

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 study the profitability of personalized pricing policies in a setting with consumer-level panel data. To compare pricing policies, we propose an inverse probability weighted estimator of profits, discuss how to handle non-random price variation, and show how to apply it in a typical consumer packaged good market with supermarket scanner data. We generate pricing policies from Bayesian hierarchical choice models, regularized regressions, neural networks, and nonparametric classifiers using different sets of data inputs. We find that the performance of machine learning methods is highly varied, ranging from a 30.7% loss to a 14.9% gain relative to a blanket couponing strategy, whereas hierarchical models generate profit gains in the range of 13-16.7%. Across all models, information on consumers' purchase histories leads to large improvements in profits, while demographic information only has a small impact. We find that out-of-sample fit statistics are uncorrelated with profit estimates 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: July 14, 2022

The authors are listed in reverse alphabetical order. We thank Rob Donnelly, Elea Feit, Jim Griffin, Bruce Hardie, Günter Hitsch, Yufeng Huang, Unnati Narang, and Max Pachali for helpful comments, as well as seminar participants at CREST, the Frankfurt School of Finance and Management, Penn State, Temple University, UCL, UC Riverside, University of Illinois at Urbana-Champaign, University of Rochester, Wirtschaftsuniversität Wien, and the European Quant Marketing Seminar, as well as participants at the 2021 Marketing Science Conference, the 26th International Panel Data Conference, the KDD '21 Workshop on Machine Learning for Consumers and Markets, the 2021 CESifo digitization area meeting, and DAGStat 2022. 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.

1 Introduction

The personalization of marketing instruments has long been a focus of marketing research and practice. Typical approaches for deriving personalized marketing policies such as targeted prices require the specification of a model of demand and data on customer characteristics to account for heterogeneous tastes. A canonical example in the marketing literature is Rossi et al. (1996) who estimate a Bayesian hierarchical choice model and use it to derive optimal targeted coupons. More recently, advances in machine learning have allowed researchers to estimate more flexible models by relaxing functional form assumptions and by allowing high-dimensional vectors of customer characteristics to enter as covariates that moderate customers’ price sensitivities. In this paper, we develop a profit evaluation framework and use it to quantify the economic value of model flexibility and data inputs when designing targeted pricing policies.

We make two contributions. First, we develop an evaluation framework based on an inverse probability weighted estimator of profits which is independent of demand-side modeling assumptions and can be used to compare expected profits across any set of targeting policies. This approach allows us to compare models directly on the decision-relevant metric (i.e., profits) rather than a statistical measure of model 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. We extend this approach to the common setting where a company has access to panel data on past interactions with customers and prices may not vary randomly. We first derive the profit estimator, highlight its main assumptions, and discuss how a company can test these assumptions in practice. We then show how the framework can be tailored to a consumer-packaged good (CPG) setting such as the one we use in our empirical application. While more narrow in focus, the CPG application shares many features with other retail markets as well as online platforms, and therefore provides value beyond traditional brick-and-mortar grocery retail.

Our second contribution is to derive targeting strategies from a variety of models and then use the proposed evaluation framework to compare the performance of these models in terms of out-of-sample profits. We use household scanner data from the mayonnaise product category and fit a large set of economic and machine learning models including Bayesian hierarchical choice models, regularized regressions, neural networks, K -nearest neighbor (KNN) classifiers, and bagged/boosted classification trees. For each model, we measure the value of different customer characteristics such as demographic and purchase history variables. We find that the profitability of the induced targeting policies varies substantially across different model specifications, but several key patterns emerge from our empirical analysis.

First, demographic variables only have a small impact on profits whereas information on consumers’ purchase behavior — either encoded in a unit-level likelihood in a Bayesian hierarchical model or as covariates in a machine learning model — tends to be more valuable for designing personalized pricing policies. In several of the machine learning models we find that only using demographic information to account for preference heterogeneity leads to pricing policies that un-

derperform relative to a blanket coupon. These patterns reaffirm a key result in Rossi et al. (1996) who show that demographic variables provide limited value with regard to the profitability of personalized pricing strategies. Our findings establish that this remains true even after including more demographic variables and allowing them to enter demand more flexibly.

Second, we find that our profit-based measure of model performance ranks models very differently than standard statistical measures of fit. Specifically, the correlation between our preferred out-of-sample profit metric and either out-of-sample hit probabilities or log-likelihood values is close to zero across all model specifications. The latter metrics assess model performance purely in terms of predictive power, do not require us to solve for the optimal pricing policy, and do not evaluate model performance in economic terms. Our findings caution against using simple statistical measures of fit to evaluate models. Instead, models should be compared on the basis of a loss function matching the decision problem at hand.

Third, we find that the most profitable pricing policy is generated by a hierarchical logit model, which outperforms the blanket coupon by 16.7%. Most machine learning models generate slightly smaller profit increases of 10–15% if purchase history variables are included as covariates, although estimated profits from several machine learning models are not statistically distinguishable from those generated based on hierarchical logit models. Despite the similarity in performance, we find that Bayesian hierarchical models offer attractive “off-the-shelf” performance — they neither require parameter tuning nor an optimal choice of data inputs to capture consumer heterogeneity. Across all specifications, hierarchical models generate profit gains in the range of 13–16.7%, whereas the performance of machine learning methods ranges from -30.7% to 14.9% depending on data inputs.

Our paper builds upon a large literature on personalized marketing strategies. Within the context of targeted pricing, existing approaches for estimating the profitability of personalized policies 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 from customized coupons (Zhang and Krishnamurthi, 2004; Pancras and Sudhir, 2007; Zhang and Wedel, 2009; Johnson et al., 2013; Howell et al., 2016; Donnelly et al., 2021; Gabel and Timoshenko, 2021) and retail prices (Besanko et al., 2003; Morozov et al., 2021). Similar approaches have also been used to evaluate gains from targeted advertising (Zantedeschi et al., 2017; Deng and Mela, 2018), detailing (Narayanan and Manchanda, 2009), and fundraising (Bumbaca et al., 2020). We deviate from this approach in the *evaluation* stage by using an inverse probability weighted estimator of profits that does not depend on any demand-side model structure and only assesses predicted profits based on observed demand 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 targeted 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 a uniform pricing benchmark. 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 “off-policy” approach for evaluating the profitability of targeting policies that were not observed in the data. The combination of randomized-by-action experiments and off-policy estimators has also been used by Ascarza (2018), Hitsch and Misra (2018), Yoganarasimhan et al. (2020), Yang et al. (2020), Cagala et al. (2021), and Liu (2022). We contribute to this literature by extending off-policy estimators — specifically inverse probability weighted estimators — to a setting with panel data and non-random treatment assignment of prices.

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 present the data and descriptive statistics. Section 4 presents details on the various demand models we estimate and Section 5 reports results regarding the performance of each model’s pricing policy. We offer concluding remarks in Section 6.

2 Evaluating the Profitability of Price Targeting Policies

In this section, we provide a framework to evaluate different pricing policies. We consider a setting where a company has access to panel data for a set of customers with non-experimental price variation and the objective is to find optimal prices for the same set of customers in order to offer personalized prices or coupons to those customers in the future.

A common 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 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. In this section we

take the pricing policy as given and show how to estimate expected profits. In Section 4 we discuss how we derive different pricing policies from different models of demand.

Our approach shares many features with the inverse probability weighted estimation framework employed by Hitsch and Misra (2018) and Yoganarasimhan et al. (2020), but extends it to a setting with non-random variation in prices and panel data on past behavior of customers. This setting is common in many markets where a seller repeatedly interacts with customers such as CPG markets and many online platforms. We first provide a general framework, discuss the assumptions required to obtain an unbiased profit estimator, and provide guidance on how the relevant assumptions can be tested in a given data set. We then show how our framework can be applied to evaluate targeting policies in CPG markets.

2.1 Targeting Policies and Profits

We define a targeting policy as the function $d : \mathcal{Z} \rightarrow \mathcal{P}$ which maps observed customer characteristics into prices. The firm offers a personalized price for each customer $i = 1, \dots, N$ and the targeting policy indicates the price $p \in \mathcal{P}$ at which a customer with attributes $\mathbf{z}_i \in \mathcal{Z}$ should be targeted.¹ The average per-customer profit of a targeting policy d is given by:

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

$$\pi_i(p) = Y_i(p) \times [p - mc], \quad (2)$$

where $\pi_i(p)$ refers to *potential* profits at price p and is equal to potential demand $Y_i(p)$ times the markup $[p - mc]$.² We aim to evaluate different targeting policies based on the expected average per customer profit $\mathbb{E}[\Pi(d)]$ of a given pricing policy.

We can calculate expected profits conditional on the characteristics of all consumers in the target population $\mathbf{Z} = \{\mathbf{z}_1, \dots, \mathbf{z}_N\}$ as follows:

$$\mathbb{E}[\Pi(d)|\mathbf{Z}] = \frac{1}{N} \sum_{i=1}^N \sum_{p \in \mathcal{P}} \mathbf{1}(d(\mathbf{z}_i) = p) \times \mathbb{E}[\pi_i(p)|\mathbf{z}_i]. \quad (3)$$

The expected profit expression in equation (3) matches the one in Hitsch and Misra (2018) who evaluate targeting policies based on cross-sectional data. The key difference in our case is not the objective function for the targeting policy, but the fact that we use panel data to construct an estimator of the expression in equation (3), which we turn to next.

¹For ease of exposition we assume that the customer characteristics used for targeting do not vary over time. All of our derivations generalize to targeting policies based on time-varying customer characteristics such as the identity of the product purchased on the last purchase occasion.

²We assume that potential profits for consumer i only depend on the price she is exposed to and not prices of other consumers.

2.2 Inverse Probability Weighted Estimator

To provide an estimator of expected profits, we use panel data on consumer purchases. We let p_{it} denote the price of the focal product for a specific consumer and time period. We also introduce an additional set of variables \mathbf{x}_t which denotes time-varying market characteristics such as prices and feature or display advertising of competing goods. Because we construct the profit estimator based on panel data, consumers will in general be exposed to different market characteristics and prices on different purchase occasions.

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 only on the 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 $\pi_i(p)$ for all possible prices through model-based extrapolation. In order to avoid making demand-side modeling assumptions, we need to confine ourselves to “usable” observations that do not require us to make 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:

$$e_p(\mathbf{z}_i, \mathbf{x}_t) = \mathbb{E}[\mathbf{1}(p_{it} = p) | \mathbf{z}_i, \mathbf{x}_t] = \mathbb{P}(p_{it} = p | \mathbf{z}_i, \mathbf{x}_t), \quad 0 < e_p(\mathbf{z}_i, \mathbf{x}_t) < 1, \quad (4)$$

where consumer characteristics \mathbf{z}_i and the characteristics of the market environment \mathbf{x}_t (e.g., competitor prices, feature or display advertising) explain the propensity with which a specific customer i is exposed to a given price in a specific week. The common support assumption $0 < e_p(\mathbf{z}_i, \mathbf{x}_t) < 1$ requires each price level $p \in \mathcal{P}$ to have positive support conditional on \mathbf{z}_i and \mathbf{x}_t . There are several important cases where the propensity score expression can be simplified. If prices are randomly allocated over both time and consumers, then the propensity score is independent of \mathbf{z}_i and \mathbf{x}_t . If prices vary over time but are not targeted at the consumer-level, then the propensity score is only a function of market characteristics \mathbf{x}_t , but not customer characteristics \mathbf{z}_i .³

The inverse probability weighted estimator is given by:

$$\hat{\Pi}(d) = \frac{1}{N} \sum_{i=1}^N \frac{1}{T_i} \sum_{t \in \mathbf{T}_i} \sum_{p \in \mathcal{P}} \mathbf{1}(d(\mathbf{z}_i) = p) \times \frac{\mathbf{1}(p_{it} = p)}{e_p(\mathbf{z}_i, \mathbf{x}_t)} \times \tilde{\pi}_{it}(p) \quad (5)$$

where \mathbf{T}_i denotes the set of time periods during which consumer i visit the store, $T_i = |\mathbf{T}_i|$ denotes

³If the firm knows the determinants of its pricing policy then the propensity score, i.e., how prices depend on other variables, is also known. If the policy is not known (as is the case in our empirical application), then propensity score needs to be estimated.

the total number of store visits by customer i , and $\tilde{\pi}_{it}(p)$ denotes per-period profits. For each consumer, we sum over all visits of that customer and the $1/T_i$ scaling ensures that all consumers are weighted equally. We assume that the distribution of \mathbf{x}_t in the data is identical to the distribution of market characteristics when the targeting policy is implemented⁴ and therefore it follows that $\mathbb{E}_{\mathbf{x}} [\mathbb{E} [\tilde{\pi}_{it}(p)|\mathbf{x}_t]] = \mathbb{E} [\pi_i(p)]$.

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:

$$\begin{aligned}
\mathbb{E}[\hat{\Pi}(d)|\mathbf{Z}, \mathbf{X}] &= \frac{1}{N} \sum_{i=1}^N \frac{1}{T_i} \sum_{t \in \mathbf{T}_i} \mathbb{E} \left[\sum_{p \in \mathcal{P}} \mathbf{1}(d(\mathbf{z}_i) = p) \times \frac{\mathbf{1}(p_{it} = p)}{e_p(\mathbf{z}_i, \mathbf{x}_t)} \times \tilde{\pi}_{it}(p) | \mathbf{z}_i, \mathbf{x}_t \right] \\
&= \frac{1}{N} \sum_{i=1}^N \frac{1}{T_i} \sum_{t \in \mathbf{T}_i} \sum_{p \in \mathcal{P}} \mathbf{1}(d(\mathbf{z}_i) = p) \times \mathbb{E} \left[\frac{\mathbf{1}(p_{it} = p)}{e_p(\mathbf{z}_i, \mathbf{x}_t)} \times \tilde{\pi}_{it}(p) | \mathbf{z}_i, \mathbf{x}_t \right] \\
&= \frac{1}{N} \sum_{i=1}^N \frac{1}{T_i} \sum_{t \in \mathbf{T}_i} \sum_{p \in \mathcal{P}} \mathbf{1}(d(\mathbf{z}_i) = p) \times \frac{e_p(\mathbf{z}_i, \mathbf{x}_t)}{e_p(\mathbf{z}_i, \mathbf{x}_t)} \times \mathbb{E} [\tilde{\pi}_{it}(p) | \mathbf{z}_i, \mathbf{x}_t] \\
&= \frac{1}{N} \sum_{i=1}^N \frac{1}{T_i} \sum_{t \in \mathbf{T}_i} \sum_{p \in \mathcal{P}} \mathbf{1}(d(\mathbf{z}_i) = p) \times \mathbb{E} [\tilde{\pi}_{it}(p) | \mathbf{z}_i, \mathbf{x}_t]. \tag{6}
\end{aligned}$$

where $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_t, \dots\}$ denotes market characteristics in all time periods. The third equality follows from the assumption that $\mathbf{1}(p_{it} = p)$ and $\pi_{it}(p)$ are conditionally independent given $(\mathbf{z}_i, \mathbf{x}_t)$.

We note that even though the targeting strategy is only based on \mathbf{z}_i we nevertheless require the conditional independence assumption to hold also with regard to \mathbf{x}_t . To see why conditioning on \mathbf{x}_t is necessary, consider the case where lower prices are correlated with advertising for the focal product, i.e., potential demand in low price periods is higher and therefore makes low prices seem more profitable than they actually are. Such a pattern would violate the conditional independence assumption and lead to a biased profit estimator.

⁴A company could decide to implement a targeted pricing strategy in time periods with a specific realization of \mathbf{x}_t such as a targeting strategy during a specific time of the year. It is also straightforward to apply our framework to such a setting by conditioning the profit estimand and estimator on the relevant realization of \mathbf{x}_t .

Next, we take expectations over market characteristics \mathbf{x}_t :

$$\begin{aligned}
\mathbb{E}[\hat{\Pi}(d)|\mathbf{Z}] &= \mathbb{E}_{\mathbf{x}} \left[\mathbb{E}[\hat{\Pi}(d)|\mathbf{Z}, \mathbf{x}] \right] \\
&= \frac{1}{N} \sum_{i=1}^N \frac{1}{T_i} \sum_{t \in \mathbf{T}_i} \sum_{p \in \mathcal{P}} \mathbf{1}(d(\mathbf{z}_i) = p) \times \mathbb{E}_{\mathbf{x}} \left[\mathbb{E}[\tilde{\pi}_{it}(p)|\mathbf{z}_i, \mathbf{x}_t] \right] \\
&= \frac{1}{N} \sum_{i=1}^N \frac{1}{T_i} \sum_{t \in \mathbf{T}_i} \sum_{p \in \mathcal{P}} \mathbf{1}(d(\mathbf{z}_i) = p) \times \mathbb{E}[\pi_i(p)|\mathbf{z}_i] \\
&= \frac{1}{N} \sum_{i=1}^N \sum_{p \in \mathcal{P}} \mathbf{1}(d(\mathbf{z}_i) = p) \times \mathbb{E}[\pi_i(p)|\mathbf{z}_i] \\
&= \mathbb{E}[\Pi(d)|\mathbf{Z}], \tag{7}
\end{aligned}$$

where the third line follows from the law of iterated expectations and the assumption that the distribution of \mathbf{x}_t in the data is identical to the distribution of market characteristics when the targeting policy is implemented and therefore $\mathbb{E}_{\mathbf{x}} [\mathbb{E}[\tilde{\pi}_{it}(p)|\mathbf{z}_i, \mathbf{x}_t]] = \mathbb{E}[\pi_i(p)|\mathbf{z}_i]$. The second-to-last line follows from the fact that $\sum_{p \in \mathcal{P}} \mathbf{1}(d(\mathbf{z}_i) = p) \times \mathbb{E}[\pi_i(p)|\mathbf{z}_i]$ 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.⁵ 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: $\frac{1}{T_i} \sum_{t \in \mathbf{T}_i} \sum_{p \in \mathcal{P}} \mathbf{1}(d(\mathbf{z}_i) = p) \times [\mathbf{1}(p_{it} = p)/e_p(\mathbf{z}_i, \mathbf{x}_t)] \times \tilde{\pi}_{it}(p)$. 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 Discussion of Assumptions

In this section, we discuss the necessary assumptions for the profit estimator in (5) to be unbiased. In doing so, we also illustrate how a firm may assess the suitability of an inverse probability weighted

⁵For 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 3.4 we show that this condition holds for the customer characteristics considered in our setting.

estimator of profits in a given price targeting setting. Lastly, we highlight possible limitations and make comparisons with model-based approaches to evaluating profits.

Conditional Independence and Common Support 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 variables $(\mathbf{z}_i, \mathbf{x}_t)$ such that price variation can be treated as-good-as-random conditional on $(\mathbf{z}_i, \mathbf{x}_t)$. The profit estimator also requires all price levels to have support in the data conditional on $(\mathbf{z}_i, \mathbf{x}_t)$. If a specific price level is never observed conditional on a specific realization of $(\mathbf{z}_i, \mathbf{x}_t)$, then $e_p(\mathbf{z}_i, \mathbf{x}_t) = 0$ and the profit estimator is not defined. A firm that aims to implement our profit evaluation framework will likely have detailed information on how prices were set in the past and can hence assess which variables enter its pricing policy. In this section, we discuss conditional independence and the common support condition separately for consumer characteristics \mathbf{z}_i and variables describing the market environment \mathbf{x}_t . The latter set of variables is unique to a panel data setting such as ours and not relevant when working with cross-sectional data.

With regard to consumer characteristics, our discussion closely follows the one in Hitsch and Misra (2018) who propose a similar profit estimator in the context of cross-sectional data. Ignoring market characteristics for a moment, if prices are randomly assigned to different customers, then price exposure is unconditionally independent of profits and $e_p(\mathbf{z}_i) = \bar{e}_p$. If the targeting policy is known, then the propensity score needs to be modeled as a function of the customer characteristics that enter the targeting policy. Moreover, the support condition needs to hold, which rules out deterministic targeting policies as a function of demographics. More specifically, conditional on \mathbf{z}_i all prices need to have a positive probability in the sample.

Next, we turn to the market environment \mathbf{x}_t and assume for the moment that prices are not targeted at the consumer-level and hence \mathbf{z}_i does not enter the propensity score. This particular constellation is likely to be relevant in many retail settings where prices in historic data vary over time due to temporary price discounts that are posted on the webpage or on a shelf in a physical store and therefore apply to all consumers. It also applies to our empirical application to purchase data in CPG markets and many studies of price targeting in similar settings (e.g., Rossi et al., 1996). Similar to the consideration with regard to customer characteristics outlined above, a firm will need to use information on how prices were set in the past to determine the set of factors that enter the propensity score. For example, if price discounts are more likely around the holidays, when demand is generally higher, then a holiday indicator would need to be included in the propensity score. Generally, either internal information on past price setting or a researcher’s institutional knowledge of a specific market will need to inform which variables are deemed necessary for the conditional independence assumption to hold and that should therefore enter the propensity score.

A second consideration with regard to market characteristics that is likely relevant in many settings (including our empirical application) is the possibility of perfect/almost-perfect correlations between different types of marketing activity which limits the applicability of our framework. For

example, our framework does not allow 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 and hence the common support assumption is violated. However, often the nature of price setting in historic data will often continue to apply to future price setting strategies the firm is contemplating. If a firm historically never offered price discounts at the same time a competitor offered a discount, this aspect of the pricing policy will likely continue in the future.

While we believe that our framework can be applied in most settings relevant to companies, the possible violations of the common support assumption illustrate the key trade-off when using a profit evaluation framework that does not impose any demand-side modeling assumptions. Namely, we are not able to extrapolate to pricing patterns that are never observed in the data. If a company wants to analyze policies for which the common support assumption is violated in historic data, the firm would need to run an A/B test that generates the requisite price variation and then apply our evaluation framework to the newly collected data.

Static Demand In order to apply the evaluation framework we also need to assume that current period profits constitute the appropriate objective function for evaluating price targeting strategies. Our approach therefore does not account for demand dynamics such as learning or stockpiling. The presence of demand dynamics would impact our evaluation framework in two ways that a researcher needs to be aware of. First, demand dynamics will lead to past prices affecting current demand and hence past prices have to be treated in a similar fashion as contemporaneous market characteristics \mathbf{x}_t . Therefore, if past prices correlate with current prices, then past prices will need to be included in the propensity score. Second, single-period profits do not capture the full profit impact of different pricing policies if future demand is impacted by current choices. Unfortunately, accounting for changes in future profits is difficult without additional assumption on the nature of demand and therefore difficult to accommodate in our framework.⁶ We therefore focus on price setting that maximizes current period profits throughout this paper, but acknowledge that not accounting for inter-temporal dependencies might be a strong assumption in some settings. In Appendix A we provide a more detailed discussion of dynamic demand in the context of our profit evaluation framework.

Comparison to Other Approaches It is instructive to consider which of the assumptions outlined above are unique to our evaluation method and which assumptions are typically made also in a model-based evaluation approach à la Rossi et al. (1996). Conditional independence is an assumption that is required for an unbiased estimator of consumer utility. In the typical demand model a relevant set of control variables would be included in the model so that price variation can be treated as-good-as-random after conditioning on the relevant set of controls, and so the price

⁶Different mechanisms such as learning, stockpiling, or brand loyalty will cause different temporal dependencies in demand which are typically modeled via a structural model of demand that is tailored to a specific empirical context. The specificity of most dynamic demand models is at odds without our evaluation framework that avoids assumptions about demand.

coefficient represents a causal effect price on demand (and hence profits). The set of variables that needs to enter the propensity score is identical to the set of variables that a researcher would need to include in the demand model. While model-based approaches can in principle accommodate dynamic demand, an assumption of static demand is often made in related work (e.g., Rossi et al., 1996; Besanko et al., 2003; Donnelly et al., 2021).

The one assumption that is usually not necessary in model-based approaches is the common support assumption. A typical model of demand would require residual variation in price after controlling for certain variables. The common support assumption instead is more restrictive and requires all price levels to occur for each value of the variables that enter the propensity score. Common support therefore constitutes the only additional assumption relative to model-based methods. We show below that in our empirical application the common support assumption does impose some restrictions on the pricing policies that we can analyze. However, we believe that in most cases the restrictions created by the common support assumption do not limit the set of pricing policies that are of interest to a firm.

2.4 Application to CPG panel data

Next, we illustrate how to operationalize the profit evaluation framework to the specific setting of CPG markets, which we study in our empirical application. In particular, we discuss the role of the common support and conditional independence assumptions outlined in the previous section. We believe that many of the institutional features and data patterns we rely on for our CPG application also apply to other retail settings including online platforms. While some of the steps we outline are motivated by patterns specific to our data set, the general ideas presented here are likely to be more broadly applicable.

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: (i) when a given product is discounted, it is also almost always featured and displayed while other competitor products in the same category tend to not be discounted, featured, or displayed at the same time; and (ii) when a given product is not discounted, it is rarely featured or displayed. We report descriptive statistics in Section 3.1 providing evidence for these type of promotion patterns in our data. As discussed in the previous section, 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.

In order to ensure that the common support assumption holds, we 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 promotion. 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. We also later test whether observable market characteristics such as competitor prices predict whether the focal product is discounted and find that only the chain-identity of the store predicts discount frequencies. We therefore model the propensity score as a function of chain identity, but not any other time-varying market variables.

3 Data

For our 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 22 months of data (January 2011 – October 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 confine our sample to households that purchase in the mayonnaise category at least three times and make at least one grocery shopping trip in each month of the sample period. We retain all trips regardless of whether a consumer purchased any product in the category. In total, we observe 581 households for 95,949 trips and 3,394 purchases in the category. For the average (median) household we observe 165 (135) trips and 5.84 (5) purchases. In our analysis, we take the perspective of a focal manufacturer, Hellmann’s, who can engage in third-degree price discrimination by issuing coupons that offer personalized prices to each customer.

3.1 Product and Market Characteristics

Prices are identical for all products in stores that belong to the same chain and do not vary across customers for a given chain and week. Figure 1 displays the time series of chain-level prices for all three brands. We find that prices at all chain/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 weekly 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

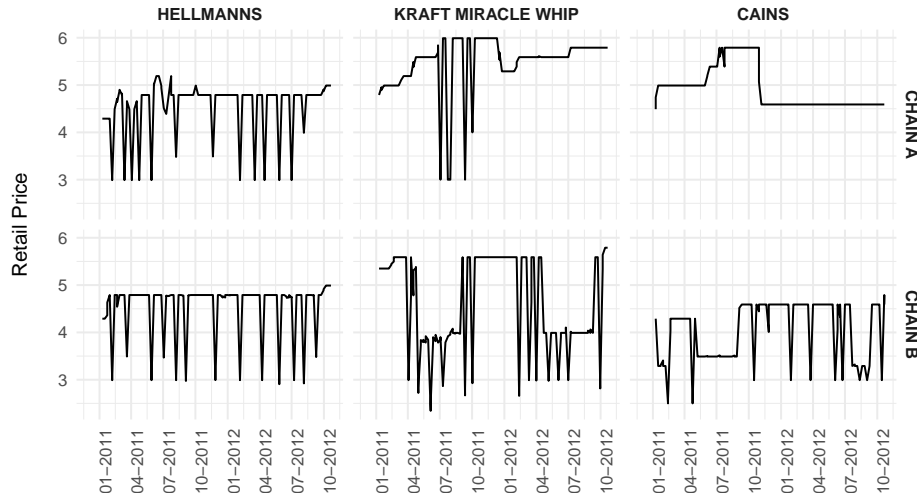


Figure 1: **Time Series of Prices.**

products. For example, Cain’s is never discounted and never featured or displayed in chain A, where chain B engages in various marketing activities for Cain’s. Both chains 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. 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. 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. In the same panel we also show 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 rarely discounted, 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. These patterns are evidence of the type of correlations between marketing activities typical in CPG markets that we discussed in Section 2.4 and form the basis for a sample selection step when generating the evaluation sample which we discuss in detail in Section 3.3.

3.2 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

Panel (I)			Median	Discount	Feature	Display	Weekly	
Brand	Chain		Price	Freq.	Freq.	Freq.	Units Sold	Market Share
Hellmann's	A		4.79	0.16	0.20	0.19	260.16	0.87
Kraft Miracle Whip	A		5.59	0.05	0.05	0.02	25.03	0.10
Cain's	A		4.59	0	0	0	10.10	0.03
Hellmann's	B		4.79	0.14	0.14	0.15	168.29	0.79
Kraft Miracle Whip	B		5.35	0.28	0.17	0.18	24.35	0.13
Cain's	B		4.50	0.16	0.09	0.15	17.67	0.08

Panel (II)		Hellmann's		Hellmann's		Other Brand		Other Brand	
		Featured		Displayed		Feat./Display		Discounted	
		no	yes	no	yes	no	yes	no	yes
Hellmann's	no	0.83	0.02	0.82	0.02	0.65	0.20	0.64	0.21
Discounted	yes	0	0.15	0.01	0.14	0.13	0.02	0.13	0.02

Table 1: **Summary of Product and Market Characteristics.** Panel (I) reports summary statistics for each brand. The unit of observation in each row is a week. Panel (II) reports the co-occurrence of Hellmann's price discounts with own and competitor promotions across chain/weeks.

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. This results in a total of 11 purchase history variables.⁷ 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. Because our analysis focuses on time-invariant customer characteristics, we do not

⁷We construct purchase history variables using all purchases that a given consumer made in the training sample (between January 2011 and March 2012).

		Mean	Std. Dev.
<u>Base Demographics:</u>	Income	52,081	29,883
	Family Size	2.54	1.18
	Retired	0.064	0.244
	Unemployed	0.251	0.434
	Single Mother	0.084	0.278
<u>Extra Demographics:</u>	Age	58.64	10.64
	High School	0.520	0.500
	College	0.451	0.498
	White Collar	0.454	0.498
	Number of Children	0.077	0.268
	Married	0.098	0.298
	Number of Dogs	0.208	0.406
	Number of Cats	0.201	0.406
	Renter	0.886	0.318
	Number of TVs	2.28	0.86
<u>Purchase History Variables:</u>	Hellmann’s Choice Share	0.787	0.380
	Kraft Choice Share	0.111	0.290
	Cain’s Choice Share	0.059	0.211
	Hellmann’s Choice Count	2.75	2.48
	Kraft Choice Count	0.489	1.548
	Cain’s Choice Count	0.231	1.209
	Discount Share	0.359	0.350
	Feature Share	0.363	0.356
	Display Share	0.479	0.361
	Number of Brands Purchased	1.08	0.42
	Total Spending	14.53	11.89

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

include variables that capture recency such as the identity of the previously purchased brand. We present descriptive statistics for all customer characteristics in Table 2.

3.3 Evaluation Sample

We form our evaluation sample based on the last 7 months of data from April to October 2012. As discussed in Sections 2.3 and 2.4 and based on the empirical patterns presented in Section 3.1, we use only a subset of chain/weeks in our evaluation sample. First, we only focus on the two most common price levels: \$4.79 and \$3.00⁸ and eliminate chain/weeks with different prices. Most other prices have relatively little support in our data and conditioning on these two price levels allows us retain 48 out of 62 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 further confine this group to chain/weeks

⁸We treat prices that are 5 cents above or below those modal prices to be identical to those two price levels.

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 displayed. After applying these selection criteria we are left with 39 chain/weeks and a total of 18,629 purchase occasions.

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.

3.4 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 (5) 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 chain is likely uncorrelated with demand shifters because manufacturers 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 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

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	0.022*** (0.007)	0.022*** (0.007)	0.021** (0.007)	0.161 (0.226)
Income (in \$10,000s)	-0.000 (0.001)	0.000 (0.001)	0.000 (0.001)	
Family Size	-0.002 (0.003)	-0.003 (0.003)	-0.004 (0.003)	
Retired	0.012 (0.012)	0.012 (0.012)	0.014 (0.012)	
Unemployed	-0.010 (0.007)	-0.011 (0.007)	-0.012 (0.007)	
Single Mother	0.002 (0.011)	0.003 (0.012)	0.006 (0.012)	
Price Kraft Miracle Whip				0.116 (0.161)
Price Cain's				0.080 (0.215)
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)	0.924	0.813	0.791	1.237
F-stat (month dummies)				1.600
F-stat (all variables)	2.753	1.506	1.203	1.102
Unit of Observation	Store Visit	Store Visit	Store Visit	Chain/ Week
Observations	18,629	18,629	18,629	39

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

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 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

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

(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 monthly dummies. 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 month dummies are jointly insignificant. We also regress the price of Hellman’s on lagged price in the previous week (not reported in the table). Past prices might impact demand in the presence of demand dynamics as discussed in Section 2.3 and Appendix A. Therefore, past prices will need to be included in the propensity score if they correlate with current prices. We find that the coefficient on past prices is not statistically significant and small in magnitude.¹¹

Based on the results from the regressions in Table 3, we only account for chain identity when constructing the propensity scores. At chains A and B, the propensity score for the regular price level of \$4.79 is equal to 0.83 and 0.81, respectively. On the remainder of store visits in the evaluation sample consumers encounter a discounted price of \$3.00.

4 Demand Estimation & Pricing Policies

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),¹² 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 and \mathbf{p}_{-jt} to denote the vector of all prices except the price of good j . 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: $\mathbb{E}[Y_{ijt}(\tilde{\mathbf{x}}_t, \mathbf{z}_i)] = \mathbb{P}(j|\tilde{\mathbf{x}}_t, \mathbf{z}_i)$. We let \mathcal{M} denote a given model chosen by the researcher to estimate choice probabilities.

¹¹When regressing current price on lagged price and a chain dummy using the evaluation sample, we find an estimated coefficient (standard error) on lagged price of -0.15 (0.17). Because the evaluation sample is relatively short, we also run the same regression for our entire sample and find an estimated coefficient (standard error) of -0.03 (0.08).

¹²We use \mathbf{x}_t to refer both to time-varying variables characterizing the market environment as well as time-invariant product characteristics. We ignored time-invariant factors in Section 2 since they are not relevant to our evaluation framework.

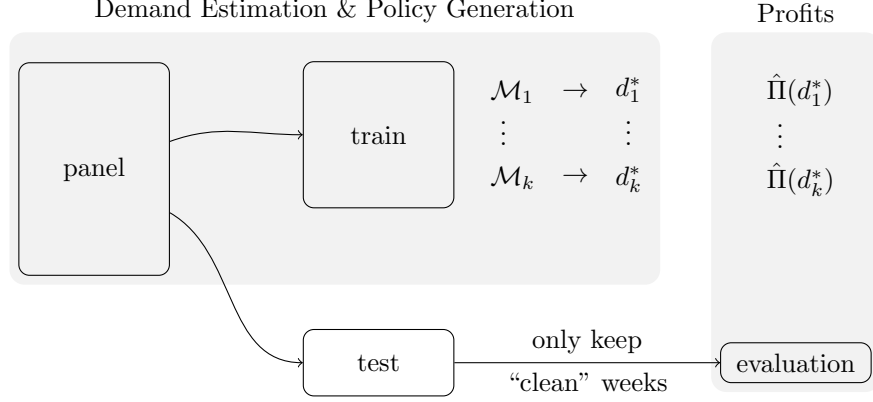


Figure 2: **Sample Selection, Demand Estimation, Pricing Policies, and Profit Evaluation.** The full panel data is first split into a training set and a test set. The training set is used to fit various models of demand \mathcal{M} and then solve for an optimal pricing policy d^* . The evaluation sample 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.

Below we first describe how pricing policies are derived given a model \mathcal{M} and then outline a set of common economic choice models and machine learning methods that can be used to estimate demand.

4.1 Constructing Price Targeting Policies

Given other market and customer characteristics $(\mathbf{p}_{-jt}, \mathbf{x}_t, \mathbf{z}_i)$ 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:

$$\pi_i(p_{ij}|\mathbf{p}_{-jt}, \mathbf{x}_t, \mathbf{z}_i, \mathcal{M}) = \mathbb{P}(j|p_{ij}, \mathbf{p}_{-jt}, \mathbf{x}_t, \mathbf{z}_i, \mathcal{M})[p_{ij} - mc_j]. \quad (8)$$

Due to uncertainty in the market environment (e.g., changes in competitor prices) at the time the coupon is issued, we integrate the profit function over the empirical distribution of demand shifters:¹³

$$\mathbb{E}_{\mathbf{p}_{-j}, \mathbf{x}}[\pi_i(p_{ij}|\mathbf{p}_{-jt}, \mathbf{x}_t, \mathbf{z}_i, \mathcal{M})] = [p_{ij} - mc_j] \int \mathbb{P}(j|p_{ij}, \mathbf{p}_{-jt}, \mathbf{x}_t, \mathbf{z}_i, \mathcal{M}) dF(\mathbf{p}_{-jt}, \mathbf{x}_t). \quad (9)$$

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 a customer with characteristics \mathbf{z}_i that maximizes expected profits:

$$d^*(\mathbf{z}_i) \equiv p_{ij}^* = \operatorname{argmax}_{p_{ij} \in \mathcal{P}} \mathbb{E}_{\mathbf{p}_{-j}, \mathbf{x}}[\pi_i(p_{ij}|\mathbf{p}_{-jt}, \mathbf{x}_t, \mathbf{z}_i, \mathcal{M})]. \quad (10)$$

¹³When deriving pricing policies from a Bayesian hierarchical model, we also integrate over the posterior distribution of individual-level parameters.

To evaluate the profitability of pricing policies, we apply the profit estimator defined in Section 2.2 on the evaluation sample which is comprised of a subset of chain-weeks (see the discussion in Section 3.3) from the last 7 months of our 22-month-long data period.¹⁴ The first 15 months constitute the training sample which is used to estimate all candidate demand models $\mathcal{M}_1, \dots, \mathcal{M}_k$. Figure 2 provides a diagram of how we construct the training and evaluation samples, and how we use both samples to estimate a series of demand models, derive pricing policies from those models, and evaluate the expected profits of each policy. Note that our sample is comprised of households that are observed in every month of the data and therefore all households are contained in both training and evaluation data sets. The temporal structure and the household overlap between training and evaluation sample mimics the decision problem the firm is facing, which is to use historic data on a given set of customers to optimize prices for the same set of customers in a future time period.¹⁵

4.2 Demand Models

Next, we describe different economic and machine learning methods that can be used to estimate demand. We focus on hierarchical choice models, regularized regressions, neural networks, nearest neighbor classifiers, and bagged/boosted classification trees. Each of these methods imposes different functional form assumptions, requires different training algorithms, and adopts different strategies to prevent overfitting — all of which can affect flexibility and predictive performance. For example, regularized regressions can only capture interactions between data inputs if they are explicitly included as covariates in the training data. In contrast, neural networks and classification trees offer more automatic procedures for discovering interactions between data inputs. In Appendix B, we provide more details on how each model structure fits within a general demand modeling framework. We also refer the reader to Hastie et al. (2009) and Murphy (2022) for comprehensive reviews on the theory and practice of machine learning, as well as Dzyabura and Yoganarasimhan (2018), Taddy (2019), and Athey and Imbens (2019) for discussions of machine learning in the context of marketing and economics.

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

¹⁴Due to the manufacturer/retailer interaction and because we do not observe marginal costs in our data, we re-write manufacturer profits as $Y_{it}(p) \times [m - (\bar{p} - p)]$ where \bar{p} denotes the regular price level, $\bar{p} - p$ denotes the face value of the coupon, 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%. We assume that any price reduction reduces the manufacturer margin 1-to-1 and does not alter the retailer margin.

¹⁵An alternative way to construct the training and evaluation sample would be to use all months up to month t to predict demand (and generate a pricing policy) for month $t + 1$. Such a procedure would maximize the data being used for prediction in each month of the evaluation sample and therefore increase power relative to our fixed sample split. For expositional purposes we refrain from such a “shifting time window” procedure which would require us to recompute purchase history variables for different training samples and which would make pricing policies consumer / time-period specific rather than prices being determined at the consumer level.

sets of customer characteristics introduced in Section 3.2: base demographics, an extended set of demographics, and extended demographics plus purchase history variables.

Bayesian Hierarchical Choice Models Bayesian hierarchical choice models are comprised of a unit-level random utility model and a random effects distribution that characterizes cross-sectional preference heterogeneity in the population. The random utility model posits that demand Y_{ijt} is the outcome of a latent utility maximization process:

$$Y_{ijt} = \begin{cases} 1 & \text{if } u_{ijt} \geq \max(u_{i0t}, u_{i1t}, \dots, u_{iJt}) \\ 0 & \text{otherwise} \end{cases} \quad (11)$$

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

$$\begin{aligned} u_{ijt} &= v_j(\tilde{\mathbf{x}}_t; \psi_i) + \varepsilon_{ijt} \\ &= \tilde{x}'_{jt} \psi_i + \varepsilon_{ijt}. \end{aligned} \quad (12)$$

Here $v_j(\tilde{\mathbf{x}}_t; \psi_i) = \tilde{x}'_{jt} \psi_i$ 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:

$$\mathbb{P}(j|\tilde{\mathbf{x}}_t, \psi_i) = \frac{\exp(\tilde{x}'_{jt} \psi_i)}{1 + \sum_{k=1}^J \exp(\tilde{x}'_{kt} \psi_i)} \quad (13)$$

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:

$$\psi_i \sim N(\Delta' \mathbf{z}_i, V_\psi). \quad (14)$$

We also estimate models with more flexible distributions of heterogeneity such as mixtures of normals (Rossi, 2014a).

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:

$$v_j(\tilde{\mathbf{x}}_t, \mathbf{z}_i; \Psi_j) = \mathbf{p}'_t \Psi_j^\alpha + \mathbf{x}'_t \Psi_j^\beta + \mathbf{z}'_i \Psi_j^\gamma + \sum_{k=1}^J (p_{kt} \cdot \mathbf{z}_i)' \Psi_j^\delta \quad (15)$$

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. Contrary to the canonical logit

model specification outlined above, 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 (15), 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 tuning parameters.

Neural Networks Neural networks are highly flexible parametric models. As classifiers, neural networks model the probability of each response class as a composition of nonlinear functions of data inputs. Model performance is affected by the choice of network architecture, regularization, and optimization algorithm. We estimate two types of feedforward neural networks: a standard single hidden layer neural network and a deep (i.e., multiple hidden layer) neural network. In both cases, we use ReLU activation functions for hidden layers and the softmax activation function for the output layer. We adopt the standard practice of choosing a relatively large number of hidden units per layer and then train the model with regularization (Hastie et al., 2009). Specifically, we use the Adam optimizer (Kingma and Ba, 2015) with a step decay learning schedule, early stopping, and dropout regularization. We let the dropout rate be 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 $\mathcal{D}_{i't'}$, $(i', t') \neq (i, t)$. Neighborhoods are defined using the Euclidian measure of distance $\|\mathcal{D}_{i't'} - \mathcal{D}_{it}\|$ and the number of neighbors K is a tuning parameter.

Classification Trees (Bagged and Boosted) 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. We consider two types of tree models. The first is a random forest (Breiman, 2001), which averages predictions across many trees grown on bootstrap replicates of the training data (“bagging”), and de-correlates predictions across trees by selecting a random subset of predictor variables at each node to use as candidates for splitting. We specify the random forest to have 500 trees and let the number of candidate predictors at each node be a tuning parameter. The second model is XGBoost (Chen and Guestrin, 2016), which represents the state-of-the-art implementation of tree “boosting.” Like bagging, boosting constructs predictions through an ensemble of trees. Unlike bagging — where separate trees are grown on subsamples of the data — boosting fits a *sequence* of trees on the complete, but modified, training data. Specifically, the residuals from the last iteration are used as the target response variable in the current iteration, which allows the model to learn slowly and gradually improve its predictions in areas where it previously did not perform well. XGBoost contains a large set of tuning parameters which together govern the number and complexity of trees, the learning rate, and the data inputs used to grow each tree.

4.3 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 personalization and therefore the profit gains of these policies crucially depend on how customer heterogeneity enters the demand function, which differs across model specifications 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 various machine learning models we estimate, Bayesian hierarchical models allow for *unobserved* heterogeneity (Rossi and Allenby, 1993; Allenby and Rossi, 1999) through a prior on ψ_i , which gives rise to individual-level posteriors. These posteriors naturally encode purchase history information via the likelihood and therefore summaries of purchase histories do not have to be constructed and included as covariates in the choice model.

To see this formally, consider the posterior distribution of individual-level parameters ψ_i obtained through Bayes' rule:

$$f(\psi_i | \mathbf{Y}_i, \tilde{\mathbf{x}}_i, \mathbf{z}_i, \Delta, V_\psi) = \frac{L(\mathbf{Y}_i | \tilde{\mathbf{x}}_i, \psi_i) f(\psi_i | \mathbf{z}_i, \Delta, V_\psi)}{\int L(\mathbf{Y}_i | \tilde{\mathbf{x}}_i, \psi_i) dF(\psi_i | \mathbf{z}_i, \Delta, V_\psi)} \quad (16)$$

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 (13), and $f(\psi_i | \mathbf{z}_i, \Delta, V_\psi)$ is the distribution of heterogeneity described in equation (14), where \mathbf{z}_i denotes demographic variables, but not purchase histories. Inferences about a customer's preference vector ψ_i will be shaped by

both their own purchase history data (the likelihood) as well as the distribution of preferences in the population (the prior), with the weights on each component being driven by the signal in the data. Intuitively, when little purchase information is available for a customer, the corresponding 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 that customer’s sample information. We re-iterate that because purchase histories are encoded in the likelihood function, we do not need to include the constructed purchase history variables as covariates in the indirect utility specification of the economic choice model.

5 Results

We fit the five different classes of models outlined in Section 4 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 lasso and elastic net penalties. We estimate a neural network with one hidden layer as well as a deep neural network where the number of hidden layers is a tuning parameter.¹⁶ This results in a total of ten different models: Bayesian hierarchical logit (with normal and mixtures of normals heterogeneity), Bayesian pooled logit, lasso regression, elastic net regression, neural network, deep neural network, KNN, random forest, and XGBoost.

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 models via cross-validation. We provide more details on the relevant tuning parameters and our choice of the tuning grid for each machine learning model in Appendix C. Because our goal is to predict demand for a fixed set of customers, we implement a cross-validation procedure where each fold includes all customers but separate time periods. In particular, we implement 5-fold cross-validation where each fold is given by a 3-month period of our 15-month-long training sample. This particular way of constructing folds — rather than randomly sampling from all purchase occasions — ensures that the same set of customers is present in each fold and thus mirrors the construction of our training and evaluation samples (which also both contain the same set of households).¹⁷

¹⁶The deep neural network is specified so that each hidden layer has the same number of units. Other network architectures such as a funnel shape (where the number of units increases/decreases with each additional hidden layer) or a diamond shape (where the number of units first increases and then decreases across hidden layers) led to similar, but marginally lower profits.

¹⁷We also trained the machine learning classifiers using a cross-validation method with different household / time-period combinations as folds and found that it performed worse on average than the temporal cross-validation procedure. In Appendix D, we describe the construction of folds in more detail and report profit results from this alternative cross-validation procedure.

5.1 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 5.50 and 6.45 respectively.¹⁸ 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 an evaluation 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 4. 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, none of the machine learning models outperform the blanket coupon. When we add in more demographic variables, profitability increases somewhat, but only the two regularized regression models and the deep neural network lead to profits higher than a blanket couponing policy. When we add in purchase histories, however, we find sizable improvements across all machine learning models. A single-layer neural network performs best among the machine learning models with policies that yield average profits of 7.41, representing a 14.9% gain relative to the blanket coupon. The standard Bayesian hierarchical logit model with normal heterogeneity and only base demographic inputs generates a profit of 7.53 and thus represents a 16.7% gain in profits relative to a blanket coupon. Adding a more flexible 5-component mixture of normals distribution of heterogeneity or adding extra demographic variables to the upper-level model leads to a slight decrease in profitability. As discussed in Section 4.3, the Bayesian hierarchical models encode purchase histories directly so we do not estimate versions of these two models that contain summaries of purchase histories as covariates.

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 the value of machine learning models is only realized when detailed purchase history information is included as covariates. This is in contrast to Bayesian hierarchical models which naturally

¹⁸Average 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 Purch. Hist.
Bayesian Hierarchical Logit			
– normal heterogeneity	7.53 (0.47)	7.48 (0.44)	
– mixtures of normals heterogeneity	7.38 (0.46)	7.29 (0.43)	
Bayesian Pooled Logit	5.74 (0.42)	5.77 (0.43)	7.14 (0.40)
Lasso	5.77 (0.41)	6.53 (0.43)	7.27 (0.41)
Elastic Net	5.50 (0.40)	6.57 (0.43)	7.27 (0.41)
Neural Network	5.87 (0.41)	6.35 (0.35)	7.41 (0.44)
Deep Neural Network	5.28 (0.40)	6.47 (0.38)	7.26 (0.42)
KNN	4.47 (0.26)	6.00 (0.30)	7.07 (0.41)
Random Forest	5.24 (0.30)	5.60 (0.27)	6.33 (0.35)
XGBoost	5.42 (0.44)	5.50 (0.42)	7.17 (0.43)
Blanket Coupon	6.45 (0.32)		
No Coupon	5.50 (0.42)		
Panel (II): Out-of-Sample Hit Probabilities			
Bayesian Hierarchical Logit			
– normal heterogeneity	0.929	0.929	
– mixtures of normals heterogeneity	0.929	0.929	
Bayesian Pooled Logit	0.927	0.928	0.929
Lasso	0.927	0.928	0.930
Elastic Net	0.927	0.927	0.930
Neural Network	0.928	0.930	0.931
Deep Neural Network	0.925	0.925	0.926
KNN	0.929	0.927	0.930
Random Forest	0.941	0.933	0.931
XGBoost	0.930	0.931	0.932

Table 4: **Targeting Policy Profits and Model Fit.** Panel (I) reports the average customer-level profits (scaled by 100) from each targeted pricing policy. Bootstrapped standard errors are reported in parentheses. Panel (II) reports the out-of-sample hit probabilities from each model.

encode purchase history information in the unit-level likelihood, 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 model. We find that this pooled logit model performs slightly worse than most of the machine learning models when all customer characteristics 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 has 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 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 statistical measures of fit. For the sake of brevity, we focus our discussion on out-of-sample hit probabilities which are reported in the lower panel of Table 4. The corresponding out-of-sample log-likelihood values are reported in Appendix E. 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, purely statistical measures of fit 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 hierarchical logit model performs relatively poorly. In terms of models that use all data inputs, the hit probability metric ranks XGBoost highest whereas based on our profit estimator the two regularized regression models and both neural networks perform better. We also calculate the correlation between profits and hit probabilities across all model specifications and find that the two metrics are almost uncorrelated with a correlation coefficient of -0.081. Similarly, the correlation between out-of-sample profits and log-likelihood values (reported in Appendix E) is equal to only 0.084. We conclude that statistical measures of fit may provide poor guidance for selecting the appropriate model for deriving personalized prices.

5.2 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 5. For ease of exposition we focus on the Bayesian hierarchical model with normal 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 40-50% of customers. XGBoost prescribes by far the lowest discount frequency of 28%. 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 agree for roughly 90% of customers. However, correlation with the logit does not fully account for the difference in performance. For example, XGBoost and the (single-layer) neural network have a similar rate of

Panel (I): Frequency of Policy Agreement									
	Discount Freq.	Logit	Lasso	Elastic Net	NNet	Deep NNet	KNN	Random Forest	XGBoost
Logit	0.414	1	0.886	0.884	0.844	0.882	0.796	0.708	0.839
Lasso	0.436		1	0.994	0.941	0.949	0.837	0.770	0.848
Elastic Net	0.439			1	0.943	0.950	0.831	0.769	0.845
NNet	0.475				1	0.923	0.831	0.745	0.802
Deep NNet	0.413					1	0.823	0.768	0.864
KNN	0.463						1	0.723	0.775
Random Forest	0.530							1	0.711
XGBoost	0.284								1

Panel (II): Predictors of Pricing Policies									
<i>Dependent Variable: Discount Dummy</i>									
Model		Logit	Lasso	Elastic Net	NNet	Deep NNet	KNN	Random Forest	XGBoost
Chain B Dummy		-0.015 (0.022)	-0.019 (0.020)	-0.017 (0.020)	-0.072*** (0.020)	-0.045* (0.020)	-0.043* (0.020)	0.01 (0.023)	0.059* (0.023)
Income (in \$10,000s)		-0.015*** (0.004)	-0.001 (0.004)	-0.001 (0.004)	0.005 (0.004)	-0.001 (0.004)	0.002 (0.004)	-0.006 (0.004)	0.001 (0.004)
Family Size		-0.019 (0.010)	-0.016 (0.009)	-0.017 (0.009)	-0.013 (0.009)	-0.017 (0.009)	-0.017 (0.009)	-0.009 (0.011)	-0.022* (0.011)
Retired		-0.154*** (0.047)	-0.001 (0.043)	0.024 (0.043)	-0.034 (0.044)	-0.108* (0.043)	-0.092* (0.043)	-0.073 (0.049)	0.027 (0.050)
Unemployed		-0.016 (0.027)	-0.01 (0.025)	-0.007 (0.024)	-0.008 (0.025)	-0.006 (0.025)	-0.019 (0.024)	0.016 (0.028)	0.002 (0.028)
Single Mother		0.061 (0.041)	-0.008 (0.038)	-0.012 (0.037)	0.009 (0.038)	-0.035 (0.037)	-0.014 (0.037)	-0.025 (0.043)	-0.035 (0.043)
Hellmann's Choice Share		-0.242*** (0.041)	-0.037 (0.039)	-0.04 (0.038)	-0.100* (0.039)	0.003 (0.038)	0.001 (0.038)	-0.182*** (0.044)	-0.013 (0.044)
Hellmann's Choice Count		0.052*** (0.009)	0.007 (0.008)	0.003 (0.008)	-0.031*** (0.008)	0.004 (0.008)	0.001 (0.008)	-0.01 (0.009)	0.028** (0.010)
Discount Share		0.783*** (0.031)	0.972*** (0.029)	0.991*** (0.029)	0.999*** (0.030)	0.926*** (0.029)	0.955*** (0.029)	0.812*** (0.033)	0.799*** (0.034)
# Brands Purchased		-0.184*** (0.029)	-0.015 (0.027)	-0.024 (0.027)	-0.041 (0.027)	0.070** (0.027)	0.037 (0.027)	0.001 (0.031)	0.088** (0.031)
Total Spending		-0.013*** (0.002)	-0.005** (0.002)	-0.004* (0.002)	0.002 (0.002)	-0.006*** (0.002)	-0.005** (0.002)	-0.005** (0.002)	-0.001 (0.002)
Observations		1,162	1,162	1,162	1,162	1,162	1,162	1,162	1,162

Table 5: **Policy Comparisons.** Panel (I) reports discount frequencies and the frequency of agreement between pricing policies. Panel (II) reports a linear probability regression of discount incidence on a selected set of customer characteristics. The unit of observation is a customer/chain combination. Significance codes: *p<0.05; **p<0.01; ***p<0.001.

agreement with the logit policy, despite a difference in average profits.

In the lower panel of Table 5 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 demographics, most 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 regard 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. Moreover, consumers with lower category spending are more likely to receive a discount in most of the models. For the other purchase history variables, the effect tends to vary across models.

5.3 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 cross-validation. However, the performance of machine learning models can suffer from various modeling pitfalls or data challenges, 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 majority class. Oversampling adds synthetically created observations to the minority classes, whereas undersampling randomly removes observations from the majority class.¹⁹ We provide additional details on our implementation of the SMOTE method in Appendix F. 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 6 (which correspond

¹⁹Other 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.

	<i>Addressing Class Imbalance</i>		<i>Feature Selection</i>			<i>Lagged Purchase Histories</i>
			Base Demos			Base Demos
			Base Demos	Extra Demos		Extra Demos
	Base Demos	Extra Demos	Purch. Hist.	Purch. Hist.	Var. Imp.	Purch. Hist.
Lasso	5.50 (0.43)	5.56 (0.41)	6.91 (0.44)	7.28 (0.42)	7.28 (0.40)	7.18 (0.42)
Elastic Net	5.50 (0.42)	5.56 (0.40)	6.91 (0.44)	7.28 (0.41)	7.28 (0.40)	6.84 (0.43)
Neural Network	5.60 (0.42)	7.00 (0.46)	7.13 (0.42)	7.53 (0.45)	5.54 (0.38)	7.11 (0.38)
Deep Neural Network	5.50 (0.40)	6.12 (0.36)	7.02 (0.42)	7.18 (0.42)	6.31 (0.42)	6.94 (0.42)
KNN	5.41 (0.37)	6.46 (0.43)	6.78 (0.42)	7.12 (0.40)	6.63 (0.38)	7.06 (0.41)
Random Forest	6.17 (0.41)	6.54 (0.45)	7.20 (0.39)	7.11 (0.41)	6.21 (0.41)	6.32 (0.32)
XGBoost	5.95 (0.43)	5.83 (0.44)	6.84 (0.43)	7.07 (0.43)	6.67 (0.42)	6.30 (0.43)

Table 6: **Targeting Policy Profits from Additional Specifications.** Each cell reports average profits per 100 customers with standard errors in parentheses. The first three columns correspond to models fit on modified data sets built using the SMOTE method to correct for class imbalance. The columns under the heading “feature selection” correspond to models fit using a selected subset of predictor variables. The final column calculates purchase history variables using data from the year preceding the training sample.

to the first three columns of results in Table 4). We find that model performance in terms of profits worsens in the majority of 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 6 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 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.²⁰ 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 discounts. The profits based on this approach are reported in the final column of Table 6. We find that using only purchase history variables leads to slight improvements in the performance for most models, whereas the second approach only improves performance of the two regularized regressions and leads to lower performance for all other models.

²⁰We report the top 10 most important variables from each machine learning model in Appendix G. Details of model-specific variable importance metrics can be found here: <https://topepo.github.io/caret/variable-importance.html>.

Lagged Purchase Histories In a final robustness check we modify the way in which we calculate purchase history variables. In our main specification, we calculate all purchase history variables over the entire sample period of the training sample. However, the following issue arises when calculating purchase histories contemporaneously: When training each of the machine learning models by cross-validation, the left-out fold contains information that was used to calculate purchase histories rather than being entirely separated from the other folds. In order to avoid this type of leakage between folds we estimate a version of each model where purchase histories variables are calculated based on data from the year prior to the training sample. Therefore, these modified purchase history variables are not a function of the data in any of the folds used for cross-validation in the training sample. We report profit estimates for models based on modified purchase histories (plus the extended set of demographics which remains unchanged) in the final column of Table 6. We find that profits remain relatively similar across all models and decrease slightly when using lagged rather than contemporaneous purchase histories.

6 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 is based on observable profits and hence does not rely on any demand-side modeling assumptions. We apply the evaluation approach to supermarket scanner data and establish several findings. First, we find that demographic information is less useful than purchase histories in terms of informing profitable price targeting strategies. This finding generalizes an earlier finding by Rossi et al. (1996). Notably, demographics have little impact on profits from targeting even when using a larger set of demographics and when allowing them to enter in a more flexible fashion. Second, we find that comparing models based on out-of-sample profits rather than statistical fit leads to a very different ranking of models. Third, we find that the most profitable policy is generated by a Bayesian hierarchical logit model, closely followed by neural networks and regularized regressions. While these profit levels are not statistically distinguishable, the performance of the Bayesian hierarchical model is nevertheless notable because it naturally encodes purchase history information and does not require parameter tuning. Taken together, our findings highlight the importance of using a decision-relevant loss function and accounting for past purchase histories when solving for targeted prices with panel data.

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. W. IMBENS (2019): “Machine Learning Methods That Economists Should Know About,” *Annual Review of Economics*, 11, 1–41.
- 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.
- BREIMAN, L. (2001): “Random Forests,” *Machine Learning*, 45, 5–32.
- 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.
- CHEN, T. AND C. GUESTRIN (2016): “XGBoost: A Scalable Tree Boosting System,” *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- 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 (2021): “Counterfactual Inference for Consumer Choice Across Many Product Categories,” *Quantitative Marketing and Economics*, 19, 369–407.

- 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): “Product Choice with Large Assortments: A Scalable Deep-Learning Model,” *Management Science*, 68, 1808–1827.
- 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.
- HASTIE, T., R. TIBSHIRANI, AND J. FRIEDMAN (2009): *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, New York: Springer, 2nd ed.
- 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.
- HOWELL, J. R., S. LEE, AND G. M. ALLENBY (2016): “Price Promotions in Choice Models,” *Marketing Science*, 35, 319–334.
- 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.
- KINGMA, D. P. AND J. BA (2015): “Adam: A Method for Stochastic Optimization,” *Proceedings of the 3rd International Conference on Learning Representations*.
- 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.
- LIU, X. (2022): “Dynamic Coupon Targeting Using Batch Deep Reinforcement Learning: An Application to Livestream Shopping,” *Marketing Science*, *Forthcoming*.
- 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*, 40, 813–1007.
- MURPHY, K. P. (2022): *Probabilistic Machine Learning: An Introduction*, MIT Press.
- 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.
- TADDY, M. (2019): *Business Data Science: Combining Machine Learning and Economics to Optimize, Automate, and Accelerate Business Decisions*, McGraw Hill.
- TIBSHIRANI, R. (1996): “Regression Shrinkage and Selection via the Lasso,” *Journal of the Royal Statistical Society: Series B (Methodological)*, 58, 267–288.
- YANG, J., D. ECKLES, P. DHILLON, AND S. ARAL (2020): “Targeting for Long-Term Outcomes,” *Working Paper*.
- 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, 2706–2728.
- 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 Dynamic Demand

Incorporating demand dynamics into our evaluation framework based on the inverse probability weighted profit estimator is difficult for several reasons. Most importantly, the key idea of using this particular profit estimator is to avoid making demand side modeling assumption and to use only observed profits when calculating expected profits. As we show in Section 2.2, in a static demand setting we can simply use observed demand (i.e. the consumer purchases the focal product or not) to calculate profits for a specific household and time period. However, if demand is dynamic, then the full profit impact of a particular price being offered in time period t would need to take into account demand in future time periods as well. Typically, a model-based approach is used to analyze dependencies in demand that arise from brand loyalty / switching costs, learning, stockpiling or other drivers of dynamic demand patterns. Due to the potentially complicated nature of the relationship between current demand and demand in future time periods, we do not see an easy way to extend an estimator without demand side assumption to a dynamic setting.

A related concern is that even if we incorrectly use static profits in the profit estimator, we need the conditional independence assumption to hold with regard to past prices on top of other market and consumer characteristics. For example, if consumers stockpile and prices alternate between high and low prices, then a high price period will be associated with lower demand because it is more likely that the consumer purchased in the previous period when the product was offered at a low price. If prices are uncorrelated over time, then consumers that face a discount or regular price will have faced the same price patterns in previous periods and therefore past prices do not lead to a difference in potential profits between consumers that face a regular or a discounted price in the current period.

In our empirical application we assess whether prices are correlated over time and find that this is not the case. We therefore do not condition on past prices in the propensity score. For the reasons outlined in the first paragraph above, we focus on targeted prices that maximize current period profits and hence our evaluation framework ignores any possible impact of pricing on future demand.

B 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.

$$\mathbb{P}(j|\tilde{\mathbf{x}}_t, \mathbf{z}_i) = \sigma_j(v(\tilde{\mathbf{x}}_t, \mathbf{z}_i; \Psi)) \quad (17)$$

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 4 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 $\mathcal{V} = \mathbb{R}^J$.

$$\mathbb{P}(j|\tilde{\mathbf{x}}_t, \mathbf{z}_i) = \frac{\exp(v_j(\tilde{x}_{jt}, \mathbf{z}_i; \Psi))}{1 + \sum_{k=1}^J \exp(v_k(\tilde{x}_{kt}, \mathbf{z}_i; \Psi))} \quad (18)$$

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

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

$$\mathbb{P}(j|\tilde{\mathbf{x}}_t, \mathbf{z}_i) = \frac{\exp(v_j(\tilde{\mathbf{x}}_t, \mathbf{z}_i; \Psi_j))}{\sum_{k=0}^J \exp(v_k(\tilde{\mathbf{x}}_t, \mathbf{z}_i; \Psi_k))} \quad (19)$$

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.²¹

$$v_j(\tilde{\mathbf{x}}_t, \mathbf{z}_i; \Psi_j) = \mathbf{p}'_t \Psi_j^\alpha + \mathbf{x}'_t \Psi_j^\beta + \mathbf{z}'_i \Psi_j^\gamma + \sum_{k=1}^J (p_{kt} \cdot \mathbf{z}_i)' \Psi_j^\delta \quad (20)$$

- (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: $\Psi = \{\psi_\ell, \ell = 1, \dots, L\}$.

$$\mathbb{P}(j|\tilde{\mathbf{x}}_t, \mathbf{z}_i) = \sigma_j(v(\tilde{\mathbf{x}}_t, \mathbf{z}_i; \Psi)) = \sigma_j(g(g(\dots g(\tilde{\mathbf{x}}_t, \mathbf{z}_i; \psi_1); \psi_{L-1}); \psi_L)) \quad (21)$$

More specifically, at each node of the first layer, we estimate a set of parameters ψ_{ln} (where n denotes a specific node of layer l) that generates an output based on a rectified linear unit

²¹This 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 .

(ReLU) activation function: $g(\tilde{\mathbf{x}}_t, \mathbf{z}_i; \psi_{ln}) = \max(0, [\tilde{\mathbf{x}}_t \ \mathbf{z}_i]'\psi_{ln})$. At each hidden layer, the inputs are given by the output values from all nodes in the previous layer.

- (iv) Nonparametric methods — namely KNN classifiers and bagged/boosted classification trees — where $v(\cdot)$ is an identity function with range $\mathcal{V} = \mathcal{X}$ and $\sigma_j(\cdot)$ is the target estimand.

$$\mathbb{P}(j|\tilde{\mathbf{x}}_t, \mathbf{z}_i) = \sigma_j(\tilde{\mathbf{x}}_t, \mathbf{z}_i) \quad (22)$$

C Tuning Parameters

Method (R package)	Tuning Parameters		Grid
Lasso (<code>glmnet</code>)	<code>alpha</code>	Mixing percentage	1
	<code>lambda</code>	Regularization parameter	$10^{\text{seq}(-8,0,\text{length}=1000)}$
Elastic Net (<code>glmnet</code>)	<code>alpha</code>	Mixing percentage	$\{0,0.05,0.1,\dots,0.95\}$
	<code>lambda</code>	Regularization parameter	$10^{\text{seq}(-8,0,\text{length}=50)}$
Neural Network (<code>keras</code>)	<code>layers</code>	Number of hidden layers	1
	<code>units</code>	Number of hidden units per layer	100
	<code>dropout</code>	Fraction of input units to use at each update during training	$\{0.2,0.3,0.4,\dots,0.8\}$
Deep Neural Network (<code>keras</code>)	<code>layers</code>	Number of hidden layers	$\{2,3,5\}$
	<code>units</code>	Number of hidden units per layer	100
	<code>dropout</code>	Fraction of input units to use at each update during training	$\{0.2,0.3,0.4,\dots,0.8\}$
KNN (<code>class</code>)	<code>k</code>	Number of neighbors	$\{11, 21, 31, \dots, 301\}$
Random Forest (<code>rf</code>)	<code>ntree</code>	Number of trees	500
	<code>mtry</code>	Number of randomly selected predictors	$\{1,2,3,\dots,15\}$
XGBoost (<code>xgboost</code>)	<code>nrounds</code>	Number of boosting iterations	$\{200,400,600,\dots,2000\}$
	<code>eta</code>	Learning rate	$\{0.3,0.2,0.1,0.01,0.001\}$
	<code>max_depth</code>	Maximum depth of a tree	$\{6,12,24\}$
	<code>gamma</code>	Minimum loss reduction required to make a split	$\{0,5,10\}$
	<code>min_child_weight</code>	Minimum sum of instance weight (hessian) needed in a child	1
	<code>subsample</code>	Fraction of observations (rows) used when growing a tree	0.5
	<code>colsample_bytree</code>	Fraction of predictors (columns) used when growing a tree	0.5

Table A1: **Tuning Parameters and Cross-Validation Grids.**

D Alternative Cross-Validation with Time-Household Blocks

In our baseline specification, we estimate all machine learning models using 5-fold cross-validation where each fold corresponds to a 3-month period. As discussed in Section 5, we believe that this way of implementing cross-validation is most appropriate given our goal of predicting demand for the same set of households in a different time period.

As a robustness check, we also implement a cross-validation procedure that splits the data into a grid of five temporal blocks and five household blocks. Based on this two-way split of the data, we then use a 25-fold cross-validation procedure where each test sample includes data from one temporal block and one household block and the corresponding training sample includes data from all other temporal blocks and all other household blocks. For example, when predicting choices for household block A in time block 5, we only use data from household blocks B, C, D, and E and time periods 1, 2, 3, and 4. We report out-of-sample profit results from ML models that are fitted using this alternative cross-validation procedure in Table A2 alongside our baseline results for comparison. We find that the alternative cross-validation method generates lower profits on average and leads to lower profits for almost all models in the specification that uses all data inputs (the final column in Table A2). The only exception is the random forest model which exhibits marginally higher profits under the alternative cross-validation procedure. Overall, we conclude that our baseline cross-validation method produces better results on average and our key substantive findings do not change when using a household / time-period based cross-validation procedure.

Our interpretation of the relative performance of the two cross-validation methods is as follows: a method that evaluates performance based on predictions for a new set of households is more appropriate when trying to generate a pricing policy for a new set of customers rather than our prediction goal of finding personalized prices for the same set of customers. While performance does not change uniformly in one direction, the results in Table A2 are broadly in line with the reasoning that a temporal cross-validation procedure is more appropriate in our setting.

	Base Demos	Base Demos Extra Demos	Base Demos Extra Demos Purch. Hist.
Main Specification:			
CV with Time Blocks			
Lasso	5.77 (0.41)	6.53 (0.43)	7.27 (0.41)
Elastic Net	5.50 (0.40)	6.57 (0.43)	7.27 (0.41)
Neural Network	5.87 (0.41)	6.35 (0.35)	7.41 (0.44)
Deep Neural Network	5.28 (0.40)	6.47 (0.38)	7.26 (0.42)
KNN	4.47 (0.26)	6.00 (0.30)	7.07 (0.41)
Random Forest	5.24 (0.30)	5.60 (0.27)	6.33 (0.35)
XGBoost	5.42 (0.44)	5.50 (0.42)	7.17 (0.43)
Alternative Specification:			
CV with Time-Household Blocks			
Lasso	5.74 (0.43)	7.08 (0.45)	7.23 (0.42)
Elastic Net	5.50 (0.42)	5.59 (0.41)	7.25 (0.42)
Neural Network	6.28 (0.41)	6.40 (0.45)	7.28 (0.43)
Deep Neural Network	5.41 (0.31)	5.00 (0.39)	7.25 (0.42)
KNN	4.47 (0.25)	6.00 (0.32)	7.06 (0.41)
Random Forest	5.36 (0.32)	5.81 (0.31)	6.40 (0.34)
XGBoost	5.50 (0.44)	5.50 (0.42)	7.14 (0.43)

Table A2: **Targeting Policy Profits with Alternative Cross-Validation Procedure.** Both panels report the average customer-level profits (scaled by 100) from each targeted pricing policy. Bootstrapped standard errors are reported in parentheses. In both specifications, models are trained using cross-validation. In the main specification (5-fold CV), folds are constructed using five time blocks. In the alternative specification (25-fold CV), folds are constructed using combinations of five time blocks and five household blocks.

E Additional Out-of-Sample Fit Results

		Base Demos	Base Demos	Base Demos
			Extra Demos	Extra Demos
		Base Demos	Extra Demos	Purch. Hist.
Bayesian Hierarchical Logit				
– normal heterogeneity	-5418	-5437		
– mixtures of normals heterogeneity	-5569	-5611		
Bayesian Pooled Logit	-5924	-6028		-5940
Lasso	-5876	-5863		-5908
Elastic Net	-5876	-5866		-5907
Neural Network	-5899	-5815		-6177
Deep Neural Network	-5859	-5790		-6629
XGBoost	-5906	-5764		-6080

Table A3: **Out-of-Sample Log-Likelihood.** The table reports log-likelihood (negative cross-entropy) values for each model. The log-likelihood values for KNN and random forest methods are negative infinity in our evaluation sample (due to predicted choice probabilities that are equal to zero) and are thus not reported.

F 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.

G Variable Importance for Machine Learning Models

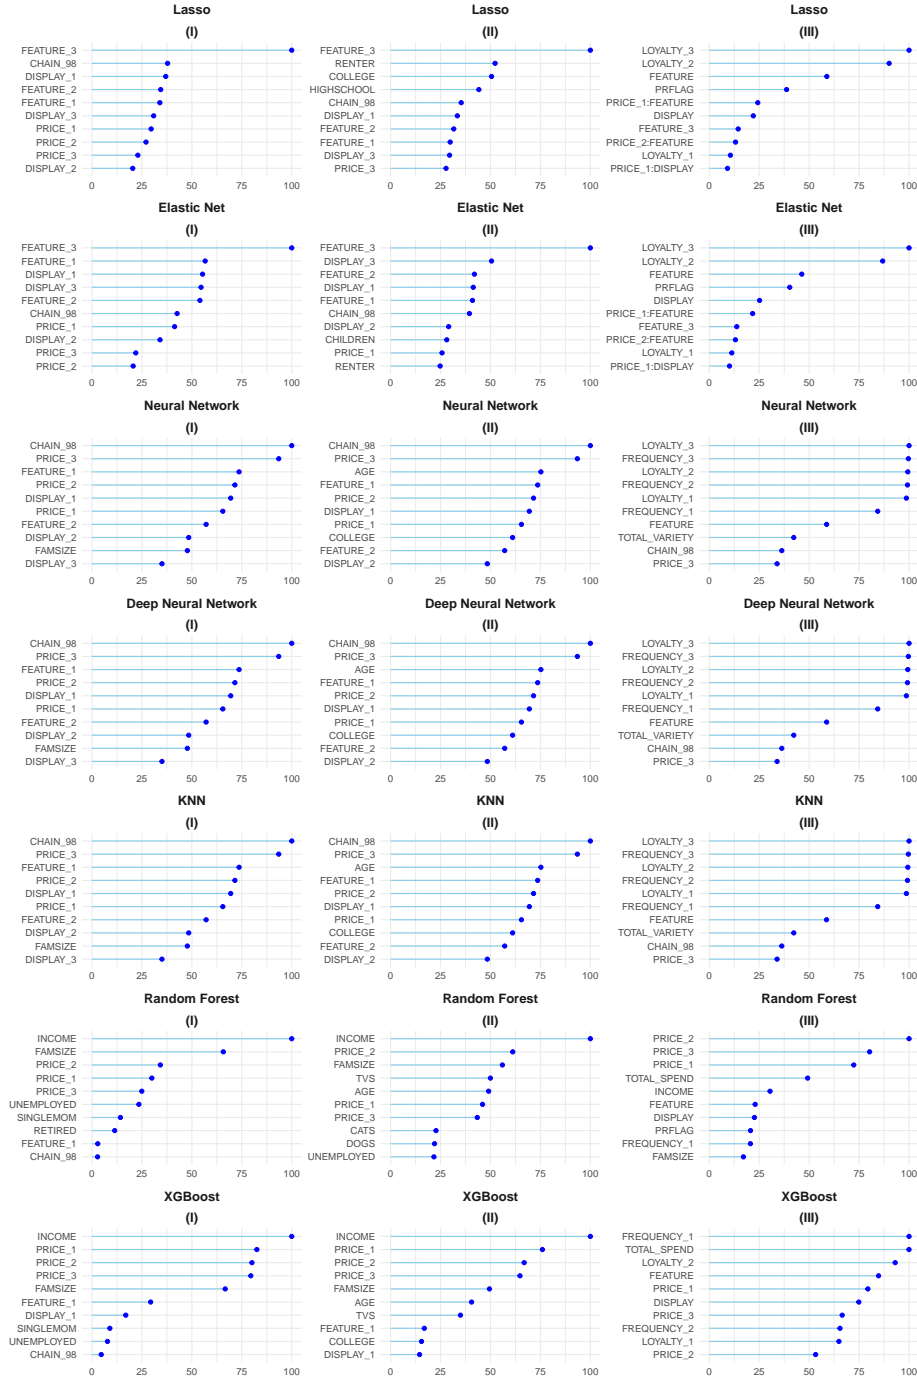


Figure A1: **Variable Importance.** The top 10 variables (based on the variable importance metric) are reported for each machine learning model. Specification (I) corresponds to base demographics, (II) corresponds to base demographics and extra demographics, and (III) corresponds to base demographics, extra demographics, and purchase histories.