, we find that the correlation between the out-of-sample hit probabilities and our preferred out-of-sample pricing policy profit estimator is close to zero across all model specifications. Our findings therefore caution against using a simple statistical measure of fit to evaluate models. Instead, models should be compared on the basis of a loss function matching the decision problem at hand.

Our paper builds upon a large literature on personalized marketing strategies. Within the context of targeted pricing, existing approaches for evaluating the effectiveness of personalized actions are typically model-based. That is, the researcher first estimates a model of demand and then uses that same model to predict demand and compute expected profits across a set of candidate prices. In this approach, the demand model is used twice: once for *estimation* and again for profit *evaluation*. This framework was implemented by Rossi et al. (1996) who measure the gains from targeted couponing under different information sets, and continues to be widely used for evaluating gains to customizing coupons (Zhang and Krishnamurthi, 2004; Pancras and Sudhir, 2007; Zhang and Wedel, 2009; Johnson et al., 2013; Donnelly et al., 2019; Gabel and Timoshenko, 2021), retail prices (Besanko et al., 2003; Morozov et al., 2021), and other marketing activities (Narayanan and Manchanda, 2009; Zantedeschi et al., 2017; Ascarza, 2018; Deng and Mela, 2018; Bumbaca et al., 2020). We deviate from this approach by using an inverse probability weighted estimator of profits that does not depend on any demand-side model structure, does not involve any extrapolation, and only assesses predicted profits based on data observed in a hold-out sample. Separating the model-based derivation of pricing policies from the evaluation framework allows us to compare pricing policies derived from different models of demand.

Our work also relates to an emerging literature on machine-learning-based targeting and tar-

geted policy evaluation. Dubé and Misra (2019) use data from a pricing experiment to estimate a logit choice model that projects demand parameters onto a high dimensional vector of characteristics, which is used to construct a targeted pricing policy. The authors then implement their targeting policy in a second experiment to assess its profitability relative to two uniform pricing benchmarks. In practice, however, researchers rarely have the ability to conduct randomized experiments over the space of possible policies. Simester et al. (2020a) emphasize the value of *randomized-by-action* instead of *randomized-by-policy* experiments, and suggest an estimation approach for evaluating the profitability of targeting policies that were not observed in the given sample. Hitsch and Misra (2018), Yoganarasimhan et al. (2020), and Cagala et al. (2021) use inverse probability weighted estimators of profits to evaluate different machine-learning-based targeting policies using cross-sectional data with randomized treatment assignment. We extend these “off-policy” profit estimators to evaluate different price targeting policies in typical CPG markets, which requires accounting for both the non-random treatment assignment of prices and the panel nature of the data. More specifically, and in contrast to most policy evaluation papers mentioned above, our data contains price variation over time, but not across individuals at a given point in time. A key step of our approach is to extrapolate from temporal price variation to a hypothetical targeting strategy that varies prices across consumers at a given point in time.

The remainder of the paper is organized as follows. In Section 2 we outline the proposed framework for comparing price targeting strategies. In Section 3 we show how to derive price targeting strategies based on different models of demand and different data inputs. In Section 4 we present the data and descriptive statistics. Section 5 presents details on the various demand models we estimate and Section 6 reports results regarding the relative performance of each model’s pricing policy. We offer concluding remarks in Section 7.

2 Evaluating the Profitability of Price Targeting Policies

In this section, we provide a general discussion of how to evaluate different pricing policies. We consider a setting where: (i) a company has access to consumer-level panel data with (non-experimental) prices that are not targeted, but vary over time; and (ii) the objective is to find optimal prices for different customers. This type of decision problem is common in CPG markets and includes the canonical manufacturer coupon targeting problem studied in Rossi et al. (1996). These features are also common in many other retail settings including online platforms (Morozov et al., 2021).

The typical approach for evaluating the profitability of price targeting policies is model-based. That is, the researcher specifies a model of demand, estimates model parameters, and solves for expected customer-level profits at candidate prices. In this framework, the demand model is used both to construct the targeting policy (i.e., solving for prices that maximize expected profits) and again to calculate expected profit levels associated with offering this policy. Therefore, such an approach does not allow us to compare targeting policies derived from different models of

demand. Counterfactual profits associated with a specific pricing policy are not observed and it is therefore not possible to separate model performance in terms of picking a pricing policy from the model’s ability to predict counterfactual profits. In what follows, we outline a framework for estimating demand and comparing profits of targeting policies which is independent of any demand-side modeling assumptions used to generate those policies. In our approach, the demand model is only used for the purpose of generating a candidate targeting policy. We then compare policies on the basis of *observed profits* in the data where the targeted price matches the observed price.

2.1 Targeting Policies and Profits

We define a targeting policy as the function $d : \mathcal{Z} \rightarrow \mathcal{P}$ which maps observed time-invariant customer characteristics into prices and therefore indicates the price $p \in \mathcal{P}$ with which a customer with attributes $\mathbf{z}_i \in \mathcal{Z}$ should be targeted. For each customer $i = 1, \dots, N$ the firm offers a personalized price in time period t in order to maximize profits $\pi_{it}(p)$, where $\pi_{it}(p)$ refers to a *potential* outcome.² The average per-customer profit of a targeting policy d implemented at time t is then given by:

$$\Pi_t(d) = \frac{1}{N} \sum_{i=1}^N \left[\sum_{p \in \mathcal{P}} \mathbf{1}(d(\mathbf{z}_i) = p) \times \pi_{it}(p) \right] \quad (1)$$

$$\pi_{it}(p) = Y_{it}(p) \times [p - mc], \quad (2)$$

where potential profits are equal to potential demand $Y_{it}(p)$ times the markup $[p - mc]$. We aim to evaluate different targeting policies based on the expected average per customer profit $\mathbb{E}[\Pi_t(d)]$ of a given pricing policy.

Given that our goal is to evaluate the profitability of targeting strategies from observational data, we must also account for the fact that potential profits $\pi_{it}(p)$ may be influenced by other variables in the market environment such as prices of competing goods, feature or display advertising, and seasonal variables, which we denote by \mathbf{x}_t . We focus on assessing the profitability of a “one-off” couponing campaign and assume that the firm does not strategically decide the timing of the campaign. We therefore take expectations over market characteristics when calculating expected profits:

²We assume that the stable unit treatment value assumption (SUTVA) holds. In other words, we assume that the potential profits are affected only by the treatment assignment (i.e., price) for customer i in time period t , but not by the treatment assignment for any other observations.

$$\mathbb{E}[\Pi_t(d)|\mathbf{Z}, \mathbf{x}_t] = \frac{1}{N} \sum_{i=1}^N \sum_{p \in \mathcal{P}} \mathbb{E}[\mathbf{1}(d(\mathbf{z}_i) = p) \times \pi_{it}(p)|\mathbf{z}_i, \mathbf{x}_t] \quad (3)$$

$$\begin{aligned} \mathbb{E}[\Pi_t(d)|\mathbf{Z}] &= \mathbb{E}_{\mathbf{x}}\{\mathbb{E}[\Pi_t(d)|\mathbf{Z}, \mathbf{x}_t]\} \\ &= \frac{1}{N} \sum_{i=1}^N \sum_{p \in \mathcal{P}} \mathbb{E}[\mathbf{1}(d(\mathbf{z}_i) = p) \times \pi_{it}(p)|\mathbf{z}_i], \end{aligned} \quad (4)$$

where $\mathbf{Z} = \{\mathbf{z}_1, \dots, \mathbf{z}_N\}$ denotes the characteristics of all consumers in the target population.

Unconditional expected average profits can be obtained by taking expectations over the population distribution of customer characteristics. We note that the expected profit expression in equation (4) matches the one in Hitsch and Misra (2018) who evaluate targeting policies based on a similar objective function with cross-sectional data. The key difference in our case is that we need to take expectations over the distribution of market characteristics \mathbf{x}_t in order to derive expected profits.

2.2 Inverse Probability Weighted Estimator

The fundamental problem in comparing the profitability of targeted pricing policies is that the price p_{it} that customer i sees on purchase occasion t in the evaluation sample will not generally coincide with the proposed targeting policy, $p_{it} \neq d(\mathbf{z}_i)$. To overcome this challenge, we specify an inverse probability weighted estimator (Robins et al., 1994; Lunceford and Davidian, 2004) of profits. The idea is that we can calculate expected profits based on the observed sample by using only the “usable” observations for which the actual price and the proposed targeted price agree, $p_{it} = d(\mathbf{z}_i)$. If we were to use a model-based evaluation approach instead (as in Rossi et al., 1996), we could compute $\pi_{it}(p)$ for all possible prices through model-based extrapolation. In order to avoid making demand-side modeling assumption, we need to confine ourselves to “usable” observation that do not require making predictions about behavior at counterfactual prices.

To account for the rate at which a given observation is usable, the estimator of expected profits weights observations by the inverse of the propensity score, i.e., the conditional probability that a specific price level is observed:

$$e_p(\mathbf{x}_t) = \mathbb{E}[\mathbf{1}(p_{it} = p)|\mathbf{x}_t] = \mathbb{P}(p_{it} = p|\mathbf{x}_t), \quad 0 < e_p(\mathbf{x}_t) < 1, \quad (5)$$

where the characteristics of the market environment \mathbf{x}_t (e.g., competitor prices, feature or display advertising, seasonal variables) explain the propensity with which a customer is exposed to a given price in a specific week. We omit customer demographics \mathbf{z}_i from the specification of propensity scores because we consider a setting where that prices in the data are not targeted to consumers.³

³In Section 4.4 we test whether demographics predict exposure to different price levels in our data and find no evidence for differential exposure to specific price levels across demographic groups.

Moreover, the common support assumption $0 < e_p(\mathbf{x}_t) < 1$ requires that each price level $p \in \mathcal{P}$ have positive support conditional on \mathbf{x}_t .

Our profit estimator is based on panel data where each customer i visits the retailer T_i times in the evaluation sample. The inverse probability weighted estimator is given by:

$$\hat{\Pi}(d) = \frac{1}{N} \sum_{i=1}^N \frac{1}{T_i} \sum_{t=1}^{T_i} \sum_{p \in \mathcal{P}} \mathbf{1}(d(\mathbf{z}_i) = p) \times \frac{\mathbf{1}(p_{it} = p)}{e_p(\mathbf{x}_t)} \times \pi_{it}(p) \quad (6)$$

where for each consumer, we sum over all store visits T_i of that customer and the $1/T_i$ scaling ensures that all consumers are weighted equally.

Under the assumption that exposure to a given price level is conditionally independent of potential profits, it follows that the profit estimator is unbiased. To show this property of the estimator, we first derive expected profits conditional on customer characteristics and market environment:

$$\begin{aligned} \mathbb{E}[\hat{\Pi}(d)|\mathbf{Z}, \mathbf{x}_t] &= \frac{1}{N} \sum_{i=1}^N \frac{1}{T_i} \sum_{t=1}^{T_i} \mathbb{E} \left[\sum_{p \in \mathcal{P}} \mathbf{1}(d(\mathbf{z}_i) = p) \times \frac{\mathbf{1}(p_{it} = p)}{e_p(\mathbf{x}_t)} \times \pi_{it}(p) | \mathbf{z}_i, \mathbf{x}_t \right] \\ &= \frac{1}{N} \sum_{i=1}^N \frac{1}{T_i} \sum_{t=1}^{T_i} \sum_{p \in \mathcal{P}} \mathbf{1}(d(\mathbf{z}_i) = p) \times \mathbb{E} \left[\frac{\mathbf{1}(p_{it} = p)}{e_p(\mathbf{x}_t)} \times \pi_{it}(p) | \mathbf{z}_i, \mathbf{x}_t \right] \\ &= \frac{1}{N} \sum_{i=1}^N \frac{1}{T_i} \sum_{t=1}^{T_i} \sum_{p \in \mathcal{P}} \mathbf{1}(d(\mathbf{z}_i) = p) \times \frac{e_p(\mathbf{x}_t)}{e_p(\mathbf{x}_t)} \times \mathbb{E}[\pi_{it}(p) | \mathbf{z}_i, \mathbf{x}_t] \\ &= \frac{1}{N} \sum_{i=1}^N \frac{1}{T_i} \sum_{t=1}^{T_i} \sum_{p \in \mathcal{P}} \mathbb{E}[\mathbf{1}(d(\mathbf{z}_i) = p) \times \pi_{it}(p) | \mathbf{z}_i, \mathbf{x}_t]. \end{aligned} \quad (7)$$

Here the third equality follows from the assumption that $\mathbf{1}(p_{it} = p)$ and $\pi_{it}(p)$ are conditionally independent given \mathbf{x}_t . Next, we take expectation over market characteristics \mathbf{x}_t :

$$\begin{aligned} \mathbb{E}[\hat{\Pi}(d)|\mathbf{Z}] &= \mathbb{E}_{\mathbf{x}} \{ \mathbb{E}[\hat{\Pi}(d)|\mathbf{Z}, \mathbf{x}_t] \} \\ &= \frac{1}{N} \sum_{i=1}^N \frac{1}{T_i} \sum_{t=1}^{T_i} \sum_{p \in \mathcal{P}} \mathbb{E}[\mathbf{1}(d(\mathbf{z}_i) = p) \times \pi_{it}(p) | \mathbf{z}_i] \\ &= \frac{1}{N} \sum_{i=1}^N \sum_{p \in \mathcal{P}} \mathbb{E}[\mathbf{1}(d(\mathbf{z}_i) = p) \times \pi_{it}(p) | \mathbf{z}_i] \\ &= \mathbb{E}[\Pi_t(d)|\mathbf{Z}], \end{aligned} \quad (8)$$

where the second line follows from the law of iterated expectations and the third line follows from the fact that $\mathbb{E}[\mathbf{1}(d(\mathbf{z}_i) = p) \times \pi_{it}(p) | \mathbf{z}_i]$ is not time-varying.

In order to gain some intuition for how the profit estimator is calculated, consider the following simple example. Each consumer can be offered either a regular or a discounted price and consumers encounter those two prices levels with probabilities of 80% and 20% in the evaluation sample.

Further assume that we are trying to calculate the profit contribution of a specific consumer who went shopping 10 times during the evaluation sample and on 8 trips the focal product was offered at its regular price whereas the product was discounted on the remaining 2 shopping trips.⁴ If the pricing policy prescribes a regular price for this specific consumer, we calculate observed profits for each of the 8 trips where the product was offered at the regular price level. We then divide trip-level profits by the propensity score associated with the regular price level of 0.8, which rescales the sum of trip-level profits over 8 trips to the same scale as the 10 total trips. Dividing by $T_i = 10$ yields the consumer-specific contribution to the average profit estimator: $\frac{1}{T_i} \sum_{t=1}^{T_i} \sum_{p \in \mathcal{P}} \mathbf{1}(d(\mathbf{z}_i) = p) \times [\mathbf{1}(p_{it} = p)/e_p(\mathbf{x}_t)] \times \pi_{it}(p)$. If a specific pricing policy prescribes a discounted price for the same consumer, we would calculate profits over the 2 trips where the product is discounted and divide profits by 0.2 to rescale them.

2.3 Conditional Independence, Propensity Scores, and Sample Selection

So far we have outlined a general profit evaluation framework based on panel data where prices are not personalized and hence customer characteristics do not enter the propensity score. We believe this scenario applies to many retail settings where a firm repeatedly interacts with a given set of customers and product prices vary over time. In this section, we discuss the conditional independence assumption by appealing to institutional features and data patterns that we believe are common in CPG markets and many related settings. While some of the steps we outline are motivated by patterns in our specific data set, the general ideas presented here are likely to be more broadly applicable.

As described in the previous section, in order for our estimator to be unbiased, we require exposure to different prices to be conditionally independent of potential profits. Hence we need to condition on the relevant market variables \mathbf{x}_t such that price variation can be treated as-good-as-random conditional on \mathbf{x}_t . The profit estimator also requires all price levels to have support in the data conditional on \mathbf{x}_t . If a specific price level is never observed conditional on a specific realization of \mathbf{x}_t , then $e_p(\mathbf{x}_t) = 0$ and the profit estimator is not defined. This required common support condition of the profit estimator motivates our approach to dealing with non-random price variation that we outline below. We take two steps in order to generate an evaluation sample where the conditional independence assumption is likely to hold.

First, we note that in typical retail markets, certain variables will be highly (sometimes perfectly) correlated with discounts for the focal product. For example, most grocery retailers adhere to a promotional calendar such that in weeks when a given product is discounted, other competitor products in the same category tend to not be discounted, featured or displayed at the same

⁴For simplicity, we assume that this individual consumer’s frequency of exposure to the two price levels coincides with the averages frequencies of the price levels in the evaluation sample. This need not be the case in the actual evaluation sample. Moreover, it is possible that an individual consumer is never exposed to a specific price level. Because we are pooling information across consumers, we do not need full support across price levels for each individual consumer. We only require that observed customer characteristics \mathbf{z}_i are uncorrelated with exposure to different price level. In Section 4.4 we show that this condition holds for the customer characteristics considered in our setting.

time. Moreover, whenever the focal product is discounted, it is also almost always featured and displayed at the same time. Conversely, when the focal product is not discounted, it is rarely featured or displayed. We report descriptive statistics in Section 4.1 providing evidence for these type of promotion patterns in our data.

The presence of market factors that have near-perfect correlations with price makes it difficult to apply the inverse probability weighted profit estimator due to the common support requirement. For example, it would not be possible for us to evaluate a price discount that occurs at the same time as a discount on a competitor product if such a scenario is never observed in the data. We therefore confine our profit evaluation sample to only include weeks where no competitor product is discounted, featured, or displayed and the focal product is either discounted and featured and displayed at the same time or offered at its regular price without an accompanying feature or display. While this sample construction is dictated by the properties of the profit estimator, we believe that it also aptly captures the nature of coupons. Like regular discounts, coupons are likely to only be issued in weeks with no other marketing activity for competitor products. Moreover, our sample construction captures the fact that a coupon is likely to also have a salience effect to the consumer that receives the coupon (similar to the salience effect of features and displays).

In a second step, we assess whether we need to condition on additional non-promotional market factors to ensure that the conditional independence assumption holds. In many retail settings, price promotions are not strategically timed by manufacturers and instead manufacturers only agree on a specific number of promotions that the retailer is required to implement in a given time span. Therefore, it is often argued (see for instance the discussion in Rossi, 2014b) that price variation over time can be treated to be as-good-as-random. We believe that this assumption is likely to be satisfied after eliminating any variation in other marketing activity through the sample selection process described above. Intuitively, our argument rests on the fact that the timing of discounts is as-good-as-random over time in a given store (after controlling for other marketing activity) and therefore the observations where observed prices are not matching those prescribed by a given targeting policy are “missing at random” and no different from usable observations in terms of potential profits. We also note that the assumptions required for our profit estimator to be unbiased are similar to the ones employed by Rossi et al. (1996) in the context of a model-based profit evaluation. Rossi et al. (1996) evaluate profits based on estimates from a demand model that controls for product fixed effects as well as feature and display advertising. Therefore, their model identifies the causal effect of price on demand (and hence on profits) if temporal price variation is as-good-as-random after controlling for other marketing activity.

In any specific setting, the researcher can empirically test whether specific factors are correlated with price. We implement such an analysis in Section 4.4 and find that, apart from the chain-identity of the store, other variables such as competitor prices do not predict whether the focal product is discounted. We therefore model the propensity score as a function of only chain identity, but not any of the other time-varying market variables. In general, we believe that after the selection step, which eliminates correlation with other marketing activity, the set of variables entering into $e_p(\cdot)$

is likely small and usually observed by the researcher.

In summary, we confine the profit evaluation to a subset of weeks that contain no variation in other promotional activity. For the remaining weeks, we take account of differential exposure to prices (as a function of \mathbf{x}_t) via the propensity score $e_p(\cdot)$. We assume that after the sample selection step, potential profits are independent of exposure to different price levels conditional on the set of variables that enter the propensity score.

3 Constructing Price Targeting Policies & Workflow

We now turn to the specification of demand and the construction of targeted pricing policies. Let $Y_{ijt}(\mathbf{p}_t, \mathbf{x}_t, \mathbf{z}_i)$ denote the demand of good $j \in \{0, 1, \dots, J\}$ by consumer i at time t , with $j = 0$ denoting the outside option. Demand depends on a complete vector of prices $\mathbf{p}_t = (p_{1t}, \dots, p_{Jt})$, a vector of non-price demand shifters \mathbf{x}_t (e.g., product characteristics or feature/display advertising),⁵ and a vector of time-invariant customer characteristics \mathbf{z}_i . When convenient, we write $\tilde{\mathbf{x}}_t = (\mathbf{p}_t, \mathbf{x}_t)$ to denote the full vector of demand shifters. We assume that demand arises from a model of discrete choice, $Y_{ijt} \in \{0, 1\}$, and so the goal is to model choice probabilities: $\mathbb{E}[Y_{ijt}(\tilde{\mathbf{x}}_t, \mathbf{z}_i)] = \mathbb{P}(j|\tilde{\mathbf{x}}_t, \mathbf{z}_i)$.⁶ We let \mathcal{M} denote a given model chosen by the researcher to estimate choice probabilities.

3.1 Policy Construction

Given data $(\tilde{\mathbf{x}}_t, \mathbf{z}_i)$ and a model \mathcal{M} , we can define profits associated with charging price p_{ij} for good j to customer i at time t as:

$$\pi_{it}(p_{ij}|\tilde{\mathbf{x}}_t, \mathbf{z}_i, \mathcal{M}) = \mathbb{P}(j|p_{ij}, \tilde{\mathbf{x}}_t, \mathbf{z}_i, \mathcal{M})[p_{ij} - mc_j]. \quad (9)$$

The optimal price targeting policy $d^*(\mathbf{z}_i)$ induced by model \mathcal{M} will then solve for the price $p_{ij} \in \mathcal{P}$ for each customer that maximizes expected profits.

$$d^*(\mathbf{z}_i) \equiv p_{ij}^* = \operatorname{argmax}_{p_{ij} \in \mathcal{P}} \frac{1}{T_i} \sum_{t=1}^{T_i} \mathbb{E}[\pi_{it}(p_{ij}|\tilde{\mathbf{x}}_t, \mathbf{z}_i, \mathcal{M})] \quad (10)$$

There are two sources of uncertainty in (9) that are relevant for computing expected profits and solving for an optimal policy. The first is uncertainty in the market environment $\tilde{\mathbf{x}}_t$ (e.g., changes in competitor prices) at the time the coupon is issued, which we account for by integrating the profit expression over the empirical distribution of demand shifters:

$$\mathbb{E}_{\tilde{\mathbf{x}}_t}[\pi_{it}(p_{ij}|\tilde{\mathbf{x}}_t, \mathbf{z}_i, \mathcal{M})] = [p_{ij} - mc_j] \int \mathbb{P}(j|p_{ij}, \tilde{\mathbf{x}}_t, \mathbf{z}_i, \mathcal{M}) dF(\tilde{\mathbf{x}}_t). \quad (11)$$

⁵We use \mathbf{x}_t to refer both to time-varying variables characterizing the market environment as well as time-invariant product characteristics (which were ignored in Section 2 since they are not relevant to our evaluation framework).

⁶The assumption of discrete choice can be easily relaxed. For example, in the case of continuous or discrete/continuous demand one can model the conditional mean function $\mathbb{E}[Y_{ijt}(\tilde{\mathbf{x}}_t, \mathbf{z}_i)]$ directly.

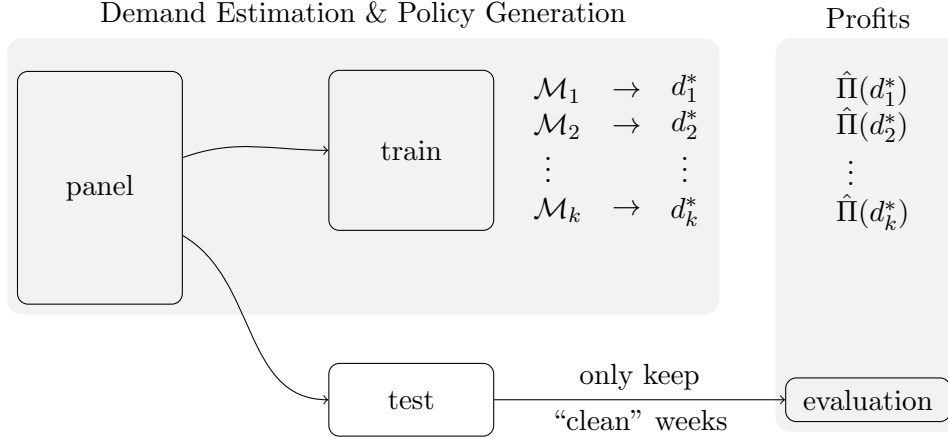


Figure 1: **Sample Selection Workflow.** The full panel data is first split into a training set and a test set. The training set is used to fit various models of demand \mathcal{M} and then solve for an optimal pricing policy d^* . The evaluation sample is constructed by only keeping “clean” weeks with no variation in feature/display activity. Profits are estimated from the evaluation sample using the inverse probability weighted estimator.

The second is uncertainty in model parameters. It is straightforward to account for parameter uncertainty in a Bayesian model by integrating (9) over the posterior distribution of individual-level parameters $\Psi = \{\psi_1, \dots, \psi_N\}$:

$$\mathbb{E}_{\tilde{\mathbf{x}}_t, \Psi}[\pi_i(p_{ij}|\tilde{\mathbf{x}}_t, \mathbf{z}_i, \mathcal{M})] = [p_{ij} - mc_j] \int \int \mathbb{P}(j|\psi_i, p_{ij}, \tilde{\mathbf{x}}_t, \mathbf{z}_i, \mathcal{M}) dF(\tilde{\mathbf{x}}_t) dF(\psi_i|\text{data}). \quad (12)$$

Accounting for parameter uncertainty is less straightforward in a machine learning model estimated through a cross-validation procedure, and we therefore resort to a “plug-in” estimator where expected profits in equation (11) are evaluated at the point estimates of parameter values.

3.2 Workflow

Given a set of candidate models $\mathcal{M}_1, \dots, \mathcal{M}_k$ we now describe our policy construction and evaluation workflow. For each household in the panel, we first randomly allocate 20% of their purchase occasions towards the test sample and the remaining 80% towards the training sample. By sampling “within households” we ensure that all households are contained in both the training and test data. We then estimate all candidate models $\mathcal{M}_1, \dots, \mathcal{M}_k$ using the training sample, and use the procedure outlined above to derive the associated set of optimal price targeting policies d_1^*, \dots, d_k^* . To evaluate the profitability of those policies, we take the test sample and apply the sample selection criteria outlined in Section 2.3. Profits are then estimated from this “evaluation sample” using the profit estimator defined in Section 2.2. Figure 1 provides a diagram of this workflow.

4 Empirical Application

We use household-level scanner panel data from the IRI academic data set (Bronnenberg et al., 2008). We focus on the mayonnaise product category in the Pittsfield, Massachusetts market and use two years of data (2011-2012) from four stores that belong to two chains (two stores per chain). Our sample includes price and purchase information for the three largest brands of mayonnaise (Hellmann's, Kraft Miracle Whip, and Cain's) and is comprised of all trips of households that visited any of the four focal stores. We also confine our sample to households that purchase in the mayonnaise category at least once in each of the two calendar years, but otherwise maintain all trips regardless of whether a consumer purchased any product in the category. In total, we observe 659 households for 100,194 trips and 3,542 purchases in the category. For the average (median) household we observe 152 (128) trips and 5.37 (4) purchases.

In our analysis, we take the perspective of a focal manufacturer, Hellmann's, who can engage in third-degree price discrimination by issuing personalized coupons that offer targeted prices to each customer. This application fits into the evaluation framework outlined above: our data contains non-random price variation over time but no targeted prices across customers. Moreover, we assume that Hellman's is not in direct control of the timing the coupon, but that coupons are part of a promotional calendar (similar to other marketing activities such as discounts of the shelf price and feature advertising) where the retailer will decide the exact timing of the coupon. Therefore, targeting is based only on customer characteristics, but not on time-varying market characteristics.⁷

4.1 Product and Market Characteristics

Prices are identical for all products in stores that belong to the same chain. Figure 2 displays the time series of chain-level prices for all three brands. We find that prices at all store/product pairs follow patterns typically observed for many consumer packaged goods and are characterized by a relatively stable regular price level and occasional discounts that lead to a sharp temporary decrease in price. Table 1 displays descriptive statistics on prices, discounts (defined as a decrease of at least 25% relative to the regular price), sales, as well as other marketing activity. Regular price levels are fairly similar across brands and chains, but the use of other marketing activity tends to differ across chains and products. For example, Cain's is never discounted and never featured or displays in chain A, where chain B engages in various marketing activities for Cain's. Both stores are similar in terms of overall sales as well as the market shares of the three brands.

Next, we present some additional descriptive statistics for the focal brand, Hellmann's. As Figure 2 and Panel (I) of Table 1 show, both the regular price level and the frequency and level of discount are relatively similar in both chains. We therefore present further descriptive statistics

⁷Due to the manufacturer / retailer interaction, we re-write manufacturer profits as $Y_{it}(p) \times [m - (p - \bar{p})]$ where \bar{p} denotes the regular price level and m denotes the manufacturer dollar margin at the regular price level. We assume that the manufacturer margin is equal to 60% and the retailer margin is 20%. Under these assumptions, the dollar margin of the manufacturer (at the regular price level) is equal to 48% of the regular price level (based on the following calculation: manufacturer dollar margin = manufacturer %-margin \times [1 - retailer %-margin] \times regular price). We assume that any price reduction reduces the manufacturer margin 1-to-1 and does not alter the retailer margin.

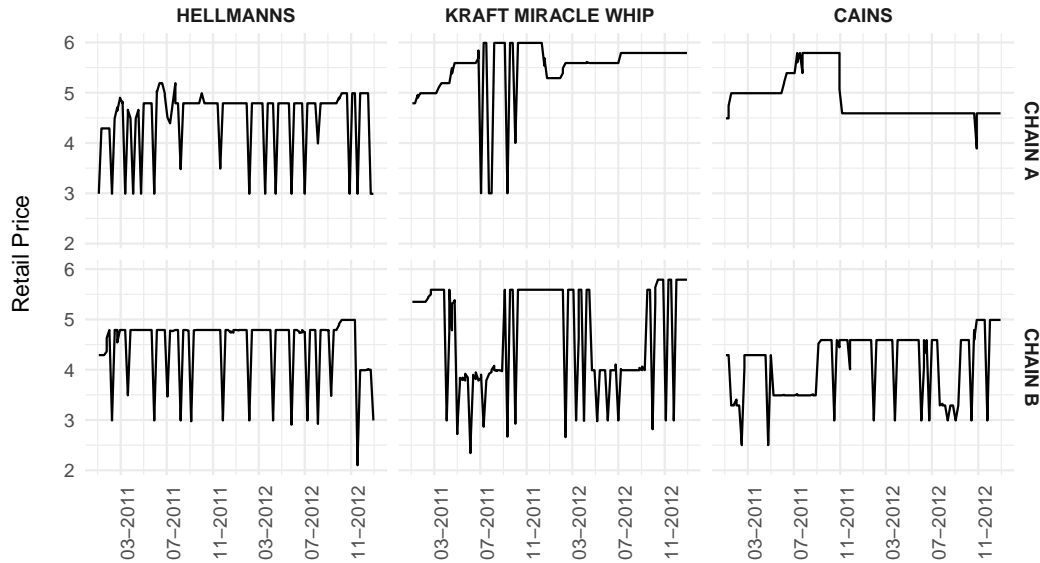


Figure 2: **Time Series of Prices.**

that pool data across chains. Panel (II) in Table 1 displays the cross-tabulation of discounts with feature advertising as well as with product displays (across chain/week combinations) for Hellmann's. We base the calculations in this panel only on chain/week combinations where the price of Hellmann's is either at \$4.79 or \$3.00, which are the two modal price levels.⁸ Those cross-tabulations show that discounts in our data are always accompanied by feature advertising and almost always by a product display. Occasionally, displays and feature advertising are used even when the product is not discounted. We also show in the same panel, how discounts for the focal brand correlate with discounts and other marketing activity for the two competing brands. We find that when the focal brand is discounted, other brands are almost never discounted, and rarely displayed or featured. The cross-tabulations are consistent with the typical pattern of accompanying discounts with other promotional activity as well as with the fact that only one brand in the category is discounted (or featured / displayed) in a given week. The latter pattern is most likely due to contractual arrangements between retailers and manufacturers that lead to an absence of discounting for multiple brands at the same time. The patterns discussed in this paragraph are evidence of the type of correlations between marketing activities discussed in Section 2.3 and form the basis for a sample selection step when generating the evaluation sample which we discuss in detail in Section 4.3.

⁸We treat prices that are 5 cents above or below those modal prices to be identical to those two price levels. Based on this selection procedure we retain 307 out of 420 chain/week pairs.

Panel (I)						Weekly	
Brand	Chain	Avg Price	Discount	Feature	Display	Units Sold	Market Share
Hellmann's	A	4.50	0.16	0.20	0.19	245.15	0.87
Kraft Miracle Whip	A	5.50	0.05	0.05	0.02	25.03	0.09
Cain's	A	4.89	0.00	0.00	0.00	10.10	0.04
Hellmann's	B	4.49	0.14	0.14	0.15	168.29	0.80
Kraft Miracle Whip	B	4.67	0.28	0.17	0.18	24.35	0.12
Cain's	B	4.10	0.16	0.09	0.15	17.67	0.08

Panel (II)		Hellmann's Featured		Hellmann's Displayed		Other Brand Featured or Displayed	Other Brand Discounted
		no	yes	no	yes	no	yes
Hellmann's	no	0.824	0.020	0.827	0.016	0.62	0.22
Discounted	yes	0.000	0.156	0.010	0.147	0.15	0.01

Table 1: **Summary of Product and Market Characteristics.** Panel (I) reports summary statistics for each brand. Panel (II) reports the co-occurrence of Hellmann's price discounts with own and competitor promotions.

4.2 Customer Characteristics

One focus of this paper is to assess the importance of different customer characteristics in terms of their ability to generate profitable targeted pricing policies. We therefore construct three sets of customer characteristics which are shown in Table 2. The first set includes 5 “base” demographic variables: income, family size, employment status, retirement status, and an indicator for a single mother. These variables are identical to the demographics in Rossi et al. (1996) and similar to demographic variables in other related research (e.g., Gupta and Chintagunta, 1994; Ainslie and Rossi, 1998; Manchanda et al., 1999; Horsky et al., 2006). We also include a set of 10 “extra” demographic variables including age, education, type of employment, number of children, marital status, number of cats and dogs, homeowner status, and number of TVs in the household. Together, these 15 total demographic variables are substantially richer than the demographics traditionally used and represent an exhaustive set of the demographic variables available in our data.

In addition to demographics, we also compute a set of customer purchase history variables, which capture brand preferences, price sensitivity, and responsiveness to other marketing instruments. Specifically, we calculate each brand's purchase share, each brand's total purchase count, the number of different brands purchased in the sample, the share of purchases where the purchased brand (regardless of which brand was chosen) was offered at a discount, the proportion of times the consumers buys a featured product, the proportion of times the consumers buys a product that is on display, and finally, the total category spend of a given consumer during the sample period.

		Mean	Std. Dev.
<i>Base Demographics:</i>	Income	51,722	29,583
	Family Size	2.52	1.15
	Retired	0.058	0.233
	Unemployed	0.232	0.423
	Single Mother	0.085	0.279
<i>Extra Demographics:</i>	Age	58.57	10.74
	High School	0.525	0.500
	College	0.451	0.498
	White Collar	0.455	0.498
	Number of Children	0.590	1.359
	Married	0.103	0.304
	Number of Dogs	0.200	0.401
	Number of Cats	0.203	0.407
	Renter	0.876	0.330
	Number of TVs	2.31	0.87
<i>Purchase History Variables:</i>	Hellmann's Choice Share	0.823	0.341
	Kraft Choice Share	0.118	0.294
	Cain's Choice Share	0.052	0.186
	Hellmann's Choice Count	3.54	2.89
	Kraft Choice Count	0.533	1.700
	Cain's Choice Count	0.263	1.387
	Discount Share	0.526	0.335
	Feature Share	0.452	0.349
	Display Share	0.447	0.34
	Number of Brands Purchased	1.15	0.40
	Total Spending	17.85	13.52

Table 2: **Summary of Customer Characteristics.**

This results in a total of 11 purchase history variables. Some of the variables capture aspects of purchase histories in a similar way to RFM (recency, frequency, monetary value) variables often used for targeting. Total spending and purchase counts capture monetary value and frequency respectively. Recency-related variables are not relevant in our setting because we are considering a targeting strategy whose timing is not under the control of the manufacturer. We present descriptive statistics for all customer characteristics in Table 2.

4.3 Evaluation Sample

As discussed in Section 2.3 and based on the empirical patterns presented in Section 4.1, we use only a subset of chain/weeks in our evaluation sample. First, we only focus on two price levels: \$4.79 and \$3.00 and eliminate chain/weeks with different prices (the same selection criterion that we used in Panel (II) of Table 1). Most other prices have relatively little support in our data

and conditioning on these two price levels allows us retain 307 out of 420 chain/weeks. Moreover, we only retain chain/weeks that fall into one of the following two groups. The first group (the “discount group”) is comprised of all chain/weeks where the focal product is discounted and priced at \$3.00. We furthermore confine this group to chain/weeks where the focal brand is featured and displayed and no other brand is discounted, featured or displayed. As the analysis in Panel (II) of Table 1 shows, these conditions are almost always met in our sample. The second group (the “regular price group”) is comprised of regular price weeks where neither the focal brand nor any competing brand is discounted, featured, or displays. After applying these selection criteria we are left with 181 chain/weeks at the regular price level and 36 chain/weeks at the discounted price level.

This particular way of constructing the evaluation sample allows us to compare a scenario where the focal product is discounted as well as featured and displayed with a scenario where the focal product is offered at the regular price and neither featured nor displayed. In both scenarios competitor brands are offered at their regular price and neither featured nor displayed. This selection procedure mimics how a couponing strategy is likely to be implemented. We documented above that retailers tend to exclusively discount (as well as feature and display) one product at a time. We assume that such exclusivity of marketing activity will also hold for a couponing campaign and thus we are considering a situation where a retailer is offering a coupon for the focal brand in a week where no other marketing activity occurs for competing brands. In this particular week, some consumers will not receive a coupon and hence the product is offered at the regular price to them. All remaining consumers receive a coupon that allows them to purchase the focal product at the discounted price. These two groups correspond to the regular price group and discount group we defined above. Furthermore, we assume that offering a coupon is equivalent to an in-store discount that is accompanied by both feature advertising and a product display. This assumption accounts for the fact that merely offering a coupon will make the product more salient to the consumer.

4.4 Propensity Scores and Independence Assumption

Because we use data from two retail chains for estimation and evaluation, we derive a separate set of optimal prices for each consumer/chain combination. We therefore compute the inverse probability weighted profit estimator in equation (6) for each chain separately and also calculate separate propensity scores for each chain.⁹ Moreover, as discussed in Section 2.3, we need to condition on the relevant market variables \mathbf{x}_t such that price variation can be treated as-good-as-random conditional on \mathbf{x}_t . Based on institutional features in CPG markets, we believe that price variation over time within a given store is likely uncorrelated with demand shifters because manufacturer are not able to strategically time discounts. Instead, discounts tend to be agreed upon as part of an annual promotional calendar that leaves the timing of discounts at the retailer’s discretion. In what follows we provide supporting evidence for the conditional independence assumption by

⁹When calculating average per-customer profits, we first compute average profits at the chain level and then take a simple average across the two chains.

	<i>Dependent Variable: Discount Dummy</i>			
	(1)	(2)	(3)	(4)
Chain B Dummy	-0.049 *** (0.003)	-0.049 *** (0.003)	-0.048 *** (0.003)	-0.157 (0.090)
Income (in \$10,000s)	0.000 (0.001)	0.001 (0.001)	0.000 (0.001)	
Family Size	0.000 (0.001)	0.001 (0.002)	0.001 (0.002)	
Retired	0.010 (0.007)	0.008 (0.008)	0.009 (0.008)	
Unemployed	-0.001 (0.004)	-0.001 (0.004)	-0.001 (0.004)	
Single Mother	0.001 (0.006)	0.001 (0.006)	0.003 (0.006)	
Year 2012 Dummy				-0.082 (0.084)
Price Kraft Miracle Whip				-0.090 (0.084)
Price Cain's				-0.009 (0.084)
Additional Demographics	No	Yes	Yes	n/a
Purchase History Variables	No	No	Yes	n/a
Month Controls	No	No	No	Yes
F-stat (excl. chain dummy)	0.357	0.572	0.831	0.996
F-stat (month dummies)				1.159
F-stat (all variables)	36.448	14.092	8.834	0.994
Unit of Observation	Store/ Visit	Store/ Visit	Store/ Visit	Chain/ Week
Observations	57,732	57,732	57,732	116

Table 3: **Determinants of Price Discounts.** Columns (1), (2), and (3) correspond to a linear probability model estimated on the full evaluation sample where the unit of analysis is the store/visit. Column (4) corresponds to a model estimated only on the time series of prices where the unit of analysis is the store-week and where the weeks match the weeks present in the evaluation sample. Significance codes: * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

showing that a large set of observed market variables are uncorrelated with the price level of our focal product.

Before turning to market variables, we first test whether the probability of encountering a regular or discounted price level is correlated with customer demographics. Such a correlation is unlikely because prices are uniform across customers on a given day and only vary over time. However, demographics could in principle be correlated with prices if some demographics groups are more likely to visit a store during a promotional period than other groups. Although, we believe such

Chain	Price	Propensity Score	# of Store Visits	Purchase Prob.
A	\$3.00	0.188	1392	0.133
A	\$4.79	0.812	6019	0.019
B	\$3.00	0.138	543	0.048
B	\$4.79	0.862	3383	0.017

Table 4: **Propensity Scores.**

selection into store visits is unlikely to occur, we nevertheless test for correlations between prices and demographics in our data. To this end, we use the evaluation sample described above and estimate a linear probability model where a discount dummy is regressed on a vector of demographics.¹⁰ Column (1) of Table 3 uses a base set of demographics as regressors, whereas column (2) use the full set of demographic variables. In both cases no individual demographic variables have a significant impact and the sets of demographics in both regressions are jointly insignificant. Finally, column (3) also adds the purchase history variables discussed to the regression. Similar to the results in the previous two columns, we find that the full set of customer characteristics is jointly insignificant. We therefore conclude that customer characteristics are uncorrelated with the price level consumers encounter on a store visit and hence we can model the propensity score $e_p(\cdot)$ as a function of \mathbf{x}_t but not \mathbf{z}_i .

Next, we implement a set of regressions to test whether other time-varying variables are correlated with prices. Specifically, we assess whether fluctuations in the regular price level of competing products predict discounts for the focal product (we do not include competitor feature/display variables since weeks with competitor advertising are excluded from our evaluation sample). Moreover, we test for seasonal effects and time trends by including a set of month dummies and a dummy for one of the two years in our sample. Finally, we also include a dummy for one of the two chains. The regression is implemented at the chain/week level based on all chain/weeks that are part of the evaluation sample. The results are reported in column (4) of Table 3. We find that competitor prices and the year dummy are insignificant and the month dummies are jointly insignificant. As we outlined earlier, we are computing the profit estimator at the chain level and therefore allow for chain-specific propensity scores. Based on the lack of correlation of other variables with price reported in Table 3, we do not construct the propensity score at a more granular level.

Table 4 displays the total number of store visits that occur within each chain/price-level combination in the evaluation sample, as well as the chain-specific propensity scores that we compute using a frequency estimator. We also display the purchase probabilities for each of the cells. As expected, we find that purchase probabilities increase significantly at both chains when the focal brand is discounted.

¹⁰We control for chain identity in all regression because chains differ in discount frequencies and chain identity might be correlated with specific customer characteristics.

5 Demand Models

In this section, we outline common economic choice models and machine learning methods that can be used to estimate demand. In Appendix A we provide a more detailed presentation of each model, including a discussion of key similarities and differences in underlying model structure. All models are based on two sets of data inputs: product and market characteristics $\tilde{\mathbf{x}}_t$ and customer characteristics \mathbf{z}_i . The product and market variables are kept constant across all models that we estimate and include product intercepts as well as price, a display dummy, and a feature advertising dummy for all three brands, and a chain dummy. We then experiment with three sets of customer characteristics introduced in Section 4.2: base demographics, extended set of demographics, extended demographics plus purchase history variables.

5.1 Bayesian Hierarchical Choice Models

Bayesian hierarchical choice models are comprised of two key ingredients: (i) a unit-level random utility model; and (ii) a random effects distribution that characterizes preference heterogeneity in the population. The random utility model posits that demand Y_{ijt} is the outcome of a latent utility maximization process:

$$Y_{ijt} = \begin{cases} 1 & \text{if } u_{ijt} \geq \max(u_{i0t}, u_{i1t}, \dots, u_{iJt}) \\ 0 & \text{otherwise} \end{cases} \quad (13)$$

where utility u_{ijt} is expressed as a linear function of product characteristics:

$$\begin{aligned} u_{ijt} &= v_j(\tilde{\mathbf{x}}_t; \psi_i) + \varepsilon_{ijt} \\ &= \tilde{\mathbf{x}}'_{jt} \psi_i + \varepsilon_{ijt}. \end{aligned} \quad (14)$$

Here $v_j(\tilde{\mathbf{x}}_t; \psi_i) = \tilde{\mathbf{x}}'_{jt} \psi_i$ represents the deterministic component of utility (with the normalization $v_0(\tilde{\mathbf{x}}_t; \psi_i) = 0$ for the outside option) while ε_{ijt} is an idiosyncratic error. We assume the ε_{ijt} 's are independent and identically distributed type I extreme value and so the choice probabilities are given by the following expression:

$$\mathbb{P}(j|\tilde{\mathbf{x}}_t, \psi_i) = \frac{\exp(\tilde{\mathbf{x}}'_{jt} \psi_i)}{1 + \sum_{k=1}^J \exp(\tilde{\mathbf{x}}'_{kt} \psi_i)} \quad (15)$$

where the denominator takes into account that the consumer can choose an outside option of not purchasing in the category on a given trip.

To complete the model, we define a hierarchical prior on unit-level parameters ψ_i that incorporates observable characteristics about each customer.

$$\psi_i = \Delta' \mathbf{z}_i + \xi_i, \quad \xi_i \sim N(0, V_\psi) \quad (16)$$

This specification induces a multivariate normal prior for ψ_i with mean $\Delta'z_i$ and covariance matrix V_ψ . We also estimate models with more flexible distributions of heterogeneity such as mixtures of normals (Rossi, 2014a). The posterior distribution of individual-level parameters ψ_i is then obtained through Bayes' rule:

$$f(\psi_i|\mathbf{Y}_i, \tilde{\mathbf{x}}_i, \mathbf{z}_i, \Delta, V_\psi) = \frac{L(\mathbf{Y}_i|\tilde{\mathbf{x}}_i, \psi_i)f(\psi_i|\mathbf{z}_i, \Delta, V_\psi)}{\int L(\mathbf{Y}_i|\tilde{\mathbf{x}}_i, \psi_i)dF(\psi_i|\mathbf{z}_i, \Delta, V_\psi)} \quad (17)$$

where $\mathbf{Y}_i = \{\mathbf{Y}_{i1}, \dots, \mathbf{Y}_{iT_i}\}$ is the full history of choices for customer i (including the outside option), $\tilde{\mathbf{x}}_i = \{\tilde{\mathbf{x}}_{i1}, \dots, \tilde{\mathbf{x}}_{iT_i}\}$ is the set of market characteristics observed at each purchase occasion, $L(\mathbf{Y}_i|\tilde{\mathbf{x}}_i, \psi_i)$ is the likelihood function induced by the model in (15), and $f(\psi_i|\mathbf{z}_i, \Delta, V_\psi)$ is the distribution of heterogeneity.

Inferences about a customer's preference vector ψ_i will be shaped by both their own purchase history data as well as the distribution of preferences in the population, with the weights on each component being driven by the "signal" in the data. Intuitively, when relatively little purchase information is available for a customer, the likelihood will be flat and the posterior will be informed by the population-level parameters (conditional on demographics). If we instead observe a long history of purchases, then the likelihood will be more peaked and the posterior will concentrate more heavily around the sample information. For example, for a consumer that always purchased the same brand even when alternative products are discounted, we will infer that she has a strong preference (i.e. a large product intercept) for the brand she always buys.

5.2 Machine Learning Models

Next, we summarize a series of flexible classification methods for estimating choice probabilities. Each model is trained to minimize either the log loss (for parametric models) or misclassification error (for nonparametric models).¹¹ Each model also has a separate set of tuning parameters which are chosen using k -fold cross validation.

Regularized Multinomial Logistic Regression Multinomial logistic regression is a generalized linear model parameterized by an inverse logit link function and an index function that is linear in consumer and market characteristics. Specifically, we write the index as:

$$v_j(\tilde{\mathbf{x}}_t, \mathbf{z}_i; \Psi_j) = \mathbf{p}'_t \Psi_j^\alpha + \mathbf{x}'_t \Psi_j^\beta + \mathbf{z}'_i \Psi_j^\gamma + \sum_{k=1}^J (p_{kt} \cdot \mathbf{z}_t)' \Psi_j^\delta \quad (18)$$

which includes main effects of all price and non-price demand shifters, and customer characteristics, as well as an interaction of price with customer characteristics. This multinomial logistic regres-

¹¹An alternative approach is to redefine the training loss function in terms of the conditional average treatment effect (CATE) rather than the outcome levels (Athey and Imbens, 2016; Hitsch and Misra, 2018). However, this "direct estimation" approach is based on the canonical setup with a binary treatment, and so the CATE is defined as the difference between two conditional expectations. Extensions to our setting are not straightforward given that the profit function is nonlinear in the price coefficient and there are more than two distinct prices in the training data.

sion parameterization departs from the random utility underpinnings of the canonical logit model outlined above in two ways. First, each product's index is a function of *all* product attributes, not just its own attributes. Second, each choice alternative has its own parameter vector Ψ_j . Given the potentially high-dimensional set of predictor variables on the right-hand side of equation (18), regularization is helpful to mitigate overfitting by encouraging sparsity in the underlying parameter vector. We implement both lasso (Tibshirani, 1996) and elastic net (Zou and Hastie, 2005) forms of regularization, with the degree of regularization being controlled by a tuning parameter.

Neural Network Neural networks are highly flexible parametric classifiers. Choice probabilities take the softmax functional form (like the logit models described above), but the conditional mean is expressed as a convolution of many nonlinear functions of the data. We estimate a neural network with one hidden layer which contains a set of nodes constructed by passing inputs through ReLU activation functions and a softmax activation function for the output layer. The number of nodes is a tuning parameter.

KNN Classifier Nearest neighbor algorithms are nonparametric classifiers with the prediction associated with a data input row $\mathcal{D}_{it} = (\tilde{\mathbf{x}}_t, \mathbf{z}_i)$ given by the majority response associated with the K nearest data points $\mathcal{D}_{i't'}$. We standardize the data and use the Euclidian measure of distance $\|\mathcal{D}_{i't'} - \mathcal{D}_{it}\|$. The number of neighbors K is a tuning parameter.

Random Forest Classification trees are nonparametric classifiers based on a partitioning of the feature space. Specifically, the data inputs $(\tilde{\mathbf{x}}_t, \mathbf{z}_i)$ are recursively split into rectangular regions such that the response classes are as homogeneous as possible within a given region. Random forests extend the standard classification tree in two ways. First, predictions are averaged across many trees. Second, to minimize correlations across tree predictions, each tree is grown using only a random subset of predictor variables. We specify the random forest to have 500 trees and let the number of predictors to use in each tree be a tuning parameter.

5.3 The Role of Heterogeneity & Data Inputs

We use the various demand models outlined above to predict profits at a set of candidate prices and then construct optimal price targeting policies. The extent of the personalization of these policies (and any subsequent profit or welfare gains therein) will crucially depend on how the shape of the demand function can change across individuals. However, the way in which this heterogeneity is operationalized across model specifications differs in a few key ways that we discuss below.

Demographics and Past Purchase Data All of the demand models we estimate account for observable heterogeneity by specifying demand as a function of customer characteristics such as demographics or past purchase histories. This approach has a long history in marketing dating back to Guadagni and Little (1983) who specify the utility index for a given brand as a function

of various past purchase and brand loyalty covariates. In brand choice models it is common to allow demographics to moderate brand preference and the sensitivity to the marketing mix (e.g., Allenby and Ginter, 1995; Ainslie and Rossi, 1998; Manchanda et al., 1999; Horsky et al., 2006).

Whether these observable characteristics alone are sufficient for capturing heterogeneity will depend on both the richness of the characteristics and the way in which they enter the demand function. For example, there is evidence that traditional demographic variables like income, age, or education have limited power in explaining brand preference heterogeneity within brand choice models (Mittal, 1994; Rossi et al., 1996; Fennell et al., 2003; Dubé et al., 2008). However, most prior papers only allow for a limited set of demographics to enter in a restrictive fashion. We revisit the role of demographics and purchase history information by estimating models with different flexible functional forms of demand as well as different sets of demographics and purchase history summaries. Notably, some of the models we estimate naturally allow for non-linear effects and interactions between different customer characteristics.

Unobserved Heterogeneity In contrast to the machine learning models outlined above, Bayesian hierarchical models not only allow for *observed* heterogeneity, but also *unobserved* heterogeneity (Rossi and Allenby, 1993; Allenby and Rossi, 1999). The unit-level parameters associated with a given customer will be informed by both the hierarchical prior $f(\psi_i|\mathbf{z}_i, \Delta, V_\psi)$ and the unit-level likelihood function $L(\mathbf{Y}_i|\tilde{\mathbf{x}}_i, \psi_i)$. If two customers share the same demographic profile but have different purchase histories, then we will infer that there are differences in their preferences through the differences in the respective likelihood functions. Therefore, we do not need to allow summary measures of purchase histories to directly enter the Bayesian hierarchical choice model. Instead, the influence of purchase histories is driven by the calculation of the posterior distribution of preference parameters as shown in (17).

6 Results

We fit the five different classes of models outlined in Section 5 to the data. For the Bayesian hierarchical models, we include both normal and mixtures of normal distributions of heterogeneity, as well as a simple “pooled” logit model that does not allow for unobserved heterogeneity. For the regularized regressions, we include both the lasso and elastic net penalties. This results in a total of eight different models: Bayesian hierarchical logit (with normal and mixtures of normals heterogeneity and without unobserved heterogeneity), regularized regressions (lasso and elastic net), neural network, KNN classifier, and a random forest. All models are fit in R (R Core Team, 2020) using standard packages. The Bayesian hierarchical models are estimated using the `bayesm` package (Rossi, 2019), which includes an MCMC sampler for hierarchical logit models with mixtures of normals heterogeneity. All machine learning models are fit using the `caret` package (Kuhn, 2020), which provides a unified syntax for training and tuning a wide class of machine learning models via k -fold cross validation (we use $k = 10$). We provide more details on the relevant tuning parameters

and our choice of the tuning grid for each machine learning model in Appendix B.

6.1 Policy Profits and Model Fit

The focus of our analysis is the expected per-customer profit associated with the targeted pricing policy induced by each model of demand. For ease of exposition, we rescale profit numbers by a factor of 100 and so the average profits can either be interpreted as dollar profits per 100 customers or profits per customer in cents. As a benchmark, we first calculate profits from a “no coupon” and a “blanket coupon” strategy. These two degenerate targeting strategies yield profits of 4.52 and 5.53 respectively.¹² We would expect any targeting strategy to outperform the blanket targeting strategy by selecting only a subset of customers for which the discount generates a large enough effect to offset the reduction in mark-up when offering a coupon. However, because we separate the derivation of the targeting strategy from the profit evaluation which is implemented on a test sample that is not used in estimation, it is not guaranteed that a candidate pricing policy outperforms a blanket coupon. We report the estimated average per-customer profit numbers across price targeting strategies derived from different models in Table 5. Block bootstrapped standard errors (using households as blocks and based on 500 bootstrap replicates) are reported in parentheses.

We find that the profitability of targeting policies varies dramatically across model specifications and data inputs. For the machine learning models, in particular, the data inputs play a key role. When customer characteristics only include base demographic variables, only the elastic net and the neural network outperform the blanket coupon. When we add in more demographic variables, profitability increases only marginally and even degrades the performance of the neural network and random forest. When we add in purchase histories, however, we find sizeable improvements across all machine learning models. In particular, the two regularized regressions and the neural network perform best with policies that yield average profits of 6.33, representing a 14.5% gain relative to the blanket coupon. The standard Bayesian hierarchical logit model with normal heterogeneity and only base demographic inputs generates a profit of 6.47 and thus represents a 17% gain in profits relative to a blanket coupon. Adding a more flexible 5-component mixture of normals distribution of heterogeneity improves the profitability slightly with a 17.5% gain relative to the blanket coupon. Similar to most of the machine learning models, performance decreases when adding additional demographic variables to the hierarchical model.

Together, these results support an earlier finding in Rossi et al. (1996) who show that purchase histories contain more valuable information than demographics in the context of generating profitable price targeting strategies. Our results show that this finding holds even after allowing for a larger set of demographic variables and allowing those variables to enter the model more flexibly. Our results also show that the value of flexible functional forms is inherently tied to the data inputs, and so the value of machine learning is only realized when detailed purchase history information is

¹²Average profits are relatively low because consumers purchase infrequently in the mayonnaise category and therefore observed profits are zero on many shopping trips.

Panel (I): Average Profits (Per 100 Customers)

	Base Demos	Base Demos + Extra Demos	Base Demos + Extra Demos + Purchase Histories
Bayesian Hierarchical Logit			
– normal heterogeneity	6.47 (0.51)	6.16 (0.50)	
– mixtures of normals heterogeneity	6.50 (0.52)	5.99 (0.49)	
Bayesian Pooled Logit	5.27 (0.45)	5.07 (0.42)	6.22 (0.52)
Lasso	5.52 (0.45)	5.71 (0.43)	6.33 (0.54)
Elastic Net	5.63 (0.45)	5.68 (0.43)	6.33 (0.52)
Neural Network	5.61 (0.48)	4.35 (0.36)	6.33 (0.49)
KNN	4.38 (0.45)	4.78 (0.41)	5.68 (0.46)
Random Forest	5.15 (0.44)	4.54 (0.31)	6.08 (0.46)
Blanket Coupon	5.53 (0.45)		
No Coupon	4.52 (0.40)		

Panel (II): Out-of-Sample Hit Probabilities

Bayesian Hierarchical Logit			
– normal heterogeneity	0.936	0.936	
– mixtures of normals heterogeneity	0.936	0.936	
Bayesian Pooled Logit	0.934	0.933	0.935
Lasso	0.934	0.934	0.935
Elastic Net	0.934	0.934	0.935
Neural Network	0.936	0.936	0.936
KNN	0.935	0.934	0.937
Random Forest	0.951	0.941	0.939

Table 5: **Targeting Policy Profits and Model Fit.** Each cell in Panel (I) reports the average customer-level profit (scaled by 100) of a targeted pricing policy. Panel (II) reports the out-of-sample hit probabilities from each model. Bootstrapped standard errors are reported in parentheses.

included as covariates. This is in contrast to Bayesian hierarchical models which naturally encode purchase history information in the unit-level likelihood, as shown in equation (17), which may be one explanation for why the Bayesian hierarchical model performs well.

For comparison, we also estimate a logit model without unobserved heterogeneity that otherwise follows the same structure as the Bayesian hierarchical logit. By turning off unobserved heterogeneity, we effectively restrict household-level parameters to be exactly equal to the population-level means and so the only differences in household-level estimates are due to differences in the observed characteristics (as in the machine learning models). We find that the pooled logit model performs worse than the two regularized regression models, however this difference is small once purchase histories are included. The regularized regressions have a similar functional form as the pooled logit, but are more flexible because they allow for all variables to enter the index of each product

rather than constraining the utility index of a particular product to only depend on characteristics of that specific product. Relaxing these utility-based restrictions seems to have only a small impact on profits. Another informative comparison is the performance of the pooled logit model with the full set of data inputs relative to the hierarchical model with unobserved heterogeneity. Both approaches allow for purchase histories to impact pricing policies, however in the model with unobserved heterogeneity, purchase information is encoded in the unit-level likelihood and driven by the overall model structure rather than being based on ad-hoc purchase history variables entering as covariates. We find that the former approach, which sidesteps the need to define a set of purchase history variables, performs better in this application.

In addition to comparing models based on profits, we also explore comparisons based on purely statistical measures of fit. In the lower panel of Table 5 we calculate the out-of-sample hit probability of each model in our evaluation sample.¹³ The hit probability is defined as the average predicted probability of chosen options in the evaluation sample and constitutes a standard statistical measure of model fit that is often used to rank model performance. Contrary to the out-of-sample profit estimator that we focus on, hit probabilities do not require us to solve for pricing policies and do not translate model performance into a decision-relevant metric such as profits. Interestingly, we find that ranking models based on hit probabilities is vastly different from our preferred ranking of models based on out-of-sample profits of the pricing policy derived from a given model. The best model in terms of hit probability is the random forest with only base demographic inputs whereas the Bayesian logit model performs relatively poorly. In terms of models that use all data inputs, the hit probability metric shows a better performance for the KNN model and the random forest, whereas based on our profit estimator the two regularized regression models perform relatively better. We also calculate a correlation coefficient of the two performance metrics across all model specifications and find that the two metrics are almost uncorrelated with a correlation coefficient of -0.148. We conclude that a fit based metrics such as the hit probability provides poor guidance with regards to selecting the appropriate model for deriving personalized prices.

6.2 Exploring Differences in Model Performance

To further explore the difference in performance across models, we report discount frequencies and the frequency of agreement between policies in Table 6. For ease of exposition we focus on the Bayesian hierarchical model with mixtures of normals heterogeneity and base demographics and the version of each machine learning model that utilizes the most comprehensive set of customer characteristics. Turning to discount frequencies first, we find that most models prescribe discounts for around 50-60% of customers. The neural network prescribes the lowest discount frequency of 50%. The remainder of the table reports the share of customers for which a pair of targeting policies prescribes the same price. We find that the logit and the two regularized regression models

¹³We are not able to calculate an out-of-sample log-likelihood for all models because the nonparametric models admit predicted choice probabilities equal to zero. Other fit statistics like the F1 score are often undefined because some models predict that at least one product will never be chosen in the test sample.

Panel (I): Frequency of Policy Agreement							
	Discount Freq.	Logit	Lasso	Elastic Net	Neural Network	KNN	Random Forest
Logit	0.558	1	0.912	0.910	0.879	0.810	0.806
Lasso	0.541		1	0.995	0.910	0.822	0.833
Elastic Net	0.543			1	0.908	0.820	0.835
Neural Network	0.502				1	0.804	0.815
KNN	0.567					1	0.807
Random Forest	0.524						1

Panel (II): Predictors of Pricing Policies							
	<i>Dependent Variable: Discount Dummy</i>						
Model	Logit	Lasso	Elastic Net	Neural Network	KNN	Random Forest	
Chain B Dummy	-0.008 (0.020)	-0.011 (0.018)	-0.003 (0.018)	0.035 (0.021)	0.038 (0.020)	0.020 (0.020)	
Income (in \$10,000s)	-0.016 *** (0.004)	-0.001 (0.003)	0.000 (0.003)	0.000 (0.004)	0.003 (0.004)	-0.005 (0.004)	
Family Size	-0.020 * (0.009)	-0.024 ** (0.008)	-0.023 ** (0.008)	-0.013 (0.010)	-0.013 (0.010)	-0.018 (0.009)	
Retired	-0.104 * (0.045)	-0.048 (0.040)	-0.043 (0.040)	-0.075 (0.048)	0.039 (0.046)	-0.027 (0.045)	
Unemployed	-0.048 (0.025)	-0.004 (0.022)	-0.014 (0.022)	-0.020 (0.027)	0.001 (0.026)	-0.043 (0.025)	
Single Mother	0.023 (0.037)	0.075 * (0.032)	0.072 * (0.033)	0.036 (0.039)	0.002 (0.037)	0.018 (0.037)	
Hellmann's Choice Share	-0.045 (0.044)	-0.017 (0.039)	-0.017 (0.039)	0.477 *** (0.047)	-0.040 (0.045)	-0.033 (0.044)	
Hellmann's Choice Count	0.016 * (0.007)	0.015 * (0.006)	0.015 * (0.006)	-0.009 (0.008)	0.018 * (0.007)	0.002 (0.007)	
Discount Share	0.882 *** (0.032)	1.066 *** (0.028)	1.068 *** (0.028)	0.786 *** (0.034)	0.989 *** (0.032)	0.991 *** (0.032)	
# Brands Purchased	-0.033 (0.026)	-0.008 (0.023)	-0.008 (0.023)	0.181 *** (0.028)	0.076 ** (0.027)	0.029 (0.026)	
Total Spending	-0.008 *** (0.001)	-0.006 *** (0.001)	-0.006 *** (0.001)	-0.002 (0.002)	0.000 (0.001)	-0.003 (0.001)	
Observations	1,318	1,318	1,318	1,318	1,318	1,318	

Table 6: **Policy Comparisons.** Panel (I) reports discount frequencies and the overlap in targeted pricing policies. Panel (II) reports a linear probability regression of discount incidence on a selected set of product, market, and customer characteristics. The unit of observation is a customer/chain combination. Significance codes: * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

agree for over 90% of customers. However, correlation with the logit does not fully account for the difference in performance. For example, KNN and the random forest have a similar rate of agreement with the logit policy, despite a difference in average profits.

In the lower panel of Table 6 we explore which customer characteristics predict whether a consumer receives a discount by regressing a discount dummy (using a linear probability model) on a set of customer characteristics separately for each pricing policy. We find that in terms of baseline demographics, there is some amount of agreement between policies generated from different models of demand. A larger family size reduces the likelihood of price discounts in all models, although the strength of the effect varies. Retirement, employment status, and single motherhood directionally affect prices similarly across most models, but the estimated effect is often not statistically significant. More generally, most demographic variables do not predict whether a discount is prescribed to a specific consumer, which is in line with the fact that demographics have a small impact on the profitability of targeting policies. With regards to the purchase history variables, we find agreement across models for some variables: consumers that frequently purchase on discount are more likely to receive a discount across all specifications of demand. Consumers with lower category spending and consumers that purchase a larger number of different brands (and are hence less loyal to a specific brand) are more likely to receive a discount, but this effect is only found in a subset of models. For the other purchase history variables, the effect tends to vary across models. We note that we regard the results from these regressions as exploratory and only as an approximation for the importance of specific variables, especially in models that are not based on linear relationships.

6.3 Machine Learning Robustness Checks

So far we have reported results from a standard implementation of machine learning methods where each model is trained and tuned through k -fold cross validation. However, the performance of machine learning models can suffer from various modeling pitfalls or data pathologies, such as including superfluous predictors or working with imbalanced data (Levin and Zahavi, 2005; Dzyabura and Yoganarasimhan, 2018; Simester et al., 2020b). In this section, we explore the extent to which addressing some of these issues can further improve the profitability of the pricing policies induced by machine learning models.

Class Imbalance We first address the fact that response classes are highly imbalanced in our data. In particular, consumers choose the outside option of “no purchase in the category” on 96% of their shopping trips. Although a high choice share for the outside good is common in many demand estimation applications, class imbalance can result in poor performance of classification algorithms trained to maximize prediction accuracy. We therefore retrain the same set of ML models discussed above on rebalanced data generated from the widely used SMOTE (Synthetic Minority Oversampling Technique) method (Chawla et al., 2002; Fernández et al., 2018). SMOTE resolves the class imbalance issue by oversampling the minority classes and undersampling the ma-

	<i>Addressing Class Imbalance</i>			<i>Feature Selection</i>	
	Base Demos	Base Demos + Extra Demos	Base Demos + Extra Demos + Purchase Histories	Only Purchase Histories	Variable Importance Criterion
Lasso	4.52 (0.41)	4.75 (0.44)	5.92 (0.46)	6.40 (0.52)	6.40 (0.54)
Elastic Net	4.52 (0.39)	4.75 (0.40)	5.92 (0.49)	6.40 (0.51)	6.48 (0.52)
Neural Network	4.52 (0.40)	4.52 (0.40)	4.57 (0.37)	6.11 (0.49)	3.62 (0.39)
KNN	4.46 (0.36)	4.63 (0.37)	5.74 (0.45)	5.59 (0.43)	5.37 (0.40)
Random Forest	5.65 (0.44)	4.92 (0.37)	5.94 (0.49)	5.95 (0.46)	4.52 (0.43)

Table 7: **Targeting Policy Profits from Additional Specifications.** Each cell reports average profits per 100 customers. The first three columns correspond to models fit on modified data sets built using the SMOTE method to correct for class imbalance. The last two columns correspond to models fit using a selected subset of predictor variables.

majority class. Oversampling adds synthetically created observations to the minority classes, whereas undersampling randomly removes observations from the majority class.¹⁴ We provide additional details on our implementation of the SMOTE method in Appendix D. In our rebalanced data, the choice shares for the outside good and inside goods are split 50/50, and the 50% choice share for the inside goods is distributed in proportion to the *conditional* choice shares in the original data. We report results of the SMOTE method in the first three columns of Table 7 (which correspond to the first three columns of results in Table 5). We find that model performance in terms of profits actually worsens in almost all cases after correcting for class imbalance.

Feature Selection We also investigate whether the machine learning models can be improved through feature selection – i.e., removing noisy or unimportant predictor variables before estimating the model. The fourth column in Table 7 reports profits for models that only include purchase history variables as customer characteristics and remove all demographic variables. This simple feature selection step is motivated by our earlier finding that demographics play a relatively small role in most models. We also implement a second, data-driven features selection approach where only the most important predictor variables (as measured by the data) are included in each model. In particular, we first compute variable importance metrics for each of the machine learning models using all three types of customer characteristics.¹⁵ For each model, we retain the top 10% of variables based on their importance and then re-run the model based on this subset of input variables. We also retain the prices of all 3 products because price is the primary variables that will determine consumer-specific sensitivity to price. The profits based on this approach are reported

¹⁴Other available techniques either focus solely on an oversampling of the minority class or an undersampling of the majority class. Due to the large imbalance in our setting, the latter would lead to a very small data set whereas the former method would create a large number of additional observations. We therefore opt for the SMOTE method which keeps the total size of the data roughly constant, but alters the share of observations with different outcomes.

¹⁵We report the top 10 most important variables from each machine learning model in Appendix C.

in the final column of Table 7. We find that both feature selection approaches lead to slight improvements in the performance of the regularized regression models, but actually degrade the performance of the KNN, random forest, and neural network models.

7 Conclusion

In this paper, we propose an approach to evaluate and compare different model-based price targeting policies. We employ an inverse probability weighted profit estimator that uses only observable profits and hence does not rely on any demand-side modeling assumptions. We apply the evaluation approach to supermarket scanner data and argue that price variation over time for a given store is as good-as-random after making adjustments to the evaluation sample. We provide supporting evidence for this assumption by showing that a rich set of observable consumer and market characteristics do not correlate with prices. We use the evaluation framework to compare the performance of different modeling approaches and data inputs in terms of generating profitable price targeting strategies.

We find large variation in performance across model specifications, ranging from a 21% reduction in profits relative to a blanket couponing strategy to a 17.5% gain in the case of the Bayesian hierarchical model. We also find that demographics at best lead to small improvements in performance and sometimes even generate lower profits relative to a blanket coupon, whereas using purchase histories as inputs leads to larger improvements across all models. Lastly, we show that the profitability of a model's targeting policy is not correlated with statistical measures of model fit, such as out-of-sample hit probabilities. We therefore caution against using statistical measures of fit to choose a "best model" and instead recommend that models be compared on the basis of a decision-relevant loss function.

References

- AINSLIE, A. AND P. E. ROSSI (1998): “Similarities in Choice Behavior Across Product Categories,” *Marketing Science*, 17, 91–106.
- ALLENBY, G. M. AND J. L. GINTER (1995): “Using Extremes to Design Products and Segment Markets,” *Journal of Marketing Research*, 32, 392–403.
- ALLENBY, G. M. AND P. E. ROSSI (1999): “Marketing Models of Consumer Heterogeneity,” *Journal of Econometrics*, 89, 57–78.
- ASCARZA, E. (2018): “Retention Futility: Targeting High-Risk Customers Might be Ineffective,” *Journal of Marketing Research*, 55, 80–98.
- ATHEY, S. AND G. IMBENS (2016): “Recursive Partitioning for Heterogeneous Causal Effects,” *Proceedings of the National Academy of Sciences*, 113, 7353–7360.
- BESANKO, D., J.-P. DUBÉ, AND S. GUPTA (2003): “Competitive Price Discrimination Strategies in a Vertical Channel Using Aggregate Retail Data,” *Management Science*, 49, 1121–1138.
- BRONNENBERG, B. J., M. W. KRUGER, AND C. F. MELA (2008): “Database Paper – The IRI Marketing Data Set,” *Marketing Science*, 27, 745–748.
- BUMBACA, F. R., S. MISRA, AND P. E. ROSSI (2020): “Scalable Target Marketing: Distributed Markov Chain Monte Carlo for Bayesian Hierarchical Models,” *Journal of Marketing Research*, 57, 999–1018.
- CAGALA, T., U. GLOGOWSKY, J. RINCKE, AND A. STRITTMATTER (2021): “Optimal Targeting in Fundraising: A Machine-Learning Approach,” *Working Paper*.
- CHAWLA, N. V., K. W. BOWYER, L. O. HALL, AND W. P. KEGELMEYER (2002): “SMOTE: Synthetic Minority Over-sampling Technique,” *Journal of Artificial Intelligence Research*, 16, 321–357.
- DENG, Y. AND C. F. MELA (2018): “TV Viewing and Advertising Targeting,” *Journal of Marketing Research*, 55, 99–118.
- DONNELLY, R., F. R. RUIZ, D. BLEI, AND S. ATHEY (2019): “Counterfactual Inference for Consumer Choice Across Many Product Categories,” *Working Paper*.
- DUBÉ, J.-P., G. J. HITSCH, P. E. ROSSI, AND M. A. VITORINO (2008): “Category Pricing with State-Dependent Utility,” *Marketing Science*, 27, 417–429.
- DUBÉ, J.-P. AND S. MISRA (2019): “Personalized Pricing and Customer Welfare,” *Working Paper*.

- DZYABURA, D. AND H. YOGANARASIMHAN (2018): “Machine Learning and Marketing,” in *Handbook of Marketing Analytics: Methods and Applications in Marketing Management, Public Policy, and Litigation Support*, ed. by N. Mizik and D. M. Hanssens, Edward Elgar Publishing.
- FENNELL, G., G. M. ALLENBY, S. YANG, AND Y. EDWARDS (2003): “The Effectiveness of Demographic and Psychographic Variables for Explaining Brand and Product Category Use,” *Quantitative Marketing and Economics*, 1, 223–244.
- FERNANDÉZ, A., S. GARCÍA, F. HERRERA, AND N. V. CHAWLA (2018): “SMOTE for Learning from Imbalanced Data: Progress and Challenges, Marking the 15-year Anniversary,” *Journal of Artificial Intelligence Research*, 16, 863–905.
- GABEL, S. AND A. TIMOSHENKO (2021): “Cross-Category Product Choice: A Scalable Deep-Learning Model,” *Management Science*, *Forthcoming*.
- GUADAGNI, P. M. AND J. D. C. LITTLE (1983): “A Logit Model of Brand Choice Calibrated on Scanner Data,” *Marketing Science*, 2, 203–238.
- GUPTA, S. AND P. K. CHINTAGUNTA (1994): “On Using Demographic Variables to Determine Segment Membership in Logit Mixture Models,” *Journal of Marketing Research*, 31, 128–136.
- HITSCH, G. J. AND S. MISRA (2018): “Heterogeneous Treatment Effects and Optimal Targeting Policy Evaluation,” *Working Paper*.
- HORSKY, D., S. MISRA, AND P. NELSON (2006): “Observed and Unobserved Preference Heterogeneity in Brand-Choice Models,” *Marketing Science*, 25, 322–335.
- JOHNSON, J., G. J. TELLIS, AND E. H. IP (2013): “To Whom, When, and How Much to Discount? A Constrained Optimization of Customized Temporal Discounts,” *Journal of Retailing*, 89, 361–373.
- KUHN, M. (2020): *caret: Classification and Regression Training*, R package version 6.0-86.
- LEVIN, N. AND J. ZAHAVI (2005): “Data Mining for Target Marketing,” in *Data Mining and Knowledge Discovery Handbook*, ed. by O. Maimon and L. Rokach, Boston, MA: Springer US, 1261–1301.
- LUNCEFORD, J. K. AND M. DAVIDIAN (2004): “Stratification and Weighting via the Propensity Score in Estimation of Causal Treatment Effects: A Comparative Study,” *Statistics in Medicine*, 23, 2937–2960.
- MANCHANDA, P., A. ANSARI, AND S. GUPTA (1999): “The “Shopping Basket”: A Model for Multicategory Purchase Incidence Decisions,” *Marketing Science*, 18, 95–114.
- MITTAL, B. (1994): “An Integrated Framework for Relating Diverse Consumer Characteristics to Supermarket Coupon Redemption,” *Journal of Marketing Research*, 31, 533–544.

- MOROZOV, I., S. SEILER, X. DONG, AND L. HOU (2021): “Estimation of Preference Heterogeneity in Markets with Costly Search,” *Marketing Science*, *Forthcoming*.
- NARAYANAN, S. AND P. MANCHANDA (2009): “Heterogeneous Learning and the Targeting of Marketing Communication for New Products,” *Marketing Science*, 28, 403–615.
- PANCRAS, J. AND K. SUDHIR (2007): “Optimal Marketing Strategies for a Customer Data Intermediary,” *Journal of Marketing Research*, 44, 560–578.
- R CORE TEAM (2020): *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria.
- ROBINS, J. M., A. ROTNITZKY, AND L. P. ZHAO (1994): “Estimation of Regression Coefficients When Some Regressors are not Always Observed,” *Journal of the American Statistical Association*, 89, 846–866.
- ROSSI, P. E. (2014a): *Bayesian Non- and Semi-parametric Methods and Applications*, Princeton University Press.
- (2014b): “Even the Rich Can Make Themselves Poor: A Critical Examination of IV Methods in Marketing Applications,” *Marketing Science*, 33, 621–762.
- (2019): *bayesm: Bayesian Inference for Marketing/Micro-Econometrics*, R package version 3.1-4.
- ROSSI, P. E. AND G. M. ALLENBY (1993): “A Bayesian Approach to Estimating Household Parameters,” *Journal of Marketing Research*, 30, 171–182.
- ROSSI, P. E., R. E. MCCULLOCH, AND G. M. ALLENBY (1996): “The Value of Purchase History Data in Target Marketing,” *Marketing Science*, 15, 321–340.
- SIMESTER, D., A. TIMOSHENKO, AND S. I. ZOUMPOULIS (2020a): “Efficiently Evaluating Targeting Policies: Improving Upon Champion vs. Challenger Experiments,” *Management Science*, 66, 3412–3424.
- (2020b): “Targeting Prospective Customers: Robustness of Machine-Learning Methods to Typical Data Challenges,” *Management Science*, 66, 2495–2522.
- TIBSHIRANI, R. (1996): “Regression Shrinkage and Selection via the Lasso,” *Journal of the Royal Statistical Society: Series B (Methodological)*, 58, 267–288.
- YOGANARASIMHAN, H., E. BARZEGARY, AND A. PANI (2020): “Design and Evaluation of Personalized Free Trials,” *Working Paper*.
- ZANTEDESCHI, D., E. M. FEIT, AND E. T. BRADLOW (2017): “Measuring Multichannel Advertising Response,” *Management Science*, 63, 2397–2771.

ZHANG, J. AND L. KRISHNAMURTHI (2004): “Customizing Promotions in Online Stores,” *Marketing Science*, 23, 561–578.

ZHANG, J. AND M. WEDEL (2009): “The Effectiveness of Customized Promotions in Online and Offline Stores,” *Journal of Marketing Research*, 46, 190–206.

ZOU, H. AND T. HASTIE (2005): “Regularization and Variable Selection via the Elastic Net,” *Journal of the Royal Statistical Society: Series B (Methodological)*, 67, 301–320.

APPENDIX

A A General Demand Modeling Framework

To highlight the variety of models that can be used to estimate demand, we write the choice probabilities as flexible functions of the data.

$$\mathbb{P}(j|\tilde{\mathbf{x}}_t, \mathbf{z}_i) = \sigma_j(v(\tilde{\mathbf{x}}_t, \mathbf{z}_i; \Psi)) \quad (19)$$

Here $v : \mathcal{X} \rightarrow \mathcal{V}$ is an index function characterized by parameters Ψ and $\sigma : \mathcal{V} \rightarrow (0, 1)^J$ maps the index into choice probabilities. Let v_j denote the j th element of $v(\cdot)$ when the range of the index \mathcal{V} is J -dimensional and let σ_j denote the j th element of $\sigma(\cdot)$. When demand is specified parametrically, we let $\sigma(\cdot)$ be the inverse logit link or “softmax” function which encompasses both economic choice models (e.g., multinomial logit) and more flexible parametric classifiers (e.g., regularized multinomial logistic regression or neural networks). Demand can also be specified nonparametrically in which case $\sigma(\cdot)$ is estimated directly and $v(\cdot)$ is the identity function.

Below, we provide a brief overview of the models we present in Section 5 and show how each model can be represented by the generalized notation outlined above.

- (i) Multinomial logit model where $\sigma(\cdot)$ is the softmax (or the inverse logit-link) function and $v(\cdot)$ maps consumer and product characteristics into utility space $\mathcal{V} = \mathbb{R}^J$.

$$\mathbb{P}(j|\tilde{\mathbf{x}}_t, \tilde{\mathbf{z}}_i) = \frac{\exp(v_j(\tilde{x}_{jt}, \mathbf{z}_i; \Psi))}{1 + \sum_{k=1}^J \exp(v_k(\tilde{x}_{kt}, \mathbf{z}_i; \Psi))} \quad (20)$$

where $v_j(\tilde{x}_{jt}, \mathbf{z}_i; \Psi) = \tilde{x}_{jt}\psi_i$ and ψ_i denotes a vector of consumer-specific parameters that are drawn from a hyper-parameter distribution that depends on \mathbf{z}_i .

- (ii) Regularized regression with an inverse logit-link function.

$$\mathbb{P}(j|\tilde{\mathbf{x}}_t, \mathbf{z}_i) = \frac{\exp(v_j(\tilde{x}_{jt}, \mathbf{z}_i; \Psi_j))}{\sum_{k=0}^J \exp(v_k(\tilde{x}_{kt}, \mathbf{z}_i; \Psi_k))} \quad (21)$$

Here each product is allowed to have a separate function with its own parameter vector Ψ_j and $v_j(\cdot)$ is a linear function of inputs. We let $\tilde{\mathbf{x}}_t$ and \mathbf{z}_i enter linearly and also allow for interactions between price and the characteristics vector \mathbf{z}_i in order to capture differential responsiveness to prices.¹⁶

$$v(\tilde{\mathbf{x}}_t, \mathbf{z}_i; \Psi_j) = p'_{jt}\Psi_j^\alpha + x'_{jt}\Psi_j^\beta + \mathbf{z}'_i\Psi_j^\gamma + (p_{jt} \cdot \mathbf{z}_i)' \Psi_j^\delta \quad (22)$$

¹⁶This specification is analogous to the Lasso regression in (Hitsch and Misra, 2018), where demographics enter on their own as well as interacted with a binary treatment indicator. In our setting, the outcome is not binary and we also allow for the influence of other demand shifters \mathbf{x}_t .

- (iii) L -hidden-layer neural network where $\sigma(\cdot)$ is the output layer activation function (which is also the softmax in our setting) and $v(\cdot)$ defines a composition of non-linear functions $g(\cdot)$. The parameters that determine the purchase probability prediction can be decomposed into the layer specific parameters: $\Psi = \{\psi_\ell, \ell = 1, \dots, L\}$.

$$\mathbb{P}(j|\tilde{\mathbf{x}}_t, \mathbf{z}_i) = \sigma_j\left(v(\tilde{\mathbf{x}}_t, \mathbf{z}_i; \Psi)\right) = \sigma_j\left(g\left(g\left(\dots g(\tilde{\mathbf{x}}_t, \mathbf{z}_i; \psi_1); \psi_{L-1}\right); \psi_L\right)\right) \quad (23)$$

More specifically, at each node of the first layer, we estimate a set of parameters ψ_{1n} (where n denotes a specific node of layer l) that generates output based on a rectified linear activation function (RELU): $g(\tilde{\mathbf{x}}_t, \mathbf{z}_i; \psi_{1n}) = \max(0, [\mathbf{p}_t \ \mathbf{x}_t \ \mathbf{z}_i]' \psi_{1n})$. At each layer beyond the first one, the inputs are given by the output values from all nodes in the previous layer. In our application we estimate a neural network with one hidden layer. The number of nodes is determined by cross-validation.

- (iv) Nonparametric models, namely classification trees and KNN classifiers where $v(\cdot)$ is an identity function with range $\mathcal{V} = \mathcal{X}$ and $\sigma_j(\cdot)$ is targeted directly.

$$\mathbb{P}(j|\tilde{\mathbf{x}}_t, \mathbf{z}_i) = \sigma_j(\tilde{\mathbf{x}}_t, \mathbf{z}_i) \quad (24)$$

B Tuning Parameters

Models	Function (<code>train=</code>)	Tuning Parameters		Tuning Grid
Lasso	<code>glmnet</code>	<code>alpha</code>	Mixing percentage	1
		<code>lambda</code>	Regularization parameter	<code>10^seq(-10, 0, len=10)</code>
Elastic Net	<code>glmnet</code>	<code>alpha</code>	Mixing percentage	{0.1, 0.2, ..., 0.9}
		<code>lambda</code>	Regularization parameter	<code>10^seq(-10, 0, len=10)</code>
Neural Network	<code>nnet</code>	<code>size</code>	Number of hidden nodes	{5, 10, 20, 25, 50}
		<code>decay</code>	Weight Decay (decay speed of optimisation)	<code>10^seq(0, -4, len=10)</code>
KNN	<code>knn</code>	<code>k</code>	Number of minimum neighbors to create association	{11, 21, 31, ..., 201}
Random Forest	<code>rf</code>	<code>ntrees</code>	Number of trees grown	500
		<code>mtry</code>	Number of randomly selected predictors (columns) per tree	{1, 2, ..., 10}

Table B.1: Overview of Tuning Parameters and Tuning Grids.

C Variable Importance for Machine Learning Models

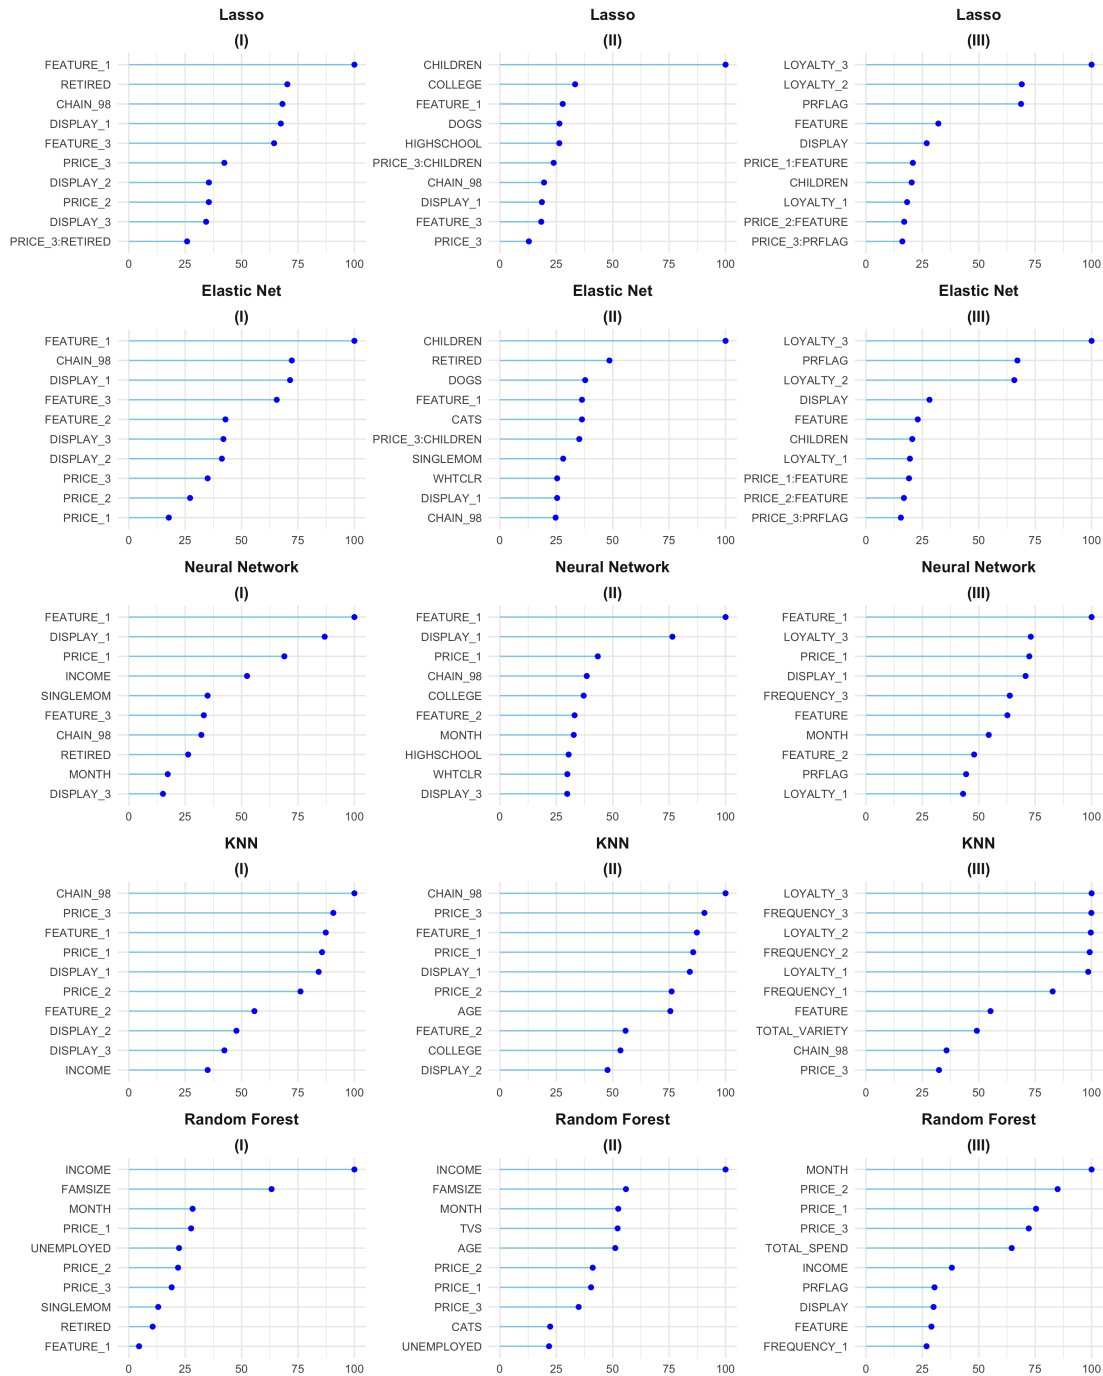


Figure C.1: **Variable Importance.** The top 10 variables (based on the variable importance metric) are reported for each machine learning model. Specification (I) corresponds to base demographics, (II) corresponds to base demographics and extra demographics, and (III) corresponds to base demographics, extra demographics, and purchase histories.

D Addressing Class Imbalance using SMOTE

In this section, we provide some additional details on the implementation of the SMOTE method to address class imbalance (Chawla et al., 2002). The general approach is described in detail in Section 2.2 of Fernández et al. (2018). The approach consists of selecting a minority class observation at random from the data. We then find its K nearest neighbors, choose N of the K instances randomly and calculate the difference between the feature vector for the observation under consideration and each of the selected neighbors. This difference is multiplied by a random number between 0 and 1, which we then add to the feature vector of the original observation. This step generates N new synthetic observations of the minority class. The step is repeated until the desired sample size for a specific class is achieved. The majority class is undersampled by randomly removing observations.

In our setting, we have 4 possible outcomes. The choice of each of the three available brands and the outside option. The latter is chosen in 96% of all observations and hence is over-represented relative to all other outcomes. We resample based on the method outline above in such a way that the conditional brand shares remain the same, but the relative importance of the outside option is diminished. In our resampled data only 50% of observations belong to the majority class (the outside option) relative to 96% in the original data.