

Doing More with Less: Overcoming Ineffective Long-Term Targeting Using Short-Term Signals

Ta-Wei Huang,^{a,*} Eva Ascarza^a

^aMarketing Unit, Harvard Business School, Boston, Massachusetts 02163

*Corresponding author

Contact: thuang@hbs.edu,  <https://orcid.org/0000-0002-6735-2954> (T-WH); eascarza@hbs.edu,  <https://orcid.org/0000-0002-4840-5344> (EA)

Received: October 20, 2022

Revised: May 2, 2023; November 21, 2023

Accepted: December 24, 2023

Published Online in Articles in Advance:
March 6, 2024

<https://doi.org/10.1287/mksc.2022.0379>

Copyright: © 2024 INFORMS

Abstract. Firms are increasingly interested in developing targeted interventions for customers with the best response. This requires identifying differences in customer sensitivity, typically through the conditional average treatment effect (CATE) estimation. In theory, to optimize long-term business performance, firms should design targeting policies based on CATE models constructed using long-term outcomes. However, we show theoretically and empirically that this method can fail to improve long-term results, particularly when the desired outcome is the cumulative result of *recurring customer actions*, like repeated purchases, due to the accumulation of unexplained individual differences over time. To address this challenge, we propose using a *surrogate index* that leverages short-term outcomes for long-term CATE estimation and policy learning. Moreover, for the creation of this index, we propose the *separate imputation* strategy, designed to reduce the additional variance caused by the inseparable nature of customer churn and purchase intensity, prevalent in marketing contexts. This involves constructing two distinct surrogate models, one for the observed last purchase time and the other for the observed purchase intensity. Our simulation and real-world application show that (i) using short-term signals instead of the actual long-term outcome significantly improves long-run targeting performance, and (ii) the separate imputation technique outperforms existing imputation approaches.

History: Catherine Tucker served as the senior editor for this article.

Supplemental Material: The online appendices and data files are available at <https://doi.org/10.1287/mksc.2022.0379>.

Keywords: long-term targeting • heterogeneous treatment effect • statistical surrogacy • customer churn • field experiments • conditional average treatment effect (CATE)

1. Introduction

Recent advancements in business experimentation combined with machine learning have transformed how companies execute targeted interventions. Through controlled experiments, businesses can infer causal relationships between their marketing offerings and customers' responses. Rather than simply measuring the average impact across all customers, companies can further identify differences in customer sensitivity based on individual characteristics, commonly quantified as the conditional average treatment effect (CATE). This approach empowers firms to focus on customers predicted to respond most favorably to their objectives (e.g., profits or purchases), particularly those with the highest predicted CATEs. This method has gained significant popularity among organizations to design effective targeted intervention, and some tech companies, such as Microsoft (Oprescu et al. 2019) and Uber (Chen et al. 2020), have taken a step further by open sourcing

their tools for CATE estimation. This has enabled more companies to adopt this approach and develop highly precise targeted marketing interventions at scale.

This *test-to-target* approach has proven effective in various marketing contexts, such as customer retention (Guelman et al. 2012, Ascarza 2018, Lemmens and Gupta 2020), membership subscription (Simester et al. 2020, Yoganarasimhan et al. 2023), and catalog mailing purchases (Hitsch et al. 2023). Despite its popularity, it remains untested whether this approach can effectively optimize long-term outcomes, such as customer lifetime value (CLV) or repeated purchases, which are typically the top-line metrics for a firm. In theory, the observation window should not alter the way firms optimize their resource allocation—if the business goal is to maximize long-run outcomes, firms should target customers based on their *long-term sensitivity* to the intervention, measured as the CATE of the intervention on the long-term business outcome.

However, our research shows that the conventional test-to-target approach can be ineffective at optimizing long-term outcomes, especially those driven by *recurring customer behaviors*, such as total purchases over an extended postintervention period. Unlike short-term outcomes (e.g., immediate purchases after the intervention), long-term outcomes accumulate individual customer behaviors, such as unobserved heterogeneity or unexplained customer attrition, that cannot be explained by the observed customer characteristics collected before the intervention. As a result, long-term outcomes not only carry information about the treatment effect (which is what CATE models aim to capture), but also accumulate unexplained variations that grow over time. This presents a significant and under-studied challenge in CATE estimation, as existing models may generate unstable and high-variance CATE predictions when unexplained variations are large. Targeting customers based on such high-variance models can result in an ineffective targeting strategy when optimizing for long-term outcomes.

This paper has two main objectives. First, we examine the challenge of estimating CATEs for *long-term, recurrent customer behaviors*. Our theoretical analysis shows that such outcomes can accumulate unexplained variations due to *unobserved heterogeneity* or *customer attrition*. Consequently, as we extend the observation window, the outcome variance—and by extension, the variance of most CATE models—tends to rise. This increased variance inadvertently heightens the likelihood of incorrect targeting. Second, we present a solution that enables firms to implement more effective targeted interventions and achieve better long-term performance. We recommend that firms create a *noise-reduced proxy* based on immediate postintervention behaviors and use this proxy for CATE estimation instead of the actual long-term outcome. Although it may seem counterintuitive, this method can potentially minimize unexplained variations and effectively capture long-term treatment effects that are reflected in short-term behavioral changes.

To construct this noise-reduced proxy, we adopt the surrogate index approach (Athey et al. 2019a, Yang et al. 2023), which uses historical data that is readily available to firms to infer the relationship between short-term behaviors and the long-term outcome. In addition to providing valid CATE estimation, we formally show that the surrogate index has smaller unexplained variations than the actual long-term outcome, leading to more accurate CATE estimation. This enables firms to effectively target customers based on their long-term sensitivity to the intervention while mitigating the noise accumulation problem.

We emphasize that the conventional method of surrogate index construction, as recommended by Athey and Wager (2019) and Yang et al. (2023), may not be

optimal in scenarios with customer attrition. In the presence of attrition, long-term repeated behaviors become the product of two critical variables: the duration of a customer's lifetime and the intensity of these behaviors during that lifetime. As both these variables are influenced by short-term behaviors and idiosyncratic variations in each period, the conventional modeling approach fails to identify the correct relationship between short-term and long-term outcomes (Brown 1983, Hoderlein and Mammen 2009, Su et al. 2019).

To address this issue, we propose a novel *separate imputation* technique. Our approach involves developing two separate models using historical data: one to predict the time of a customer's last observed purchase (i.e., the proxy for lifetime) and another to predict the average purchases per period when a customer is still active. We then combine the predictions of both models to estimate expected future purchases. This approach stands out from other surrogate models by effectively mitigating the issue of increased variance that arises from the inseparable nature of customer attrition and purchase intensity. Consequently, this technique enables firms to construct more accurate and robust surrogate indices, leading to improved CATE estimation and more effective targeting strategies for optimizing long-term outcomes.

Through simulation analyses and a real-world marketing campaign, we demonstrate that our proposed approach is significantly more effective than relying on the actual long-term outcome. Specifically, we show that targeting rules based on immediate, short-term signals—either directly from short-term outcomes or through surrogate indices—consistently yield better results than those using CATE models estimated using the long-term outcome. This is surprising because it suggests that, for optimal long-term performances, firms should rely on short-term outcomes and historical information rather than on the long-term outcomes themselves. Furthermore, we demonstrate that our separate imputation approach achieves the best targeting performance. In the real-world application, targeting customers using the proposed solution yields a 6% increase in profits (compared with directly rolling out the best action to all customers), while targeting based on the long-term outcome results in a 3% profit loss.

There are several compelling reasons for firms to implement our proposed solution. First, it uses existing historical data, obviating the need for new experimental data to reduce variance. This benefit translates into cost savings, as firms can bypass the expenses associated with expanding experimental samples. Second, the solution integrates seamlessly with standard machine learning algorithms and existing software packages, facilitating a swift and effective deployment geared toward higher profits. Third, it is applicable to

a wide variety of business settings, including retailers, e-commerce, apparel, and nonprofit organizations. Last, our method expedites decision-making processes. Firms no longer have to endure the typical delays associated with waiting for long-term outcomes, thereby accelerating the implementation of targeted interventions (Athey et al. 2019a, Yang et al. 2023).

Our research contributes to the literature in four strands. First, we address practical challenges in designing and implementing targeting policies. We underscore the limitations of current best practices (Ascarza 2018, Simester et al. 2020, Ellickson et al. 2022, Yoganarayanan et al. 2023), particularly in their ineffectiveness for optimizing noisy long-term outcomes. We contribute to this literature by proposing a new targeting paradigm where firms *reduce noise* in the outcome variable *before* estimating any CATE model. Although ignoring the actual outcome of interest may seem counterintuitive, we demonstrate how creating the “right” proxy using the surrogate index approach with proper imputation methods results in more effective targeting.

Second, our work highlights a significant challenge in estimating the CATE and effective targeting in scenarios characterized by a low signal-to-noise ratio. Considerable research has been dedicated to developing methods for CATE estimation (Imai and Strauss 2011, Imai and Ratkovic 2013, Guelman et al. 2015, Grimmer et al. 2017, Chernozhukov et al. 2018, Athey et al. 2019b, Künzel et al. 2019, Nie and Wager 2021, Kennedy 2023) and policy learning (Manski 2004, Kitagawa and Tetenov 2018, Athey and Wager 2021, Mbakop and Tabord-Meehan 2021). To the best of our knowledge, this paper is the first to theoretically explore how unexplained variations affect the predictive accuracy of CATE models and their targeting performance. Moreover, we incorporate behavioral insights from marketing literature to illustrate why the issue of high noise is common in many marketing contexts. Our study provides important insights into the limitations of existing CATE models and highlights the need for robust solutions to estimate CATEs.

Third, we contribute to the literature on statistical surrogacy and long-term treatment effect estimation (Prentice 1989, Athey et al. 2019a, Qian et al. 2021, Imbens et al. 2022, Yang et al. 2023). This literature has traditionally assumed that firms use short-term signals because of the cost of waiting to observe long-term outcomes. We further demonstrate, using formal theory and empirical evidence, the value of leveraging short-term proxies even when the actual long-term outcome is observed. Moreover, past research has primarily constructed surrogate indices using standard regression models, whereas our work highlights the importance of considering the data generating process for surrogate indices construction.

Last, our work contributes to the literature on treatment effect estimation for low-sensitivity experiments

(i.e., experiments with outcome variance much larger than the treatment effect). Our approach differs from previous research (Deng et al. 2013, Guo et al. 2021, Jin and Ba 2023) in a critical way—We do not reduce variance by eliminating variations that can be explained by customer observables. Instead, we advocate for the use of information explainable by short-term signals and pretreatment covariates for CATE estimation and targeting. This strategy is pivotal, as it targets the core issue of high variance in CATE estimates arising from unexplained variations, rather than from the observable heterogeneity among customers.

2. Data and Motivation

We use a field experiment from a retail technology company in Taiwan to highlight the challenges of targeting with the goal of maximizing repeated purchases over an extended period of time. This company manages a network of self-service vending machines at various locations within a city. Customers can grab food and beverages from a vending machine, and the machine automatically counts the items using Radio Frequency Identification technology and charges the customers through their preregistered payment methods. The company uses a third-party messaging platform (like WhatsApp) to manage customer profiles and send marketing messages. To use this service, customers must join the company’s messaging app channel and register their payment methods through the messaging app.

2.1. Marketing Intervention

As part of the customer activation process, the company sends a 15% discount coupon to every new customer following their first purchase. The coupon is automatically applied to the next purchase made within 14 days, after which it expires. The company considered offering additional coupons to some newly acquired customers, but only if doing so would increase their total purchases in the subsequent months. To develop a targeted approach, they conducted a randomized controlled experiment to identify customers who would increase their purchases as a result of such an intervention.

In the experiment, the company randomly assigned customers who had just made their first purchase to one of two groups: a control group ($W_i = 0$) that received one coupon (i.e., the “business-as-usual” case) or a treatment group ($W_i = 1$) that received three coupons. All coupons offered a 15% discount and expired after 14 days. The experiment included 1,853 customers, with 889 assigned to the treatment group and 964 to the control group. The company collected several pretreatment covariates on customers to design personalized interventions, such as acquisition channels and information about their first purchase. (Further details on these

covariates and randomization checks can be found in Table 2 in Section 6.)

2.2. Average Treatment Effect on Repeated Purchasing Behaviors

Before delving into customer-specific impacts, we first assess the average treatment effect of the intervention on overall customer purchases. In the week immediately after the intervention, the total number of purchases, represented as $Y_{i,1}$, increase by 10% over the control group, but this effect is not statistically significant ($p=0.62$). When we extend the observation horizon to 10 weeks,¹ the average treatment effect on 10-week purchases (denoted as $Y_{i,10}$) is 0.3153 (with $p=0.03$), corresponding to a 30% increase (where the control group had an average of $Y_{i,10}$ of 0.99). Clearly, the company's intervention had a long-lasting impact on repeated customer purchases, whereas the short-term impact was relatively small.

To gain deeper insights into the intervention's effects on repeated purchasing behaviors, we further investigate the differences in customer attrition and purchase frequency between the two treatment groups conditional on being "alive" (Figure 1). The leftmost figure illustrates the percentage of "alive" customers, defined as those making a purchase in a given week or afterward (up to 50 weeks after their initial purchase). The data suggests that the intervention reduced customer churn, with the treatment group consistently exhibiting higher percentages of "alive" customers than the control group. The rightmost figure in Figure 1 shows the weekly average number of purchases per active customer. For the first seven weeks, retained customers in the treatment group made more purchases on average than those in the control group.

2.3. Designing Targeted Interventions Through CATE Estimation

The primary objective of the focal firm was to identify customer segments with the most favorable responses

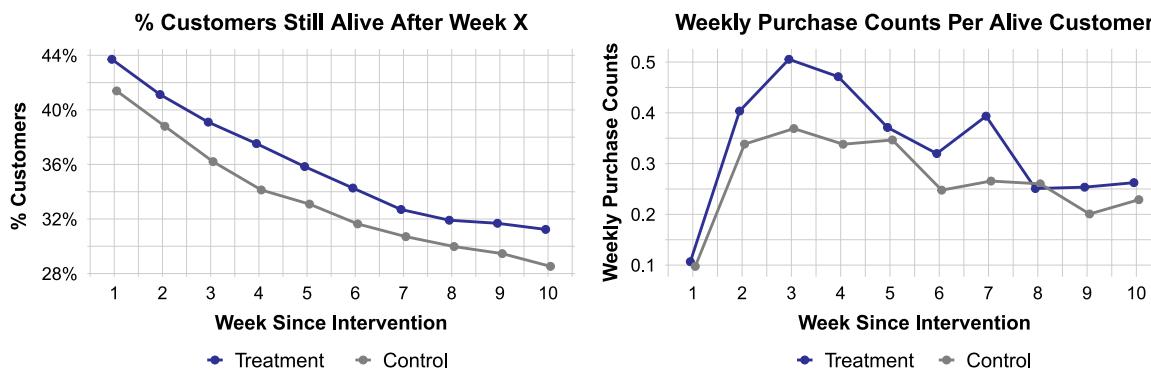
to the intervention with the goal of exclusively targeting these segments in future activation campaigns. With this in mind, the company aims to create a *treatment prioritization rule*, optimizing outcomes through targeted treatment assignments (Athey 2017, Ascarza 2018, Hitsch et al. 2023). We now demonstrate how the firm can use the experiment to achieve this aim.

Let's begin by examining a scenario where the focal firm's goal is to maximize customer purchases within the first week postintervention (i.e., $Y_{i,1}$). This approach aligns with prevalent coupon targeting literature (Dubé et al. 2017, Gubela et al. 2017), where the primary aim is enhancing immediate purchases after the intervention. To identify which customers should be offered additional coupons, we construct a CATE model for $Y_{i,1}$. Specifically, this model is designed to estimate the quantity $\tau_{Y_1}(\mathbf{X}_i) \equiv \mathbb{E}[Y_{i,1}(1)|\mathbf{X}_i] - \mathbb{E}[Y_{i,1}(0)|\mathbf{X}_i]$, where $Y_{i,1}(W_i)$ is the potential outcome (Rubin 1974) of customer i 's first-week purchase given the treatment condition W_i , and \mathbf{X}_i includes the pretreatment customer covariates mentioned previously.

To evaluate the model's performance, we apply a bootstrap validation method similar to that used in Ascarza (2018) (see Section 6.3 for details). Briefly, we first estimate a CATE model using the training set and predict CATEs for the validation customers. We then sort validation customers based on their predicted CATEs and group them into quintiles, with Q_1 containing customers with the highest predicted CATEs (i.e., those who are predicted to increase purchases the most because of the intervention), and Q_5 containing those with the lowest predicted CATEs. Finally, we evaluate the model's ability to identify the "right targets" by computing the *group average treatment effect* (GATE) for each quintile. We compute two measures: (i) the predicted CATEs (Prediction) and (ii) the actual outcome $Y_{i,1}$ (Data).

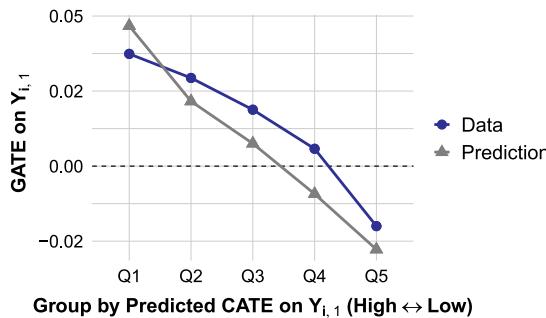
Figure 2 presents the predicted vs. actual GATEs for customers in the validation data set.² The closeness between predicted and actual GATEs indicates the

Figure 1. (Color online) Percentage of Alive Customers and Average Weekly Purchases After the Intervention



Note. Customers are labeled as "alive" if they made at least one purchase in that week or later (up to week 50).

Figure 2. (Color online) Predicted and Actual GATEs by Predicted CATE Levels When the Outcome Variable Is $Y_{i,1}$



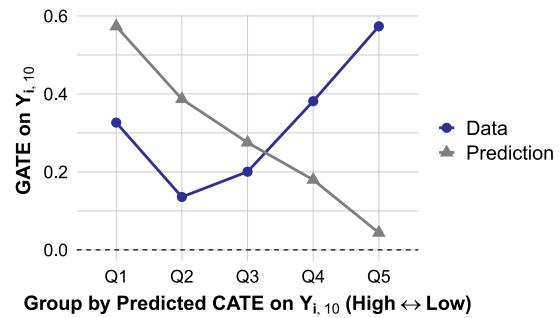
Notes. Groups Q_1, \dots, Q_5 are categorized based on the decreasing order of treatment effects predicted by the CATE model for $Y_{i,1}$. Hence, the predicted GATEs (triangles line) are monotonically decreasing by definition. Actual GATEs (circles line) are computed from $Y_{i,1}$. For example, the predicted and actual GATE on $Y_{i,1}$ for Q_1 are 0.046 and 0.037, respectively.

CATE model's accuracy in estimating the true treatment effect. Specifically, the target segment recommended by the model (Q_1) includes customers for whom the intervention was the most beneficial (with an actual GATE of 0.037, four times the ATE), whereas the do-not-touch segment (Q_5) includes customers for whom the intervention did not generate additional purchases (with an actual GATE of -0.0196). This suggests that the model can effectively rank customers according to their responsiveness to the intervention, enabling the firm to design targeted policies that maximize customer transactions within one week following the intervention.

However, the firm's primary goal is to stimulate purchases across a longer time frame, particularly in the 10 weeks following the intervention, rather than merely increasing immediate purchases. In theory, the same targeting approach should be applicable, with the *only difference* being the use of $Y_{i,10}$ as the outcome for estimating CATEs, that is, $\tau_{Y_{10}}(\mathbf{X}_i) \equiv \mathbb{E}[Y_{i,10}(1)|\mathbf{X}_i] - \mathbb{E}[Y_{i,10}(0)|\mathbf{X}_i]$, and comparing predicted and actual GATEs. Therefore, we replicate the same analysis, this time using the total purchases made within the 10 weeks following the intervention as the dependent variable.

Figure 3 presents the predicted vs. actual GATEs on $Y_{i,10}$ for the validation customers. The U-shaped trajectory of the actual GATEs underscores the inability of the CATE model to accurately rank customers by their treatment effects on $Y_{i,10}$. For instance, if the company chooses to target customers within Q_1 (those anticipated to show the strongest effect), the actual uplift from targeting this segment would be an increase of 0.41 purchases, contrasting the 0.58 predicted by the model. This divergence is even more apparent for customers predicted to benefit the least from the treatment (those in Q_5). Although the model predicts zero impact

Figure 3. (Color online) Predicted and Actual GATEs by Predicted CATE Levels When the Outcome Variable Is $Y_{i,10}$



Notes. Groups Q_1, \dots, Q_5 are categorized based on the decreasing order of treatment effects predicted by the CATE model for $Y_{i,10}$. Actual GATEs (circles line) are computed from $Y_{i,10}$.

for this group, in reality, targeting them would result in an uplift of 0.63 purchases, notably outperforming the outcome of targeting Q_1 . Consequently, targeting strategies based on this CATE model would be ineffective.

Why does the test-to-target approach succeed for optimizing short-term outcome ($Y_{i,1}$) but struggle with long-term outcome ($Y_{i,10}$)? Is this a general phenomenon that extends beyond this particular example? If so, how can marketers design targeted interventions to optimize long-term outcomes? We address these questions in the remaining of this article. Section 3 examines common consumer behaviors that drive the accumulation of unexplained variations and analyzes their consequences for CATE estimation and targeting when optimizing long-term outcomes. Section 4 introduces a general solution that uses the less noisy short-term behaviors to predict the long-term treatment effect while reducing unexplained variations, along with the proposed strategy to address customer attrition. In Section 5, we validate our solution through simulation analyses and explore the trade-off between information gain and noise accumulation. We demonstrate the superiority of our approach in a real-world marketing campaign in Section 6. Finally, we conclude in Section 7 and suggest several research directions for future work.

3. Problem: Unexplained Variations in Long-Term Repeated Purchases

The difficulty of targeting for long-term outcomes arises from two interrelated issues. First, long-term outcomes—particularly those involving repeated consumer interactions—tend to have high levels of unexplained variation due to the accumulation of unexplained customer behavior. Second, this noise accumulation problem can significantly undermine the precision of popular CATE models, resulting in sub-optimal targeting approaches. Our theoretical analysis

formalizes and generalizes this problem by examining two critical aspects: (i) identifying *when* and *why* the unexplained variance for long-term outcomes increases as the observation window expands and (ii) understanding the impact of noise accumulation on the accuracy of state-of-the-art CATE models and the consequent targeting performance. (Detailed proofs of all theoretical results are available in the Online Appendix A.)

3.1. Characterizing Long-Term Repeated Purchasing Outcomes

We examine a particular type of outcome variable: the *cumulative sum of recurring behaviors* over time. Such outcomes are commonly observed in marketing, ranging from repeated transactions in retail companies, to total engagement time for social media platforms, and customer lifetime value for Software-as-a-service businesses. For illustration throughout this paper, let's take the example of a retail company aiming to maximize total purchases over T periods. The firm plans to achieve this by targeting a marketing intervention, like distributing additional coupons to specific types of customers. Let $W_i \in \{0, 1\}$ denote the treatment assigned to customer i , and let \mathbf{X}_i represent the pretreatment covariates used for determining this assignment. The total purchases made by customer i over T periods is represented by $Y_{i,T}$. This is equivalent to the aggregate of transactions across all periods, expressed as $Y_{i,T}(W_i) \equiv \sum_{t=1}^T S_{i,t}(W_i)$. Here, $S_{i,t}(W_i)$ denotes the purchases by customer i in period t given the treatment W_i .

Our first objective is to investigate how the unexplained variations in $Y_{i,T}(W_i)$ —specifically the variation in $Y_{i,T}(W_i)$ that cannot be attributed to \mathbf{X}_i and W_i —evolves as T increases. To achieve this, we break down the long-term outcome in the following manner:

$$Y_{i,T}(W_i) = \underbrace{\sum_{t=1}^T \mathbb{E}[S_{i,t}(W_i) | \mathbf{X}_i]}_{= \mathbb{E}[Y_{i,T}(W_i) | \mathbf{X}_i]} + \underbrace{\sum_{t=1}^T \varepsilon_{i,t}^S}_{\equiv \varepsilon_i^{Y_T}},$$

where $\mathbb{E}[S_{i,t}(W_i) | \mathbf{X}_i]$ is the expected purchases in period t for customer i given treatment W_i and covariates \mathbf{X}_i , and $\varepsilon_{i,t}^S$ represents the mean-zero unexplained variations of purchases in period t . Then, the variance of unexplained variations in the outcome of interest is given by

$$\begin{aligned} \text{Var}[\varepsilon_i^{Y_T}] &= \text{Var} \left[\sum_{t=1}^T \varepsilon_{i,t}^S \right] = \sum_{t=1}^T \text{Var}[\varepsilon_{i,t}^S] \\ &+ 2 \sum_{1 \leq t_1 < t_2 \leq T} \text{Cov}[\varepsilon_{i,t_1}^S, \varepsilon_{i,t_2}^S]. \end{aligned}$$

This decomposition reveals that when there is nonnegative serial correlation in the per-period unexplained

variations (specifically, when $\text{Cov}[\varepsilon_{i,t_1}^S, \varepsilon_{i,t_2}^S] \geq 0$ for any $t_1 < t_2 \leq T$), the variance of these unexplained variations in $Y_{i,T}$ increases with the length of the observation period T . We next argue that common factors in marketing contexts, such as unobserved heterogeneity and customer attrition, often lead to positive serial correlation in these unexplained variations (Guadagni and Little 1983, Fader and Lattin 1993, Roy et al. 1996). Consequently, the existence of these factors in customer behavior typically results in a progressive accumulation of noise over time.

3.1.1. Unobserved Heterogeneity. When there are variations in customers' intrinsic preference toward the company—commonly known as *unobserved heterogeneity* or *individual fixed effects*—positive serial correlation in unexplained variations arises (Jones and Landwehr 1988, Gonul and Srinivasan 1993). To illustrate that, assume that the unexplained variation in each period can be expressed as

$$\varepsilon_{i,t}^S = \bar{\varepsilon}_i^S + \eta_{i,t}^S,$$

where $\bar{\varepsilon}_i^S$ represents the time-invariant individual purchase tendency (that is not captured by \mathbf{X}_i), and $\eta_{i,t}^S$ denote the independent per-period shock that is independent of $\bar{\varepsilon}_i^S$. Then, the serial correlation of $\varepsilon_{i,t}^S$ is positive because

$$\begin{aligned} \text{Cov}[\varepsilon_{i,t_1}^S, \varepsilon_{i,t_2}^S] &= \text{Var}[\bar{\varepsilon}_i^S] + \text{Cov}[\bar{\varepsilon}_i^S, \eta_{i,t_1}^S] + \text{Cov}[\bar{\varepsilon}_i^S, \eta_{i,t_2}^S] \\ &\quad + \text{Cov}[\eta_{i,t_1}^S, \eta_{i,t_2}^S] = \text{Var}[\bar{\varepsilon}_i^S] > 0. \end{aligned}$$

This result emphasizes that, when observable characteristics (\mathbf{X}_i) are insufficient to capture customer heterogeneity in the long-term outcome, the unexplained variations will exhibit positive serial correlations over time.

3.1.2. Customer Attrition. Also known as customer churn, customer attrition represents the progressive loss of customers over time. This phenomenon exists in many business contexts and has been extensively analyzed in the context of recurring purchase behaviors and customer relationship management (Schmittlein et al. 1987, Fader et al. 2005, Neslin et al. 2006, Ascarza et al. 2018b, Bachmann et al. 2021). In the following discussion, we demonstrate how customer attrition, whether observed or latent, results in a positive serial correlation in unexplained variations.

To get the intuition why customer attrition implies positive serial correlation, let's first consider the scenario where a customer churns in period t . In such a case, all future unexplained variations for this customer would be negative, given that their actual purchases reduce to zero. Conversely, if a customer remains active at time t , the unexplained variations from all earlier periods for this person would likely skew positive. This

occurs because the expected per-period purchases across all customers (comprising both churned and active individuals) tend to be lower than the realized purchases of a customer who remains active. Therefore, we can infer that the unexplained variation in per-period purchases exhibits positive serial correlations.

To formally demonstrate that customer attrition causes positive serial correlation of unexplained variations, we need to unpack the dynamics of how customer attrition affects $\varepsilon_{i,t}^S$. When a customer is still “alive” in period t , the actual purchase for that period might diverge from the expected purchase, and we denote this deviation as $\eta_{i,t}^S$. For simplicity, let’s assume these deviations are independent across distinct periods (i.e., there is no unobserved heterogeneity). Conversely, for a customer who has churned in period t (or before), the unexplained variation is the negative expectation of their purchase in that period, that is, $-\mathbb{E}[S_{i,t}(W_i)|\mathbf{X}_i]$. Consequently, the unexplained variation for period t can be expressed as

$$\begin{aligned}\varepsilon_{i,t}^S &= \mathbb{1}[i \text{ remains active at } t]\eta_{i,t}^S \\ &\quad - \mathbb{1}[i \text{ has churned at } t] \cdot \mathbb{E}[S_{i,t}(W_i)|\mathbf{X}_i].\end{aligned}$$

Given this formulation, the expected value of $\eta_{i,t}^S$ is positive, ensuring that $\mathbb{E}[\varepsilon_{i,t}^S] = 0$. This is consistent with the earlier intuition: Because the expected short-term purchase, denoted as $\mathbb{E}[S_{i,t}(W_i)|\mathbf{X}_i]$, is an average taken across both alive and churned customers, it follows that the expected purchase from an “alive customer” should surpass this average.

Proceeding to the correlation dynamics, we show in Online Appendix A.1 that the covariance between the unexplained variations for periods $t_1 < t_2$ is

$$\begin{aligned}\text{Cov}[\varepsilon_{i,t_1}^S, \varepsilon_{i,t_2}^S] &= [1 - \theta_{t_1}(W_i|\mathbf{X}_i)]\{\mathbb{E}[S_{i,t_1}(W_i)|\mathbf{X}_i] + \mathbb{E}[\eta_{i,t_1}^S]\} \\ &\quad \times \theta_{t_2}(W_i|\mathbf{X}_i)\{\mathbb{E}[S_{i,t_2}(W_i)|\mathbf{X}_i] + \mathbb{E}[\eta_{i,t_2}^S]\} \geq 0, \quad (1)\end{aligned}$$

where $\theta_t(W_i|\mathbf{X}_i)$ denotes the probability that the customer is still alive at t . Essentially, this covariance represents the comovement attributed to customer attrition at time t_1 , as it is the product of (i) the likelihood of customer churn at t_1 multiplied by the expected impact if the customer churns at this time (i.e., $[1 - \theta_{t_1}(W_i|\mathbf{X}_i)]\{\mathbb{E}[S_{i,t_1}(W_i)|\mathbf{X}_i] + \mathbb{E}[\eta_{i,t_1}^S]\}$), and (ii) the probability of the customer remaining active at t_2 combined with the expected purchase when alive at t_2 (i.e., $\theta_{t_2}(W_i|\mathbf{X}_i)\{\mathbb{E}[S_{i,t_2}(W_i)|\mathbf{X}_i] + \mathbb{E}[\eta_{i,t_2}^S]\}$). Given that every term in (1) is nonnegative, the resulting covariance is also nonnegative.

3.1.3. Other Customer Behaviors. Certainly, other behavioral factors can also contribute to positive (or negative) serial correlation. For instance, the presence of

state dependence, habit persistence, and psychological switching costs can lead consumers’ past consumption to positively influence their future consumption (Roy et al. 1996, Keane 1997, Seetharaman 2004, Dubé et al. 2010). Under such circumstances, a notable increase in previous purchases can boost subsequent purchases, resulting in positively correlated unexplained variations. Conversely, in settings where consumers exhibit stockpiling behavior (Tulin et al. 2002), there might exist a negative correlation between unexplained variations across different periods. It will depend on whether the strength of the stockpiling behavior is stronger than that of the unexplained variations driven by unobserved heterogeneity and attrition.

To summarize, the degree of noise accumulation is determined by two main factors: (i) the presence and significance of behavioral drivers in the data set and (ii) the effectiveness of using observable characteristics to predict these drivers. For instance, if the observable characteristics fail to comprehensively capture customers’ inherent preferences toward the company, this can lead to a significant accumulation of unexplained variations in long-term outcomes. Similarly, in situations where customer attrition is prevalent, the unexpected churn can lead to significant shocks in total purchases, resulting in increased unexplained variations in the long-term outcomes.

In our empirical application, we expect to have unobserved heterogeneity since the information from the initial purchase is unlikely to account for all the variation in individual preferences. Besides, Figure 1 demonstrates significant customer attrition, contributing to the accumulation of unexplained variation due to unpredictable churn. Moreover, stockpiling is unlikely in this context given the perishable nature of the goods sold in their vending machines. Taking all these factors into consideration, the focal firm is indeed facing the problem of increasing noise in their outcome of interest (i.e., long-term total purchases) as they increase the observation window. We present evidence supporting noise accumulation and positive serial correlation of unexplained variations in Section 6.2.

3.2. Implications for CATE Estimation and Targeting

The second objective of our theoretical analyses is to examine how noise accumulation in the outcome variable impacts the predictive accuracy of CATE models and their effectiveness in targeting. Within this context, we establish two key results: first, the variance in predicted CATEs escalates with increasing unexplained variations; and second, as this variance grows, the likelihood of making incorrect targeting decisions also increases.

In our analysis, we assume the firm has conducted a randomized controlled experiment such that the

complete randomization, overlap, and no interference assumptions hold (Imbens and Rubin 2015). Additionally, we focus on a broad spectrum of CATE models, as detailed in Assumption App-2 in Online Appendix A.2. A defining feature of these CATE models is that the prediction for a new individual can be represented as a weighted average of residualized outcomes from the training data, where the weights are determined by the degree of similarity in covariates between individuals in the training set and the new individual and estimated using honest estimation (Athey et al. 2019b). Besides, the residualization function is constructed using cross-fitting (Newey and Robins 2018). Essentially, this class includes popular models such as Causal Forest (Wager and Athey 2018), S-learners and T-learners with different outcome models (Künzel et al. 2019), and R-learners with a variety of second-stage estimators (Nie and Wager 2021, Kennedy 2023). Online Appendix A.2 provides a formal characterization of these models.

The following theorem formally establishes the relationship between the magnitude of unexplained variations and the variance of state-of-the-art CATE models.

Theorem 1 (Variance of CATE Prediction). *Assume that the CATE model (denoted as $\hat{\tau}_{Y_T}$) belongs to the general class described previously. Then, the variance of the predicted CATE for an individual with covariates \mathbf{x}_{new} scales with the amount of unexplained variations in the outcome variable. Mathematically, for given the same level of observed heterogeneity in the training set (i.e., $\{\mathbf{X}_i, W_i\}_{i=1}^N$), there exists two constants C_1 and C_2 such that*

$$C_1 \text{Var}[\varepsilon_i^{Y_T}] \leq \text{Var}[\hat{\tau}_{Y_T}(\mathbf{x}_{\text{new}}) | \{\mathbf{X}_i, W_i\}_{i=1}^N] \leq C_2 \text{Var}[\varepsilon_i^{Y_T}], \quad (2)$$

when $\text{Var}[\varepsilon_i^{Y_T}] \geq \sigma_0^2$ for some σ_0^2 .

Theorem 1 shows that the upper and lower bounds of the variance of the predicted CATE are proportional to the variance of the unexplained variation in the outcome variable when it is nontrivial, with both bounds increasing monotonically in $\text{Var}[\varepsilon_i^{Y_T}]$. Consequently, the unexplained variation in the outcome variable, irrespective of its impact on the consistency of CATE estimators, leads to instability in CATE models. This insight is especially crucial for practitioners focusing on optimizing long-term outcomes, as accumulating unexplained variations can significantly heighten the variance of the CATE estimator.

Next, we assess how often targeting decisions informed by the CATE model diverge from the optimal targeting strategy, which involves targeting customers with positive true CATEs. The result in Theorem 1 leads to the following.

Corollary 1 (Mistargeting Probability). *Suppose the company has an unbiased CATE model $\hat{\tau}_{Y_T}$ and targets only customers predicted to have a positive treatment effect.³*

Then, the probability of the model deviating from the optimal policy, $\mathbb{P}[\hat{\tau}_{Y_T}(\mathbf{x}_{\text{new}}) \cdot \hat{\tau}_{Y_T}(\mathbf{x}_{\text{new}}) < 0]$, increases with the variance of the predicted CATE, that is, $\text{Var}[\hat{\tau}_{Y_T}(\mathbf{x}_{\text{new}})]$.

Corollary 1 reveals a fundamental challenge in targeting for long-term outcomes: Instability in CATE predictions caused by unexplained variations in the outcome variable can result in erroneous targeting decisions. Therefore, regardless of the application domain, companies will invariably encounter challenges in accurate targeting when faced with substantial unexplained variations in the outcome variable.

A possible solution to the high variance problem in optimizing long-term outcomes is the use of CATE estimators specifically tailored to reduce prediction variance. This can be achieved by constraining the complexity of the CATE model, such as applying lasso regularization during the estimation process of R-learners (Nie and Wager 2021). Although these estimators are effective at reducing variance, they often introduce substantial underfitting bias. Our empirical analyses indicate that in cases with limited sample sizes, this bias can significantly impair targeting accuracy (refer to Table 1 for empirical results). Therefore, the effective use of regularization necessitates large sample sizes.

Finally, there are alternative approaches for developing targeting policies. These methods, known as policy learning, directly leverage experimental data to learn the optimal allocation of treatments. They typically derive a proxy variable for the true CATE, such as the inverse probability weighting (Kitagawa and Tetenov 2018) or doubly robust scores (Athey and Wager 2021) and then estimate a policy (often using machine learning models) to determine which customers the firm should target. Because these methods depend on a proxy variable derived from experimental data, they are conceptually similar to the targeting approach examined in this section. Therefore, they encounter a similar challenge: substantial unexplained variations in the outcome variable can increase the variance of the CATE proxy, resulting in ineffective targeting policies. We provide additional empirical evidence in Section 6.4.2.

3.3. Summary

Our theoretical analysis reveals that for outcomes based on repeated customer behaviors, factors like unobserved heterogeneity and customer churn are key contributors to unexplained per-period variations. Over time, these variations accumulate, resulting in increased variance in outcomes as the observation period lengthens. As a result, the accuracy of conventional CATE models is compromised due to this noise accumulation, leading to less effective targeting decisions. These insights highlight a critical challenge in estimating CATEs and developing effective targeting strategies for long-term business outcomes.

4. Solution: Surrogate Index with Separate Imputation

Thus far, we identified that the primary issue with ineffective long-term targeting arises from the accumulation of unexplained variations in the outcome of interest. Therefore, a potential solution to improve targeting effectiveness is to reduce the unexplained variations in long-term outcomes that are not attributable to the intervention. To achieve this, we suggest using a *noise-reduced proxy* in place of the actual outcome variable when estimating the CATE model. This proxy aims to capture the long-term treatment effect heterogeneity while excluding the variations unrelated to the intervention.

4.1. Solution Overview

We use the *surrogate index* (Athey and Wager 2019, Yang et al. 2023) to construct the noise-reduced proxy. This index represents the expected long-term outcome, derived from a set of *observed short-term outcomes*—such as immediate postintervention purchases in our case—and pretreatment covariates. To construct a surrogate index, companies can use historical data to build a model that identifies the relationship between short-term behaviors (along with customer covariates) and the actual long-term outcome. After establishing such a model, companies can collect short-term data, use it to predict long-term outcomes, and then assess the impact of their intervention based on these predicted long-term results. This approach offers three major advantages:

1. *Noise Reduction*: By design, the surrogate index incorporates only unexplained variations from the short-term outcomes in the surrogate model. Therefore, the predicted long-term outcome (or proxy) excludes the unexplained variations of behaviors occurring after these short-term outcomes were collected. As a result, the surrogate index includes *fewer unexplained variations* (i.e., those unrelated to the intervention) compared with the actual long-term outcome.

2. *Long-Term Orientation*: The surrogate index effectively represents the long-term effect of the intervention. Many marketing interventions result in long-term shifts in consumer behaviors by initiating immediate behavioral modifications. In such scenarios, short-term outcomes frequently account for a significant portion, if not all, of the long-term treatment effect. For example, the coupon promotion in our empirical application can reduce customer churn in the initial week following the intervention, as illustrated in Figure 1, and this enhanced retention eventually results in increased aggregate long-term purchases due to the extended lifetime of customers who received more promotions. The short-term purchase data in this scenario reflects the enhanced retention, enabling the surrogate index to use this information to forecast the sustained, long-term retention improvements.

3. *Acceleration of Decision Time*: For constructing a surrogate index, companies require: (i) a model that links short-term signals with long-term outcomes, which is estimated using historical data; (ii) preintervention covariates that are available before implementing the intervention; and (iii) short-term outcomes that are observed immediately following the intervention. As a result, companies can develop targeting strategies for optimizing long-term outcomes soon after the intervention (i.e., as soon as the short-term signals are observed), thus bypassing the delay associated with waiting for observing the actual long-term outcomes.

The first benefit of the surrogate index method highlights its potential to improve long-term targeting effectiveness. By reducing unexplained variations in the outcome variable, it enhances CATE estimation precision (as per Theorem 1), leading to more accurate targeting decisions (Corollary 1). The second benefit confirms the validity of using the surrogate index for estimating long-term treatment effects: although such targeting policies mainly depend on short-term behaviors, they can be effective if these behaviors explain a significant proportion of the long-term treatment effect. The third benefit, although not essential in our case since the focal company can postpone the implementation of new targeting rules, might be crucial in situations where rapid decision-making can lead to substantial savings in organizational costs.

We next provide the formal definition of a surrogate index, specify the conditions required for it to capture long-term treatment effects (benefit 2), and illustrate its variance reduction property (benefit 1).

4.2. Identification and Variance Reduction Using Surrogate Index

Assume that the company has access to two data sets: the experimental data with the intervention (denoted as \mathcal{E}), and the historical data without the intervention (denoted as \mathcal{H}). The *surrogate index* is defined as follows.

Definition 1 (Surrogate Index). The surrogate index is the expected long-term outcome ($Y_{i,T}$) of customers in \mathcal{H} , conditioned on their short-term behaviors ($\mathbf{S}_{i,T_0} = \{S_{i,1}, \dots, S_{i,T_0}\}$ for some $T_0 < T$) and pretreatment covariates (\mathbf{X}_i). Mathematically, it can be represented as $\tilde{Y}_T(\mathbf{S}_{T_0}, \mathbf{X}_i) \equiv \mathbb{E}_{\mathcal{H}}[Y_{i,T} | \mathbf{S}_{T_0}, \mathbf{X}_i]$.

To ensure identification of CATEs using the short-term signals, the following assumptions are made (Athey et al. 2019a).

Assumption 1 (Identification Assumptions for Long-Term CATEs). For the identification of the long-term treatment effect through a surrogate index, the following assumptions are made:

1. (*Surrogacy*) *The short-term outcomes can fully mediate the treatment effect of W_i on $Y_{i,T}$; that is, $W_i \perp\!\!\!\perp Y_{i,T} | \mathbf{S}_{i,T_0}, \mathbf{X}_i, \forall i \in \mathcal{E}$.*

2. (Comparability) The experimental and historical data are comparable in distribution; that is, $Y_{i,T} | \mathbf{S}_{i,T_0}, \mathbf{X}_i, i \in \mathcal{E} \stackrel{d}{\sim} Y_{i,T} | \mathbf{S}_{i,T_0}, \mathbf{X}_i, i \in \mathcal{H}$.

The surrogacy assumption states that the variations in short-term behaviors induced by the intervention can fully reflect its causal impact on the long-term outcome. Therefore, if this assumption holds, and we can accurately estimate the influence of \mathbf{S}_{i,T_0} on $Y_{i,T}$, then the long-term treatment effect can be deduced simply by observing the short-term outcomes of the experimental units. The comparability assumption ensures that the impact of \mathbf{S}_{i,T_0} on $Y_{i,T}$ is the same for both the experimental data and the historical data, which implies that one can use the historical data to infer the impact of \mathbf{S}_{i,T_0} on $Y_{i,T}$ and then extrapolate this relationship to the experimental data.

Altogether, when both of these assumptions hold true, the surrogate index representation becomes a reliable tool for estimating the CATE on the long-term outcome. Furthermore, because the surrogate index is based on short-term outcomes, it contains less unexplained variation compared with the actual long-term outcome. Formally, we state the theorem as follows.

Theorem 2 (Identification and Variance Reduction Using Surrogate Index). *Suppose that Assumption 1 holds.*

1. *The CATE of the intervention on the long-term outcome is equal to the CATE on the surrogate index, that is, $\tau_{Y_T}(\mathbf{X}_i) = \tilde{Y}_T(\mathbf{S}_{i,T_0}(1), \mathbf{X}_i) - \tilde{Y}_T(\mathbf{S}_{i,T_0}(0), \mathbf{X}_i)$, where $\mathbf{S}_{i,T_0}(W_i)$ denotes the potential outcome of the short-term outcomes.*

2. *The variance of the surrogate index is smaller than the variance of the actual long-term outcome, that is, $\text{Var}[\tilde{Y}_T(\mathbf{S}_{i,T_0}(W_i), \mathbf{X}_i)] < \text{Var}[Y_{i,T}(W_i) | \mathbf{X}_i]$.*

Theorem 2(1) suggests that the surrogate index offers an unbiased estimation of the CATE for the long-term outcome, assuming we can correctly estimate $\mathbb{E}_{\mathcal{H}}[Y_{i,T} | \mathbf{S}_{T_0}, \mathbf{X}_i]$. Meanwhile, Theorem 2(2) demonstrates that the surrogate index has lower variance compared with the actual long-term outcome. Combining these insights with those from Theorem 1 and Corollary 1, it becomes evident that employing the surrogate index method for CATE estimation can (i) reduce variance in CATE predictions and (ii) decrease the likelihood of mistargeting when the objective is optimizing the long-term outcome.

Importantly, the realization of these benefits depends primarily on the validity of Assumption 1, as discussed in Section 4.4. It also relies on the firm's capability to accurately estimate $\mathbb{E}_{\mathcal{H}}[Y_{i,T} | \mathbf{S}_{T_0}, \mathbf{X}_i]$ using available historical data, which we will discuss in the following section.

4.3. Separate Imputation Approach

Conventionally, researchers have constructed the surrogate index by using a regression model that associates

long-term outcomes with short-term results and preintervention covariates (Athey and Wager 2019, Yang et al. 2023). However, in many marketing scenarios, this method may encounter a challenge of high variance, particularly in the presence of *customer attrition*, as unexplained customer churn hinders the model's ability to distinguish between variations that can and cannot be explained by the short-term outcomes. In this section, we delve into the reasons for this shortcoming and introduce an novel solution that firms can seamlessly adopt when constructing the surrogate index.

4.3.1. Challenge: Inseparable Unexplained Variations

Arising from Attrition. When customer attrition exists, the total purchases of customer i over T periods ($Y_{i,T}$) are determined by two factors: the customer's active lifetime up to period T (denoted as \mathcal{T}_i^T), and their expected purchase per period while active (denoted as Λ_i^T). We can further break down these factors into variations that can be explained by \mathbf{S}_{i,T_0} and \mathbf{X}_i , and unobserved variations arising from individual preferences toward the firm not captured in the data (e.g., unobserved heterogeneity, random shocks, etc.). Hence, the aggregate purchase counts for customer i can be expressed as

$$\begin{aligned} Y_{i,T} &= \mathcal{T}_i^T \times \Lambda_i^T = \{\mathbb{E}[\mathcal{T}_i^T | \mathbf{S}_{i,T_0}, \mathbf{X}_i] + \varepsilon_i^T\} \\ &\quad \times \{\mathbb{E}[\Lambda_i^T | \mathbf{S}_{i,T_0}, \mathbf{X}_i] + \varepsilon_i^\Lambda\} \\ &= \mathbb{E}[\mathcal{T}_i^T | \mathbf{S}_{i,T_0}, \mathbf{X}_i] \mathbb{E}[\Lambda_i^T | \mathbf{S}_{i,T_0}, \mathbf{X}_i] \\ &\quad + \underbrace{\mathbb{E}[\Lambda_i^T | \mathbf{S}_{i,T_0}, \mathbf{X}_i] \varepsilon_i^T + \mathbb{E}[\mathcal{T}_i^T | \mathbf{S}_{i,T_0}, \mathbf{X}_i] \varepsilon_i^\Lambda}_{\equiv \xi_i(\mathbf{S}_{i,T_0}, \mathbf{X}_i), \text{ additional variations}} + \varepsilon_i^T \varepsilon_i^\Lambda, \end{aligned} \quad (3)$$

where ε_i^T and ε_i^Λ denote the unexplained variations in lifetime and purchase intensity, respectively.

The previous formulation reveals a key challenge in estimating the relationship between short-term and long-term outcomes: There is an additional term, $\xi_i(\mathbf{S}_{i,T_0}, \mathbf{X}_i)$, which connects explained variations in customer lifetime (attributable to short-term outcomes and observed covariates) with unexplained variations in per-period purchase intensity, and vice versa. This implies that directly modeling $Y_{i,T}$ using \mathbf{S}_{i,T_0} (and \mathbf{X}_i) becomes problematic, as the model may not effectively differentiate whether variations in $Y_{i,T}$ are driven by \mathbf{S}_{i,T_0} (and \mathbf{X}_i) or by ε_i^T and ε_i^Λ . This inseparability results in *high variance* in the surrogate model, as it struggles to separate the explainable and unexplainable variations present in the historical data. (Further discussion and an empirical example of this phenomenon are available in the Online Appendix B.)

Previous research has addressed the inseparability problem by introducing additional assumptions (Brown

1983, Roehrig 1988, Chesher 2003, Hoderlein and Mammen 2007, Imbens and Newey 2009). For example, a common method to tackle the multiplicative noise structure is to assume that the unexplained part in the outcome variable (e.g., $\xi_i(\mathbf{S}_{i,T_0}, \mathbf{X}_i)$ and $\varepsilon_i^T \varepsilon_i^\Lambda$ in (3)) is independent of the target variables (\mathbf{S}_{i,T_0} in our case) after controlling for observed covariates (Hoderlein and Mammen 2007, 2009; Su et al. 2019). However, this assumption does not hold in our scenario because $\xi_i(\mathbf{S}_{i,T_0}, \mathbf{X}_i)$ is clearly associated with \mathbf{S}_{i,T_0} even after controlling for \mathbf{X}_i . Another solution proposed in the literature is to use instrumental variables that correlate with \mathbf{S}_{i,T_0} but remain unassociated with the unexplained variations (Chernozhukov et al. 2007). Yet, this approach is rarely feasible in practice since it necessitates historical data that includes exogenous shocks capable of serving as instrumental variables when modeling the relationship between short-term and long-term outcomes.

Consequently, we propose a new imputation technique for creating the surrogate index, specifically designed to tackle the inseparability problem stemming from customer attrition.

4.3.2. Solution: Separate Models for Customer Attrition and Purchase Intensity. Upon examining (3), it becomes clear that the issue of inseparable unexplained variations can be addressed by separately estimating two relationships: first, between customer lifetime and short-term outcomes (i.e., $\mathbb{E}[\mathcal{T}_i^T | \mathbf{S}_{i,T_0}, \mathbf{X}_i]$), and another between purchase intensity and short-term outcomes (i.e., $\mathbb{E}[\Lambda_i^T | \mathbf{S}_{i,T_0}, \mathbf{X}_i]$), and combine the predictions of these two models afterward. We refer to this as the *separate imputation* approach as we determine the long-term outcome by combining predictions from two distinct surrogate models rather than making a direct prediction.

To be more specific, the separate imputation capitalizes on the measurable nature of both customer lifetime and purchase intensity. The first model predicts customer lifetime using \mathbf{S}_{i,T_0} and \mathbf{X}_i (denoted by $\widehat{\mathcal{T}}_T(\mathbf{X}_i | \mathbf{S}_{i,T_0})$), whereas the second one estimates the purchase intensity using \mathbf{S}_{i,T_0} and \mathbf{X}_i (denoted by $\widehat{\Lambda}_T(\mathbf{X}_i | \mathbf{S}_{i,T_0})$). After constructing the two surrogate models using the historical data, we predict the lifetime and purchase intensity for customers in the experimental data using their observed short-term outcomes ($\mathbf{S}_{i,T_0}(W_i)$) and pretreatment characteristics (\mathbf{X}_i). We then combine these predicted values by multiplication to create the surrogate index that will be used for CATE estimation and policy learning, that is,

$$\begin{aligned} \widehat{Y}_T^{\text{Sep}}(\mathbf{X}_i | \mathbf{S}_{i,T_0}(W_i)) &= \widehat{\mathcal{T}}_T(\mathbf{X}_i | \mathbf{S}_{i,T_0}(W_i)) \\ &\times \widehat{\Lambda}_T(\mathbf{X}_i | \mathbf{S}_{i,T_0}(W_i)), \quad i \in \mathcal{E}. \end{aligned}$$

By design, the estimation of $\widehat{Y}_T^{\text{Sep}}(\mathbf{X}_i | \mathbf{S}_{i,T_0}(W_i))$ is unaffected by $\xi_i(\mathbf{S}_{i,T_0}, \mathbf{X}_i)$, circumventing the high-variance

issue stemming from the multiplicative noise pattern seen in (3).

4.3.3. Implementing Separate Imputation in Practice. Our approach can be applied to various business contexts as long as the firm has access to historical data \mathcal{H} that can be used to calibrate the two surrogate models. In contractual settings, where customer churn is observable, creating two separate models is straightforward, as both customer lifetime and average purchases per alive period can be directly calculated from the available data. In noncontractual settings, such as the one in our empirical application, customer churn is not immediately apparent, which means that the actual customer lifetime cannot be directly obtained from the data. In such contexts, we utilize established metrics like *recency* and *frequency* (derived from observed customer behaviors) to approximate customer lifetime and purchase intensity. In our application, we employ the *last observed purchase time* up to T as a proxy for \mathcal{T}_i^T , and the *average number of purchases per period* up to $T_{i,T}$ as a surrogate for the average purchase intensity Λ_i^T .

While \mathcal{T}_i^T and Λ_i^T may not flawlessly represent the true customer lifetime or purchase intensity (as customers might churn after the last period of purchase), these metrics are useful proxies because (i) together, they provide sufficient information to infer customer lifetime and purchase intensity (Fader et al. 2005, 2010), and (ii) \mathcal{T}_i^T is largely reflective of customer lifetime (for example, customers who churn early typically show low recency), whereas Λ_i^T correlates strongly with purchase intensity (such as high average purchases per period indicating consistent frequent buyers). Therefore, using these proxies for estimating surrogate index is a justified and practical approach to address the challenge of inseparability in noncontractual settings.

To conclude, using $\widehat{Y}_T^{\text{Sep}}(\mathbf{X}_i | \mathbf{S}_{i,T_0}(W_i))$ as a substitute for the actual outcome in CATE estimation presents significant benefits for addressing the challenges of long-term targeting. First, this approach results in less noise compared with using the actual long-term outcome, as it omits unpredictable variations that occur between $T_0 + 1$ and T . Second, it enables valid inferences about the long-term treatment effect under Assumption 1. Last, using separate models for churn and purchase helps to more accurately assess the influence of short-term behavioral changes on long-term purchasing patterns. As illustrated in Sections 5 and 6, this separate imputation method outperforms other strategies commonly used by firms, including existing imputation techniques.

4.4. Potential Limitations and Implementation Considerations

Although the surrogate index method presents significant advantages, it is important to acknowledge its

limitations and the potential challenges encountered during implementation. Understanding these factors is essential for assessing the generalizability of the proposed approach and its applicability in different situations.

4.4.1. Existence of Less Noisy Surrogate Variables. Similar to Athey et al. (2019a), we use immediate outcomes per period as the surrogate variables in our empirical study. These variables can reflect potential long-term impact in two ways. First, being integral parts of $Y_{i,T}$, they inherently capture parts of the long-term treatment effect. Second, changes in short-term behaviors can be indicative of shifts in long-term purchasing trends. For example, an intervention that enhances customer retention, thereby increasing long-term purchases, would likely result in a greater proportion of customers in the treatment group making nonzero short-term purchases due to reduced churn. Similarly, if the intervention encourages sustainable purchasing habits, an increase in short-term purchase frequency among treated customers would be observed, signaling the habit formation influenced by the intervention.

Surrogate variables of this kind are particularly beneficial in contexts where the outcome involves recurring activities such as repeat purchases or engagement, commonly seen in industries like retail, subscription media, food service, and transportation. However, for sectors characterized by long purchase cycles, like the automotive industry, finding valid surrogate variables can be challenging. In these cases, our method may be less applicable, as it can be difficult to detect short-term indicators. Furthermore, when short-term outcomes are also noisy (e.g., when the data are obscured for privacy protection), the benefits of a surrogate index may not be significant. In such situations, the surrogate index method is not applicable given the lack of effective surrogate variables.

4.4.2. Choice of Surrogates. The surrogacy assumption requires that short-term outcomes *fully mediate* the long-term treatment effect. When this condition is not met, the CATEs identified using the surrogate index may diverge from the actual CATEs, leading to potential mistargeting errors (Yang et al. 2023). One way to mitigate the risk of surrogacy violation is to include a large number of surrogates (Athey et al. 2019a). For example, when addressing long-term outcomes such as repeated transactions, incorporating more periods into the surrogate index reduces the chances of violating this assumption. However, adding more periods also increases unexplained variations, which may decrease the effectiveness of targeting policies. This issue is formally characterized in the following.

Corollary 2 (Noise Accumulation of Surrogate Indices). *Suppose that Assumption 1 holds. When we build two surrogate indices based on different periods of short-term outcomes, where $T_0 > T'_0$, it follows that*

$$\text{Var}[\tilde{Y}_T(\mathbf{S}_{i,T_0}(W_i), \mathbf{X}_i)] > \text{Var}[\tilde{Y}_T(\mathbf{S}_{i,T'_0}(W_i), \mathbf{X}_i)].$$

Corollary 2 and the surrogacy assumption highlight the tradeoff between information gain and noise accumulation for optimal targeting performance. Specifically, including more periods of short-term outcomes improves the validity of the surrogacy assumption. However, it also introduces unexplained variations which in turn increase the mistargeting probability. Determining the optimal number of periods for estimating surrogate models requires empirical investigation, which is challenging due to the difficulty of directly testing the surrogacy assumption. In this paper, we address this question by comparing holdout targeting performance across surrogate indices constructed using varying numbers of periods. Our simulations and real-world data analysis indicate that surrogate models with fewer periods may significantly improve targeting effectiveness, even though they might risk violating the surrogacy assumption.

4.4.3. Violation of Comparability Assumption. In this work, we construct surrogate indices using historical data. This method is based on the assumption that the relationships between short-term and long-term outcomes remain consistent across both historical and experimental data sets. This assumption is expected to be valid in our empirical analysis, as our experiment includes *all* newly acquired customers, who share characteristics with those in the historical data set, and the company's product offerings and acquisition strategies remained unchanged. However, if the experiment targets a specific subset of customers (like those acquired through a particular channel), additional adjustments may be necessary to ensure the surrogate models, developed from historical data, remain generalizable (Miratrix et al. 2018, Sahoo et al. 2022).

5. Empirical Performance: Simulation Evidence

5.1. Simulation Setting

We first validate our solution using synthetic data. Our simulations consider a company implementing a marketing intervention with the aim to maximize total purchases ($Y_{i,T}$) over a 10-week period ($T=10$) after the intervention. We generate experimental data (\mathcal{E}) with accumulated unexplained variations over time. Specifically, we assume that the intervention has a direct and heterogeneous impact on churn probability and purchase intensity during the first three weeks, and no further impact after the fourth week. For simplicity, we

assume that the unexplained variations in purchase intensity (while alive) are independent across time. Therefore, the noise accumulation is mainly driven by unexplained customer attrition. More details regarding the simulation setting and the noise accumulation behaviors can be found in Online Appendices C.1 and C.2.

We also generate historical data (\mathcal{H}) for 5,000 customers using the data generating process for the control group, which reflects the company not implementing the intervention in the past. These data will be used to construct the surrogate models. For the main analysis, we select $T_0 = 3$ periods for constructing the surrogate index, ensuring compliance with the surrogacy assumption (Assumption 1). Additionally, we vary the number of periods in constructing the surrogacy index to explore the balance between gaining information and accumulating noise. The model specifications for creating surrogate indices are detailed in the Online Appendix C.5.

5.2. Comparison Methods

5.2.1. Alternative Imputation Methods. We compare our separate imputation approach to various imputation methods for surrogate index construction. The first technique we examine is the single imputation approach proposed by Athey et al. (2019a) and Yang et al. (2023). In this method, the outcome variable ($Y_{i,T}$) is regressed directly onto the short-term signals ($S_{i,1}, \dots, S_{i,T_0}$) and the preintervention covariates (X_i). However, this imputation approach does not account for the inseparable nature of unexplained variations stemming from customer attrition. As a result, this model offers a less accurate estimation for the impact of short-term outcomes on $Y_{i,T}$ compared with the separate imputation technique.

We also examine the beta-geometric (BG)/NBD model with covariates (Fader and Hardie 2007) as an alternative imputation approach. This model has not been traditionally suggested as an imputation technique in the surrogate index literature, but it could be a reasonable candidate in our context because it has been shown to be effective in capturing unobserved heterogeneity in customer attrition and purchase intensity. The accuracy of the BG/NBD model, along with similar variants like the Pareto/NBD model, heavily depends on the distributional assumptions and the assumed functional forms of the relationship between customer covariates. Consequently, in situations where the relationship between observed covariates and the heterogeneity of treatment effects is intricate, we expect the BG/NBD model to underperform when contrasted with more flexible frameworks, such as nonlinear regressions or machine learning models.

5.2.2. Alternative Variance Reduction Methods. In addition to using alternative imputation methods to

obtain the surrogate index, we investigate other techniques to reduce the variance in CATE estimation. One such technique is to explicitly regularize the CATE function during the estimation process. Here, we use R-learner with lasso regularization when estimating CATE function (Nie and Wager 2021) to study the effectiveness of regularization as a potential solution. Although regularization can reduce the variance of CATE models, it can also introduce significant underfitting bias (Hastie et al. 2009)—The penalty term from regularization may cause the model to overlook crucial data patterns, which can be particularly problematic when dealing with a small training sample with large unexplained variations.

Another obvious alternative to reduce variance is to increase the sample size of the experiment data. Although this is straightforward to evaluate when using synthetic data, this is not a feasible or optimal solution for enhancing targeting policies in practice. First, the number of customers who qualify for the intervention is often limited, which limits the sample size available to most firms.⁴ Second, even when companies can increase the experimental sample size, the rate at which the variance of CATE decreases with respect to sample size can be considerably slow when the noise level is high. Nonetheless, we incorporate this alternative approach in our simulation analysis to evaluate the potential advantages of scaling the experiment, as opposed to employing different dependent variables for CATE estimation.

5.2.3. Baseline Approaches. Finally, we examine two baseline methods commonly used in practice: the *default* approach and the *myopic* approach. The default approach targets customers based on their actual long-term outcome, $Y_{i,T}$. This approach is the simplest and most common, but it is expected to be ineffective due to the substantial unexplained variations in $Y_{i,T}$. The myopic approach targets customers based on their short-term performance, $Y_{i,T_0} = \sum_{t=1}^{T_0} S_{i,t}$ (i.e., based on their behavior only a few periods right after the intervention). This approach can avoid noise accumulation (as there is less unexplained variations in behavior up to T_0), but it may not yield optimal performance because it disregards the disparities between short-term and long-term treatment effects.

5.3. Evaluation Procedure

We assess targeting performance through 200 bootstrap replications and report the mean and standard deviation of key metrics. In each replication, we first create a training and a validation set. For each of the approaches considered, we use the training set to construct a CATE model $\hat{\tau}_{\hat{Y}}(X_i)$ using \hat{Y} as the dependent variable (e.g., $\hat{Y} = Y_{10}$ for the default approach,

$\hat{Y} = \hat{Y}_T^{\text{Sep}}$ for the proposed approach, etc.). We then calculate the area under the targeting operating characteristic curve (AUTOC) (Yadlowsky et al. 2021) using the actual long-term outcome ($Y_{i,T}$) on the validation set.

Specifically, AUTOC is constructed as follows. Given the predicted CATEs $\hat{\tau}_{\hat{Y}}(\mathbf{X}_i)$, the *targeting operator characteristic* (TOC) for $Y_{i,T}$ is defined as

$$\text{TOC}(\phi; \hat{\tau}_{\hat{Y}}) = \mathbb{E}[Y_{i,T}(1) - Y_{i,T}(0) | F_{\hat{\tau}_{\hat{Y}}}(\hat{\tau}_{\hat{Y}}(\mathbf{X}_i)) \geq 1 - \phi] - \mathbb{E}[Y_{i,T}(1) - Y_{i,T}(0)],$$

where $F_{\hat{\tau}_{\hat{Y}}}$ is the cumulative distribution function of the predicted CATEs. The TOC measures the incremental gains from targeting the top $\phi \times 100\%$ customers, as the difference in ATE between customers in the top $\phi \times 100\%$ CATE group and all customers. Then, the AUTOC is defined as

$$\text{AUTOC}(\hat{\tau}) = \int_0^1 \text{TOC}(\phi; \hat{\tau}) d\phi.$$

A model $\hat{\tau}_{\hat{Y}}$ is better than another model $\hat{\tau}_{\hat{Y}'}$ in identifying customers in the top $\phi \times 100\%$ CATE group if $\text{TOC}(\phi; \hat{\tau}_{\hat{Y}}) > \text{TOC}(\phi; \hat{\tau}_{\hat{Y}'})$. The AUTOC is a useful metric for evaluating the effectiveness of a CATE model because it quantifies how well the model ranks units based on their treatment effect, with a higher AUTOC indicating a more effective targeting (or treatment prioritization) rule.

5.4. Results

Table 1 shows the AUTOC values of different approaches. Each row corresponds to CATE models for a specific outcome variable, and the columns indicate the two methods used for CATE estimation and the two training sample sizes ($n = 1,000$ and $n = 50,000$). In Online Appendix C.6, we also provide the results from other CATE models (including S-learner and T-learner) to corroborate that our findings are not driven by a particular CATE model.

In the scenario of a small sample size ($n = 1,000$), there are several important findings. First, all the methods

that use short-term proxies as the dependent variable for CATE estimation have superior AUTOC performance (both higher mean values and smaller standard deviations) than the default approach. This implies that relying on short-term signals rather than the actual long-term outcome can significantly improve the effectiveness of targeting. Second, the separate imputation method consistently achieves the highest performance (highest AUTOC means with smallest standard deviations), regardless of models being used to estimate CATEs. In contrast, the single imputation performs worse than other short-term approaches. This finding highlights the importance of separating churn and purchase when creating a surrogate index. Third, although the BG/NBD technique fares better than using the actual outcome, it falls short of the performance achieved by the separate imputation method. This can be attributed to the BG/NBD approach being less efficient when the relationship between observed characteristics and key parameters of interest is complex, as it is the case in our simulation.

Next, we highlight the efficiency of the proposed solution with respect to sample size. In situations where separate imputation is applied to small experimental data ($n = 1,000$), its performance (AUTOC = 0.88 for both methods) nearly matches the performance observed with 50 times larger data sets (AUTOC = 0.92 without regularization and AUTOC = 0.93 with regularization). Notably, the difference in AUTOCs between the small and large data sets is lower for the separate imputation method compared with other strategies, which highlights the sample size efficiency of our proposed solution. Furthermore, when using a nonregularized CATE model, using short-term proxies for targeting within a small experimental data set yields better performance with AUTOC scores between 0.83 and 0.88, compared with the standard approach trained on a significantly larger data set, which achieves an AUTOC of 0.73. Taken together, these results make a strong case for the separate imputation technique, highlighting its ability to achieve near-optimal performance even with significantly fewer data points compared with larger datasets.

Table 1. Comparison of AUTOC Values for Different Outcomes and CATE Models

| Outcome \hat{Y} | $N = 1,000$ | | $N = 50,000$ | |
|------------------------|------------------------|---------------------|------------------------|---------------------|
| | Without regularization | With regularization | Without regularization | With regularization |
| Separate imputation | 0.88 (0.04) | 0.88 (0.13) | 0.92 (0.02) | 0.93 (0.02) |
| Single imputation | 0.83 (0.09) | 0.63 (0.40) | 0.91 (0.02) | 0.91 (0.02) |
| BG/NBD imputation | 0.84 (0.07) | 0.78 (0.30) | 0.91 (0.02) | 0.93 (0.02) |
| Myopic ($Y_{i,3}$) | 0.85 (0.06) | 0.81 (0.28) | 0.91 (0.02) | 0.93 (0.02) |
| Default ($Y_{i,10}$) | 0.55 (0.30) | 0.31 (0.45) | 0.73 (0.04) | 0.92 (0.02) |

Notes. Higher AUTOC reflects better prioritization rule. We average the results over 200 replications and show in parentheses the standard deviation. We use Causal Forest as the CATE model without regularization and R-learner with lasso regularization as the model with regularization. The performance of different CATE models is provided in Online Appendix C.6.

Finally, it is noteworthy that the performance of the R-lasso remains consistent across all methods when working with a large experimental data set ($n = 50,000$). This result suggests that when there is an abundance of samples, regularization can effectively address the noise accumulation challenge. However, in smaller sample settings, applying regularization to CATE models for the default approach can lead to a significant drop in targeting performance, with the AUTOC value plummeting to 0.31. Such an outcome is driven by R-lasso's inclination to underestimate treatment effect heterogeneity, particularly in contexts with limited data.⁵ Hence, firms should not depend exclusively on regularization for mitigating noise accumulation—Its effectiveness may realize only when dealing with large-scale experiments.

5.5. Tradeoff Between Information Gain and Noise Accumulation

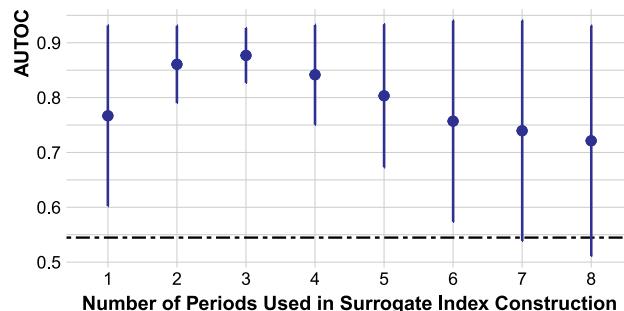
In the previous section, we emphasize cases where we incorporate only the minimal set of short-term outcomes ($T_0 = 3$) to meet the surrogacy assumption (Assumption 1(1)). However, verifying the surrogacy assumption in real-world applications is challenging, and companies must balance between information capture and noise accumulation when determining the number of periods to include in the surrogate index. Although incorporating more periods into the surrogate model could help fulfill the surrogacy assumption, it might also increase unexplained variations, potentially diminishing the targeting effectiveness (as discussed in Section 4.4.2).

We explore this tradeoff with our simulated data, simulating a real-world situation where firms determine the number of surrogates to include in their surrogate model. Specifically, we construct eight distinct surrogate models, each using a different number of short-term outcomes (ranging from $T_0 = 1$ to $T_0 = 8$) using the separate imputation approach. Following this, we generated 200 bootstrap replications and report the average and standard deviations of the AUTOC for each model. Figure 4 presents the results corresponding to the number of periods ranging from $T_0 = 1$ to $T_0 = 8$.

The inverted U-shaped relationship in Figure 4 reflects the tradeoff between information gain and noise accumulation. As we increase T_0 from one to three periods, the AUTOC improves because the intervention has a direct impact until the third week. However, the AUTOC starts to decline once we include behaviors beyond the third period (after satisfying surrogacy). This pattern is consistent with the guidance provided by Athey et al. (2019a) and Yang et al. (2023), recommending that companies should use the smallest set of short-term outcomes to create surrogate models, provided that the surrogacy assumption holds.

Interestingly, models using only one or two periods of information (where the surrogacy assumption is

Figure 4. (Color online) Tradeoff Between Information Gain and Noise Accumulation: An Analysis of Causal Forest AUTOCS with Surrogate Index Constructed Using Different Periods



Notes. Each point reports the average over 200 simulation replications together with the one standard deviation interval. The dashed line represents the mean AUTOC of the default approach. We used 1,000 customers in the training set. We present here the results of using Causal Forest, but our findings are robust across different CATE models. See Online Appendix C.8 for the results of different CATE models.

violated) outperform the default approach. This suggests that the benefits of noise reduction can outweigh the drawbacks of information loss. In other words, if the outcome of managerial interest (in our case, long-term cumulative purchases) has significant unexplained variations, violating the surrogacy assumption might not be the major concern. In turn, firms can improve targeting performance by using short-term outcomes in the surrogate models, even if these short-term behaviors do not capture the full impact of the intervention.

6. Empirical Performance: Real-World Application

This section evaluates the effectiveness of our proposed method using data from a retail technology company in Taiwan, as detailed in Section 2. We begin by offering additional information about the customer covariates, followed by an empirical demonstration of the noise accumulation issue. Last, we showcase how our proposed solution effectively identifies the most responsive customers, thereby enhancing profitability.

6.1. Observed Customer Covariates

The company gathered a collection of customer covariates at the time of their initial purchase. This set included information about their first transaction, such as total sales, item count, and product categories, as well as the location of the purchase (for example, a public-access vending machine) and whether the customer was referred by a friend. Given that these data are accessible to the company prior to the execution of the intervention, it can be used to estimate CATEs and determine whom to target. Table 2 presents the summary statistics

Table 2. Pretreatment Covariates and Comparison Across Two Experimental Conditions

| Variable | Treatment (N = 889) | Control (N = 964) | Difference p value |
|--|---------------------|-------------------|--------------------|
| $\log(\text{Sales})$ in the first transaction | -0.0205 | 0.0182 | 0.601 |
| $\log(\text{Quantity})$ in the first transaction | 0.0034 | -0.0031 | 0.930 |
| Was the first-visit fridge open to public? | 0.0035 | -0.0031 | 0.859 |
| Did the first purchase include any side dish item? | 0.0023 | -0.0020 | 0.865 |
| Did the first purchase include any dessert item? | 0.0002 | -0.0002 | 0.990 |
| Did the first purchase include any beverage item? | 0.0191 | -0.0169 | 0.332 |
| Did the first purchase include any lunch box item? | 0.0128 | -0.0114 | 0.461 |
| Did the first purchase include any item from other categories? | 0.0036 | -0.0032 | 0.697 |
| Was the customer referred by another customer? | -0.0016 | 0.0015 | 0.916 |

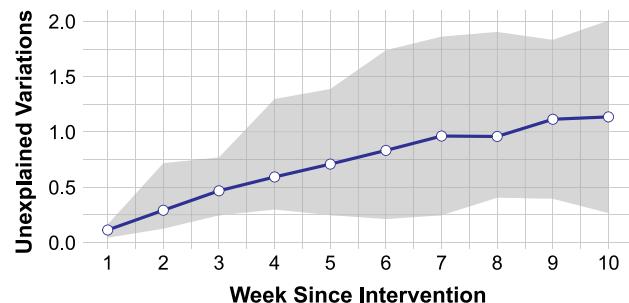
Notes. All continuous variables were first standardized with mean zero and variance one. All binary variables were first subtracted by the mean of all customers. We use the log scale for sales and quantity to create CATE models, as outliers in these variables may impact the performance of CATE models. However, there is no significant difference for the two variables in the original scale.

of these variables for both the treatment and control groups. To maintain the company's privacy, all data points have been standardized. The results confirm successful randomization, as no significant differences exist between the groups in any of the provided variables.

6.2. Evidence of Noise Accumulation

As discussed in Section 3.1, long-term outcomes often accumulate unexplained variations that deteriorate the accuracy of CATE models, leading to suboptimal targeting. In this section, we provide empirical evidence of the noise accumulation behavior in our specific context.

First, to determine the portion of the variability in $Y_{i,T}$ that can be explained by X_i and W_i , we construct a regression forest model that correlates $Y_{i,T}$ with X_i and W_i . We then estimate the unexplained variations by calculating the absolute differences between the actual purchases and their predictions from the model. Figure 5 shows the unexplained variations in total purchases over T periods, ranging from $T=1$ through $T=10$. In the figure, the dot indicates the median, and

Figure 5. (Color online) Unexplained Variations in $Y_{i,T}$ for $T = 1, \dots, 10$ 

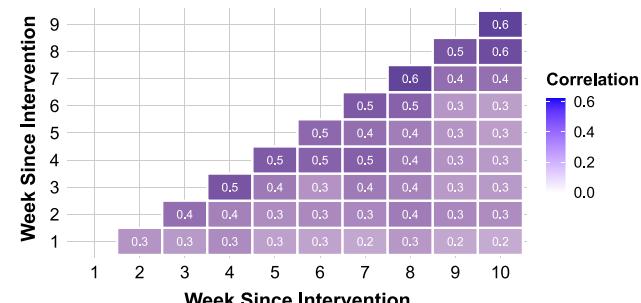
Notes. Each dot illustrates the median of unexplained variations for all customers in the experimental data. The gray shaded area presents the range between the highest 10% and the lowest 10% of these variations.

the gray shaded region covers the top and bottom 10% of values of the distribution of unexplained variations. Notably, the unexplained variations increase as the duration of the observation period extends.

Next, we demonstrate the positive serial correlations of unexplained variations across different periods. Using a similar approach, we create a regression forest model, $\widehat{\mathbb{E}}[S_{i,t} | X_i, W_i]$, to predict the number of purchases for each period. The residuals, $\widehat{\varepsilon}_{i,t}^S = S_{i,t} - \widehat{\mathbb{E}}[S_{i,t} | X_i, W_i]$, indicate the unexplained variations in $S_{i,t}$. We then examine the cross-correlation of these residuals, $\text{Cor}(\widehat{\varepsilon}_{i,t_1}^S, \widehat{\varepsilon}_{i,t_2}^S)$. Figure 6 presents a consistent positive correlation between the residuals for each $1 \leq t_1 < t_2 \leq 10$ highlighting a significant noise accumulation issue in our empirical context. As highlighted in Section 3.1, the observed positive correlation is typically attributed to unobserved heterogeneity and attrition, both of which are highly likely to occur in this empirical context.

6.3. Empirical Analysis

Similar to our analysis in Section 5, we assess the effectiveness of our targeting approach against various alternatives. Specifically, we compare it to the default and

Figure 6. (Color online) Cross-Correlation Matrix of Unexplained Variations in Each Period

Note. All correlation coefficients reported are significantly nonzero with $p < 0.01$.

myopic approaches, as well as two other imputation methods: single imputation and the BG/NBD model.⁶

We created surrogate indices using historical data from customers who were acquired at least 10 weeks before the experiment began ($n = 4,031$ customers), ensuring that sufficient historical data are available to model the relationship between short- and long-term outcomes.⁷ In all approaches, we use the first-week short-term outcome (i.e., $S_{i,1}$) to estimate the surrogate indices. We determine the number of short-term periods empirically. (See Online Appendix D.7 for details and complete set of results.)

6.3.1. Validation Approach and Key Metrics. To assess the targeting performance of each approach, we use a bootstrap validation scheme similar to that of Ascarza (2018). Specifically, we generate $B = 500$ data splits consisting of training (70%) and validation (30%) sets. In each split, we estimate CATE models using the training set, with distinct outcome variables (\tilde{Y}) serving as the dependent variable. Then, we predict the corresponding CATEs ($\hat{\tau}_{\tilde{Y}}$) for customers in the validation set.

Using the predictions for validation customers, we evaluate the effectiveness of each targeting approach based on two widely used metrics: the GATEs across predicted CATE quintile groups and the expected profit gained by targeting customers with positive predicted CATEs.

6.3.1.1. GATEs by Predicted CATE Levels. Similar to the analyses presented in Section 2, we start by dividing validation customers into quintile groups based on their predicted CATEs ($\hat{\tau}_{\tilde{Y}}$), with $Q_{1\tilde{Y}}$ having the highest predicted CATEs and $Q_{5\tilde{Y}}$ having the lowest predicted CATEs. Next, we calculate the GATE for each quintile group using the *actual* long-term outcome ($Y_{i,10}$):

$$\widehat{\text{GATE}}_{Y_{10}}(Q_k^{\tilde{Y}}) = \frac{\sum_{i: i \in Q_{k\tilde{Y}}, W_i=1} Y_{i,10}}{|\{i: i \in Q_{k\tilde{Y}}, W_i=1\}|} - \frac{\sum_{i: i \in Q_{k\tilde{Y}}, W_i=0} Y_{i,10}}{|\{i: i \in Q_{k\tilde{Y}}, W_i=0\}|}.$$

6.3.1.2. Expected Profitability of Targeting Policies. To compute expected profitability, we consider a policy that targets customers with positive predicted CATEs (i.e., $\pi^{\tilde{Y}}(\mathbf{X}_i) = \mathbb{1}\{\hat{\tau}_{\tilde{Y}}(\mathbf{X}_i) > 0\}$) and calculate the expected purchase counts in the next 10 weeks using the inverse-probability-weighted (IPW) estimator (Horvitz and Thompson 1952). Specifically,

$$\widehat{V}(\pi^{\tilde{Y}}) = \frac{1}{|\text{Validation Set}|} \cdot \sum_{i \in \text{Validation Set}} \left(\frac{\mathbb{1}[W_i = \pi^{\tilde{Y}}(\mathbf{X}_i)]}{\widehat{\mathbb{P}}[\pi^{\tilde{Y}}(\mathbf{X}_i) = W_i]} \right) Y_{i,10}, \quad (4)$$

where $\widehat{\mathbb{P}}[\pi^{\tilde{Y}}(\mathbf{X}_i) = W_i]$ is the (estimated) propensity score for customers who are assigned the same treatment by $\pi^{\tilde{Y}}$ as in the actual data.⁸ Although the treatment assignment in the data are random and independent of the derived targeting policy, we use IPW adjustment to account for any possible imbalances between treated and nontreated customers because the sample used for profit evaluation (i.e., validation customers who were assigned the same treatment in the actual data as $\pi^{\tilde{Y}}$ assigns for policy evaluation) is relatively small. The IPW adjustment is also frequently used in other marketing literature that uses randomized controlled experiments for policy evaluation (Hitsch et al. 2023, Yoganarasimhan et al. 2023).

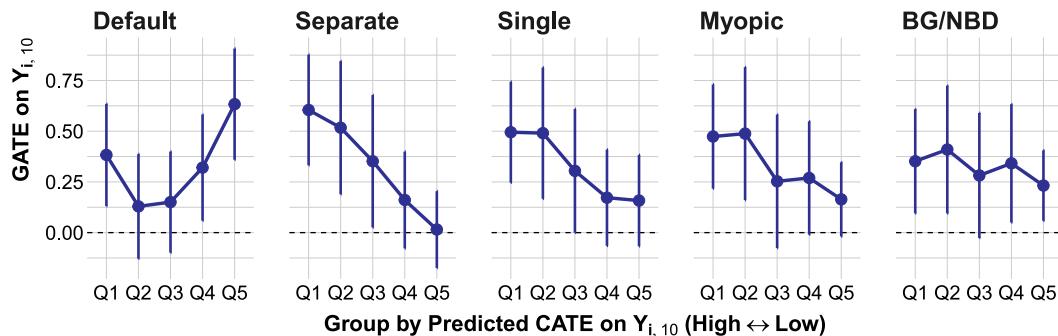
We then calculate the expected profit under policy $\pi^{\tilde{Y}}$ using the following formula:

$$\text{Profit}(\pi^{\tilde{Y}}) = \text{AOV} \cdot \left[p \cdot \widehat{V}(\pi^{\tilde{Y}}) - d \cdot \left(\frac{\sum_{i=1}^N \pi^{\tilde{Y}}(\mathbf{X}_i)}{N} \right) \cdot U^{W=1} - d \cdot \left(1 - \frac{\sum_{i=1}^N \pi^{\tilde{Y}}(\mathbf{X}_i)}{N} \right) \cdot U^{W=0} \right],$$

where AOV is the average order value,⁹ p is the average profit margin, $d = 15\%$ is the discount the coupon provided, U^W is the average number of coupons being used under the treatment condition W , and $\frac{\sum_{i=1}^N \pi^{\tilde{Y}}(\mathbf{X}_i)}{N}$ calculates the proportion of customers being treated under policy $\pi^{\tilde{Y}}$.

When assessing the profitability of different approaches, we also investigate the *policy learning* approach for determining targeting policies. As discussed in Section 3.2, using the actual long-term outcome for policy construction could lead to issues with noise accumulation. To address this, one could incorporate our proposed method into policy learning. This involves replacing the actual long-term outcome with a proxy that has reduced variance and subsequently applying policy learning using this proxy as the outcome variable.

Specifically, we use the doubly robust policy learning technique introduced by Athey and Wager (2021). Our method involves several steps: First, we estimate the doubly robust (DR) scores using various outcome variables (\tilde{Y}) as the proxy for CATE. We then develop targeting policies ($\pi_{\tilde{Y}}$) by creating cost-sensitive classifiers that predict which customers would have positive DR scores, using the DR score itself as the misclassification cost. Finally, we compute the expected profit improvement for each policy. Detailed information on the implementation is available in Online Appendix D.3.

Figure 7. (Color online) Actual GATEs by Predicted CATE Levels for Different Outcome Variables

Notes. Each point represents the mean of bootstrap results on the validation customers together with the one standard deviation interval. Groups $Q_1^{\tilde{Y}}, \dots, Q_5^{\tilde{Y}}$ are categorized based on the decreasing order of treatment effects predicted by CATE models for various outcome variables. GATEs are computed on the actual long-term outcome ($Y_{i,10}$). We present the results from T-learner as it gives the best targeting profitability. Our findings are robust across different CATE models. See Online Appendix D.5 for the results from other CATE models.

6.4. Empirical Results

6.4.1. GATEs by Predicted CATE Levels. Figure 7 shows the GATEs by predicted CATE groups. As discussed in Section 2, the U-shaped curve generated by the default approach indicates that the CATE model for $Y_{i,10}$ is unable to identify customers with the highest or lowest incremental effects. In contrast, all models that use short-term signals to estimate CATEs are more effective at ranking customers' long-term treatment effect than the default method.

Among all these models, the separate imputation method produces the best targeting performance as it generates the steepest curve. Specifically, the GATE for $Q_1^{\tilde{Y}}$ (representing the most sensitive customers identified by the method) is much larger than that of $Q_2^{\tilde{Y}}$, larger than that of $Q_3^{\tilde{Y}}$, and so on. Conversely, the BG/NBD model yields the least favorable result among all the proxies. This finding aligns with our intuition that the BG/NBD approach is likely to be ineffective when the parametric specifications of key parameters are different from the actual relationships, which is likely to be the case in this empirical application.

6.4.2. Profitability of Targeting Policies. We compare the expected profit of each targeting policy, denoted by

$\pi^{\tilde{Y}}$ (as described in Section 6.3.1), with the profit the company would obtain if it ran a uniform policy, π_0 , which applies the best intervention uniformly to all customers. In our scenario, given that the average treatment effect is positive, π_0 involves sending three coupons to every customer. Specifically, we define the profit improvement (PI) for each targeting approach as: $PI(\pi^{\tilde{Y}}) = \frac{\text{Profit}(\pi^{\tilde{Y}})}{\text{Profit}(\pi_0)} - 1$, where $\text{Profit}(\pi_0)$ is calculated in the same way as $\text{Profit}(\pi^{\tilde{Y}})$.

Table 3 presents the results of different approaches. The first column, labeled "Predicted CATEs," illustrates the profit improvement achieved when targeting rules are based on predicted CATEs.¹⁰ The second column, labeled "Policy Learning," details the profit improvement attained by using the doubly robust policy learning approach with various outcome variables.

Several key results are worth highlighting. First, consistent with our simulation results, the separate imputation approach (shown in the first row) consistently delivers the highest expected profits, whether we use predicted CATEs or policy learning for targeting. In contrast, the default method (last row) consistently leads to a loss in profit. This suggests that basing targeting decisions on predicted CATEs derived from actual long-term outcomes can detrimentally affect the long-

Table 3. Expected Profit Improvement: Targeting Based on Predicted CATEs or Targeting Using Policy Learning Based on Different Outcome Variables

| Outcome variable \tilde{Y} | Predicted CATEs | Policy learning |
|------------------------------|-----------------|-----------------|
| Separate Imputation | 5.81% (4.19%) | 4.57% (3.90%) |
| Single Imputation | 4.06% (4.51%) | 3.66% (4.75%) |
| BG/NBD Imputation | 0.95% (2.77%) | -0.26% (4.36%) |
| Myopic | 3.34% (3.60%) | 3.51% (4.58%) |
| Default | -3.52% (3.17%) | -3.44% (3.37%) |

Notes. We average the profit improvement over 200 replications and show in parentheses the standard deviations. For the "Predicted CATEs" approaches, we present the results from T-learner as it gives the best targeting profitability. Our findings are robust to using different CATE models.

term profitability, highlighting the substantial impact of noise accumulation. Interestingly, the myopic approach (fourth row) also mitigates much of the profit loss caused by high noise levels. This suggests that in certain scenarios, using less information—such as fewer observed periods—can, paradoxically, lead to better outcomes for long-term profitability.

Second, the separate and single imputation methods (first and second rows, respectively) outperform the myopic strategy in terms of profit. This highlights the importance of connecting short-term and long-term outcomes. However, the BG/NBD model (third row) falls short in performance compared with other short-term proxies, suggesting that flexible machine learning methods may be more advantageous in our empirical context. Last, the results obtained from policy learning are marginally less effective than those based on predicted CATEs, potentially due to the lower accuracy of the doubly robust scores compared with the predicted CATEs from the CATE model.

7. Conclusion and Future Directions

Firms often leverage targeted interventions to improve their business outcomes. An increasingly popular approach is to combine experimentation (or A/B testing) and customer data to predict each customer's sensitivity to an intervention. However, our research—both from theoretical and empirical perspectives—shows that this method can be ineffective in situations where outcomes are the result of recurrent behaviors that accumulate unexplained variations over time. To address this challenge, we propose a new targeting strategy that emphasizes *reducing the noise in outcome variables* prior to estimating CATE models. This method enhances the precision in estimating customers' long-term sensitivity to interventions, thereby significantly improving targeting efficacy.

Specifically, our proposed solution involves developing a *surrogate index* using *separate imputation* to effectively address the challenge of estimating long-term CATE. This method uses short-term behavioral changes to infer long-term treatment effects while distinctively accounting for the dynamics of customer attrition and purchase intensity. As a result, it effectively reduces the impact of unexplained variations when estimating CATE for long-term outcomes, offering significant performance improvements over current methodologies. Our solution is readily applicable using commonly accessible machine learning algorithms, rendering it practical for a broad range of businesses, spanning industries with both contractual and noncontractual customer relationships. By capitalizing on their existing historical purchase data, companies can improve their targeting strategies without the need for expanding the size of their experiments, thus avoiding extra costs.

We evaluate our proposed solution using both simulation analyses and a real-world marketing campaign, demonstrating superior targeting performance compared with existing practices. Our results also highlight the tradeoff between information gain and noise accumulation, emphasizing the importance of balancing these factors when determining the optimal number of short-term outcomes to include in a surrogate index model. Our findings indicate that when the long-term outcome is notably noisy, using a smaller set of short-term outcomes can outperform targeting strategies based on predicted CATEs of the actual long-term outcome, even in scenarios where the short-term outcomes do not entirely explain the long-term treatment effect. In practice, companies can conduct empirical testing to determine the most effective number of short-term outcome periods for inclusion in their surrogate models, thus maximizing their targeting performance.

Although our research provides valuable insights and solutions, there are limitations that suggest directions for future research. First, our proposed solution directly addresses the issue of *unobserved heterogeneity* in customer attrition and purchase intensity, which is prevalent in various marketing contexts (Fader and Hardie 2010, Ascarza et al. 2018a). However, other dynamics can cause more unexplained variations in the outcome variable, such as customer inertia and variety-seeking (Bawa 1990), state dependence (Roy et al. 1996, Dubé et al. 2010), or consumer learning (Erdem and Keane 1996). Incorporating these behaviors explicitly into surrogate models may further mitigate unexplained variations and enhance targeting performance. Furthermore, there are different modeling approaches available to connect the relationship between short-term and long-term outcomes, especially when we have multiple points in time for interventions. For example, Mazoure et al. (2021) proposes an innovative reinforcement learning framework that optimizes long-term customer engagement by combining immediate rewards with an estimate of residual value derived from future product usage. Thus, future research could explore the integration of these dynamics and develop new modeling approaches for surrogate index construction to enhance targeting performance.

Second, when the long-term outcome is a repeated purchase measure, it is natural to use short-term purchases after the intervention for surrogate index construction. However, when firms have different long-term objectives, there may not exist obvious short-term signals to use as surrogates. Therefore, it is essential to develop a general surrogate selection procedure and document potential surrogate outcomes for various marketing applications. For example, Han et al. (2021) proposes an estimation method to quantify the percentage of the long-term treatment effect that short-term surrogates can explain. Additionally, Yoganarasimhan et al. (2023) provides evidence that short-term conversion on subscription can be an

effective low-variance proxy for long-run revenue, and Wang et al. (2022) documents potential surrogate outcomes for the long-term user experience in the context of content recommendation. Future research could focus on identifying appropriate surrogate outcomes for different marketing contexts and developing methods to evaluate the effectiveness of these surrogates in improving targeting performance.

Third, there may be scenarios where no surrogate outcome is available for noise reduction, such as when the objective is to directly optimize a short-term outcome with significant unexplained variations. In such cases, future research could explore the development of new CATE models that are more resilient to noise in the outcome variable. For example, Huang and Ascarza (2023) proposes an iterative error correction procedure to improve CATE estimation when the data are intentionally masked by large noise for privacy protection. Furthermore, it would be worthwhile to investigate how to incorporate the estimation uncertainty of CATEs into the targeting strategy and determine whether it can further enhance the profitability of a marketing campaign.

Finally, the proposed imputation strategy relies on state-of-the-art machine learning methods to predict future purchases based on observed short-term behaviors. However, machine learning models may also overfit large unexplained variations in historical data, resulting in inaccurate long-term outcome predictions. Future research could explore alternative imputation strategies that are more robust to data noise. For instance, Padilla et al. (2023) proposes a Bayesian approach to predict purchase likelihood by incorporating information from intermediate stages in the customer journey. It would be worthwhile to investigate whether their approach can further mitigate the impact of unexplained variations in historical data.

Acknowledgments

The authors thank an anonymous company for providing the data used in this research; the review team, Bruce Hardie, Duncan Simester, Jeremy Yang, Liangzong Ma, and faculty and students in the Marketing Unit at Harvard Business School for valuable feedback; and participants of the Harvard Business School Doctoral Digital Workshop, 2022 American Causal Inference Conference, 2022 Marketing Science Conference, European Quantitative Marketing Seminar (EQMS), 2023 Marketing Dynamics Conference, and seminars at Harvard Business School, Massachusetts Institute of Technology, Cornell, Carnegie Mellon University, Michigan, Yale, Copenhagen Business School, Esade, WU Vienna, Northwestern, Stanford, and Washington University in St. Louis for helpful comments. T.-W. Huang has served on the advisory board for the focal company (who chose to stay anonymous) described in Section 6.

Endnotes

¹ We use 10 weeks to align with the time frame used by the focal firm when considering future purchases for newly acquired customers.

² Figures 2 and 3 present the results of using the X-learner to estimate CATE. The results are consistent across various CATE models, including Causal Forest, T-learner, and S-learner, as detailed in Online Appendix D.4.

³ This proposition also holds in the case where the firm aims to target customers with treatment effects larger than a certain threshold. See Online Appendix A for details.

⁴ Simester et al. (2022) propose an approach to calculate the sample size required to train and certify targeting policies.

⁵ To provide additional context, in approximately 60% of the bootstrap replications, the default method of R-lasso generated the same CATE prediction (which was equal to the ATE) for all consumers. This suggests a notable underestimation of the heterogeneity in treatment effects.

⁶ Unlike in simulations, we cannot substantially increase the sample size in a real-world context due to the finite number of customers the firm could acquire over time. Furthermore, we omit the results of the R-learner with lasso regularization, as it yielded identical CATE predictions for all customers. This result was expected, considering the limited sample size and the substantial unexplained variations present in the data.

⁷ We use random forest and BG/NBD to construct those surrogate models, which are described in detail in Online Appendix D.1.

⁸ We estimate this quantity in each iteration using the probability forest implemented by the grf package.

⁹ We did not observe a significant difference in AOV between the treatment and control groups (mean difference = \$0.05 with a *p* value of 0.88).

¹⁰ We present the results when using a T-learner to compute CATEs. For robustness, results from alternative CATE estimation methods are provided in Online Appendix D.6, which show consistent findings.

References

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

Ascarza E, Netzer O, Hardie BG (2018b) Some customers would rather leave without saying goodbye. *Marketing Sci.* 37(1):54–77.

Ascarza E, Neslin SA, Netzer O, Anderson Z, Fader PS, Gupta S, Hardie BG, et al. (2018a) In pursuit of enhanced customer retention management: Review, key issues, and future directions. *Customer Needs Solutions* 5:65–81.

Athey S (2017) Beyond prediction: Using big data for policy problems. *Science* 355(6324):483–485.

Athey S, Wager S (2019) Estimating treatment effects with causal forests: An application. *Observational Stud.* 5(2):37–51.

Athey S, Wager S (2021) Policy learning with observational data. *Econometrica* 89(1):133–161.

Athey S, Tibshirani J, Wager S (2019b) Generalized random forests. *Ann. Statist.* 47(2):1148–1178.

Athey S, Chetty R, Imbens GW, Kang H (2019a) The surrogate index: Combining short-term proxies to estimate long-term treatment effects more rapidly and precisely. Technical report, National Bureau of Economic Research, Cambridge, MA.

Bachmann P, Meierer M, Näf J (2021) The role of time-varying contextual factors in latent attrition models for customer base analysis. *Marketing Sci.* 40(4):783–809.

Bawa K (1990) Modeling inertia and variety seeking tendencies in brand choice behavior. *Marketing Sci.* 9(3):263–278.

Brown BW (1983) The identification problem in systems nonlinear in the variables. *Econometrica* 51(1):175–196.

Chen H, Harinen T, Lee JY, Yung M, Zhao Z (2020) CausalML: Python package for causal machine learning. Preprint, submitted 25, <https://arxiv.org/abs/2002.11631>.

Chernozhukov V, Imbens GW, Newey WK (2007) Instrumental variable estimation of nonseparable models. *J. Econometrics* 139(1): 4–14.

Chernozhukov V, Demirer M, Duflo E, Fernandez-Val I (2018) Generic machine learning inference on heterogeneous treatment effects in randomized experiments, with an application to immunization in India. Technical report, National Bureau of Economic Research, Cambridge, MA.

Chesher A (2003) Identification in nonseparable models. *Econometrica* 71(5):1405–1441.

Deng A, Xu Y, Kohavi R, Walker T (2013) Improving the sensitivity of online controlled experiments by utilizing pre-experiment data. *Proc. 6th ACM Internat. Conf. Web Search Data Mining (WSDM '13)* (Association for Computing Machinery, New York), 123–132.

Dubé JP, Hitsch GJ, Rossi PE (2010) State dependence and alternative explanations for consumer inertia. *RAND J. Econom.* 41(3): 417–445.

Dubé JP, Fang Z, Fong N, Luo X (2017) Competitive price targeting with smartphone coupons. *Marketing Sci.* 36(6):944–975.

Ellickson PB, Kar W, Reeder JC (2022) Estimating marketing component effects: Double machine learning from targeted digital promotions. *Marketing Sci.* 42(4):704–728.

Erdem T, Keane MP (1996) Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets. *Marketing Sci.* 15(1):1–20.

Fader PS, Hardie BGS (2007) Incorporating time-invariant covariates into the Pareto/NBD and BG/NBD models, <http://www.brucehardie.com/notes/019>.

Fader PS, Hardie BG (2010) Customer-base valuation in a contractual setting: The perils of ignoring heterogeneity. *Marketing Sci.* 29(1): 85–93.

Fader PS, Lattin JM (1993) Accounting for heterogeneity and nonstationarity in a cross-sectional model of consumer purchase behavior. *Marketing Sci.* 12(3):304–317.

Fader PS, Hardie BG, Lee KL (2005) “Counting your customers” the easy way: An alternative to the Pareto/NBD model. *Marketing Sci.* 24(2):275–284.

Fader PS, Hardie BG, Shang J (2010) Customer-base analysis in a discrete-time noncontractual setting. *Marketing Sci.* 29(6):1086–1108.

Gonul F, Srinivasan K (1993) Modeling multiple sources of heterogeneity in multinomial logit models: Methodological and managerial issues. *Marketing Sci.* 12(3):213–229.

Grimmer J, Messing S, Westwood SJ (2017) Estimating heterogeneous treatment effects and the effects of heterogeneous treatments with ensemble methods. *Political Anal.* 25(4):413–434.

Guadagni PM, Little JD (1983) A logit model of brand choice calibrated on scanner data. *Marketing Sci.* 2(3):203–238.

Gubela RM, Lessmann S, Haupt J, Baumann A, Radmer T, Gebert F (2017) Revenue uplift modeling. *Machine Learning for Marketing Decision Support*.

Guelman L, Guillén M, Pérez-Marín AM (2012) Random forests for uplift modeling: An insurance customer retention case. Engemann KJ, Gil-Lafuente AM, Merigó JM, eds. *Modeling and Simulation in Engineering, Economics and Management*, MS 2012: Lecture Notes in Business Information Processing, vol. 115 (Springer, Berlin, Heidelberg), 123–133.

Guelman L, Guillén M, Pérez-Marín AM (2015) Uplift random forests. *Cybernetic Systems* 46(3–4):230–248.

Guo Y, Coey D, Konutgan M, Li W, Schoener C, Goldman M (2021) Machine learning for variance reduction in online experiments. *Adv. Neural Inform. Processing Systems* 34:8637–8648.

Han L, Wang X, Cai T (2021) On the evaluation of surrogate markers in real world data settings. Preprint, submitted April 12, 2021, <https://arxiv.org/abs/2104.05513>.

Hastie T, Tibshirani R, Friedman JH, Friedman JH (2009) *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, vol. 2 (Springer, Berlin).

Hitsch GJ, Misra S, Walter Z (2023) Heterogeneous treatment effects and optimal targeting policy evaluation. Preprint, submitted November 6, <https://dx.doi.org/10.2139/ssrn.3111957>.

Hoderlein S, Mammen E (2007) Identification of marginal effects in nonseparable models without monotonicity. *Econometrica* 75(5):1513–1518.

Hoderlein S, Mammen E (2009) Identification and estimation of local average derivatives in non-separable models without monotonicity. *Econom. J.* 12(1):1–25.

Horvitz DG, Thompson DJ (1952) A generalization of sampling without replacement from a finite universe. *J. Amer. Statist. Assoc.* 47(260):663–685.

Huang TW, Ascarza E (2023) Debiasing treatment effect estimation for privacy-protected data: A model audition and calibration approach. Preprint, September 18, <https://dx.doi.org/10.2139/ssrn.4575240>.

Imai K, Ratkovic M (2013) Estimating treatment effect heterogeneity in randomized program evaluation. *Ann. Appl. Statist.* 7(1): 443–470.

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

Imbens GW, Newey WK (2009) Identification and estimation of triangular simultaneous equations models without additivity. *Econometrica* 77(5):1481–1512.

Imbens GW, Rubin DB (2015) *Causal Inference in Statistics, Social, and Biomedical Sciences* (Cambridge University Press, Cambridge, UK).

Imbens G, Kallus N, Mao X, Wang Y (2022) Long-term causal inference under persistent confounding via data combination. Preprint, submitted February 15, <https://arxiv.org/abs/2202.07234>.

Jin Y, Ba S (2023) Toward optimal variance reduction in online controlled experiments. *Technometrics* 65(2):231–242.

Jones JM, Landwehr JT (1988) Removing heterogeneity bias from logit model estimation. *Marketing Sci.* 7(1):41–59.

Keane MP (1997) Modeling heterogeneity and state dependence in consumer choice behavior. *J. Bus. Econom. Statist.* 15(3):310–327.

Kennedy EH (2023) Toward optimal doubly robust estimation of heterogeneous causal effects. *Electronic J. Statist.* 17(2): 3008–3049.

Kitagawa T, Tetenov A (2018) Who should be treated? Empirical welfare maximization methods for treatment choice. *Econometrica* 86(2):591–616.

Künzel SR, Sekhon JS, Bickel PJ, Yu B (2019) Metalearners for estimating heterogeneous treatment effects using machine learning. *Proc. Natl. Acad. Sci. USA* 116(10):4156–4165.

Lemmens A, Gupta S (2020) Managing churn to maximize profits. *Marketing Sci.* 39(5):956–973.

Manski CF (2004) Statistical treatment rules for heterogeneous populations. *Econometrica* 72(4):1221–1246.

Mazoure B, Mineiro P, Srinath P, Sede RS, Precup D, Swaminathan A (2021) Improving long-term metrics in recommendation systems using short-horizon reinforcement learning. Preprint, submitted June 1, <https://arxiv.org/abs/2106.00589>.

Mbakop E, Tabord-Meehan M (2021) Model selection for treatment choice: Penalized welfare maximization. *Econometrica* 89(2): 825–848.

Miratrix LW, Wager S, Zubizarreta JR (2018) Shape-constrained partial identification of a population mean under unknown probabilities of sample selection. *Biometrika* 105(1):103–114.

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

Newey WK, Robins JR (2018) Cross-fitting and fast remainder rates for semiparametric estimation. Preprint, submitted January 27, <https://arxiv.org/abs/1801.09138>.

Nie X, Wager S (2021) Quasi-oracle estimation of heterogeneous treatment effects. *Biometrika* 108(2):299–319.

Oprescu M, Syrgkanis V, Battocchi K, Hei M, Lewis G (2019) EconML: A Python package for ML-based heterogeneous treatment effects estimation. <https://github.com/py-why/EconML>.

Padilla N, Ascarza E, Netzer O (2023) The customer journey as a source of information. Preprint, submitted November 21, <https://dx.doi.org/10.2139/ssrn.4612478>.

Prentice RL (1989) Surrogate endpoints in clinical trials: Definition and operational criteria. *Statist. Medicine* 8(4):431–440.

Qian T, Yoo H, Klasnja P, Almirall D, Murphy SA (2021) Estimating time-varying causal excursion effects in mobile health with binary outcomes. *Biometrika* 108(3):507–527.

Roehrig CS (1988) Conditions for identification in nonparametric and parametric models. *Econometrica* 56(2):433–447.

Roy R, Chintagunta PK, Haldar S (1996) A framework for investigating habits, “the hand of the past,” and heterogeneity in dynamic brand choice. *Marketing Sci.* 15(3):280–299.

Rubin DB (1974) Estimating causal effects of treatments in randomized and nonrandomized studies. *J. Ed. Psych.* 66(5):688.

Sahoo R, Lei L, Wager S (2022) Learning from a biased sample. Preprint, submitted September 5, <https://arxiv.org/abs/2209.01754>.

Schmittlein DC, Morrison DG, Colombo R (1987) Counting your customers: Who are they and what will they do next? *Management Sci.* 33(1):1–24.

Seetharaman P (2004) Modeling multiple sources of state dependence in random utility models: A distributed lag approach. *Marketing Sci.* 23(2):263–271.

Simester D, Timoshenko A, Zoumpoulis SI (2020) Targeting prospective customers: Robustness of machine-learning methods to typical data challenges. *Management Sci.* 66(6):2495–2522.

Simester D, Timoshenko A, Zoumpoulis SI (2022) A sample size calculation for training and certifying targeting policies. Preprint, submitted October 06, <https://dx.doi.org/10.2139/ssrn.4228297>.

Su L, Ura T, Zhang Y (2019) Non-separable models with high-dimensional data. *J. Econometrics* 212(2):646–677.

Tulin E, Susumu I, Keane Michael P (2002) A model of consumer brand and quantity choice dynamics under price uncertainty. *Quant. Marketing Econom.* 1:5–64.

Wager S, Athey S (2018) Estimation and inference of heterogeneous treatment effects using random forests. *J. Amer. Statist. Assoc.* 113(523):1228–1242.

Wang Y, Sharma M, Xu C, Badam S, Sun Q, Richardson L, Chung L, et al. (2022) Surrogate for long-term user experience in recommender systems. *Proc. 28th ACM SIGKDD Conf. Knowledge Discovery Data Mining (KDD '22)* (Association for Computing Machinery, New York), 4100–4109.

Yadlowsky S, Fleming S, Shah N, Brunskill E, Wager S (2021) Evaluating treatment prioritization rules via rank-weighted average treatment effects. Preprint, submitted November 15, <https://arxiv.org/abs/2111.07966>.

Yang J, Eckles D, Dhillon P, Aral S (2023) Targeting for long-term outcomes. *Management Sci.*

Yoganarasimhan H, Barzegary E, Pani A (2023) Design and evaluation of optimal free trials. *Management Sci.* 69(6):3220–3240.