Working Paper 23-023

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

Ta-Wei Huang
Eva Ascarza

Harvard Business School

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

Ta-Wei Huang
Harvard Business School

Eva Ascarza
Harvard Business School

**Working Paper 23-023**



Funding for this research was provided in part by Harvard Business School.

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

Ta-Wei Huang
Harvard Business School, thuang@hbs.edu

Eva Ascarza
Harvard Business School, eascarza@hbs.edu

Firms are increasingly interested in developing targeted interventions for customers with the best response, which 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 that such an approach may fail to improve long-term performance, and can even harm it, when the outcome of interest (e.g. repeated purchases or CLV) accumulates unobserved individual differences over time. Our theoretical analysis demonstrates that unexplained variations in the outcome variable can lead to inaccurate CATE estimates and incorrect targeting policies. To address this issue, we propose using a *surrogate index* that leverages less noisy short-term purchases for long-term CATE estimation and policy learning. Furthermore, we introduce the *separate imputation* strategy to handle the non-separable nature of churn and purchase in marketing contexts. This involves constructing two distinct surrogate models, one for the observed last purchase time and the other for purchase frequency. 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.

*Keywords*: long-run targeting, heterogeneous treatment effect, statistical surrogacy, customer churn, field experiments

*History*: First Version: October 19th, 2022. This version: April 28th, 2023.

---

## 1. Introduction

Recent advancements in business experimentation and artificial intelligence have revolutionized the way companies perform targeted interventions. By leveraging 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, commonly referred to as the conditional average treatment effect (CATE). This allows firms to target customers who are predicted to respond in a way that aligns the most with their goals (such as profits or purchases) — those with the highest predicted CATEs — and achieve their objectives more effectively. This approach has gained significant popularity among organizations, and some tech companies, such as Microsoft (Oprescu et al. 2019) and Uber (Chen et al. 2020), have taken a step further by making their software packages for CATE estimation open-source. This has enabled more companies to adopt this approach and develop highly precise targeted marketing interventions at scale.

The *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. 2022), 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 customer values, firms should target customers based on their *long-term sensitivity* to the intervention, measured as the causal impact of the intervention on the long-term business outcome.

However, our research shows that the conventional test-to-target approach can be ineffective and even harmful for optimizing long-term customer response, such as repeated transactions over an extended period following the intervention. Unlike short-term outcomes (e.g., immediate purchases after the intervention), long-term outcomes accumulate individual customer behaviors, such as unobserved heterogeneity, state dependence, or habit persistence, that cannot be explained by the observed customer characteristics. As a result, long-term outcomes not only carry information about the treatment effect (which is what CATE models aimed to capture), but also accumulate *unexplained variations* that persist over time. This presents a significant and understudied 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 and potentially harmful targeting strategy for desired long-term outcomes.

This paper has two main objectives. First, we examine the issue of CATE estimation when the outcome variable is noisy. We provide theoretical analyses that demonstrate the relationship between unexplained variations in the outcome variable and the predictive accuracy of a CATE model. Our analyses demonstrate that the variance of most existing CATE models increases when there is greater unexplained variation in the outcome variable. This problem is particularly severe when the outcome variable captures repeated customer behaviors over an extended period of time, which is often the case in marketing contexts. As we demonstrate, long-term marketing outcomes tend to accumulate unexplained customer behaviors over time. This may appear to contradict the expectation that long-term outcomes are less noisy because idiosyncratic shocks would eventually cancel out over time. However, unexplained variations in customer behaviors are often serially correlated due to unobserved heterogeneity and customer attrition, making long-term outcomes, such as repeated purchases or customer lifetime value, intrinsically noisy.

Secondly, we present a solution that enables firms to implement more effective targeted interventions and achieve better long-term performance. We suggest that firms use a *noise-reduced proxy* as the outcome variable for CATE estimation, instead of relying solely on the actual long-term outcome. This approach, although counter-intuitive, has the potential to reduce unexplained variations while capturing the heterogeneous impact of the intervention from short-term signals. To construct this proxy, we adopt the *surrogate index* approach (Athey et al. 2019a, Yang et al. 2022), which leverages historical data that is readily available to the firm to infer the relationship between short-term purchases and the actual long-term outcome. In addition to valid CATE estimation as described in Yang et al. (2022), we formally show that the unexplained variations in the surrogate index are smaller than those in the actual long-term outcome, leading to more accurate CATE estimation. As a result, this approach enables firms to effectively target customers based on their long-term sensitivity to the intervention and mitigate the problem caused by the unexplained variation in the long-term outcome.

Furthermore, we highlight that the conventional approach to constructing surrogate indices, as proposed in Athey and Wager (2019), Yang et al. (2022), may not be the most effective in situations when customer attrition exists, which is often encountered in marketing contexts. Therefore, we propose a novel imputation technique to address this issue. Our approach involves developing two separate models using historical data: one for predicting the time of a customer's last observed purchase (i.e., the proxy for lifetime) and another for predicting the average purchase frequency when a customer is still alive. We then combine the predictions of both models to estimate expected future purchases. This approach, which we call the *separate imputation* strategy, is distinct from other imputation methods in the literature on statistical surrogacy as it addresses the challenge of estimating long-term outcomes resulting from the non-separable nature of churn and purchase. 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 the effectiveness of leveraging "short-term signals" for CATE estimation rather than relying on the actual outcome. We compare the performance of CATE models based on different outcome variables and show that targeting using short-term signals, whether using short-term outcomes or surrogate indices, consistently outperforms the approach of directly targeting based on CATE models for the long-term outcome. Our findings are surprising as they suggest that to maximize long-term outcomes, firms should ignore the actual long-term outcome and instead rely on short-term outcomes and historical data. Furthermore, we demonstrate that our separate imputation approach achieves the best targeting performance. In the real-world application, targeting customers using the proposed solution can yield a 6% increase in profits compared to 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. Firstly, in addition to its targeting effectiveness, this solution is ideal because firms no longer need to wait for the long-term results to be observed, which can significantly delay decision-making (Athey et al. 2019a, Yang et al. 2022). With our approach, firms can simultaneously improve the profitability of targeted marketing and make faster decisions. Secondly, it can be easily implemented using standard machine learning models and existing software packages, enabling firms to quickly and efficiently deploy the model that generates the highest profit. Thirdly, the separate imputation approach is widely applicable to a variety of business settings, including retailers, e-commerce, apparel, and non-profit organizations. Finally, our approach does not require collecting additional experimental data and instead leverages readily available historical customer data. This means that firms can adopt our solutions without incurring the additional costs of increasing experiment size, making it a cost-effective solution for enhancing the effectiveness of targeted interventions.

Our research contributes to the literature in four stands. First, we identify practical challenges in designing and implementing targeting policies, and highlight how existing test-to-target approaches (e.g., Ascarza 2018, Simester et al. 2020, Yoganarasimhan et al. 2022, Ellickson et al. 2022) may be ineffective when firms aim to optimize customers' long-term outcomes. Our real-world application shows that interventions can have a combined effect on both churn and purchase frequency, and thus, it is crucial to consider both the short-term impact (increased purchase frequency) and potential long-term benefits (improved customer retention). However, unexplained variations in long-term purchases can significantly impact the predictive performance of CATE models, resulting in ineffective targeting. To tackle this issue, we propose a new targeting paradigm where firms reduce noise in the outcome variable *before* estimating any CATE models. Although ignoring the actual outcome of interest may seem counter-productive, we demonstrate how creating the "right" proxy using the surrogate index approach with proper imputation methods results in more effective targeting.

Second, we address a critical issue in estimating heterogeneous treatment effects. While significant work has focused on developing methods for estimation (Imai and Strauss 2011, Imai and Ratkovic 2013, Guelman et al. 2015, Grimmer et al. 2017, Chernozhukov et al. 2018, Künzel et al. 2019, Athey et al. 2019b, Kennedy 2020, Nie and Wager 2021) and policy learning (Manski 2004, Kitagawa and Tetenov 2018, Athey and Wager 2021, Mbakop and Tabord-Meehan 2021), this paper is the first, to the best of our knowledge, to formally establish the relationship between unexplained variations and the predictive accuracy of CATE models. Additionally, we bring behavioral insights from the marketing literature to understand why the high variance problem is prevalent in many different contexts. Our study provides important insights into the limitations of existing CATE models and highlights the need for robust solutions to estimate heterogeneous treatment effects.

Third, our study contributes to the literature on statistical surrogacy and long-term treatment effect estimation (Prentice 1989, Athey et al. 2019a, Yang et al. 2022, Qian et al. 2021, Imbens et al. 2022). 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. Furthermore, past research has primarily constructed surrogate indices using standard regression models, while our work highlights the importance of considering the data generating process for such indices.

Lastly, 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 (e.g., Deng et al. 2013, Guo et al. 2021, Jin and Ba 2021) in a critical way — we do not reduce variance by eliminating variations that can be explained by customer observables. Instead, we suggest only using the information that short-term signals and pre-treatment covariates can explain for targeting. This is essential to our solution as our objective is to overcome inaccurate CATE estimates caused by unexplained variations rather than high variance due to observed heterogeneity.

The paper is organized as follows. In the next section, we provide a motivating example that highlights the issue of ineffective long-term targeting. Section 3 provides a theoretical analysis of the impact of unexplained variations on CATE estimation and targeting and explores the behavioral mechanisms that drive the high levels of unexplained variations encountered in many long-term marketing outcomes. Section 4 introduces the surrogate index approach as a solution to identify the long-term treatment effect and reduce unexplained variations, along with the proposed strategy to address customer attrition when building the surrogate index. In Section 5, we validate our solution through simulation analyses and explore the potential 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.

## 2. Motivating Example

### 2.1. Identifying Sensitive Customers

We first use the data from our empirical application (described in detail in Section 6) to illustrate the challenge of long-term targeting. The data is from a food-tech company that sells fresh, pre-cooked meals through their vending machine network. The company distributes promotional coupons to newly acquired customers, with the goal of encouraging repeat purchases. The company is considering sending extra coupons to *specific* customers who respond positively to the intervention, meaning they will (hopefully) end up purchasing more due to the additional coupons. All coupons are sent at the same time, right after a newly acquired customer makes their first purchase, and have an effective period of two weeks, automatically applied to the customer's subsequent purchase. To develop a targeting policy that achieves this goal, the company conducted an A/B test. A randomly selected group of users received the additional incentive, while the control customers were in the "business-as-usual" group.

Let us examine a scenario in which the focal firm seeks to find a targeting policy that optimizes the average purchase counts of a customer within the first week after the intervention ($Y_{i,1}$). This situation is comparable to most literature on coupon targeting (e.g., Gubela et al. 2017, Dubé et al. 2017), where the objective is to maximize the immediate purchase after the intervention. To identify the customers who should receive extra coupons, we construct a CATE model based on $Y_{i,1}$ and evaluate the model's performance using a bootstrap validation method similar to that used in Ascarza (2018).[1] Briefly, we first estimate a CATE model ($\widehat{\tau}_{Y_1}$) 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 $\mathcal{Q}_1^{\widehat{\tau}_{Y_1}}$ containing customers with the highest predicted CATEs (i.e., those who are predicted to increase purchases the most because of the intervention), and $\mathcal{Q}_5^{\widehat{\tau}_{Y_1}}$ 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 GATEs using two measures: (i) the predicted CATEs (Prediction), and (ii) the actual outcome $Y_{i,1}$ (Data). If these two values are similar, a targeting policy derived from the CATE model would be effective for the firm. Specifically, the target segment recommended by the model ($\mathcal{Q}_1^{\widehat{\tau}_{Y_1}}$) would include customers for whom the intervention is beneficial (i.e., high actual GATE), while the do-not-touch segment ($\mathcal{Q}_5^{\widehat{\tau}_{Y_1}}$) would include customers for whom the intervention does not generate additional transactions. Figure 1a reports both predicted and actual GATEs for customers in the validation data.[2] The proximity between predicted and actual GATEs suggests that the prediction provided by the CATE model is close to the actual treatment effect. This implies that the model can effectively rank customers according to their sensitivity to the intervention, allowing the firm to create targeted policies that maximize customer transactions within one week of the intervention.

[1] See Section 6.2 for the details.

[2] Figure 1 shows the results when using X-learner for CATE estimation. All results in this paper are replicated if using different CATE models, including Causal Forest, T-learner, and S-learner. See Appendix C.3

**Figure 1** **Predicted and Actual GATEs by Predicted CATE Levels When (a) the Outcome Variable is $Y_{i,1}$ and (b) the Outcome Variable is $Y_{i,10}$.**

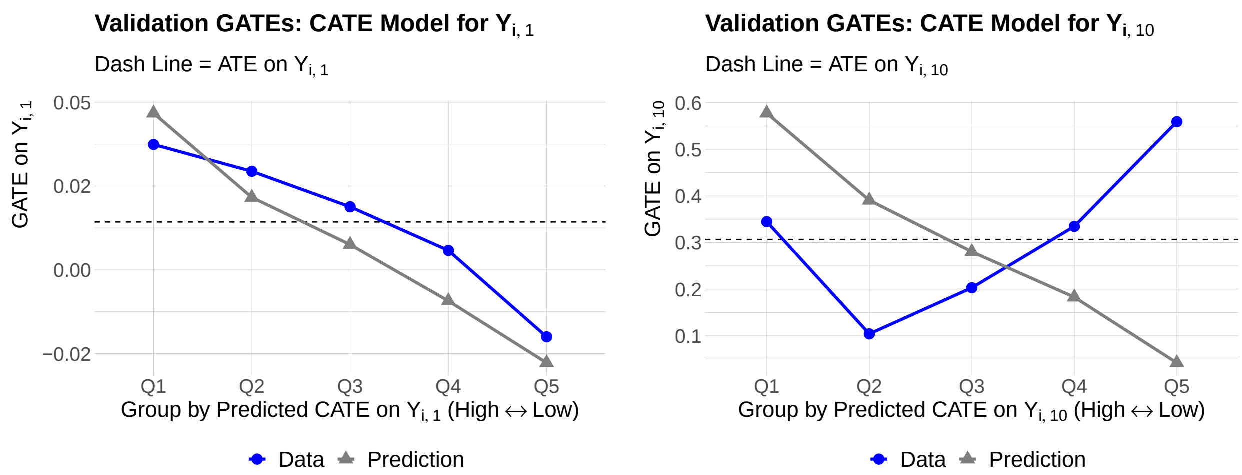


(a) CATE Model for $Y_{i,1}$

(b) CATE Model for $Y_{i,10}$

*Note.* Groups $\mathcal{Q}_1^{\hat{\tau}_\cdot}, \cdots, \mathcal{Q}_5^{\hat{\tau}_\cdot}$ are categorized based on the decreasing order of treatment effect predicted by CATE models for (a) $Y_{i,1}$ and (b) $Y_{i,10}$. Hence, the predicted GATEs (gray line) are monotonically decreasing by definition. Actual GATEs (blue line) are computed from the observed outcomes. For example, the predicted and actual GATE on $Y_{i,1}$ for $\mathcal{Q}_1^{\hat{\tau}_{Y_1}}$ in Figure 1a are 0.046 and 0.038, respectively.

### 2.2. Targeting for the Long-term Outcome

The primary objective of the focal firm, however, is not to increase transactions in the immediate future but to encourage purchases over a more extended time frame, such as ten weeks after the intervention.[3] In theory, the same approach for targeting should work, with the *only difference* being that $Y_{i,10}$ (i.e., cumulative number of purchases during the next 10 weeks) is used as the outcome when estimating CATEs and when comparing predictions with actual GATEs. Consequently, we repeat the same analysis, but now focus on purchase counts over the ten weeks following the intervention.

Figure 1b reports the predicted and actual GATEs on $Y_{i,10}$ for validation customers. The U-shaped curve of actual GATEs indicates that the CATE model cannot identify the correct rankings of treatment effects on $Y_{i,10}$. Specifically, if the company decided to target customers in $\mathcal{Q}_1^{\hat{\tau}_{Y_{10}}}$ (i.e., those predicted to have the greatest effect), the impact of targeting that group would only be 0.41 additional purchases, instead of the 0.58 predicted by the model. The discrepancy between predicted and actual effects is even more pronounced for

[3] We use ten weeks to align with the time frame used by the focal firm when considering future purchases for newly acquired customers. However, this duration may differ among companies.

customers with the least favorable predicted CATEs (i.e., those in $\mathcal{Q}_5^{\widehat{\tau}_{Y_{10}}}$), the difference between predicted and actual effects is even more pronounced — the predicted treatment effect for this group is almost zero, whereas targeting this group would actually increase transactions by 0.63, a much higher value than the result of targeting $\mathcal{Q}_1^{\widehat{\tau}_{Y_{10}}}$.

Why is the test-to-target approach ineffective for $Y_{i,10}$ while it works well for short-term outcomes? We argue that long-term outcomes, especially those that involve repeated interactions with users/consumers, are consistently impacted by high levels of noise, which can significantly reduce the accuracy of existing CATE models. Therefore, targeting policies based on the predictions from these models can lead to suboptimal outcomes and even reduce the profitability for the firm than not using any targeting at all.

## 3. The Problem: Unexplained Variations in Long-Term Outcomes

We now delve into the challenges associated with estimating CATEs for marketing outcomes over a long-term horizon. Our theoretical investigation first examines the influence of unexplained variations in the outcome variable on the effectiveness of standard CATE estimators. Specifically, we highlight how the presence of high levels of noise in long-term outcomes can significantly reduce the accuracy of existing CATE models, rendering them ineffective for targeting. Next, we characterize the data generating process of typical long-term marketing outcomes and demonstrate that the unexplained variance of such outcomes is increasing in the length of the observational period. (The proofs for all theoretical results are available in Online Appendix A).

### 3.1. The Firm's Problem

In targeting, companies aim to construct a *treatment prioritization rule* that optimizes a desired outcome, such as maximizing long-term transactions, through personalized treatment assignment. The most recent best practices (Athey 2017, Ascarza 2018, Hitsch et al. 2023) suggest companies to target customers based on the "incremental effects" of an intervention, which is characterized as the CATE on the outcome variable. Specifically, assume that the company performed a marketing intervention on their customers with two treatment conditions ($W_i \in \{0,1\}$). Then, the CATE is defined as

$$\tau_Y(\mathbf{X}_i) \equiv \mathbb{E}[Y_i(1)|\mathbf{X}_i] - \mathbb{E}[Y_i(0)|\mathbf{X}_i],$$

where $Y_i(W_i)$ is the potential outcome (Rubin 1974, Holland 1986) of customer $i$'s response given the treatment condition $W_i$, and $\mathbf{X}_i$ includes the pre-treatment customer covariates capturing the treatment effect heterogeneity.

To identify the CATE from an experimental data, we impose the following assumptions on the treatment assignment mechanism.

ASSUMPTION 1. ***[Intervention]*** *The treatment assignment in the experimental data satisfies the following assumptions:*

1. (Randomization) *The treatment assignment is independent of the potential outcome, i.e.,* $Y_i(1), Y_i(0) \perp\!\!\!\perp W_i$.
2. (Overlap) *The probability of receiving a treatment should be positive for all individuals in the population regardless of their covariate values, i.e.,* $0 < \mathbb{P}[W_i = 1|\mathbf{X}_i] < 1$.
3. (No Interference) *The potential outcome of a customer is not affected by the treatment assigned to other customers.*

In this paper, we assume complete randomization in the treatment assignment for simplicity. However, the theoretical results presented herein are readily extended to observational studies under the unconfoundedness assumption.

### 3.2. The Role of Unexplained Variations in CATE Estimation

We now present a theoretical framework to investigate the impact of *unexplained variations* in the outcome variable on the predictive accuracy of common CATE models. We define *unexplained variations* as the residual variations in the outcome variable that cannot be explained by $\mathbf{X}_i$ and $W_i$, i.e.,

$$\varepsilon_i(\mathbf{X}_i, W_i) \equiv Y_i(W_i) - \mathbb{E}[Y_{i,T}(W_i)|\mathbf{X}_i].$$

We assume a zero mean and finite variance for $\varepsilon_i(\mathbf{X}_i, W_i)$, and by construction its variance is identical to $\mathrm{Var}[Y_i(W_i)|\mathbf{X}_i]$.

In the theoretical analysis, we consider a wide range of CATE estimators including Causal Forests and other learners widely used in practice. Specifically, we assume the the following class of CATE estimators:

ASSUMPTION 2. ***[Class of CATE Estimator]*** *For a given test customer with covariates* $\mathbf{x}_{\text{test}}$*, the CATE predictor* $\widehat{\tau}_Y(\mathbf{x}_{\text{test}}) = \widehat{\mu}_Y^1(\mathbf{x}_{\text{test}}) - \widehat{\mu}_Y^0(\mathbf{x}_{\text{test}})$ *induces two weighted estimators of the residualized outcomes* ($\widehat{\mu}_Y^0$ *and* $\widehat{\mu}_Y^1$) *of the following form:*

$$\widehat{\mu}_Y^w(\mathbf{x}_{\text{test}}) = \sum_{i \in \mathcal{D}; W_i = w} \widehat{\ell}_i^w(\mathbf{x}_{\text{test}})[Y_i - \widehat{m}_Y^w(\mathbf{X}_i)],$$

*where $\mathcal{D}$ denotes the set of customers used to generate the prediction. We further impose the following assumptions about the estimation process:*

1. ***[Honest Estimation]*** *The weight function $\widehat{\ell}^w(\mathbf{x}_{\text{test}})$ is independent of $Y_j,\ \forall j \in \mathcal{D}$. This assumption suggests that $\widehat{\ell}^w(\mathbf{x}_{\text{test}})$ is constructed either only using the covariates or using a sample that is independent of $\mathcal{D}$.*
2. ***[Cross Fitting]*** *The adjustment function $\widehat{m}_Y^w(\cdot)$ is either equal to zero (i.e., no residualization) or created using a different sample that is independent of both $\mathcal{D}$ and the data employed to construct the weight function.*

Essentially, we examine a class of CATE estimators that predicts CATEs by computing the weighted average of (residualized) outcomes of customers in the training set, with the weight and residual functions constructed using the state-of-the-art sample splitting techniques such as cross-fitting (Newey and Robins 2018) and honest estimation (Athey et al. 2019b). Note that, as we show in Online Appendix A.1, this class of estimators include a wide range of popular models, such as Causal Forest (Wager and Athey 2018), S-learners and T-learners with different outcome models (Künzel et al. 2019), as well as R-learners with a variety of second-stage estimators (Nie and Wager 2021, Kennedy 2020).

The following theorem formally establishes the relationship between the amount of unexplained variation and the bias and variance of common CATE estimators:

Theorem 1. ***[Bias and Variance of CATE Prediction]*** *Suppose Assumption 1 holds and the CATE estimator $\widehat{\tau}_Y(\mathbf{x}_{\text{test}})$ belongs to the class described in Assumption 2. Then, for a test customer with covariates $\mathbf{x}_{\text{test}}$,*

1. *The bias of the predicted CATE,* $\text{Bias}\left[\widehat{\tau}_Y(\mathbf{x}_{\text{test}})\right] \equiv \mathbb{E}[\widehat{\tau}_Y(\mathbf{x}_{\text{test}}) - \tau(\mathbf{x}_{\text{test}})]$, *is not affected by the unexplained variations in the outcome variable (i.e., $\varepsilon_i(\mathbf{X}_i, W_i)$).*
2. *The variance of the predicted CATE is increasing in the variance of unexplained variations in the training set. Specifically, the variance is*

$$\text{Var}\left[\widehat{\tau}_Y(\mathbf{x}_{\text{test}})\right] = \mathbb{E}\left(\sum_{i\in\mathcal{D}}\left[\widehat{\ell}_i^{W_i}(\mathbf{x}_{\text{test}})\right]^2 \text{Var}[Y_i(W_i)|\mathbf{X}_i]\right). \tag{1}$$

Theorem 1.1 implies that consistent estimators, such as Causal Forest (Wager and Athey 2018) or R-learners with linear regression as the second-stage estimator (Semenova and Chernozhukov 2021), can be used to achieve zero bias for a large experimental sample, even when there exist substantial unexplained variations in the outcome variable. However, Theorem 1.2 highlights the significant impact that unexplained variations on the variance of a CATE estimator. This finding is particularly relevant for practitioners who aim to optimize long-term outcomes, as the accumulation of noise over time can dramatically increase the variance of the CATE estimator, leading to suboptimal targeting policies.

To illustrate the impact of unexplained variation on targeting performance, we compare the targeting decisions based on the CATE model to the optimal targeting policy, which targets customers with positive true CATEs. The result in Theorem 1 leads to the following proposition:

PROPOSITION 1. ***[Mistargeting]*** *Consider a scenario in which a company obtains a consistent CATE model $\widehat{\tau}_Y(\cdot)$ (e.g., Causal Forest) and implements a policy whereby only customers with positive (expected) treatment effect will receive the intervention.*[4] *Then, the probability of the learned CATE model deviating from the optimal policy, i.e.,* $\mathbb{P}\left[\tau_Y(\mathbf{x}_{\text{test}}) \cdot \widehat{\tau}_Y(\mathbf{x}_{\text{test}}) < 0\right]$, *is increasing in the variance of unexplained variations for customers in the training set (i.e.,* $\operatorname{Var}[Y_i(W_i)|\mathbf{X}_i]$*).*

Proposition 1 shows that the discrepancy between the learned policy and the optimal targeting policy increases as the amount of unexplained variations in the training set grows. Consequently, regardless of the business context, organizations are bound to encounter difficulties in CATE estimation when the outcome variable is noisy.

It is important to note that alternative methods exist for targeting customers beyond directly targeting those with positive CATEs. For instance, rather than building a CATE model and targeting individuals with positive predicted CATEs, Kitagawa and Tetenov (2018) and Athey and Wager (2021) first derive a proxy variable for CATE using inverse probability weighting or a doubly robust estimator, then develop a policy (usually employing machine learning models) to determine which customers the firm should target. These methods are conceptually similar to the targeting approach examined in our study and encounter the same issue, as large unexplained variations in the outcome variable also increase the variance of the CATE proxy variable. (We provide additional empirical evidence in Section 6.3.2.)

[4] 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.

### 3.3. Unexplained Variations in Long-term Marketing Outcomes

Having established the relationship between unexplained variations and targeting performance, we now examine the prevalence and severity of this problem in situations where the objective is to maximize long-term marketing outcomes such as customer lifetime value or repeated purchases. At first glance, one might expect that aggregating purchasing patterns over time would result in a less noisy measurement compared to a single observation for the same customer, as it would cancel out unexpected behaviors across multiple periods. However, as we demonstrate in this section, the presence of factors such as unobserved heterogeneity or customer attrition introduces strong temporal correlation in purchase behaviors. In such scenarios, aggregating purchase outcomes over time amplifies the unexplained variations caused by these factors, rather than cancelling out idiosyncratic shocks. To better understand and characterize this phenomenon, we now present a formal theory on long-term marketing outcomes .

**3.3.1. Formalizing Transaction Behavior.** Let us assume that a company aims to design a targeted intervention with the objective of increasing customer purchases, denoted as $Y_{i,T}$, within a specific time frame $T$ following the intervention. Based on the literature on probabilistic models for repeated purchasing behaviors (Schmittlein et al. 1987, Fader et al. 2005, Abe 2009, Bachmann et al. 2021), we make the following assumptions on the customer behaviors.

Assumption 3. ***[Data Generating Process]*** *Let $W_i$ denote the treatment assignment and $\mathbf{X}_i$ the customer covariates. We denote $t=0$ as the time when the intervention takes place and $\mathbf{X}_i$ is recorded. The customer behavior is characterized by the following assumptions:*

1. *A customer $i$ is in the "alive" state with the firm during a period of time (namely the customer lifetime) and transitions to a "permanently dead" state thereafter.*
2. *At the beginning of each period $t$, each customer decides to stay alive with the company or not. The relationship state $\delta_{i,t} \equiv \mathbb{1}(\text{Lifetime}_i \geq t)$ can be written as*

$$\delta_{i,t} = \theta_t(\mathbf{X}_i, W_i) + \varepsilon^{\theta}_{i,t},$$

*where* $\theta_t(\mathbf{X}_i, W_i) \equiv \mathbb{E}[\delta_{i,t}|\mathbf{X}_i, W_i] = \mathbb{P}[\text{Lifetime}_i \geq t|W_i, \mathbf{X}_i]$ *is the conditional mean alive probability*[5] *for period* $t$*, and* $\varepsilon^{\theta}_{i,t} \equiv \delta_{i,t} - \mathbb{E}[\delta_{i,t}|\mathbf{X}_i, W_i]$ *represents the mean-zero variation that cannot be explained by* $\mathbf{X}_i$ *and* $W_i$*.*

3. *The realized number of transactions made by customer* $i$ *during an alive period* $t$ *can be characterized as*

$$\mathcal{S}_{i,t} = \lambda_t(\mathbf{X}_i, W_i) + \varepsilon^{\lambda}_{i,t},$$

*where* $\lambda_t(\mathbf{X}_i, W_i) \equiv \mathbb{E}[\mathcal{S}_{i,t}|W_i, \mathbf{X}_i]$ *is the conditional mean transaction counts in period* $t$*, and* $\varepsilon^{\lambda}_{i,t}$ *denotes the mean-zero unexplained variation. After becoming inalive, the customer no longer makes any purchase with the firm, and the realized transaction counts is zero for all the following periods.*

4. *Unexplained variations in churn and purchase have non-negative serial correlations, i.e.,* $\text{Cor}\left(\varepsilon^{\theta}_{i,t_1}, \varepsilon^{\theta}_{i,t_2}\right) \geq 0$ *and* $\text{Cor}\left(\varepsilon^{\lambda}_{i,t_1}, \varepsilon^{\lambda}_{i,t_2}\right) \geq 0$*, for all* $t_1, t_2 > 0$*, and the unexplained variations in the two processes are independent, i.e.,* $\{\varepsilon^{\theta}_{i,t}\}_{t=1}^{T} \perp\!\!\!\perp \{\varepsilon^{\lambda}_{i,t}\}_{t=1}^{T}$*.*

Note that our framework is highly general and can be applied to both contractual and non-contractual settings (Fader and Hardie 2010, Fader et al. 2010). The distinction between the two lies in whether $\delta_{i,t}$ is directly observable or not. Furthermore, our assumptions correspond to the setting of a generalized linear model (if $\theta_t$ and $\lambda_t$ are linear functions) and a generalized nonlinear model (if $\theta_t$ and $\lambda_t$ are non-linear function) for the relationship state $\delta_{i,t}$ and realized transactions $\mathcal{S}_{i,t}$. While many probabilistic models for purchasing behaviors do not explicitly include additive errors (e.g., the Beta-Geometric model for customer lifetime and the Gamma-Poisson model for purchase counts), our framework can be applied to such cases by assuming specific distributions for $\varepsilon^{\theta}_{i,t}$ and $\varepsilon^{\lambda}_{i,t}$.

**3.3.2. Unexplained Variations in Long-term Marketing Outcomes.** We now examine the quantity the company aims to maximize, namely the number of actual transactions in the first $T$ periods following the intervention, $Y_i^T(W_i) = \sum_{t=1}^{T} S_{i,t}(W_i)$ . Building on our prior assumptions, the realized transaction counts in a particular period $t$, denoted as $S_{i,t}(W_i)$, can be expressed as:

$$S_{i,t}(W_i) = \delta_{i,t} \cdot \mathcal{S}_{i,t} + (1 - \delta_{i,t}) \cdot 0 = \delta_{i,t} \cdot \mathcal{S}_{i,t} = \left[\theta_t(\mathbf{X}_i, W_i) + \varepsilon^{\theta}_{i,t}\right] \cdot \left[\lambda_t(\mathbf{X}_i, W_i) + \varepsilon^{\lambda}_{i,t}\right].$$

[5] Note that we can write $\theta_t(\mathbf{X}_i, W_i) = \prod_{k=1}^{t}[1 - p_t(\mathbf{X}_i, W_i)]$, where $p_t(\mathbf{X}_i, W_i)$ is the expected churn probability of customer $i$ in period $t$ given the treatment $W_i$ and covariates $\mathbf{X}_i$.

Therefore, the total transactions from $t=1$ to $T$ can be written as

$$Y_i^T(W_i) = \sum_{t=1}^{T} \left[\theta_t(\mathbf{X}_i, W_i) + \varepsilon_{i,t}^{\theta}\right] \cdot \left[\lambda_t(\mathbf{X}_i, W_i) + \varepsilon_{i,t}^{\lambda}\right]. \tag{2}$$

The formulation in Equation (2) highlights when and why the noise of a long-term marketing outcome can be substantially large. For example, unexpected customer churn (whether silent or not, as discussed in Ascarza et al. (2018)) can result in large negative values of $\varepsilon_{i,t}^{\theta} = -\theta_t(W_i, \mathbf{X}_i)$ for all periods after churn, leading to significant negative unexplained variations in $Y_{i,T}(W_i)$. Similarly, "binge-buying" behavior (Zhang et al. 2013, Lu et al. 2019) can cause large positive values of $\varepsilon_{i,t}^{\lambda}$ for multiple periods, resulting in substantial positive unexplained variations in $Y_{i,T}$.

Consequently, when individual behaviors accumulate over time, the unexplained variations in the outcome variable will be larger as the firm extends the observation length for the business outcome. This property is formally shown in the following theorem:

Theorem 2. ***[Unexplained Variation in Long-term Outcomes]*** *Assume the data generating process described in Assumption 3. Then, the unexplained variations in $Y_{i,T}(W_i)$ is increasing in $T$, i.e.,*

$$\mathrm{Var}\left[Y_{i,T_2}(W_i)|\mathbf{X}_i\right] - \mathrm{Var}\left[Y_{i,T_1}(W_i)|\mathbf{X}_i\right] > 0, \quad \forall\ T_2 > T_1.$$

Combining Theorem 2 with Theorem 1, it can be inferred that the variance of standard CATE estimators will increase as the observation window becomes longer, resulting in larger prediction errors and less effective targeting policies. This observation leads to the following corollary:

Corollary 1. ***[Mistargeting for Long-Term Outcomes]*** *Suppose that Assumption 1, 2, and 3 hold. Also, assume that the intervention directly affects the purchases up to period $T_1$. Then, for a consistent CATE estimator, the mistargeting probability is higher for the policy based on $\widehat{\tau}_{Y_{T_2}}$ than $\widehat{\tau}_{Y_{T_1}}$ if $T_2 > T_1$, i.e.,*

$$\mathbb{P}\left[\tau_{Y_{T_2}}(\mathbf{x}_{\text{test}}) \cdot \widehat{\tau}_{Y_{T_2}}(\mathbf{x}_{\text{test}}) < 0\right] > \mathbb{P}\left[\tau_{Y_{T_1}}(\mathbf{x}_{\text{test}}) \cdot \widehat{\tau}_{Y_{T_1}}(\mathbf{x}_{\text{test}}) < 0\right].$$

Corollary 1 highlights the challenge associated with targeting based on long-term CATE models. Specifically, the probability of mistargeting increases when using a long-term

CATE model compared to a short-term CATE model, particularly when there is no additional signal of treatment effects in later periods. This is due to the higher uncertainty in CATE estimation due to the noise associated with long-term outcomes, which can lead to less accurate CATE estimation.

**3.3.3. Discussion of the Data Generating Assumptions in Marketing Contexts.** Before we develop a solution to this problem, it is important to examine to what extent the assumptions of our theoretical framework hold in marketing applications. A key condition for Theorem 2 to hold is the presence of non-negative serial correlation in unexplained variations (i.e., Assumption 3.4). Hence, a natural question to ask is, how realistic are these data patterns in marketing contexts?

In line with prior research on choice persistence over time (e.g., Guadagni and Little 1983, Fader and Lattin 1993, Roy et al. 1996), we argue that this phenomenon is common in real-world marketing contexts. In particular, the prevalence of *unobserved heterogeneity* and *customer attrition* drives positive serial correlation in the unexplained variations, leading to an accumulation of noise over time.

First, positive serial correlation in unexplained variations in the purchase processes occurs due to variations in customers' intrinsic purchase tendencies towards the company (i.e., *unobserved heterogeneity*), commonly accounted for via individual fixed effects (e.g., Jones and Landwehr 1988, Gonul and Srinivasan 1993). The following proposition shows that individual fixed effects result in a positive serial correlation in purchase tendencies:

Proposition 2. ***[Autocorrelation of Unexplained Purchase Variations]*** *Let us assume that the unexplained variation in the purchase process can be expressed as*

$$\varepsilon_{i,t}^{\lambda} = \overline{\varepsilon}_{i}^{\lambda} + \eta_{i,t}^{\lambda},$$

*where* $\overline{\varepsilon}_{i}^{\lambda}$ *represents the time-invariant individual purchase tendency, and* $\eta_{i,t}^{\lambda}$ *denote the i.i.d per-period shock that is independent of* $\overline{\varepsilon}_{i}^{\lambda}$. *Then, the serial correlation of* $\varepsilon_{i,t}^{\lambda}$ *is positive, i.e.,* $\mathrm{Cov}(\varepsilon_{i,t_1}^{\lambda}, \varepsilon_{i,t_2}^{\lambda}) > 0$.

This proposal emphasizes that when observable characteristics $(X_i, W_i)$ are insufficient to capture customer heterogeneity, unobserved variations in the purchase process will exhibit positive autocorrelations. Other customer behaviors, such as state dependence and habit persistence (Roy et al. 1996, Seetharaman 2004), can also contribute to positive correlations. Conversely, stockpiling behavior (Tulin et al. 2002) may lead to negative correlation. The severity of noise accumulation will depend on the extent to which these behavioral drivers are present in the data and not captured by the observables. In our empirical study, unobserved heterogeneity is likely present since pre-treatment behaviors collected by the company may not fully account for outcome variations driven by customer preferences. On the other hand, stockpiling behavior is unlikely given that the company mainly sells fresh meals with a short shelf life.

Second, the serial correlation of unexplained variations in the churn process is, by definition, positive. When a customer churns in period $t$ (resulting in a negative value of $\varepsilon^{\theta}_{i,t}$), all subsequent unexplained variations for that customer should also be negative. Conversely, if a customer remains alive at time $t$ (resulting in a positive value of $\varepsilon^{\theta}_{i,t}$), then all of the prior unexplained variations should also be positive for this customer. The following proposition provides formal proof for this property:

PROPOSITION 3. ***[Autocorrelation of Unexplained Churn Variations]*** *Under the data generating process as described in Assumption 3, we have*

$$\mathrm{Cov}\left(\varepsilon^{\theta}_{i,t_1}, \varepsilon^{\theta}_{i,t_2}\right) = \theta_{t_2}(\mathbf{X}_i, W_i)[1 - \theta_{t_1}(\mathbf{X}_i, W_i)] \geq 0.$$

Finally, our theoretical framework assumes that the unexplained variations in churn and purchase processes are independent, i.e., they reflect distinct types of individual preferences — one for retention and another for purchase frequency. This assumption simplifies our proofs and is supported by previous studies (Abe 2009) that found no significant correlation between the noises in churn and purchase processes. However, even if there were a positive correlation between these noises, our findings would remain unchanged. We acknowledge that our framework's robustness to those assumptions should be further explored in future research.

In conclusion, the state-of-the-art "test-to-target" approach is inadequate when there is noise accumulation in the outcome of interest, a problem that exacerbates as the length of the observation period increases. Consequently, it is imperative to develop a new approach that enables firms to design targeting policies that can effectively maximize long-term outcomes.

## 4. The Solution: Surrogate Index with Separate Imputation

Our proposed solution involves using a "noise-reduced proxy" as a substitute for the outcome variable when estimating CATEs. This approach aims to improve targeting policy by reducing unexplained variations that are not relevant to the intervention, without altering the firm's objective (i.e., the firm is still optimizing the long-term outcome). Although using a proxy instead of the actual outcome may seem counterproductive as it contains less information, we demonstrate that a "well-specified proxy" can effectively eliminate irrelevant variations, resulting in superior targeting policy.

To be effective, this proxy should meet three key requirements. Firstly, it should eliminate unexplained variations in the outcome variable to reduce the variance of the CATE model. Secondly, it should capture long-run treatment effect heterogeneity to ensure that the resulting targeting strategy is optimized for the desired long-term outcome. Finally, it should account for behavioral patterns present in the data, such as customer attrition. We propose a novel imputation approach that satisfies these requirements and improves long-term targeting precision.

### 4.1. Identification and Variance Reduction Using Surrogate Index

Our proposed solution builds upon the *surrogate index* approach (Athey and Wager 2019, Yang et al. 2022) for long-term treatment effect estimation. The surrogate index refers to the predicted value of the long-term outcome, based on the observed short-term outcomes and pre-treatment covariates. To construct a surrogate index, we model the relationship between short-term and long-term outcomes using historical data obtained from previously acquired customers prior to the start of the experiment. One desirable feature of the proposed solution is that these data are readily available for most firms.

More formally, let us assume that the company has access to two datasets: 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. The surrogate index is the expected long-term outcome ($Y_{i,T}$) of customers in $\mathcal{H}$, conditioned on their short-term behaviors ($\mathbf{S}_{i,T_0} = \{S_{i,1}, \cdots, S_{i,T_0}\}$ for some $T_0 < T$) and pre-treatment covariates ($\mathbf{X}_i$). Mathematically, it can be represented as:

$$\widetilde{Y}_T(\mathbf{S}_{T_0}, \mathbf{X}_i) \equiv \mathbb{E}_{\mathcal{H}}\left[Y_{i,T} | \mathbf{S}_{T_0}, \mathbf{X}_i\right].$$

There are two additional requirements for the identification of CATE on $Y_{i,T}$ through $\widetilde{Y}_T(\mathbf{S}_{T_0}, \mathbf{X}_i)$ (Athey et al. 2019a):

Assumption 4. ***[Surrogacy]*** *The short-term outcomes can fully mediate the treatment effect of $W_i$ on $Y_{i,T}$; that is, $W_i \perp\!\!\!\perp Y_{i,T} \mid \mathbf{S}_{i,T_0}, \mathbf{X}_i$ , $\forall i \in \mathcal{E}$.*

Assumption 5. ***[Comparability]*** *The experimental and historical data are comparable in distribution; that is, $Y_{i,T} \mid \mathbf{S}_{i,T_0}, \mathbf{X}_i$ , $i \in \mathcal{E} \overset{d}{\sim} Y_{i,T} \mid \mathbf{S}_{i,T_0}, \mathbf{X}_i$ , $i \in \mathcal{H}$.*

When these conditions are met, the surrogate index representation enables identification of the CATE. As a result, a targeting policy based on the surrogate index is effective in optimizing the long-term outcome. Moreover, since the surrogate index is based on short-term outcomes, it contains less unexplained variation compared to the actual long-term outcome and therefore reduces the variance of CATE predictions. More formally, we state the theorem as follows:

Theorem 3. ***[Identification and Variance Reduction Using Surrogate Index]*** *Suppose that Assumption 1, Assumption 4, and 5 hold.*

1. *The CATE of the intervention on the long-term outcome is equal to the CATE on the surrogate index, that is,*

$$\tau_{Y_T}(\mathbf{X}_i) = \widetilde{Y}_T(\mathbf{S}_{i,T_0}(1), \mathbf{X}_i) - \widetilde{Y}_T(\mathbf{S}_{i,T_0}(0), \mathbf{X}_i),$$

   *where $\mathbf{S}_{i,T_0}(W_i)$ denotes the potential outcome of the short-term outcomes given the treatment status $W_i$.*

2. *The variance of the surrogate index is smaller than the variance of the actual long-term outcome, that is,*

$$\mathrm{Var}\left[\widetilde{Y}_T(\mathbf{S}_{i,T_0}(W_i), \mathbf{X}_i)\right] < \mathrm{Var}[Y_{i,T}(W_i)|\mathbf{X}_i].$$

Altogether, Theorem 3.1 and Theorem 3.2 imply that estimating the CATE on the surrogate index leads to smaller variances for the same class of CATE model compared to estimating it on the long-term outcome while still maintaining valid estimation of the long-term treatment effect.

**4.1.1. Discussion of the Surrogacy Assumption in Marketing Contexts.** The surrogate index approach is a powerful tool, but its effectiveness relies on two key prerequisites. Firstly, it requires the existence of short-term signals that can accurately predict long-term treatment effects. This condition is easily met for long-term outcomes associated with relatively frequent behaviors (e.g., repeated purchases or continuous product usage) that accumulate or recur over time, as seen in the retail, subscription-based media, food service, or transportation industries. However, it can be challenging to find valid short-term surrogates in industries with long purchase cycles (e.g., the automotive sector). In such cases, our solution may be less relevant as short-term signals may be more difficult to identify.

Secondly, the surrogacy assumption requires that short-term outcomes can *fully mediate* the long-term treatment effect. If this condition is not met, discrepancies may arise between the actual CATEs and those identified using the surrogate index (Yang et al. 2022), resulting in significant mistargeting errors. One way to ensure that the surrogacy assumption is satisfied is to include a large number of surrogates (Athey et al. 2019a). When dealing with long-term outcomes such as CLV or repeated transactions, including more periods in the surrogate index may be the most natural way to satisfy the assumption. However, adding more periods can also increase unexplained variations, which may decrease the effectiveness of targeting policies. This issue is formally characterized in the following proposition:

PROPOSITION 4. ***[Noise Accumulation of Surrogate Indices]*** *If we construct two surrogate indices using different periods of short-term outcomes $T_0$ and $T_0'$, then we have*

$$\mathrm{Var}\left[\widetilde{Y}_T(\mathbf{S}_{i,T_0}(W_i), \mathbf{X}_i)\right] > \mathrm{Var}\left[\widetilde{Y}_T(\mathbf{S}_{i,T_0'}(W_i), \mathbf{X}_i)\right].$$

Proposition 4, along with the surrogacy assumption, highlights the trade-off that firms must balance between *information gain* and *noise accumulation* to achieve optimal targeting performance. Essentially, including more (short-term) periods as surrogates helps satisfy the surrogacy assumption but also introduces unexplained variations that increase mistargeting errors. Determining the optimal number of periods for constructing surrogate models requires empirical investigation. In this paper, we address this question through simulations and real-world data. Our findings suggest that in cases where there are substantial unexplained variations in the long-term outcome, using fewer periods in the surrogate model, despite violating the surrogacy assumption, can still result in better targeting performance.

### 4.2. Separate Imputation Approach

Creating a surrogate index in situations with customer attrition can be a challenging task. To illustrate this, consider a scenario where a firm observes short-term outcomes from $t = 1$ to $T_0$ and pre-treatment covariates, which will be used to construct the surrogate index. Let us consider two quantities: the posterior mean lifetime ($\mathcal{T}_T(\mathbf{X}_i|\mathbf{S}_{i,T_0})$) capped at period $T$, and the purchase rate per alive period ($\Lambda(\mathbf{X}_i|\mathbf{S}_{i,T_0})$) over the $\mathcal{T}_T(\mathbf{X}_i|\mathbf{S}_{i,T_0})$ periods.[6] Similar to Assumption 3, we express the observed lifetime and purchase rate as the posterior mean plus the unexplained variation. Then, the total purchase counts over $T$ periods for customers in the historical data:$\mathcal{H}$ can be written as follows:

$$
\begin{aligned}
Y_{i,T} &= \left[\mathcal{T}_T(\mathbf{X}_i|\mathbf{S}_{i,T_0}) + \varepsilon^{\mathcal{T}}_{i,T}\right] \cdot \left[\Lambda_T(\mathbf{X}_i|\mathbf{S}_{i,T_0}) + \varepsilon^{\Lambda}_{i,T}\right] \\
&= \underbrace{\mathcal{T}_T(\mathbf{X}_i|\mathbf{S}_{i,T_0}) \cdot \Lambda_T(\mathbf{X}_i|\mathbf{S}_{i,T_0})}_{=\mathbb{E}_{\mathcal{H}}\left[Y_{i,T}|\mathbf{S}_{T_0},\mathbf{X}_i\right] \equiv \widetilde{Y}_T(\mathbf{S}_{T_0},\mathbf{X}_i)} \quad + \quad \underbrace{\mathcal{T}_T(\mathbf{X}_i|\mathbf{S}_{i,T_0}) \cdot \varepsilon^{\Lambda}_{i,T} + \Lambda_T(\mathbf{X}_i|\mathbf{S}_{i,T_0}) \cdot \varepsilon^{\mathcal{T}}_{i,T} + \varepsilon^{\mathcal{T}}_{i,T} \cdot \varepsilon^{\Lambda}_{i,T}}_{=\varepsilon_{i,T}(\mathbf{X}_i,\mathbf{S}_{i,T_0})},
\end{aligned}
\tag{3}
$$

where $\varepsilon^{\mathcal{T}}_{i,T}$ and $\varepsilon^{\Lambda}_{i,T}$ denotes the unexplained variations in lifetime and purchase rate, respectively. In this case, regressing $Y_{i,T}$ on $\mathbf{S}_{i,T_0}$ and $\mathbf{X}_i$, as done by Athey et al. (2019a) and Yang et al. (2022), could result in inaccurate estimation of $\widetilde{Y}_T(\mathbf{S}_{T_0}, \mathbf{X}_i)$ due to the dependence of the unexplained variation term ($\varepsilon_{i,T}(\mathbf{X}_i, \mathbf{S}_{i,T_0})$) on $\mathbf{S}_{i,T_0}$ and $\mathbf{X}_i$.

We propose a *separate imputation* strategy to address this issue. Our method involves developing two distinct surrogate indices: one for customer lifetime ($\mathcal{T}_T(\mathbf{X}_i|\mathbf{S}_{i,T_0})$) and another for transactions ($\Lambda_T(\mathbf{X}_i|\mathbf{S}_{i,T_0})$), and combining them to produce the proxy that will serve as the outcome variable for CATE estimation or policy learning. The separate imputed outcome possesses the key characteristics of a proper surrogate index. First, it is less noisy than the actual long-term outcome as it excludes unexplained variations that occur during $T_0 + 1$ to $T$. Second, it provides valid inference for the long-term treatment effect when Assumption 4 holds. Finally, the use of separate models for churn and purchase processes addresses the estimation challenge arising from the non-separable nature of churn and purchase in the long-term outcome.[7]

[6] For simplicity, we focus on the average churn and purchase behaviors over the first $T$ periods. However, our argument can be easily extended to cases with time-varying survival probability and purchase rate.

[7] Many other behavioral processes generate a non-separable nature of the error terms. For example, if a customer needs to open an email in order to purchase, or to preview a product page before deciding how much to spend, these behaviors might also present challenges if generating the proxy by directly regressing the final outcome (e.g., purchase or not or total spend) on customer covariates and surrogate variables. Our separate imputation approach is also valid for those cases.

Constructing two separate models is straightforward in a contractual setting where customer lifetime and average purchase rate are directly observed. However, in non-contractual settings like the one we examine in this empirical application, customer attrition is no longer observable. Therefore, we address this issue by leveraging the information from widely-used *recency* and *frequency* measures. Specifically, we use *observed last purchase time* until $T$ (denoted as $\mathcal{T}_{i,T}$) as a proxy for customer lifetime, and the *average purchase counts per period* until $\mathcal{T}_{i,T}$ (denoted as $\Lambda_{i,T}$) as a proxy for average purchase rate per alive period. By using the time of the last purchase as an approximation for the time at which the customer churns, we can construct two surrogate models from the historical data $\mathcal{H}$:

$$\widehat{\mathcal{T}}_T\left(\mathbf{X}_i|\mathbf{S}_{i,T_0}\right)=\widehat{\mathbb{E}}_{i\in\mathcal{H}}\left[\mathcal{T}_{i,T}|\mathbf{S}_{i,T_0},\mathbf{X}_i\right], \qquad \widehat{\Lambda}_T\left(\mathbf{X}_i|\mathbf{S}_{i,T_0}\right)=\widehat{\mathbb{E}}_{i\in\mathcal{H}}\left[\Lambda_{i,T}|\mathbf{S}_{i,T_0},\mathbf{X}_i\right],$$

where $\widehat{\mathbb{E}}$ can be any regression or machine learning models. In practice, we can compare the targeting performance of different models empirically and select the one that provides the best result.

After constructing two surrogate models using historical data, we predict the lifetime and purchase rate for customers in the experimental data using their observed short-term outcomes $\mathbf{S}_{i,T_0}(W_i)$ and $\mathbf{X}_i$. We then combine these predicted values by multiplication to create the surrogate index for CATE estimation and policy learning, i.e.,

$$\widehat{\widehat{Y}}_T^{\mathrm{Sep}}\left(\mathbf{X}_i|\mathbf{S}_{i,T_0}(W_i)\right)=\widehat{\mathcal{T}}_T\left(\mathbf{X}_i|\mathbf{S}_{i,T_0}(W_i)\right)\times\widehat{\Lambda}_T\left(\mathbf{X}_i|\mathbf{S}_{i,T_0}(W_i)\right),\ i\in\mathcal{E}.$$

In summary, when there is customer attrition or other behavioral patterns that create a multiplicative error structure, regressing the long-term outcome on the observed short-term outcome and pre-treatment covariates directly may lead to inaccurate estimation of the surrogate index. To address this issue, we propose a separate imputation strategy that involves constructing two distinct surrogate models for customer lifetime and average purchase rate. As demonstrated in Section 5 and Section 6, the proposed surrogate index yields superior targeting performance compared to other alternatives used by firms, including existing imputation methods.

# 5. Empirical Performance: Simulation Evidence

## 5.1. Simulation Setting

We first conduct simulations to validate our argument and solution. Our simulations involve a company executing a marketing intervention aimed at maximizing the total purchase count ($Y_{i,T}$, where $T = 10$) over a ten-week period post-intervention. Using the customer behavior assumptions outlined in Assumption 1 abd Assumption 3, we generate experimental data ($\mathcal{E}$) with unexplained variations in both attrition and purchase behaviors. We provide details about the simulation setup in Online Appendix B.1. In particular, we assume that the intervention has a direct and heterogeneous impact on churn probability and purchase frequency during the first three weeks, and no impact after the fourth week. The unexplained variations in the churn process have a positive autocorrelation by design, as customers cannot become alive again once they have churned. For simplicity, we assume that the unexplained variations in the purchase process are independent across time.[8] More details regarding the simulation setting can be found in Online Appendix B.1.

We also generate historical data ($\mathcal{H}$) of 5,000 customers using the data generating process for the control group, which reflects the company not implementing the intervention in the past. The data will be used to construct the surrogate indices. We select $T_0 = 3$ for the main analysis to ensure that the surrogacy assumption (Assumption 4) is satisfied. We then adjust the number of periods used in the analysis to examine the trade-off between information gain and noise accumulation. The model specifications of surrogate indices are described in detail in Online Appendix B.3.

### 5.2. Comparison Methods

**5.2.1. Alternative Imputation Methods.** We compare our solution to various imputation methods for constructing surrogate indices. The first technique we examine is the *single imputation* approach proposed by Athey et al. (2019a) and Yang et al. (2022). This method directly regresses the outcome variable ($Y_{i,T}$) on the pre-treatment covariates ($\mathbf{X_i}$) and the short-term signals ($S_{i,1}, \cdots, S_{i,T_0}$) to reduce unexplained variations in the long-term outcome. This imputation technique does not consider the multiplicative noise structure that results from customer churn, making it less precise for CATE estimation than the separate imputation method.

The second imputation method we consider is based on the BG/NBD model (Fader and Hardie 2007). While this model has not been previously proposed used as imputation method in surrogacy models, we believe that it can be a reasonable candidate in our context for two reasons. First, the model assumes a similar data generating process to

[8] Essentially, we assume that there is no individual fixed effect for the purchase rate, which means that there is no time-invariant individual preference towards the company that affect the purchase process.

the one described in Section 3.3 and utilizes flexible probability distributions to capture unobserved heterogeneity in churn and purchase behavior. Second, the model has been shown to generate accurate transaction forecasts at the aggregate level across a variety of contexts.

We utilize historical data to estimate the BG/NBD model (with covariates) and utilize it to predict expected future transactions for each customer in the experimental dataset based on their pre-treatment covariates and short-term purchases. However, the accuracy of the model, along with similar variants like the Pareto/NBD model, heavily depends on the likelihood specification, the distributional assumptions, and the assumed functional forms of the relationship between customer covariates and key parameters of those distributions. Thus, when the relationship between the observed covariates and the treatment effect heterogeneity is complex, we anticipate that it may be less effective compared to the separate imputation approach.

**5.2.2. Alternative Variance Reduction Methods.** Aside from utilizing alternative imputation methods to obtain the surrogate index, we also investigate other techniques to reduce the variance in CATE estimation. One such technique is to simplify the CATE model by regularizing the CATE function. In our simulation, we utilize R-lasso as the regularized CATE model (Nie and Wager 2021). 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 size and/or high noise level.

Another (obvious) alternative to reduce variance is to increase the sample size. While this could be an effective solution, in principle, we argue that it is not practical and not the most efficient way to improve targeting policies. First, the number of customers who qualify for the intervention is often limited, and therefore, the sample size available to most firms is also limited.[9] Second, even in cases where the company can increase the experimental sample size, the rate at which the variance of CATE decreases with respect to sample size may be considerably slower compared to when employing imputation methods.

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

**5.2.3. Baseline Methods.** Finally, we examine two baseline methods commonly used in practice. First, we consider targeting based on the actual long-term outcome $Y_{i,T}$ (*default* approach). As demonstrated earlier, this approach is likely to be ineffective due to the substantial unexplained variations in $Y_{i,T}$. Second, we consider a *myopic* approach, which involves targeting customers based on short-term performance $Y_{i,T_0} = \sum_{t=1}^{T_0} S_{i,t}$ (i.e., based on their behavior right after the intervention). While this approach can avoid noise accumulation (as there is less unexplained variation in behavior up to $T_0$), 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. Then, for each of the approaches considered, we construct a CATE model $\widehat{\tau}_{\ddot{Y}}(\mathbf{X}_i)$ using $\ddot{Y}$ as outcome variable (e.g., $\ddot{Y} = Y_{10}$ for the default approach, $\ddot{Y} = \widehat{\widehat{Y}}{}^{\text{Sep}}_T$ for the proposed approach, etc.), and calculate the area under the targeting operating characteristic curve (AUTOC) (Yadlowsky et al. 2021) using the actual long-term outcome ($Y_{i,T}$). 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 high AUTOC indicating an effective sorting mechanism.

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

$$\mathrm{TOC}(\phi; \widehat{\tau}_{\ddot{Y}}) = \mathbb{E}\left[Y_{i,T}(1) - Y_{i,T}(0) | F_{\widehat{\tau}_{\ddot{Y}}}(\widehat{\tau}_{\ddot{Y}}(\mathbf{X}_i)) \geq 1 - \phi\right] - \mathbb{E}\left[Y_{i,T}(1) - Y_{i,T}(0)\right], \tag{4}$$

where $F_{\widehat{\tau}_{\ddot{Y}}}$ 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

$$\mathrm{AUTOC}(\widehat{\tau}) = \int_0^1 \mathrm{TOC}(\phi; \widehat{\tau}) d\phi. \tag{5}$$

Note that a model $\widehat{\tau}_{\ddot{Y}}$ is better than another model $\widehat{\tau}_{\ddot{Y}'}$ in identifying customers in the top $\phi \times 100\%$ CATE group if $\mathrm{TOC}(\phi; \widehat{\tau}_{\ddot{Y}}) > \mathrm{TOC}(\phi; \widehat{\tau}_{\ddot{Y}'})$. Thus, a higher AUTOC value suggests that the CATE model is more effective in identifying customers who exhibit the strongest response to the intervention.[10]

[10] In the empirical application we also evaluate targeting on the basis of expected profitability.

### 5.4. Results

Table 1 shows the AUTOC values of different CATE models. Each row corresponds to CATE models for a specific outcome variable, while the columns indicate different methods for CATE estimation (e.g., S-learner, Causal Forest).[11] By default, we train the models using 1,000 customers (500 per condition) and evaluate the AUTOC using 10,000 validation customers (5,000 per condition). To compute the results of increasing sample size (last row if Table 1), we increase the training sample from 1,000 to 50,000 customers.

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

| Sample Size | Outcome $\ddot{Y}$ | CATE Model | | | |
|---|---|---|---|---|---|
| | | S-GRF | T-GRF | Causal Forest | R-lasso |
| $N = 1,000$ | Separate Imputation | **0.76 (0.15)** | **0.73 (0.18)** | **0.87 (0.10)** | 0.28 (0.41) |
| | Myopic | 0.71 (0.14) | 0.69 (0.13) | 0.82 (0.07) | 0.27 (0.42) |
| | Single Imputation | 0.68 (0.15) | 0.64 (0.15) | 0.82 (0.11) | 0.24 (0.4) |
| | BG/NBD Imputation | 0.67 (0.19) | 0.63 (0.14) | 0.79 (0.14) | 0.25 (0.41) |
| | Default (Small Sample) | 0.37 (0.25) | 0.36 (0.25) | 0.53 (0.31) | 0.26 (0.41) |
| $N = 50,000$ | Default (Large Sample) | 0.62 (0.05) | 0.56 (0.04) | 0.71 (0.04) | **0.90 (0.02)** |

Higher AUTOC reflects better prioritization rule and therefore superior targeting performance. We average the results over 200 replications and show in parentheses the standard deviation. The performance of different outcomes when $N = 50,000$ is provided in Online Appendix B.4

There are several findings to highlight. First, all approaches that use short-term proxies as outcome variables in the CATE estimation exhibit higher AUTOCs than those that use the actual long-term outcome $Y_{i,T}$. This result supports that using short-term signals can considerably enhance the targeting performance. Notably, the separate imputation method consistently achieves the highest AUTOC values across all CATE models. This finding highlights the value of separating churn and purchase when creating a surrogate index. On a related note, while the BG/NBD method performs better than using the actual outcome, it performs worse than that of the other methods. As previously discussed, this approach may be less effective in situations where the relationship between the observed characteristics and the key parameters is complex, which is the case in our simulation.

Second, we highlight the sample size efficiency of our proposed solution. Specifically, we compare the performance of using short-term proxies (top rows in Table 1) to that of using the actual outcome, but trained on a much larger sample — 50 times more customers. Despite being trained on a much smaller sample, CATE models for short-term

[11] We use several CATE models to corroborate that our findings are not driven by a particular method for CATE estimation.

outcomes (except R-lasso) exhibit higher mean AUTOC and lower standard deviations than models for $Y_{i,T}$ that are trained using a much larger sample. This finding highlights the cost-effectiveness of our solution in terms of sample size, as companies do not need to conduct large-scale experiments to improve targeting performance. Instead, they can rely on existing information on historical customer behaviors in their database.

Finally, we find that applying regularization to CATE models for $Y_{i,T}$, specifically with R-lasso and hyperparameters selected through cross-validation, results in poorer targeting performance when the sample size is small. This result appears to be driven by R-lasso's tendency to underestimate treatment effect heterogeneity, especially for smaller sample sizes. [12] On the other hand, we find that the R-lasso outperforms other CATE models for $Y_{i,T}$ when the sample size is sufficient. However, it is worth noting that even with a large sample size, utilizing separate imputation *in conjunction with* R-lasso still produces higher AUTOC values. (See Online Appendix B.4 for the complete results.)

### 5.5. The Trade-off between Information Gain and Noise Accumulation

As discussed in section 4.1.1, our method assumes that the short-term outcomes can fully mediate the treatment effect of the intervention on the desired long-term outcome. While increasing the number of periods in the surrogate model is one way to meet this assumption, it can also increase unexplained variations and therefore reduce the targeting performance. Therefore, firms that implement our solution would need to balance the trade-off between acquiring more information and accumulating more noise. To investigate this trade-off, we expand upon the analyses presented in Table 1 by varying the number of periods used to construct surrogate indices (ranging from $T_0 = 1$ to $T_0 = 6$). Note that the intervention in our simulation only affects customer behavior directly for the first three periods, which means that the surrogacy assumption holds when $T_0 \geq 3$. We create 200 bootstrap replications and present the means and standard deviations of the AUTOC for surrogate indices with separate imputation, constructed using varying numbers of periods of short-term signals, in Figure 2.

The inverted U-shaped relationship in Figure 2 reflects the inherent trade-off between information gain and noise accumulation. As we increase $T_0$ from one to three periods,

[12] In fact, we find that R-lasso produces the same predicted CATE for all customers (indicating no treatment effect heterogeneity) in about 60% of the bootstrap replications, regardless of the outcome variables used for CATE estimation.

**Figure 2** **Trade-off between Information Gain and Noise Accumulation: An Analysis of Causal Forest AUTOCs with Surrogate Index Constructed Using Different Periods.**

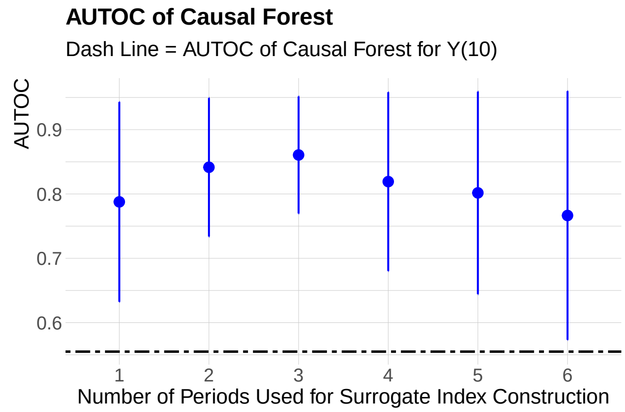


*Note.* Each point reports the average over 200 simulation replications together with the one standard deviation interval. We used 1,000 customers in the training set. The larger the AUTOC, the better targeting performance. We present here the results of using Causal Forest but our findings are robust across different CATE models. See Online Appendix B.5 for the results of different CATE models.

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 (i.e., after satisfying surrogacy). This pattern is anticipated by and aligns with the advice of Athey et al. (2019a) and Yang et al. (2022), which suggest 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 two periods of information (where the surrogacy assumption is *violated*) outperform those utilizing four or more periods of information. This result suggests that the benefits of noise reduction (proposed in this research) 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 Assumption 4 may not be a major concern because firms can still improve targeting performance by using fewer short-term outcomes in the surrogate models.

## 6. Empirical Performance: Real-world Application

In this section, we demonstrate the effectiveness of our proposed solution in a real-world marketing application with the data from a retail-technology company in Taiwan. This company deploys self-serving fridges in various locations within a city, including department stores and office buildings. Customers can conveniently grab food and beverages from the fridge, and the machine will automatically count items using RFID technology and charge customers through their pre-registered payment methods. The company utilizes a third-party messaging app (i.e., LINE) to manage customer profiles and send marketing messages. To use the service, customers must join the company's messaging app channel and register their payment methods using the app.

### 6.1. The Marketing Intervention

As part of their customer activation process, the company sends a coupon offering a 15% discount in the next purchase to every newly acquired customer. The coupon is automatically applied to the (first) next purchase made within 14 days and expires after then. The company is considering increasing the number of coupons offered to *some* newly acquired customers, but only if doing so would increase total purchases in the following months. In order to develop a targeting policy, the company conducted a randomized controlled experiment to identify customers for whom the new intervention would lead to increased purchases during their first ten weeks.[13] The variable of interest, $Y_{i,10}$, is defined as the sum of customer $i$'s transactions during the first 10 weeks, where $S_{i,t}$ denotes the number of transactions made by customer $i$ during week $t$.

In the experiment, the company selected customers who had just made their first purchase and randomly assigned them to one of two groups. The treatment group ($W_i = 1$) received three coupons offering a 15% discount for each their next three purchases, while the control group ($W_i = 0$) received only one coupon (the business-as-usual case). All coupons expired after 14 days. The experiment involved 1,853 customers, with 889 in the treatment group and 964 in the control group. Pre-treatment covariates, including purchase behaviors and referral status, were used to control for differences between the two groups. Online Appendix C.1 shows that the randomization was executed properly.

The study found that the average treatment effect on total purchases over the ten-week period was 0.3153, with a p-value of 0.028. This corresponds to a 15% increase in the average number of purchases made by customers in the treatment group compared to the control group (the average $Y_{i,10}$ for the control group was 1.99). Furthermore, the intervention had a lasting impact on customer churn and purchase frequency that extended beyond the coupon effective period (Figure 3). The leftmost figure depicts the weekly

[13] No below-the-line campaigns were conducted within the first ten weeks after acquisition, so there are no post-treatment confounders until this point.

retention rate across the two experimental groups, with the retention rate representing the percentage of customers who made a purchase on a given week or after, up to a maximum of fifty weeks after their initial acquisition. The findings reveal that the intervention reduced customer churn as the retention rate was consistently higher for the treatment group than the control group. Furthermore, the difference in retention rates between the two groups was 2.3% in the first week and increased to 2.7% in the tenth week. These results suggest that the intervention had both short-term and long-term effects on customer attrition.

**Figure 3** **Retention and Purchase Rates After the Intervention.**

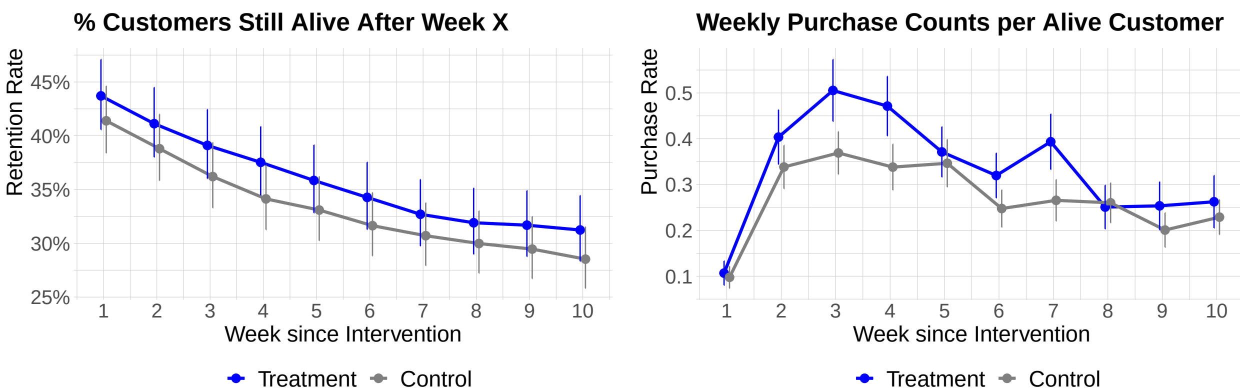


*Note.* Customers are labeled as "alive" in a week if they made at least one purchase in that week or later (cap in week 50). The purchase rate when alive is the total purchases in a specific week divided by the number of alive customers. We report the mean and two standard errors confidence intervals. The retention measure used here reflects a lower bound estimate of the true retention rate.

The rightmost figure in Figure 3 shows the weekly purchase counts per alive customer, which is measured as the average purchase counts per alive customer in a week. The results suggest that, for the first seven weeks, the retained customers in the treatment group made more purchases, on average, than the retained customers in the control group. This implies that the intervention had a long-term impact on purchase frequency, as the retained customers in the treatment group continued to make more purchases than the retained customers in the control group beyond the effective period of the coupons.

## 6.2. Empirical Analysis

**6.2.1. Comparison Methods** The primary objective of our study is to evaluate the effectiveness of the proposed targeting approach in comparison to several alternatives. Unlike the simulation study (Section 5), it is not feasible to increase the sample size in this scenario, and the size of the intervention is limited by the number of customers obtained over time. Besides, we do not present the outcomes of R-lasso since it invariably produces identical CATE predictions for all customers, which is anticipated given the limited training sample size and significant unexplained variations in the data. As a result, we compare our approach to alternatives that utilize the same sample size, including the default and myopic approaches, as well as two alternative imputation methods (i.e., single and BG/NBD). To construct the surrogate indices, we use historical data from customers who were acquired at least ten weeks before the start of the experiment, which totals 4,031 customers, and use the first-week short-term outcome (i.e., $S_{i,1}$) to estimate the surrogate indices.[14] Online Appendix C.8 conducts an empirical analysis to determine the optimal number of periods to be used in this application and provides the complete set of results.

**6.2.2. Validation Approach and Key Metrics** To assess the targeting performance of each approach, we utilize a bootstrap validation scheme similar to that of Ascarza (2018). Specifically, we generate $B = 200$ data splits consisting of training (70%) and validation (30%) sets. For each split, we estimate CATE models that use distinct outcome variables ($\ddot{Y}$) as the dependent variable using the training set. We then predict CATEs ($\widehat{\tau}_{\ddot{Y}}$) on different outcomes 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 group average treatment effects (GATEs) across predicted CATE quintile groups, and the expected profit gained by targeting customers with positive predicted CATEs.

*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 ($\widehat{\tau}_{\ddot{Y}}$), with $\mathcal{Q}_1^{\widehat{\tau}_{\ddot{Y}}}$ having the highest predicted CATEs and $\mathcal{Q}_5^{\widehat{\tau}_{\ddot{Y}}}$ having the lowest predicted CATEs. Next, we calculate the group average treatment effect (GATE) for each quintile group using the actual long-term outcome ($Y_{i,10}$):

$$\widehat{\text{GATE}}_{Y_{10}}(\mathcal{Q}_k^{\widehat{\tau}_{\ddot{Y}}}) = \frac{\sum_{i:\ i \in \mathcal{Q}_k^{\widehat{\tau}_{\ddot{Y}}}, W_i = 1} Y_{i,10}}{\left|\{i:\ i \in \mathcal{Q}_k^{\widehat{\tau}_{\ddot{Y}}}, W_i = 1\}\right|} - \frac{\sum_{i:\ i \in \mathcal{Q}_k^{\widehat{\tau}_{\ddot{Y}}}, W_i = 0} Y_{i,10}}{\left|\{i:\ i \in \mathcal{Q}_k^{\widehat{\tau}_{\ddot{Y}}}, W_i = 0\}\right|}.$$

[14] We use Random Forest and BG/NBD to construct those surrogate models, which are described in detail in Online Appendix C.4

*Expected Profitability of Targeting Policies.* To evaluate the profitability of each CATE model, we consider targeting customers with positive predicted CATEs (policy $\pi^{\widehat{\tau}_{\ddot{Y}}}$) and calculate the expected purchase counts in the next ten weeks using the inverse-probability-weighted (IPW) estimator (Horvitz and Thompson 1952):

$$\widehat{V}(\pi^{\widehat{\tau}_{\ddot{Y}}}) = \frac{1}{|\text{Validation Set}|} \cdot \sum_{i \in \text{Validation Set}} \left( \frac{\mathbb{1}[W_i = \pi^{\widehat{\tau}_{\ddot{Y}}}(\mathbf{X}_i)]}{\widehat{\mathbb{P}}[\pi^{\widehat{\tau}_{\ddot{Y}}}(\mathbf{X}_i) = W_i]} \right) Y_{i,10}, \tag{6}$$

where $\widehat{\mathbb{P}}[\pi^{\widehat{\tau}_{\ddot{Y}}}(\mathbf{X}_i) = W_i]$ is the (estimated) propensity score for customers who are assigned the same treatment by $\pi^{\widehat{\tau}_{\ddot{Y}}}$ as in the actual data.[15] While the treatment assignment in the data is random and independent of the derived targeting policy, we utilize IPW adjustment to account for any possible imbalances between treated and non-treated customers, as the sample used for profit evaluation (i.e., validation customers who were assigned the same treatment in the actual data as $\pi^{\widehat{\tau}_{\ddot{Y}}}$ assigns for policy evaluation) is relatively small. The IPW adjustment is also frequently employed in other marketing literature that utilizes randomized controlled experiments for learning targeting policies (e.g., Hitsch et al. 2023, Yoganarasimhan et al. 2022).

We then calculate the expected profit under policy $\pi^{\widehat{\tau}_{\ddot{Y}}}$ using the following formula:

$$\begin{aligned} \text{Profit}\left(\pi^{\widehat{\tau}_{\ddot{Y}}}\right) = \text{AOV} \cdot p \cdot \widehat{V}(\pi^{\widehat{\tau}_{\ddot{Y}}}) - \text{AOV} \cdot d \cdot \left( \frac{\sum_{i=1}^{N} \pi^{\widehat{\tau}_{\ddot{Y}}}(\mathbf{X}_i)}{N} \right) \cdot U^{W=1} \\ - \text{AOV} \cdot d \cdot \left( 1 - \frac{\sum_{i=1}^{N} \pi^{\widehat{\tau}_{\ddot{Y}}}(\mathbf{X}_i)}{N} \right) \cdot U^{W=0}, \end{aligned} \tag{7}$$

where AOV is the average order value (in dollars)[16], $p$ 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 $\frac{\sum_{i=1}^{N} \pi^{\widehat{\tau}_{\ddot{Y}}}(\mathbf{X}_i)}{N}$ calculates the proportion of customers being treated under policy $\pi^{\widehat{\tau}_{\ddot{Y}}}$.

### 6.3. Empirical Results

**6.3.1. GATEs by Predicted CATE Levels.** Figure 4 presents the GATEs by predicted CATE groups. As discussed in Section 2, the U-shaped curve generated by the default method indicates that the CATE model for $Y_{i,10}$ is unable to identify customers with the

[15] We estimate this quantity in each iteration using the probability forest implemented by the `grf` package.

[16] Note that 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).

highest or lowest incremental effects. In contrast, all models that employ short-term signals to estimate CATEs are more effective in ranking customers' long-term treatment effect than the default method.

**Figure 4** **Actual GATEs by Predicted CATE Levels for Different Outcome Variables.**

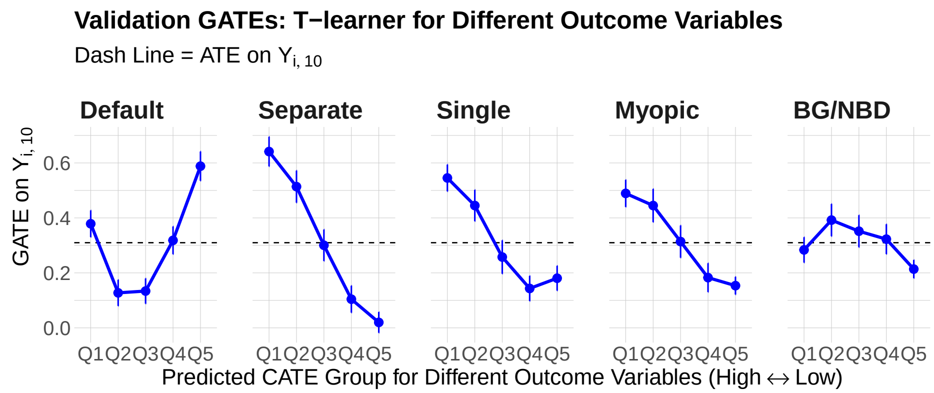


*Note.* Each point represents the mean of bootstrap results on the validation customers together with the two standard error confidence interval. Groups $\mathcal{Q}_1^{\widehat{\tau}_{\dot{Y}}}, \cdots, \mathcal{Q}_5^{\widehat{\tau}_{\dot{Y}}}$ are categorized based on the decreasing order of treatment effect predicted by the CATE model 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 C.7 for the results from other CATE models.

Of all these models, the separate imputation method produces the best targeting performance as it generates the steepest curve. Specifically, the GATE for $\mathcal{Q}_1^{\widehat{\tau}_{\dot{Y}}}$ (representing the most sensitive customers identified by the method) is much larger than that of $\mathcal{Q}_2^{\widehat{\tau}_{\dot{Y}}}$, larger than that of $\mathcal{Q}_3^{\widehat{\tau}_{\dot{Y}}}$, 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, as seems to be the case in this empirical application.

**6.3.2. Profitability of Targeting Policies.** To evaluate the profitability, we compare the expected profits of a targeting policy that targets customers with positive predicted CATEs ($\pi^{\widehat{\tau}_{\dot{Y}}}$ described in Section 6.2.2) with a uniform policy (denoted as $\pi_0$) that rolls out the best intervention to all customers. In our case, since the ATE is positive, $\pi_0$ corresponds to sending three coupons to all customers. We calculate the profit improvement (PI) as $\mathrm{PI}(\pi^{\widehat{\tau}_R}) = \frac{\mathrm{Profit}\left(\pi^{\widehat{\tau}_R}\right)}{\mathrm{Profit}(\pi_0)} - 1$, where $\mathrm{Profit}(\pi_0)$ is calculated in the same way as $\mathrm{Profit}(\pi^{\widehat{\tau}_{\dot{Y}}})$.

Table 2 presents the profit improvement from targeting policies based on different outcomes. The first column demonstrates the outcomes when we target customers based on the predicted CATEs from T-learner, expanding on the results exhibited in Figure 4. The second to fourth columns show the results when using other CATE models. In the last column, we extend our analysis to include the emerging field of policy learning (e.g., Swaminathan and Joachims 2015, Kitagawa and Tetenov 2018) and focus on the doubly robust policy learning technique (DR-PL) introduced by Athey and Wager (2021), which has strong theoretical guarantees. We apply this approach to our analysis as follows: first, we estimate the doubly robust (DR) scores to proxy the CATEs using different outcome variables ($\ddot{Y}$). We then create policies ($\pi_{\ddot{Y}}$) by building cost-sensitive classifiers based on the DR scores using the Probability Forest algorithm implemented in the `grf` package. Finally, we calculate the expected profit improvement using Equation (7). (Please refer to Online Appendix C.6 for details on the implementation of this approach.)

**Table 2** **Expected Profit Improvement: Targeting Based on Different Models.**

| Outcome Variable | T-learner | S-learner | X-learner | Causal Forest | DR-PL |
|---|---|---|---|---|---|
| Separate Imputation | 5.81% (0.58%) | 5.66% (0.58%) | 3.98% (0.54%) | 1.64% (0.57%) | 4.57% (0.58%) |
| Single Imputation | 4.06% (0.62%) | 4.45% (0.62%) | 2.29% (0.59%) | 1.06% (0.61%) | 3.66% (0.59%) |
| Myopic | 3.34% (0.57%) | 3.08% (0.54%) | 2.96% (0.54%) | 1.09% (0.57%) | 3.51% (0.60%) |
| BG/NBD Imputation | 0.95% (0.49%) | 1.61% (0.47%) | 1.93% (0.42%) | −1.44% (0.53%) | −0.26% (0.50%) |
| Default | −3.52% (0.46%) | −1.75% (0.40%) | −0.89% (0.41%) | −4.21% (0.46%) | −3.44% (0.46%) |

We average the profit improvement over 200 replications and show in parentheses the bootstrapped standard errors.

Several key results are worth highlighting. Firstly, consistent with our simulation analyses, the separate imputation method (first row) produces the highest expected profit, regardless of the models used for CATE estimation. In contrast, the default approach (last row) results in negative profit improvement regardless of the targeting method. These results highlight that the noise accumulation problem have severe consequences in terms of profitability, as personalizing the intervention can cause the firm to lose money.

Secondly, the profit loss incurred by high noise levels can be alleviated by using a myopic approach (the third row). This finding is both important and counterintuitive: using *less* information (i.e., fewer observed periods) can be more beneficial for the firm to maximize long-term outcomes. In turn, looking at all results collectively, targeting based on short-term proxies improves profitability in almost all cases (except when using the Causal Forest for BG/NBD imputation).

Thirdly, the relatively poor performance of the Causal Forest is noteworthy. While the Causal Forest performs relatively well in the simulation, this is not the case in the empirical application. This result is consistent with previous work finding that meta-learners can be superior to direct methods in targeting (Yoganarasimhan et al. 2022). Finally, we found that the profit improvement achieved by the doubly robust approach is lower than that achieved by targeting policies based on T-learner and S-learner. This could be because the doubly robust scores are less accurate than the predicted CATEs obtained from T-learner and S-learner.

## 7. Conclusion and Future Directions

Firms often use targeted interventions to increase the long-term profitability. An increasingly popular approach to designing targeting policies combines experimentation (or A/B testing) and customer data to estimate heterogeneous responses to the intervention using CATE models. This paper demonstrates, both theoretically and empirically, that this approach can become ineffective, and even harmful, when the outcome variable accumulates unexplained variations. Therefore, we propose a new targeting paradigm where firms are encouraged to *reduce the noise in the outcome variable*, particularly when it is a long-term outcome, before estimating any CATE model.

Specifically, we present the *separate imputation* approach as a solution to overcome the challenge of long-term CATE estimation in the presence of unexplained variations in the outcome variable. By utilizing short-term behavioral changes to predict long-term responses, this method provides substantial improvements over existing solutions and effectively reduces the impact of unexplained variations when estimating CATE for long-term outcomes. Our solution can be easily implemented using widely available machine learning algorithms, making it practical for businesses across various industries, including those with both contractual and non-contractual relationships. By capitalizing on their historical purchase data, businesses can enhance their marketing efforts and boost long-term profitability without incurring additional costs of increasing the experiment size.

Our proposed solution has been rigorously evaluated using both simulation analyses and real-world marketing data, demonstrating superior targeting performance compared to existing methods. Additionally, our results highlight the trade-off 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 model. In particular, we find that when substantial noise exists in the long-term outcome, utilizing fewer short-term outcomes, even if that violates the surrogacy assumption, can still yield superior targeting performance compared to targeting based on predicted CATE on the actual long-term outcome or not targeting at all. In practice, companies can perform empirical validation to identify the optimal number of short-term outcome periods to incorporate into surrogate models, thereby achieving the best targeting performance.

While our research provides valuable insights and solutions, there are limitations that suggest directions for future research. Firstly, our proposed solution directly addresses the issue of *unobserved heterogeneity* in customer attrition and purchase propensity, which is prevalent in various marketing contexts. 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), 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.

Secondly, in situations where 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. Hence, it is essential to develop a general surrogate selection procedure and document potential surrogate outcomes for various marketing applications. For instance, 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. (2022) provides evidence that short-run conversion on subscription can be an effective low-variance proxy for long-run revenue. Furthermore, 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.

Thirdly, there may be scenarios where no surrogate variable 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. Additionally, 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. (2019) 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 variation in historical data.

## Funding and Competing Interests

Ta-Wei Huang has served on the advisory board for the focal company (who chose to stay anonymous) described in Section 6.

## Acknowledgments

The authors are grateful to an anonymous company for providing the data used in this research. They also appreciate the valuable feedback from Bruce Hardie, Duncan Simester, Jeremy Yang, and faculty and students in the HBS Marketing Unit. Additionally, the authors thank participants of the HBS Doctoral Digital Workshop, 2022 American Causal Inference Conference, 2022 Marketing Science Conference, European Quantitative Marketing Seminar (EQMS), and seminars at HBS, MIT, Cornell, CMU, Michigan, Yale, and Copenhagen Business School for their helpful comments.

## References

Abe M (2009) "Counting your customers" one by one: A hierarchical bayes extension to the Pareto/NBD model. *Marketing Science* 28(3):541–553.

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

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

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

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.

Athey S, Tibshirani J, Wager S (2019b) Generalized random forests. *The Annals of Statistics* 47(2):1148–1178.

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

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

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

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

Chen H, Harinen T, Lee JY, Yung M, Zhao Z (2020) CausalML: Python package for causal machine learning.

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.

Deng A, Xu Y, Kohavi R, Walker T (2013) Improving the sensitivity of online controlled experiments by utilizing pre-experiment data. *Proceedings of the sixth ACM international conference on Web search and data mining*, 123–132.

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

Ellickson PB, Kar W, Reeder JC (2022) Estimating marketing component effects: Double machine learning from targeted digital promotions. *Marketing Science* 0(0).

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

Fader PS, Hardie BG (2007) Incorporating time-invariant covariates into the Pareto/NBD and BG/NBD models. *Technical Note* .

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

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

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

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

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

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

Guadagni PM, Little JD (1983) A logit model of brand choice calibrated on scanner data. *Marketing Science* 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. *International Conference on Modeling and Simulation in Engineering, Economics and Management*, 123–133 (Springer).

Guelman L, Guillén M, Pérez-Marín AM (2015) Uplift random forests. *Cybernetics and 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. *Advances in Neural Information Processing Systems* 34.

Han L, Wang X, Cai T (2021) On the evaluation of surrogate markers in real world data settings. *arXiv preprint arXiv:2104.05513* .

Hastie T, Tibshirani R, Friedman JH, Friedman JH (2009) *The elements of statistical learning: data mining, inference, and prediction*, volume 2 (Springer).

Hitsch GJ, Misra S, Walter Z (2023) Heterogeneous treatment effects and optimal targeting policy evaluation. *Available at SSRN 3111957* .

Holland PW (1986) Statistics and causal inference. *Journal of the American Statistical Association* 81(396):945–960.

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

Imai K, Ratkovic M (2013) Estimating treatment effect heterogeneity in randomized program evaluation. *The Annals of Applied Statistics* 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 Analysis* 19(1):1–19.

Imbens G, Kallus N, Mao X, Wang Y (2022) Long-term causal inference under persistent confounding via data combination. *arXiv preprint arXiv:2202.07234* .

Jin Y, Ba S (2021) Towards optimal variance reduction in online controlled experiments. *arXiv preprint arXiv:2110.13406* .

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

Kennedy EH (2020) Optimal doubly robust estimation of heterogeneous causal effects. *arXiv preprint arXiv:2004.14497* .

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. *Proceedings of the National Academy of Sciences* 116(10):4156–4165.

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

Lu J, Karmarkar U, Venkatraman V (2019) Planning to binge: How consumers choose to allocate media viewing time. *Available at SSRN 3493755* .

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

Mazoure B, Mineiro P, Srinath P, Sedeh RS, Precup D, Swaminathan A (2021) Improving long-term metrics in recommendation systems using short-horizon reinforcement learning. *arXiv preprint arXiv:2106.00589* .

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

Newey WK, Robins JR (2018) Cross-fitting and fast remainder rates for semiparametric estimation. *arXiv preprint arXiv: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 (2019) The customer journey as a source of information. *Working Paper* .

Prentice RL (1989) Surrogate endpoints in clinical trials: definition and operational criteria. *Statistics in 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.

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

Rubin DB (1974) Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology* 66(5):688.

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

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

Semenova V, Chernozhukov V (2021) Debiased machine learning of conditional average treatment effects and other causal functions. *The Econometrics Journal* 24(2):264–289.

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

Simester D, Timoshenko A, Zoumpoulis SI (2022) A sample size calculation for training and certifying targeting policies. *Available at SSRN 4228297* .

Swaminathan A, Joachims T (2015) Batch learning from logged bandit feedback through counterfactual risk minimization. *The Journal of Machine Learning Research* 16(1):1731–1755.

Tulin E, Susumu I, Keane Michael P (2002) A model of consumer brand and quantity choice dynamics under price uncertainty. *Quantitative Marketing and Economics* 1(1).

Wager S, Athey S (2018) Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association* 113(523):1228–1242.

Wang Y, Sharma M, Xu C, Badam S, Sun Q, Richardson L, Chung L, Chi EH, Chen M (2022) Surrogate for long-term user experience in recommender systems. *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 4100–4109.

Yadlowsky S, Fleming S, Shah N, Brunskill E, Wager S (2021) Evaluating treatment prioritization rules via rank-weighted average treatment effects. *arXiv preprint arXiv:2111.07966* .

Yang J, Eckles D, Dhillon P, Aral S (2022) Targeting for long-term outcomes. *arXiv preprint arXiv:2010.15835* .

Yoganarasimhan H, Barzegary E, Pani A (2022) Design and evaluation of optimal free trials. *Management Science* .

Zhang Y, Bradlow ET, Small DS (2013) New measures of clumpiness for incidence data. *Journal of Applied Statistics* 40(11):2533–2548.

# Online Appendix

## A. Proofs

In this appendix we present the proofs of the theoretical results presented in Section 3 of the main document.

### A.1. Class of Weighted-outcome CATE Estimators

First, we show that the class of CATE estimators studied in Section 3 include a wide variety of commonly-used CATE models, such as S-learner, T-learner, Causal Forest, and R-learner.

PROPOSITION 5. *The following CATE estimators belong to the class described in Assumption 2:*

1. *S-learners and T-learners Künzel et al. (2019) with the outcome models being OLS regressions, ridge regressions, $k$-nearest neighbors, or random forests with honest estimation*
2. *R-learners (Nie and Wager 2021) with second-stage models being OLS regressions, ridge regressions, $k$-nearest neighbors, or random forests with honest estimation*
3. *Causal Forest with honest estimation*

*Proof:* Let $\mathcal{D} = \{(\mathbf{X}_i, W_i, Y_i)\}_{i=1}^N$ be the training set for CATE prediction.

1. We provide a proof for S-learners using (a) high-dimensional ridge regression (the proof for OLS estimators is similar), (b) $k$-nearest neighbors, and (c) random forests with honest estimation as the outcome model. We omit the proofs for T-learner as they are similar to the proofs for S-learner.

   *(a) High-dimensional Ridge Regression:*

   Consider the ridge regression model $Y_i = \phi(\mathbf{X}_i, W_i)'\boldsymbol{\beta} + \varepsilon_i$, where $\phi(\mathbf{X}_i, W_i)$ is a high-dimensional feature transformation function used to construct a ridge regression estimator.

   Let $\Phi = [\phi(\mathbf{X}_1, W_1) \cdots \phi(\mathbf{X}_N, W_N)]'$ denote the feature matrix. Then, the closed-form solution of the ridge coefficient with regularization term $\lambda$ can be written as:

$$\widehat{\boldsymbol{\beta}} = (\mathbf{\Phi}'\mathbf{\Phi} + \lambda I)^{-1}\mathbf{\Phi}'\mathbf{y}.$$

Denote $\mathbf{p}_i$ be the $i$-th row of the matrix $(\mathbf{\Phi}'\mathbf{\Phi} + \lambda I)^{-1}\mathbf{\Phi}'$. Then, the predicted CATE for $\mathbf{x}_{\text{test}}$ can be written as

$$\begin{aligned}\widehat{\tau}_Y(\mathbf{x}_{\text{test}}) &= \sum_{i\in\mathcal{D}} [\phi(\mathbf{x}_{\text{test}},1) - \phi(\mathbf{x}_{\text{test}},0)]' \mathbf{p}_i Y_i \\ &= \sum_{i\in\mathcal{D}:\ W_i=1} \underbrace{[\phi(\mathbf{x}_{\text{test}},1) - \phi(\mathbf{x}_{\text{test}},0)]' \mathbf{p}_i}_{\widehat{\ell}_i^1(\mathbf{x}_{\text{test}})} Y_i - \\ &\qquad \sum_{i\in\mathcal{D}:\ W_i=0} \underbrace{[\phi(\mathbf{x}_{\text{test}},0) - \phi(\mathbf{x}_{\text{test}},1)]' \mathbf{p}_i}_{\widehat{\ell}_i^0(\mathbf{x}_{\text{test}})} Y_i.\end{aligned}$$

*(b) $k$-nearest neighbors:*

Let $N_k(\mathbf{x}_{\text{test}}, w)$ represent the set of the $k$-nearest neighboring customers in the training set for a test customer with covariates and treatment assignment $(\mathbf{x}_{\text{test}}, w)$. Then, we can express the predicted CATE as:

$$\begin{aligned}\widehat{\tau}_Y(\mathbf{x}_{\text{test}}) &= \sum_{\mathbf{X}_i\in N_k(\mathbf{x}_{\text{test}},1)} \frac{Y_i}{k} - \sum_{\mathbf{X}_i\in N_k(\mathbf{x}_{\text{test}},0)} \frac{Y_i}{k} \\ &= \sum_{i\in\mathcal{D}:\ W_i=1} \underbrace{\frac{\mathbb{1}[\mathbf{X}_i\in N_k(\mathbf{x}_{\text{test}},1)] - \mathbb{1}[\mathbf{X}_i\in N_k(\mathbf{x}_{\text{test}},0)]}{k}}_{\widehat{\ell}_i^1(\mathbf{x}_{\text{test}})} Y_i - \\ &\qquad \sum_{i\in\mathcal{D}:\ W_i=0} \underbrace{\frac{\mathbb{1}[\mathbf{X}_i\in N_k(\mathbf{x}_{\text{test}},0)] - \mathbb{1}[\mathbf{X}_i\in N_k(\mathbf{x}_{\text{test}},1)]}{k}}_{\widehat{\ell}_i^0(\mathbf{x}_{\text{test}})} Y_i.\end{aligned}$$

*(c) Random Forest with Honest Estimation:*

$\mathcal{D}_1^b, \mathcal{D}_2^b$ be the divided samples for the $b$-th tree $(b = 1, \cdots, B)$, where $\mathcal{D}_1^b$ is used to construct the regression tree and $\mathcal{D}_2^b$ is used to generate predictions. Define $L^b(\mathbf{x}_{\text{test}}, W)$ be the leaf in the $b$-th tree to which customer $(\mathbf{x}_{\text{test}}, W)$ belongs. Using this notation, we can express the predicted CATE as follows:

$$\widehat{\tau}_Y(\mathbf{x}_{\text{test}}) = \sum_{i=1}^{N} \frac{1}{B} \sum_{b=1}^{B} \left[ \frac{\mathbb{1}[\mathbf{X}_i\in L^b(\mathbf{x}_{\text{test}},1),\ i\in\mathcal{D}_2^b]}{|\{i\in\mathcal{D}_2^b:\ \mathbf{X}_i\in L^b(\mathbf{x}_{\text{test}},1)\}|} - \frac{\mathbb{1}[\mathbf{X}_i\in L^b(\mathbf{x}_{\text{test}},0),\ i\in\mathcal{D}_2^b]}{|\{i\in\mathcal{D}_2^b:\ \mathbf{X}_i\in L^b(\mathbf{x}_{\text{test}},0)\}|} \right] Y_i$$

$$= \sum_{i\in\mathcal{D}:W_i=1} \underbrace{\frac{1}{B} \sum_{b=1}^{B} \left( \frac{\mathbb{1}[\mathbf{X}_i\in L^b(\mathbf{x}_{\text{test}},1),\ i\in\mathcal{D}_2^b]}{|\{i\in\mathcal{D}_2^b:\ \mathbf{X}_i\in L^b(\mathbf{x}_{\text{test}},1)\}|} - \frac{\mathbb{1}[\mathbf{X}_i\in L^b(\mathbf{x}_{\text{test}},0),\ i\in\mathcal{D}_2^b]}{|\{i\in\mathcal{D}_2^b:\ \mathbf{X}_i\in L^b(\mathbf{x}_{\text{test}},0)\}|} \right)}_{\widehat{\ell}_i^1(\mathbf{x}_{\text{test}})} Y_i -$$

$$\sum_{i\in\mathcal{D}:W_i=0}\underbrace{\frac{1}{B}\sum_{b=1}^{B}\left(\frac{\mathbb{1}[\mathbf{X}_i\in L^b(\mathbf{x}_{\text{test}},0),\ i\in\mathcal{D}_2^b]}{|\{i\in\mathcal{D}_2^b:\ \mathbf{X}_i\in L^b(\mathbf{x}_{\text{test}},0)\}|}-\frac{\mathbb{1}[\mathbf{X}_i\in L^b(\mathbf{x}_{\text{test}},1),\ i\in\mathcal{D}_2^b]}{|\{i\in\mathcal{D}_2^b:\ \mathbf{X}_i\in L^b(\mathbf{x}_{\text{test}},0)\}|}\right)}_{\widehat{\ell}_i^0(\mathbf{x}_{\text{test}})}Y_i.$$

Note that $\widehat{\ell}_i^W(\mathbf{x}_{\text{test}})$ is independent of $Y_i$ for $i\in\mathcal{D}_2^b$ as the honest estimation only uses customers in $\mathcal{D}_1^b$ to construct the tree.

2. We only present the scenario where high-dimensional ridge regression is used in the second-stage estimation, as the derivations for $k$-nearest neighbors and random forests with honest estimation are similar, using the weights derived in 1.(b) and 1.(c).

   Let $\mathcal{D}_{\text{train}}^1,\cdots,\mathcal{D}_{\text{train}}^Q$ be the sample splits of the training set. Define $\mathcal{D}_{\text{train}}^{(-i)}=\cup_{q:\ i\notin\mathcal{D}_{\text{train}}^q}\mathcal{D}_{\text{train}}^q$ as the training set that excludes the subsample that includes the $i$-th sample. Also, denote $\widehat{e}^{\mathcal{D}_{\text{train}}^{(-i)}}(\mathbf{X}_i)$ (for propensity scores) and $\widehat{m}^{\mathcal{D}^{(-i)}}(\mathbf{X}_i)$ (for conditional means) as the nuisance models trained using the data set $\mathcal{D}_{\text{train}}^{(-i)}$. Then, the Robinson's transformation for each customer is

$$\widetilde{\tau}_i=\frac{Y_i-\widehat{m}^{\mathcal{D}_{\text{train}}^{(-i)}}(\mathbf{X}_i)}{W_i-\widehat{e}^{\mathcal{D}_{\text{train}}^{(-i)}}(\mathbf{X}_i)}.$$

   Now, consider the ridge regression model

$$\widetilde{\tau}_i=\phi(\mathbf{x}_i)'\boldsymbol{\beta}+\varepsilon_i,$$

   where $\phi(\mathbf{x})$ is the high-dimensional feature transformation function. Let $\Phi=[\phi(\mathbf{x}_1)\ \cdots\ \phi(\mathbf{x}_N)]'$ be the feature matrix. Then, the closed-form solution for the ridge coefficient is

$$\widehat{\beta}=(\mathbf{\Phi}'\mathbf{\Phi}+\lambda I)^{-1}\mathbf{\Phi}'\widetilde{\boldsymbol{\tau}}.$$

   Denote $\mathbf{p}_i$ be the $i$-th row of the matrix $(\mathbf{\Phi}'\mathbf{\Phi}+\lambda I)^{-1}\mathbf{\Phi}'$. Then, the predicted CATE for $\mathbf{x}_{\text{test}}$ can be written as

$$\begin{aligned}\widehat{\tau}_Y(\mathbf{x}_{\text{test}})&=\sum_{i\in\mathcal{D}}\phi(\mathbf{x}_{\text{test}})'\mathbf{p}_i\left(\frac{Y_i-\widehat{m}^{\mathcal{D}_{\text{train}}^{(-i)}}(\mathbf{X}_i)}{W_i-\widehat{e}^{\mathcal{D}_{\text{train}}^{(-i)}}(\mathbf{X}_i)}\right)\\&=\sum_{i\in\mathcal{D}:W_i=1}\underbrace{\left(\frac{\phi(\mathbf{x}_{\text{test}})'\mathbf{p}_i}{1-\widehat{e}^{\mathcal{D}_{\text{train}}^{(-i)}}(\mathbf{X}_i)}\right)}_{\widehat{\ell}_i^1(\mathbf{x}_{\text{test}})}[Y_i-\underbrace{\widehat{m}^{\mathcal{D}_{\text{train}}^{(-i)}}(\mathbf{X}_i)}_{\widehat{m}^1(\mathbf{X}_i)}]-\\&\qquad\sum_{i\in\mathcal{D}:W_i=0}\underbrace{\left(\frac{\phi(\mathbf{x}_{\text{test}})'\mathbf{p}_i}{1-\widehat{e}^{\mathcal{D}_{\text{train}}^{(-i)}}(\mathbf{X}_i)}\right)}_{\widehat{\ell}_i^0(\mathbf{x}_{\text{test}})}[Y_i-\underbrace{\widehat{m}^{\mathcal{D}_{\text{train}}^{(-i)}}(\mathbf{X}_i)}_{\widehat{m}^0(\mathbf{X}_i)}].\end{aligned}$$

Note that the cross-fitting and honest estimation assumptions imply that (i) $\widehat{\ell}_i^{W_i}(\mathbf{x}_{\text{test}})$ is independent of $Y_i$ and (ii) $\widehat{m}^w(\mathbf{X}_i)$ is independent of $Y_i$ and $\widehat{\ell}_i^{W_i}(\mathbf{x}_{\text{test}})$.

3. Let $\mathcal{D}_1^b, \mathcal{D}_2^b$ be the divided samples for the $b$-th tree $(b=1,\cdots,B)$, where $\mathcal{D}_1^b$ is used to construct the causal tree and $\mathcal{D}_2^b$ is used to generate predictions. Define $L^b(\mathbf{x}_{\text{test}}$ be the leaf in the $b$-th tree to which customer $(\mathbf{x}_{\text{test}}, W)$ belongs. Using this notation, we can express the predicted CATE as follows:

$$\widehat{\tau}_Y(\mathbf{x}_{\text{test}}) = \sum_{i\in\mathcal{D}:W_i=1} \underbrace{\frac{1}{B}\sum_{b=1}^{B}\frac{\mathbb{1}[\mathbf{X}_i \in L^b(\mathbf{x}_{\text{test}}),\ i\in\mathcal{D}_2^b]}{|\{i\in\mathcal{D}_2^b:\ \mathbf{X}_i\in L^b(\mathbf{x}_{\text{test}})\}|}}_{\widehat{\ell}_i^1(\mathbf{x}_{\text{test}})} Y_i -$$

$$\sum_{i\in\mathcal{D}:W_i=0} \underbrace{\frac{1}{B}\sum_{b=1}^{B}\frac{\mathbb{1}[\mathbf{X}_i \in L^b(\mathbf{x}_{\text{test}}),\ i\in\mathcal{D}_2^b]}{|\{i\in\mathcal{D}_2^b:\ \mathbf{X}_i\in L^b(\mathbf{x}_{\text{test}})\}|}}_{\widehat{\ell}_i^0(\mathbf{x}_{\text{test}})} Y_i.$$

Note that $\widehat{\ell}_i^{W_i}(\mathbf{x}_{\text{test}})$ is independent of $Y_i$ due to honest estimation.

### A.2. Proof for Theorem 1

Let $\mathcal{X}$ be the set of customer covariates, $\mathcal{W}$ the set of treatment assignments, and $\mathcal{Y}$ the set of outcomes for customers in the training set, which includes the data used to construct weights, adjustment functions, and predictions. We denote $\mathcal{D}$ as the index set for customers being used to generate CATE predictions. Define $\mu_Y(\mathbf{X}_i, W_i) = \mathbb{E}[Y_i(W_i)|\mathbf{X}_i]$ and $\sigma_Y^2(\mathbf{X}_i, W_i) = \text{Var}[\varepsilon_i(\mathbf{X}_i, W_i)] = \text{Var}[Y_i(W_i)|\mathbf{X}_i]$.

*Proof for Theorem 1.1:*

First, the conditional bias for each induced outcome model, given $\mathcal{X}$, $\mathcal{W}$, and $\mathcal{D}$, can be expressed as:

$$\begin{aligned}
&\text{Bias}_{\mathcal{Y}|\mathcal{X},\mathcal{W},\mathcal{D}}[\widehat{\mu}_Y^w(\mathbf{x}_{\text{test}})] = \mathbb{E}_{\mathcal{Y}|\mathcal{X},\mathcal{W},\mathcal{D}}[\widehat{\mu}_Y^w(\mathbf{x}_{\text{test}})] - \mu_Y(\mathbf{x}_{\text{test}}, w)\\
&= \mathbb{E}_{\mathcal{Y}|\mathcal{X},\mathcal{W},\mathcal{D}}\left[\sum_{i\in\mathcal{D}:\ W_i=w}\widehat{\ell}_i^w(\mathbf{x}_{\text{test}})\left[Y_i - \widehat{m}_Y^w(\mathbf{X}_i)\right]\right] - \mu(\mathbf{x}_{\text{test}}, w) =\\
&= \sum_{i\in\mathcal{D}:\ W_i=w}\mathbb{E}_{\mathcal{Y}|\mathcal{X},\mathcal{W},\mathcal{D}}\left[\widehat{\ell}_i^w(\mathbf{x}_{\text{test}})\left[\mu(\mathbf{X}_i, w) + \varepsilon_i(\mathbf{X}_i, w) - \widehat{m}_Y^w(\mathbf{X}_i)\right]\right] - \mu(\mathbf{x}_{\text{test}}, w).
\end{aligned} \tag{App-1}$$

Since $\widehat{\ell}_i^w(\mathbf{x}_{\text{test}})$ is independent of $Y_i$ (by the honest estimation assumption) and $\widehat{m}_Y^w(\mathbf{X}_i)$ (by the cross-fitting assumption), we have

$$\mathbb{E}_{\mathcal{Y}|\mathcal{X},\mathcal{W},\mathcal{D}}\left[\widehat{\ell}_i^w(\mathbf{x}_{\text{test}})\left[\mu\left(\mathbf{X}_i, w\right)+\varepsilon_i\left(\mathbf{X}_i, w\right)-\widehat{m}_Y^w(\mathbf{X}_i)\right]\right]=$$
$$\mathbb{E}_{\mathcal{Y}|\mathcal{X},\mathcal{W},\mathcal{D}}\left[\widehat{\ell}_i^w(\mathbf{x}_{\text{test}})\right]\cdot\mathbb{E}_{\mathcal{Y}|\mathcal{X},\mathcal{W},\mathcal{D}}\left[\mu\left(\mathbf{X}_i, w\right)+\varepsilon_i\left(\mathbf{X}_i, w\right)-\widehat{m}_Y^w(\mathbf{X}_i)\right].$$

Replacing the above term into Equation (App-1) gives

$$\text{Bias}_{\mathcal{Y}|\mathcal{X},\mathcal{W},\mathcal{D}}\left[\widehat{\mu}_Y^w(\mathbf{x}_{\text{test}})\right]=\sum_{i\in\mathcal{D}:\ W_i=w}\mathbb{E}_{\mathcal{Y}|\mathcal{X},\mathcal{W},\mathcal{D}}[\widehat{\ell}_i^w(\mathbf{x}_{\text{test}})][\mu\left(\mathbf{X}_i, w\right)-\widehat{m}_Y^w(\mathbf{X}_i)]-\mu(\mathbf{x}_{\text{test}}, w).$$

As a result, the unconditional bias becomes

$$\text{Bias}\left[\widehat{\mu}_Y^w(\mathbf{x}_{\text{test}})\right]=\mathbb{E}\left[\sum_{i\in\mathcal{D}:\ W_i=w}\widehat{\ell}_i^w(\mathbf{x}_{\text{test}})[\mu\left(\mathbf{X}_i, w\right)-\widehat{m}_Y^w(\mathbf{X}_i)]\right]-\mu(\mathbf{x}_{\text{test}}, w).$$

Finally , the bias of the predicted CATE is

$$\begin{aligned}\text{Bias}\left[\widehat{\tau}_Y(\mathbf{x}_{\text{test}})\right]&=\text{Bias}\left[\widehat{\mu}^1(\mathbf{x}_{\text{test}})\right]-\text{Bias}\left[\widehat{\mu}^0(\mathbf{x}_{\text{test}})\right]\\&=\mathbb{E}\left(\sum_{i\in\mathcal{D}:\ W_i=1}\widehat{\ell}_i^1(\mathbf{x}_{\text{test}})[\mu\left(\mathbf{X}_i, 1\right)-\widehat{m}^1(\mathbf{X}_i)]\right)-\\&\qquad\mathbb{E}\left(\sum_{i\in\mathcal{D}:\ W_i=0}\widehat{\ell}_i^0(\mathbf{x}_{\text{test}})[\mu\left(\mathbf{X}_i, 1\right)-\widehat{m}^0(\mathbf{X}_i)]\right)-\tau_Y^{\mu}(\mathbf{x}_{\text{test}}).\end{aligned}$$

It is worth noting that the bias does not depend on $\varepsilon_i(\mathbf{X}_i, W_i)$. As a result, the bias of the predicted CATE is not affected by the variance of unexplained variations.

*Proof for Theorem 1.2:*

First, the conditional variance of $\widehat{\tau}_Y(\mathbf{x}_{\text{test})}$ given $\mathcal{X}$, $\mathcal{W}$, and $\mathcal{D}$ is

$$\begin{aligned}\text{Var}_{\mathcal{Y}|\mathcal{X},\mathcal{W},\mathcal{D}}\left[\widehat{\tau}_Y(\mathbf{x}_{\text{test}})\right]&=\text{Var}_{\mathcal{Y}|\mathcal{X},\mathcal{W},\mathcal{D}}\left[\widehat{\mu}^1(\mathbf{x}_{\text{test}})-\widehat{\mu}^0(\mathbf{x}_{\text{test}})\right]\\&=\text{Var}_{\mathcal{Y}|\mathcal{X},\mathcal{W},\mathcal{D}}\left[\widehat{\mu}^1(\mathbf{x}_{\text{test}})\right]+\text{Var}_{\mathcal{Y}|\mathcal{X},\mathcal{W},\mathcal{D}}\left[\widehat{\mu}^0(\mathbf{x}_{\text{test}})\right]-2\text{Cov}_{\mathcal{Y}|\mathcal{X},\mathcal{W},\mathcal{D}}\left[\widehat{\mu}^1(\mathbf{x}_{\text{test}}),\widehat{\mu}^0(\mathbf{x}_{\text{test}})\right].\end{aligned}$$

Under Assumption 1, the covariance $\text{Cov}_{\mathcal{Y}|\mathcal{X},\mathcal{W},\mathcal{D}}\left[\widehat{\mu}^1(\mathbf{x}_{\text{test}}),\widehat{\mu}^0(\mathbf{x}_{\text{test}})\right]$ is equal to zero as $\widehat{\mu}^1(\mathbf{x}_{\text{test}})$ and $\widehat{\mu}^1(\mathbf{x}_{\text{test}})$ are derived from two different samples (i.e., the treatment and control group, respectively).

Now, consider the variance term for each outcome model. We first write it as:

$$\text{Var}_{\mathcal{Y}|\mathcal{X},\mathcal{W},\mathcal{D}}\left[\widehat{\mu}_Y^w(\mathbf{x}_{\text{test}})\right]=\mathbb{E}_{\mathcal{Y}|\mathcal{X},\mathcal{W},\mathcal{D}}\left[\left(\widehat{\mu}_Y^w(\mathbf{x}_{\text{test}})\right)^2\right]-\mathbb{E}^2_{\mathcal{Y}}\left[\widehat{\mu}_Y^w(\mathbf{x}_{\text{test}})\right]. \tag{App-2}$$

The first term in Equation (App-2) is

$$\begin{aligned}
&\mathbb{E}_{\mathcal{Y}|\mathcal{X},\mathcal{W},\mathcal{D}}\left\{\left[\widehat{\mu}_Y^w(\mathbf{x}_{\text{test}})\right]^2\right\} \\
&= \mathbb{E}_{\mathcal{Y}|\mathcal{X},\mathcal{W},\mathcal{D}}\left\{\left(\sum_{i\in\mathcal{D}:\ W_i=w}\widehat{\ell}_i^w(\mathbf{x}_{\text{test}})\left[Y_i(w)-\widehat{m}_Y^w(\mathbf{X}_i)\right]\right)^2\right\} \\
&= \mathbb{E}_{\mathcal{Y}|\mathcal{X},\mathcal{W},\mathcal{D}}\left\{\sum_{i,j\in\mathcal{D}:\ W_i=W_j=w}\widehat{\ell}_i^w(\mathbf{x}_{\text{test}})\widehat{\ell}_j^w(\mathbf{x}_{\text{test}})\left[Y_i(w)-\widehat{m}_Y^w(\mathbf{X}_i)\right]\left[Y_j(w)-\widehat{m}_Y^w(\mathbf{X}_j)\right]\right\}.
\end{aligned}$$

Note that we have the following decomposition for the term in the expectation:

$$\begin{aligned}
&\sum_{i\in\mathcal{D}:\ W_i=W_j=w}\widehat{\ell}_i^w(\mathbf{x}_{\text{test}})\widehat{\ell}_j^w(\mathbf{x}_{\text{test}})\left[Y_i(w)-\widehat{m}_Y^w(\mathbf{X}_i)\right]\left[Y_j(w)-\widehat{m}_Y^w(\mathbf{X}_j)\right] \\
&= \sum_{i,j\in\mathcal{D}:\ W_i=W_j=w}\widehat{\ell}_i^w(\mathbf{x}_{\text{test}})\widehat{\ell}_j^w(\mathbf{x}_{\text{test}})\left[\mu\left(\mathbf{X}_i,w\right)-\widehat{m}_Y^w(\mathbf{X}_i)\right]\left[\mu\left(\mathbf{X}_j,w\right)-\widehat{m}_Y^w(\mathbf{X}_j)\right]+ \\
&\qquad 2\cdot\sum_{i,j\in\mathcal{D}:\ W_i=W_j=w}\widehat{\ell}_i^w(\mathbf{x}_{\text{test}})\widehat{\ell}_j^w(\mathbf{x}_{\text{test}})\left[\mu\left(\mathbf{X}_i,w\right)-\widehat{m}_Y^w(\mathbf{X}_i)\right]\cdot\varepsilon_j\left(\mathbf{X}_j,w\right)+ \\
&\qquad \sum_{i,j\in\mathcal{D}:\ W_i=W_j=w}\widehat{\ell}_i^w(\mathbf{x}_{\text{test}})\widehat{\ell}_j^w(\mathbf{x}_{\text{test}})\cdot\varepsilon_i\left(\mathbf{X}_i,w\right)\varepsilon_j\left(\mathbf{X}_i,w\right).
\end{aligned}$$

The conditional expectation of the first term in the above decomposition is

$$\begin{aligned}
&\mathbb{E}_{\mathcal{Y}|\mathcal{X},\mathcal{W},\mathcal{D}}\left\{\sum_{i,j\in\mathcal{D}:\ W_i=W_j=w}\widehat{\ell}_i^w(\mathbf{x}_{\text{test}})\widehat{\ell}_j^w(\mathbf{x}_{\text{test}})\left[\mu\left(\mathbf{X}_i,w\right)-\widehat{m}_Y^w(\mathbf{X}_i)\right]\left[\mu\left(\mathbf{X}_j,w\right)-\widehat{m}_Y^w(\mathbf{X}_j)\right]\right\} \\
&= \mathbb{E}^2_{\mathcal{Y}|\mathcal{X},\mathcal{W},\mathcal{D}}\left[\widehat{\mu}_Y^w(\mathbf{x}_{\text{test}})\right].
\end{aligned}$$

By the assumption of honest estimation and cross-fitting, the conditional expectation of the second term in the above decomposition is zero

$$\begin{aligned}
&\mathbb{E}_{\mathcal{Y}|\mathcal{X},\mathcal{W},\mathcal{D}}\left[\sum_{i,j\in\mathcal{D}:\ W_i=W_j=w}\widehat{\ell}_i^w(\mathbf{x}_{\text{test}})\widehat{\ell}_j^w(\mathbf{x}_{\text{test}})\left[\mu\left(\mathbf{X}_i,w\right)-\widehat{m}_Y^w(\mathbf{X}_i)\right]\cdot\varepsilon_j\left(\mathbf{X}_j,w\right)\right] \\
&= \mathbb{E}_{\mathcal{Y}|\mathcal{X},\mathcal{W},\mathcal{D}}\left\{\sum_{i,j\in\mathcal{D}:\ W_i=W_j=w}\widehat{\ell}_i^w(\mathbf{x}_{\text{test}})\widehat{\ell}_j^w(\mathbf{x}_{\text{test}})\left[\mu\left(\mathbf{X}_i,w\right)-\widehat{m}_Y^w(\mathbf{X}_i)\right]\right\}\cdot\mathbb{E}_{\mathcal{Y}|\mathcal{X},\mathcal{W},\mathcal{D}}\left[\varepsilon_j\left(\mathbf{X}_j,w\right)\right] \\
&= 0.
\end{aligned}$$

The conditional expectation of the third term in the above decomposition can be written as

$$\begin{aligned}
&\mathbb{E}_{\mathcal{Y}|\mathcal{X},\mathcal{W},\mathcal{D}}\left[\sum_{i,j\in\mathcal{D}:\ W_i=W_j=w} \widehat{\ell}_i^w(\mathbf{x}_{\text{test}})\widehat{\ell}_j^w(\mathbf{x}_{\text{test}}) \cdot \varepsilon_i\left(\mathbf{X}_i, w\right)\varepsilon_j\left(\mathbf{X}_i, w\right)\right] \\
&= \sum_{i\in\mathcal{D}:\ W_i=w} \mathbb{E}_{\mathcal{Y}|\mathcal{X},\mathcal{W},\mathcal{D}}\left[\widehat{\ell}_i^w(\mathbf{x}_{\text{test}})^2\varepsilon_i^2\left(\mathbf{X}_i, w\right)\right] + \\
&\qquad \mathbb{E}_{\mathcal{Y}|\mathcal{X},\mathcal{W},\mathcal{D}}\left[\sum_{i\neq j\in\mathcal{D}:\ W_i=W_j=w} \widehat{\ell}_i^w(\mathbf{x}_{\text{test}})\widehat{\ell}_j^w(\mathbf{x}_{\text{test}}) \cdot \varepsilon_i\left(\mathbf{X}_i, w\right)\varepsilon_j\left(\mathbf{X}_i, w\right)\right] \\
&= \sum_{i\in\mathcal{D}:\ W_i=w} \mathbb{E}_{\mathcal{Y}|\mathcal{X},\mathcal{W},\mathcal{D}}\left[\widehat{\ell}_i^w(\mathbf{x}_{\text{test}})^2\right]\sigma^2\left(\mathbf{X}_i, w\right)
\end{aligned}$$

since (i) $\widehat{\ell}_i^w(\mathbf{x}_{\text{test}}), \widehat{\ell}_j^w(\mathbf{x}_{\text{test}})$ are independent of $\varepsilon_i\left(\mathbf{X}_i, w\right), \varepsilon_j\left(\mathbf{X}_i, w\right)$ (by the assumption of honest estimation and cross-fitting) and (ii) $\varepsilon_i\left(\mathbf{X}_i, w\right), \varepsilon_j\left(\mathbf{X}_i, w\right)$ are independent of each other.

Combining all the things together, we can write the unconditional variance as

$$\mathrm{Var}\left[\widehat{\mu}_Y^w(\mathbf{x}_{\text{test}})\right] = \sum_{i\in\mathcal{D}:\ W_i=w} \mathbb{E}\left[\widehat{\ell}_i^w(\mathbf{x}_{\text{test}})^2\right]\sigma^2\left(\mathbf{X}_i, w\right).$$

As a result, the unconditional variance of $\widehat{\tau}_Y(\mathbf{x}_{\text{test}})$ is

$$\begin{aligned}
\mathrm{Var}\left[\widehat{\tau}_Y(\mathbf{x}_{\text{test}})\right] &= \mathrm{Var}\left[\widehat{\mu}^1(\mathbf{x}_{\text{test}})\right] + \mathrm{Var}\left[\widehat{\mu}^0(\mathbf{x}_{\text{test}})\right] \\
&= \mathbb{E}\left(\sum_{i\in\mathcal{D}:\ W_i=1} \widehat{\ell}_i^1(\mathbf{x}_{\text{test}})^2\sigma^2\left(\mathbf{X}_i, 1\right) + \sum_{i\in\mathcal{D}:\ W_i=0} \widehat{\ell}_i^0(\mathbf{x}_{\text{test}})^2\sigma^2\left(\mathbf{X}_i, 0\right)\right).
\end{aligned}$$

### A.3. Proof for Proposition 1

The probability that the learned policy makes a different decision than the optimal targeting policy is

$$\begin{aligned}
&\mathbb{P}\left[\tau_Y(\mathbf{x}_{\text{test}}) \cdot \widehat{\tau}_Y(\mathbf{x}_{\text{test}}) < 0\right] = \\
&\quad\begin{cases}
\mathbb{P}\left[\frac{\widehat{\tau}_Y(\mathbf{x}_{\text{test}}) - \tau_Y(\mathbf{x}_{\text{test}})}{\sqrt{\mathrm{Var}[\widehat{\tau}_Y(\mathbf{x}_{\text{test}})]}} < \frac{-|\tau_Y(\mathbf{x}_{\text{test}})|}{\sqrt{\mathrm{Var}[\widehat{\tau}_Y(\mathbf{x}_{\text{test}})]}}\right], & \text{if } \tau_Y(\mathbf{x}_{\text{test}}) > 0, \\
1 - \mathbb{P}\left[\frac{\widehat{\tau}_Y(\mathbf{x}_{\text{test}}) - \tau_Y(\mathbf{x}_{\text{test}})}{\sqrt{\mathrm{Var}[\widehat{\tau}_Y(\mathbf{x}_{\text{test}})]}} < \frac{|\tau_Y(\mathbf{x}_{\text{test}})|}{\sqrt{\mathrm{Var}[\widehat{\tau}_Y(\mathbf{x}_{\text{test}})]}}\right], & \text{if } \tau_Y(\mathbf{x}_{\text{test}}) < 0.
\end{cases}
\end{aligned}$$

Then, we can see that the probability increases as $\mathrm{Var}\left[\widehat{\tau}_Y(\mathbf{x}_{\text{test}})\right]$ increases. By Theorem 1, the larger the variance of the outcome variable, the larger $\mathrm{Var}\left[\widehat{\tau}_Y(\mathbf{x}_{\text{test}})\right]$. Therefore, the mistargeting probability is also increasing in the variance of unexplained variations.

Now, let us consider the case when the firm aims to target customers with CATEs larger than the threshold $c$. The mistargeting probability is

$$\begin{aligned}
&\mathbb{P}\left[(\tau_Y(\mathbf{x}_{\text{test}})-c)\cdot(\widehat{\tau}_Y(\mathbf{x}_{\text{test}})-c)<0\right]= \\
&\quad\begin{cases}
\mathbb{P}\left[\frac{\widehat{\tau}_Y(\mathbf{x}_{\text{test}})-\tau_Y(\mathbf{x}_{\text{test}})}{\sqrt{\mathrm{Var}[\widehat{\tau}_Y(\mathbf{x}_{\text{test}})]}}<\frac{-|\tau_Y(\mathbf{x}_{\text{test}})-c|}{\sqrt{\mathrm{Var}[\widehat{\tau}_Y(\mathbf{x}_{\text{test}})]}}\right], & \text{if } \tau_Y(\mathbf{x}_{\text{test}})-c>0, \\
1-\mathbb{P}\left[\frac{\widehat{\tau}_Y(\mathbf{x}_{\text{test}})-\tau_Y(\mathbf{x}_{\text{test}})}{\sqrt{\mathrm{Var}[\widehat{\tau}_Y(\mathbf{x}_{\text{test}})]}}<\frac{|\tau_Y(\mathbf{x}_{\text{test}})-c|}{\sqrt{\mathrm{Var}[\widehat{\tau}_Y(\mathbf{x}_{\text{test}})]}}\right], & \text{if } \tau_Y(\mathbf{x}_{\text{test}})-c<0.
\end{cases}
\end{aligned}$$

Then, we can see that the mistargeting probability increases as $\mathrm{Var}\left[\widehat{\tau}_Y(\mathbf{x}_{\text{test}})\right]$ increases.

### A.4. Proof for Theorem 2

Note that we can write the long-term outcome as

$$\begin{aligned}
Y_i^T(W_i) &= \sum_{t=1}^{T}\left[\theta_t(\mathbf{X}_i,W_i)+\varepsilon_{i,t}^{\theta}\right]\cdot\left[\lambda_t(\mathbf{X}_i,W_i)+\varepsilon_{i,t}^{\lambda}\right] \\
&= \sum_{t=1}^{T}\theta_t(\mathbf{X}_i,W_i)\lambda_t(\mathbf{X}_i,W_i)+\sum_{t=1}^{T}\varepsilon_t^{S}(W_i),
\end{aligned}$$

where $\varepsilon_t^S(W_i)=\theta_t(\mathbf{X}_i,W_i)\varepsilon_{i,t}^{\lambda}+\lambda_t(\mathbf{X}_i,W_i)\varepsilon_{i,t}^{\theta}+\varepsilon_{i,t}^{\theta}\varepsilon_{i,t}^{\lambda}$. Then, we can write the variance of $Y_{i,T}(W_i)$ as follows:

$$\begin{aligned}
\mathrm{Var}[Y_{i,T}(W_i)|\mathbf{X}_i] &= \mathrm{Var}\left[\sum_{t=1}^{T}\varepsilon_{i,t}^{S}(W_i)\middle|\mathbf{X}_i\right] \\
&= \sum_{t=1}^{T}\mathrm{Var}\left[\varepsilon_{i,t}^{S}(W_i)|\mathbf{X}_i\right]+2\sum_{1\le t_1<t_2<T}\mathrm{Cov}\left[\varepsilon_{i,t_1}^{S}(W_i),\varepsilon_{i,t_2}^{S}(W_i)|\mathbf{X}_i\right].
\end{aligned} \tag{App-3}$$

The first element in Equation (App-3), $\mathrm{Var}\left[\varepsilon_{i,t}^{S}(W_i)\right]$, is always positive. The second element, $\mathrm{Cov}\left[\varepsilon_{i,t_1}^{S}(W_i),\varepsilon_{i,t_2}^{S}(W_i)|\mathbf{X}_i\right]$, can be written as

$$\begin{aligned}
&\mathrm{Cov}\left[\varepsilon_{i,t_1}^{S}(W_i),\varepsilon_{i,t_2}^{S}(W_i)|\mathbf{X}_i\right] \\
&\quad=\theta_{t_1}(\mathbf{X}_i,W_i)\theta_{t_2}(\mathbf{X}_i,W_i)\mathrm{Cov}\left(\varepsilon_{i,t_1}^{\lambda},\varepsilon_{i,t_2}^{\lambda}\right)+\theta_{t_1}(\mathbf{X}_i,W_i)\lambda_{t_2}(\mathbf{X}_i,W_i)\mathrm{Cov}\left(\varepsilon_{i,t_1}^{\lambda},\varepsilon_{i,t_2}^{\theta}\right)+ \\
&\quad\quad\theta_{t_1}(\mathbf{X}_i,W_i)\mathrm{Cov}\left(\varepsilon_{i,t_1}^{\lambda},\varepsilon_{i,t_2}^{\lambda}\varepsilon_{i,t_2}^{\theta}\right)+\lambda_{t_1}(\mathbf{X}_i,W_i)\lambda_{t_2}(\mathbf{X}_i,W_i)\mathrm{Cov}\left(\varepsilon_{i,t_1}^{\theta},\varepsilon_{i,t_2}^{\theta}\right)+ \\
&\quad\quad\lambda_{t_1}(\mathbf{X}_i,W_i)\theta_{t_2}(\mathbf{X}_i,W_i)\mathrm{Cov}\left(\varepsilon_{i,t_1}^{\theta},\varepsilon_{i,t_2}^{\lambda}\right)+\lambda_{t_1}(\mathbf{X}_i,W_i)\mathrm{Cov}\left(\varepsilon_{i,t_1}^{\theta},\varepsilon_{i,t_2}^{\lambda}\varepsilon_{i,t_2}^{\theta}\right)+ \\
&\quad\quad\theta_{t_2}(\mathbf{X}_i,W_i)\mathrm{Cov}\left(\varepsilon_{i,t_2}^{\lambda},\varepsilon_{i,t_1}^{\lambda}\varepsilon_{i,t_1}^{\theta}\right)+\lambda_{t_2}(\mathbf{X}_i,W_i)\mathrm{Cov}\left(\varepsilon_{i,t_2}^{\theta},\varepsilon_{i,t_1}^{\lambda}\varepsilon_{i,t_1}^{\theta}\right)+\mathrm{Cov}\left(\varepsilon_{i,t_1}^{\lambda}\varepsilon_{i,t_1}^{\theta},\varepsilon_{i,t_2}^{\lambda}\varepsilon_{i,t_2}^{\theta}\right).
\end{aligned}$$

By the independence assumption ($\{\varepsilon^{\theta}_{i,t}\}_{t=1}^{T} \perp\!\!\!\perp \{\varepsilon^{\lambda}_{i,t}\}_{t=1}^{T}$), we have

$$\mathrm{Cov}\left(\varepsilon^{\lambda}_{i,t_1}, \varepsilon^{\lambda}_{i,t_2}\varepsilon^{\theta}_{i,t_2}\right) = \mathbb{E}\left(\varepsilon^{\lambda}_{i,t_1}\varepsilon^{\lambda}_{i,t_2}\varepsilon^{\theta}_{i,t_2}\right) = \mathbb{E}\left(\varepsilon^{\lambda}_{i,t_1}\varepsilon^{\lambda}_{i,t_2}\right)\mathbb{E}\left(\varepsilon^{\theta}_{i,t_2}\right) = 0,$$

$$\mathrm{Cov}\left(\varepsilon^{\theta}_{i,t_1}, \varepsilon^{\lambda}_{i,t_2}\varepsilon^{\theta}_{i,t_2}\right) = \mathbb{E}\left(\varepsilon^{\theta}_{i,t_1}\varepsilon^{\theta}_{i,t_2}\varepsilon^{\lambda}_{i,t_2}\right) = \mathbb{E}\left(\varepsilon^{\theta}_{i,t_1}\varepsilon^{\theta}_{i,t_2}\right)\mathbb{E}\left(\varepsilon^{\lambda}_{i,t_2}\right) = 0,$$

$$\mathrm{Cov}\left(\varepsilon^{\lambda}_{i,t_1}\varepsilon^{\theta}_{i,t_1}, \varepsilon^{\lambda}_{i,t_2}\varepsilon^{\theta}_{i,t_2}\right) = \mathbb{E}\left(\varepsilon^{\lambda}_{i,t_1}\varepsilon^{\lambda}_{i,t_2}\right)\mathbb{E}\left(\varepsilon^{\theta}_{i,t_1}\varepsilon^{\theta}_{i,t_2}\right) = \mathrm{Cov}\left(\varepsilon^{\lambda}_{i,t_1}, \varepsilon^{\lambda}_{i,t_2}\right)\mathrm{Cov}\left(\varepsilon^{\theta}_{i,t_1}, \varepsilon^{\theta}_{i,t_2}\right).$$

Therefore, the covariance term becomes

$$\begin{aligned}\mathrm{Cov}\left[\varepsilon^{S}_{i,t_1}(W_i), \varepsilon^{S}_{i,t_2}(W_i)|\mathbf{X}_i\right] =& \theta_{t_1}(\mathbf{X}_i, W_i)\theta_{t_2}(\mathbf{X}_i, W_i)\mathrm{Cov}\left(\varepsilon^{\lambda}_{i,t_1}, \varepsilon^{\lambda}_{i,t_2}\right) + \\ & \lambda_{t_1}(\mathbf{X}_i, W_i)\lambda_{t_2}(\mathbf{X}_i, W_i)\mathrm{Cov}\left(\varepsilon^{\theta}_{i,t_1}, \varepsilon^{\theta}_{i,t_2}\right) + \mathrm{Cov}\left(\varepsilon^{\lambda}_{i,t_1}, \varepsilon^{\lambda}_{i,t_2}\right)\mathrm{Cov}\left(\varepsilon^{\theta}_{i,t_1}, \varepsilon^{\theta}_{i,t_2}\right).\end{aligned}$$

Since $\mathrm{Cov}\left(\varepsilon^{\lambda}_{i,t_1}, \varepsilon^{\lambda}_{i,t_2}\right)$ and $\mathrm{Cov}\left(\varepsilon^{\theta}_{i,t_1}, \varepsilon^{\theta}_{i,t_2}\right)$ are non-negative by assumptions, $\mathrm{Cov}\left(\varepsilon^{S}_{i,t_1}, \varepsilon^{S}_{i,t_2}\right)$ is also non-negative.

Finally, we have the variance increase property:

$$\begin{aligned}&\mathrm{Var}[Y_{i,T+1}(W_i)] - \mathrm{Var}[Y_{i,T}(W_i)] \\ &= \sum_{t=1}^{T}\mathrm{Var}\left[\varepsilon^{S}_{i,t}(W_i)|\mathbf{X}_i\right] + 2\sum_{1\le t_1<t_2<T}\mathrm{Cov}\left[\varepsilon^{S}_{i,t_1}(W_i), \varepsilon^{S}_{i,t_2}(W_i)|\mathbf{X}_i\right] > 0.\end{aligned}$$

### A.5. Proof for Proposition 2

The serial correlation of the unexplained variations is positive because

$$\begin{aligned}\mathrm{Cov}(\varepsilon^{\lambda}_{i,t_1}, \varepsilon^{\lambda}_{i,t_2}) &= \mathrm{Cov}\left(\bar{\varepsilon}^{\lambda}_{i} + \eta^{\lambda}_{i,t_1}, \bar{\varepsilon}^{\lambda}_{i} + \eta^{\lambda}_{i,t_2}\right) \\ &= \mathrm{Var}(\bar{\varepsilon}^{\lambda}_{i}) + \mathrm{Cov}\left(\bar{\varepsilon}^{\lambda}_{i}, \eta^{\lambda}_{i,t_2}\right) + \mathrm{Cov}\left(\eta^{\lambda}_{i,t_1}, \eta^{\lambda}_{i,t_2}\right) + \mathrm{Cov}\left(\eta^{\lambda}_{i,t_1}, \eta^{\lambda}_{i,t_2}\right) \\ &= \mathrm{Var}(\bar{\varepsilon}^{\lambda}_{i}) > 0,\end{aligned}$$

provided $\eta^{\lambda}_{i,t_1}$ and $\eta^{\lambda}_{i,t_2}$ are assumed to be independent of $\bar{\varepsilon}^{\lambda}_{i}$ and each other.

### A.6. Proof for Proposition 3

The covariance of unexplained variations in churn is positive because

$$\begin{aligned}\mathrm{Cov}\left(\varepsilon^{\theta}_{i,t_1}, \varepsilon^{\theta}_{i,t_2}\right) &= \mathrm{Cov}\left[\theta_{t_1}(\mathbf{X}_i, W_i) + \varepsilon^{\theta}_{i,t_1}, \theta_{t_2}(\mathbf{X}_i, W_i) + \varepsilon^{\theta}_{i,t_2}\right] \\ &= \mathrm{Cov}\left(\delta_{i,t_1}, \delta_{i,t_2}\right) \\ &= \mathbb{E}\left[\delta_{i,t_1}\delta_{i,t_2}\right] - \mathbb{E}\left[\delta_{i,t_1}\right]\mathbb{E}\left[\delta_{i,t_2}\right] \\ &= \theta_{t_2}(\mathbf{X}_i, W_i) - \theta_{t_1}(\mathbf{X}_i, W_i)\theta_{t_2}(\mathbf{X}_i, W_i) \\ &= \theta_{t_2}(\mathbf{X}_i, W_i)[1 - \theta_{t_1}(\mathbf{X}_i, W_i)] > 0.\end{aligned}$$

### A.7. Proof for Theorem 3

1. By the comparability assumption, we have

$$\widetilde{Y}_T(\mathbf{S}_{T_0}, \mathbf{X}_i) \equiv \mathbb{E}_{\mathcal{H}}\left[Y_{i,T} | \mathbf{S}_{T_0}, \mathbf{X}_i\right] = \mathbb{E}_{\mathcal{E}}\left[Y_{i,T} | \mathbf{S}_{T_0}, \mathbf{X}_i\right].$$

   By the surrogacy assumption, we have

$$\mathbb{E}_{\mathcal{E}}\left[Y_{i,T}(W_i) | \mathbf{X}_i\right] = \mathbb{E}_{\mathcal{E}}\left[Y_{i,T} | \mathbf{S}_{T_0}(W_i), \mathbf{X}_i\right].$$

   Combining these two observations gives

$$\begin{aligned}\tau_{Y_T}(\mathbf{X}_i) &= \mathbb{E}_{\mathcal{E}}\left[Y_{i,T}(1) | \mathbf{X}_i\right] - \mathbb{E}_{\mathcal{E}}\left[Y_{i,T}(0) | \mathbf{X}_i\right] \\ &= \widetilde{Y}_T(\mathbf{S}_{i,T_0}(1), \mathbf{X}_i) - \widetilde{Y}_T(\mathbf{S}_{i,T_0}(0), \mathbf{X}_i).\end{aligned}$$

2. By the law of total variance, we have

$$\begin{aligned}\mathrm{Var}[Y_{i,T}(W_i)|\mathbf{X}_i] &= \mathbb{E}\left(\mathrm{Var}[Y_{i,T}(W_i)|\mathbf{X}_i, \mathbf{S}_{i,T_0}(W_i)]\right) + \mathrm{Var}\left(\mathbb{E}[Y_{i,T}(W_i)|\mathbf{X}_i, \mathbf{S}_{i,T_0}(W_i)]\right) \\ &= \mathbb{E}\left(\mathrm{Var}[Y_{i,T}(W_i)|\mathbf{X}_i, \mathbf{S}_{i,T_0}(W_i)]\right) + \mathrm{Var}[\widetilde{Y}_T(\mathbf{S}_{i,T_0}(W_i), \mathbf{X}_i)] \\ &> \mathrm{Var}[\widetilde{Y}_T(\mathbf{S}_{i,T_0}(W_i), \mathbf{X}_i)]\end{aligned}$$

   since $\mathbb{E}\left(\mathrm{Var}[Y_{i,T}(W_i)|\mathbf{X}_i, \mathbf{S}_{i,T_0}(W_i)]\right) > 0$.

### A.8. Proof for Proposition 4

By the law of total variance, we have

$$\begin{aligned}&\mathrm{Var}[\widetilde{Y}_T(\mathbf{S}_{i,T_0}(W_i), \mathbf{X}_i)] \\ &= \mathbb{E}\left(\mathrm{Var}[\widetilde{Y}_T(\mathbf{S}_{i,T_0}(W_i), \mathbf{X}_i)|\mathbf{S}_{i,T_0'}(W_i)]\right) + \mathrm{Var}\left(\mathbb{E}[\widetilde{Y}_T(\mathbf{S}_{i,T_0}(W_i), \mathbf{X}_i)|\mathbf{S}_{i,T_0'}(W_i)]\right) \\ &= \mathbb{E}\left(\mathrm{Var}[\widetilde{Y}_T(\mathbf{S}_{i,T_0}(W_i), \mathbf{X}_i)|\mathbf{S}_{i,T_0'}(W_i)]\right) + \mathrm{Var}\left(\mathbb{E}[Y_i(W_i)|\mathbf{S}_{i,T_0'}(W_i)]\right) \\ &= \mathbb{E}\left(\mathrm{Var}[\widetilde{Y}_T(\mathbf{S}_{i,T_0}(W_i), \mathbf{X}_i)|\mathbf{S}_{i,T_0'}(W_i)]\right) + \mathrm{Var}\left[\widetilde{Y}_T(\mathbf{S}'_{i,T_0}(W_i), \mathbf{X}_i)\right] \\ &> \mathrm{Var}[\widetilde{Y}_T(\mathbf{S}'_{i,T_0}(W_i), \mathbf{X}_i)]\end{aligned}$$

since $\mathbb{E}\left(\mathrm{Var}[\widetilde{Y}_T(\mathbf{S}_{i,T_0}(W_i), \mathbf{X}_i)|\mathbf{S}_{i,T_0'}(W_i)]\right) > 0$.

## B. Further Details about the Simulation Analyses

In this appendix, we provide details about the simulation analyses described in Section 5 of the main document.

### B.1. Simulation Setting

For the simulation, we consider a company that conducts a marketing intervention and aims to maximize the total purchase counts ($Y_{i,10}$) over a ten-week time-frame following the intervention. We simulate the data based on the customer behaviors described in Section 3.3:

1. There are two pre-treatment covariates which are drawn i.i.d. from the standard normal distribution, i.e., $X_{i,1}, X_{i,2} \sim_{i.i.d.} \mathcal{N}(0,1)$.
2. At the end of each period, a customer churns with probability $p_{i,t}$.
3. The realized purchase counts in each period when a customer is alive follows a Poisson distribution with mean purchase rate $\lambda_{i,t}$, i.e., $\widetilde{S}_{i,t} \sim \mathsf{Poisson}(\lambda_{i,t})$.
4. The intervention reduces customers' churn probability ($p_{i,t}$) in the *first three periods*. In particular, the churn probability is

$$\text{(Treatment)} \quad p_{i,t}(W_i = 1) = \begin{cases} \frac{1}{\exp(1.5+0.5X_{i,1}+0.4X_{i,2})}, & \text{if } t \leq 3, \\ \frac{1}{\exp(1.4+0.5X_{i,1}+0.4X_{i,2})}, & \text{if } t > 3. \end{cases}$$

$$\text{(Control)} \quad p_{i,t}(W_i = 0) = \frac{1}{\exp(1.4+0.5X_{i,1}+0.4X_{i,2})}, \quad \forall t = 1, \cdots, 10.$$

5. The intervention increases customers' purchase rates in the *first three periods* and has no impact on later periods. The purchase rate follows:

$$\text{(Treatment)} \quad \lambda_{i,t}(W_i = 1) = \begin{cases} \exp(1+0.5X_{i,1}+0.5X_{i,2}), & \text{if } t \leq 3, \\ \exp(0.9+0.5X_{i,1}+0.4X_{i,2}), & \text{if } t > 3. \end{cases}$$

$$\text{(Control)} \quad \lambda_{i,t}(W_i = 0) = \exp(0.9+0.5X_{i,1}+0.4X_{i,2}), \quad \forall t = 1, \cdots, 10.$$

Note that we choose the logit link function for the churn probability to ensure that it lies between 0 and 1. Additionally, we use the exponential link function for the purchase rate to ensure that it is non-negative.

### B.2. Evaluation Procedure

The following procedure is performed to evaluate the performance of different approaches:

1. Derive the outcome variable $\ddot{Y}$

2. Generate one training set (with $N/2$ treated customers and $N/2$ non-treated customers) and one validation set (with $5,000$ customers for each condition)
3. Construct CATE models ($\widehat{\tau}_{\ddot{Y}}$) using the training set
4. Calculate the AUTOCs of $\widehat{\tau}_{\ddot{Y}}$ using the validation set

We generate 200 bootstrap samples and report the mean and standard deviations of each quantity for performance evaluation.

### B.3. Specification of Surrogate Indices

As discussed in Section 5 of the main document, we utilize historical data ($\mathcal{H}$) to generate surrogate indices. Here, we present our model specifications for different imputation methods:

- *(Separate Imputation)* We constructed two linear regressions to predict the observed last purchase time and average purchase rate per active period:

$$\begin{aligned}\mathcal{T}_{i,10} = \alpha_0^{\mathcal{T}} + \beta_1^{\mathcal{T}} X_{i,1} + \beta_2^{\mathcal{T}} X_{i,2} + \beta_3^{\mathcal{T}} X_{i,1} \cdot X_{i,2} + \\ \sum_{t=1}^{T_0} \left( \gamma_t^{\mathcal{T}} \cdot S_{i,t} + \xi_t^{\mathcal{T}} \cdot X_{i,1} \cdot S_{i,t} + \eta_t^{\mathcal{T}} \cdot X_{i,1} \cdot S_{i,t} \right) + \varepsilon_i^{\mathcal{T}}, \\ \Lambda_{i,10} = \alpha_0^{\Lambda} + \beta_1^{\Lambda} X_{i,1} + \beta_2^{\Lambda} X_{i,2} + \beta_3^{\Lambda} X_{i,1} \cdot X_{i,2} + \\ \sum_{t=1}^{T_0} \left( \gamma_t^{\Lambda} \cdot S_{i,t} + \xi_t^{\Lambda} \cdot X_{i,1} \cdot S_{i,t} + \eta_t^{\Lambda} \cdot X_{i,2} \cdot S_{i,t} \right) + \varepsilon_i^{\Lambda},\end{aligned}$$

  where $\mathcal{T}_{i,10}$ is the observed last transaction time and $\Lambda_{i,10}$ denotes the average per-period purchase counts until the observed last transaction.

- *(Single Imputation)* We fit the following linear regression to predict $Y_{i,T}$:

$$\begin{aligned}Y_{i,10} = \alpha_0^{Y} + \beta_1^{Y} X_{i,1} + \beta_2^{Y} X_{i,2} + \beta_3^{Y} X_{i,1} \cdot X_{i,2} + \\ \sum_{t=1}^{T_0} \left( \gamma_t^{Y} \cdot S_{i,t} + \xi_t^{Y} \cdot X_{i,1} \cdot S_{i,t} + \eta_t^{Y} \cdot X_{i,2} \cdot S_{i,t} \right) + \varepsilon_i^{Y}.\end{aligned}$$

- *(BG/NBD)* We employ a BG/NBD model with time-invariant covariates $X_{i,1}$ and $X_{i,2}$. Linear specifications were used for all key parameters. After generating expected future purchase counts after $T_0$, we add it with the observed purchase counts until $T_0$ to capture the short-term treatment effects.

### B.4. Sample Size Efficiency

We investigate the influence of training sample size on the efficiency of CATE model estimation. Figure App-1 displays the AUTOC values for various CATE models and training sample sizes. The results indicate that CATE models utilizing short-term proxies consistently outperform the default approach, regardless of the training sample size. Moreover, the separate imputation method persistently surpasses other methods in terms of AUTOC for the same CATE model and training sample size. These findings suggest that incorporating short-term outcomes is a viable strategy for enhancing targeting performance, as it enables more accurate CATE estimation without requiring substantially larger sample sizes.

**Figure App-1 Area-under-TOC Curves: CATE Models with Different Training Sample Size.**

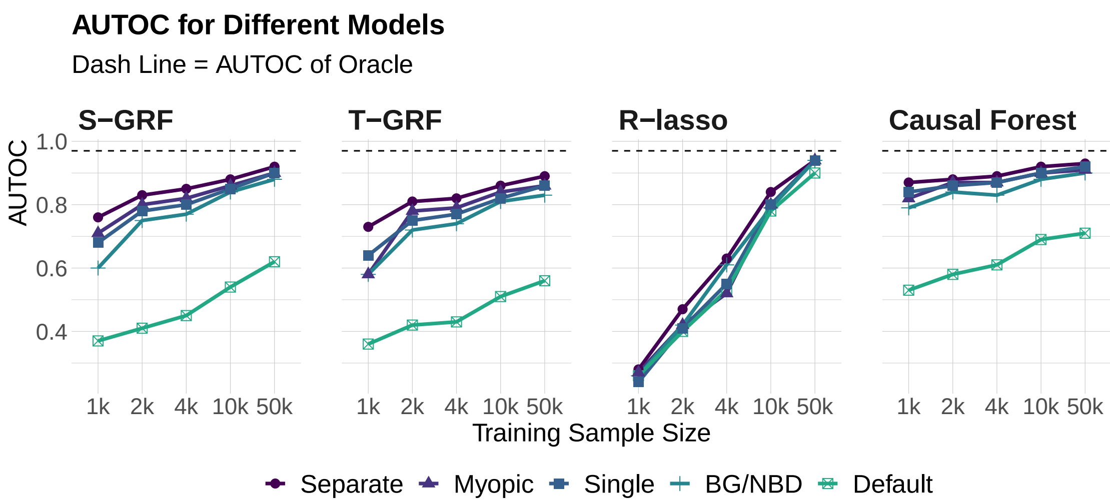


*Note.* Each point reports the average over 200 simulation replications. The larger the AUTOC, the better targeting performance.

### B.5. Trade-off between Information and Noise Accumulation

We reproduce the analyses described in Section 5.5 for different approaches and CATE models. Figure App-2 shows the results of AUTOCs for various CATE models and the number of periods used for surrogacy construction. Our results indicate that (i) using short-term proxies consistently leads to higher AUTOC compared to using the actual long-term outcome (except for R-lasso), (ii) the separate imputation method outperforms other short-term proxies regardless of the CATE models used for estimation, (iii) the separate imputation method shows higher robustness to noise, as the AUTOC declines more slowly, and (iv) using $T_0 = 3$ (i.e., the minimal number of periods such that the surrogacy assumption is satisfied) generally yields the most effective targeting performance.

**Figure App-2** **Area-under-TOC Curve: CATE Models for Outcomes Using Different Periods of Information.**

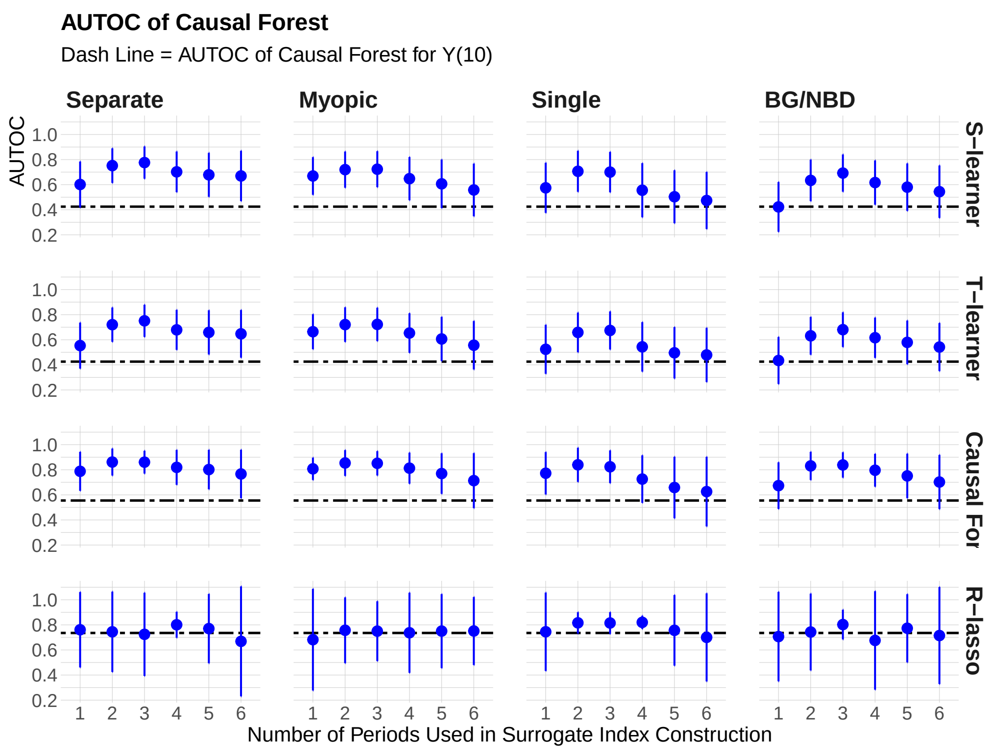


*Note.* Each point reports the average over 200 simulation replications together with the two standard error interval. The larger the AUTOC, the better targeting performance.

# C. Further Details for the Empirical Application

In this section, we provide additional analyses for the empirical application presented in Section 6 of the main document.

## C.1. Randomization Checks

Table App-1 compares treated and non-treated customers on a set of pre-treatment covariates — variables easily observed by the focal company, which are used for CATE estimation and targeting. We standardized all variables to preserve the company's confidentiality. Table App-1 suggests that the randomization was correctly executed as there is no statistically significant difference detected across the two groups of customers.

**Table App-1** **Randomization Check: Comparison of Pre-treatment Covariates under Two Experimental Conditions.**

| Variable | Three Coupons ($N = 889$) | One Coupon ($N = 964$) | Difference $p$-value |
|---|---|---|---|
| log(Sales) in the first transaction | −0.021 | 0.018 | 0.601 |
| log(Quantity) in the first transaction | 0.018 | 0.040 | 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 main-dish 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 |

All continuous variables were first standardized then summarized across conditions. 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 tree-based models. However, there is no significant differences for the two variables in the original scale as well.

## C.2. Specification of CATE Models

Given an outcome variable $\ddot{Y}_i$, we construct four types of CATE models to show robustness of our findings, including:

1. *S-learner*: we predict $\mathbb{E}[\ddot{Y}_i|W_i, \mathbf{X}_i]$ by regressing $\ddot{Y}_i$ on $W_i, \mathbf{X}_i$ using random forest and perform the automatic hyperparameter tuning using the method implemented in the `grf` package.
2. *T-learner*: we construct two random forests of $\ddot{Y}_i$ on $\mathbf{X}_i$, one for treated customers and another for non-treated customers. We perform the automatic hyperparameter tuning using the method implemented in the `grf` package.

3. *X-learner* (Künzel et al. 2019): all the outcome models in X-learner are estimated using the random forest with automatic hyperparameter tuning. We estimate the propensity score using the probability forest implemented in the `grf` package.
4. *Causal Forest* (Wager and Athey 2018): we use the causal forest function implemented in the `grf` package with automatic hyperparameter tuning.

### C.3. Replication results for the motivating example

In this appendix, we reproduce Figure 1a and Figure 1b in the main document using the CATE models outlined in Appendix C.2.

**C.3.1. Targeting for Short-term Outcome.** In the motivating example (Section 2 of the main document), we show that the focal company can develop an effective CATE model when the outcome variable is $Y_{i,1}$. Figure App-3 presents the Group Average Treatment Effects (GATEs) on $Y_{i,1}$ across quintiles based on predicted CATE. This chart is constructed similarly to Figure 1a, now presenting the results for the different CATE models. While the actual CATE curves are not perfectly decreasing for all models, they are still effective in distinguishing customers with high CATEs from those with low CATEs, as $\mathcal{Q}_1^{\widehat{\tau}_{Y_1}}, \mathcal{Q}_2^{\widehat{\tau}_{Y_1}}, \mathcal{Q}_3^{\widehat{\tau}_{Y_1}}$ have higher treatment effects than $\mathcal{Q}_4^{\widehat{\tau}_{Y_1}}$ and $\mathcal{Q}_5^{\widehat{\tau}_{Y_1}}$ have.

**Figure App-3 CATE Models for $Y_{i,1}$.**

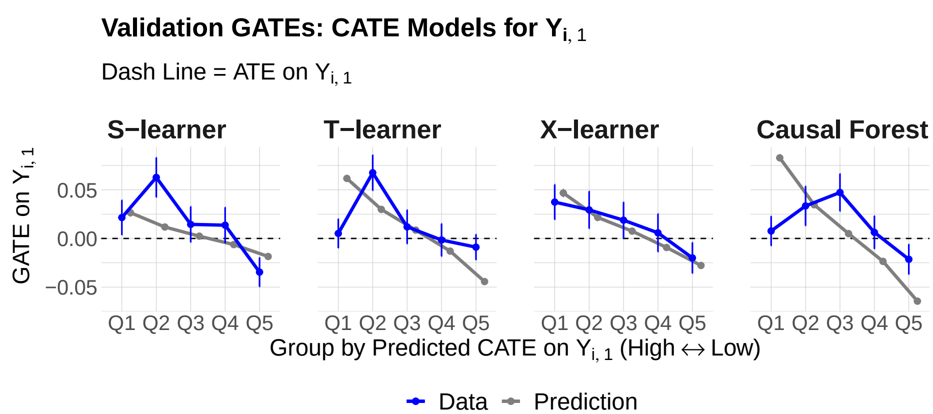


*Note.* Each point represents the mean of bootstrap results. Groups $\mathcal{Q}_1^{\widehat{\tau}_{Y_1}}, \cdots, \mathcal{Q}_5^{\widehat{\tau}_{Y_1}}$ are defined by decreasing order of predicted treatment effects on $Y_{i,1}$. Hence, the predicted GATEs (the gray line) are, by definition, monotonically decreasing. Actual GATEs (the blue line) are computed from the actual response.

**C.3.2. CATE Models for Long-term Outcome.** Similarly, Figure App-4 illustrates the predicted and actual GATEs when using $Y_{i,10}$ as outcome variable. Notably, all CATE models produce the same V-shaped curve, indicating that the firm would overlook a significant proportion (e.g., the bottom quintile group $\mathcal{Q}_5^{\widehat{\tau}Y10}$) of the "should-target" customers when the targeting policy is based on these models.

**Figure App-4 Targeting for $Y_{i,10}$.**

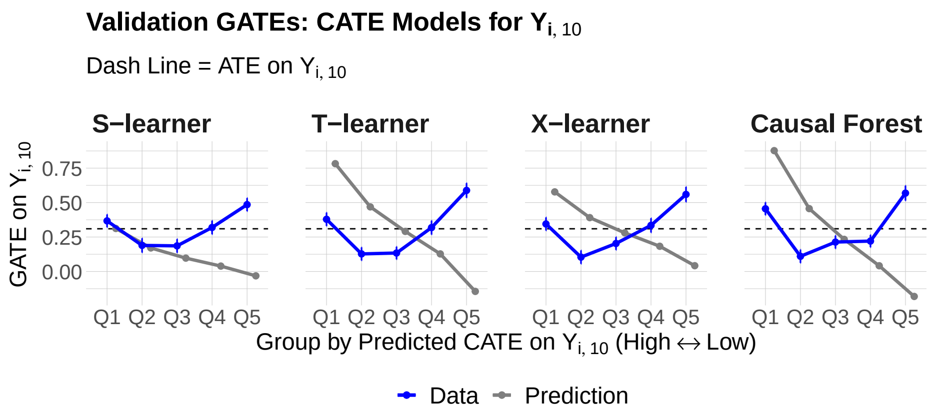


*Note.* Each point represents the mean of bootstrap results. Groups $\mathcal{Q}_1^{\widehat{\tau}_{Y10}}, \cdots, \mathcal{Q}_5^{\widehat{\tau}_{Y10}}$ are defined by decreasing order of predicted treatment effects on $Y_{i,10}$. Hence, the predicted GATEs (the gray line) are, by definition, monotonically decreasing. Actual GATEs (the blue line) are computed from the actual response.

## C.4. Model Specification of Surrogate Indices

To construct the surrogate indices, we gathered historical data of customers who were acquired at least ten weeks before the experiment started (4,031 in total) and imputed different outcome variables as follows:

- $\widetilde{Y}_{10}^{\text{Single}}(S_{i,1}, \mathbf{X}_i)$: we fit a random forest model (Athey et al. 2019) of $Y_{i,10}$ on the first-week purchase ($S_{i,1}$) and customer covariates ($\mathbf{X}_i$), reported in Appendix C.1. We perform automatic parameter tuning using the function provided by the `grf` package.
- $\widetilde{Y}_{10}^{\text{Sep}}(S_{i,1}, \mathbf{X}_i)$: we fit two random forests, one for the observed last purchase week ($\mathcal{T}_{i,10}$) and another for the purchase counts per period until the last purchase week ($\Lambda_{i,10} \equiv Y_{i,10}/\mathcal{T}_{i,10}$), on $S_{i,1}$ and $\mathbf{X}_i$. We perform automatic parameter tuning using the function provided by the `grf` package.
- $\widetilde{Y}_{10}^{\text{BTYD}}(S_{i,1}, \mathbf{X}_i)$: we fit a BG/NBD model with $\mathbf{X}_i$ as time-invariant covariates.

## C.5. Forecasting Accuracy of Surrogate Indices

Table App-2 reports the Pearson's correlation coefficients between the surrogate indices and the actual long-term outcome. Note that the correlation increases as $T_0$ increases, but the targeting performance becomes worse (as shown in Appendix C.8) since $\mathbf{S}_{T_0}$ accumulates more unexplained variations as $T_0$ increases.

**Table App-2** **Pearson's Correlation Between Short-term Proxies and Long-term Outcome.**

| Outcome Variable | $T_0 = 1$ | $T_0 = 2$ | $T_0 = 3$ | $T_0 = 4$ |
|---|---|---|---|---|
| $\widetilde{Y}_{10}^{\text{Sep}}(S_{i,T_0}, \mathbf{X}_i)$ | 0.4943 | 0.6714 | 0.7740 | 0.8521 |
| $\widetilde{Y}_{10}^{\text{Single}}(S_{i,T_0}, \mathbf{X}_i)$ | 0.4920 | 0.6687 | 0.7665 | 0.8579 |
| $\widetilde{Y}_{10}^{\text{BTYD}}(S_{i,T_0}, \mathbf{X}_i)$ | 0.4455 | 0.6097 | 0.7526 | 0.8419 |
| $Y_{i,T_0}$ | 0.5075 | 0.7276 | 0.8241 | 0.8860 |

## C.6. Details for Policy Learning Using Doubly Robust Scores

In this section, we provide a detailed explanation of how we implement doubly robust policy learning, as proposed by Athey and Wager (2021). Specifically, for each training-validation split, we learn the policy by the following steps:

1. Compute the outcome variable $\ddot{Y}$ for the training set ($\mathcal{D}_{\text{train}}$).
2. Compute the (honest) doubly robust score for $i$ in the training set:

$$\widehat{\Gamma}_i = \left[\widehat{m}^1(\mathbf{X}_i) - \widehat{m}^0(\mathbf{X}_i)\right] + \frac{W_i - \widehat{e}(\mathbf{X}_i)}{\widehat{e}(\mathbf{X}_i)} \left[Y_i - \widehat{m}^{W_i}(\mathbf{X}_i)\right].$$

   We use `grf` and `policytree` packages (Sverdrup et al. 2020) to derive the doubly robust scores for each customer.

3. Derive the targeting policy $\widehat{\pi}$ by solving the optimization problem:

$$\widehat{\pi} = \arg\max_{\pi} \sum_{i \in \mathcal{D}_{\text{train}}} \left[2\pi(\mathbf{X}_i) - 1\right] \widehat{\Gamma}_i,$$

   where we constrain $\pi$ in the class of probability forest in the `grf` package.

### C.7. Replication Results for the GATE Analysis

Figure App-5 displays the GATEs by predicted CATE levels of different outcome variables. Notably, regardless of the methods used, the separate imputation method consistently produces the steepest CATE curve, suggesting the superiority of targeting based on our proposed solution. Note that both Single and Myopic approaches also result in reasonably good performance when compared to models based on $Y_{i,10}$.

**Figure App-5** **Actual group average treatment effects by predicted CATE levels on different outcome variables.**

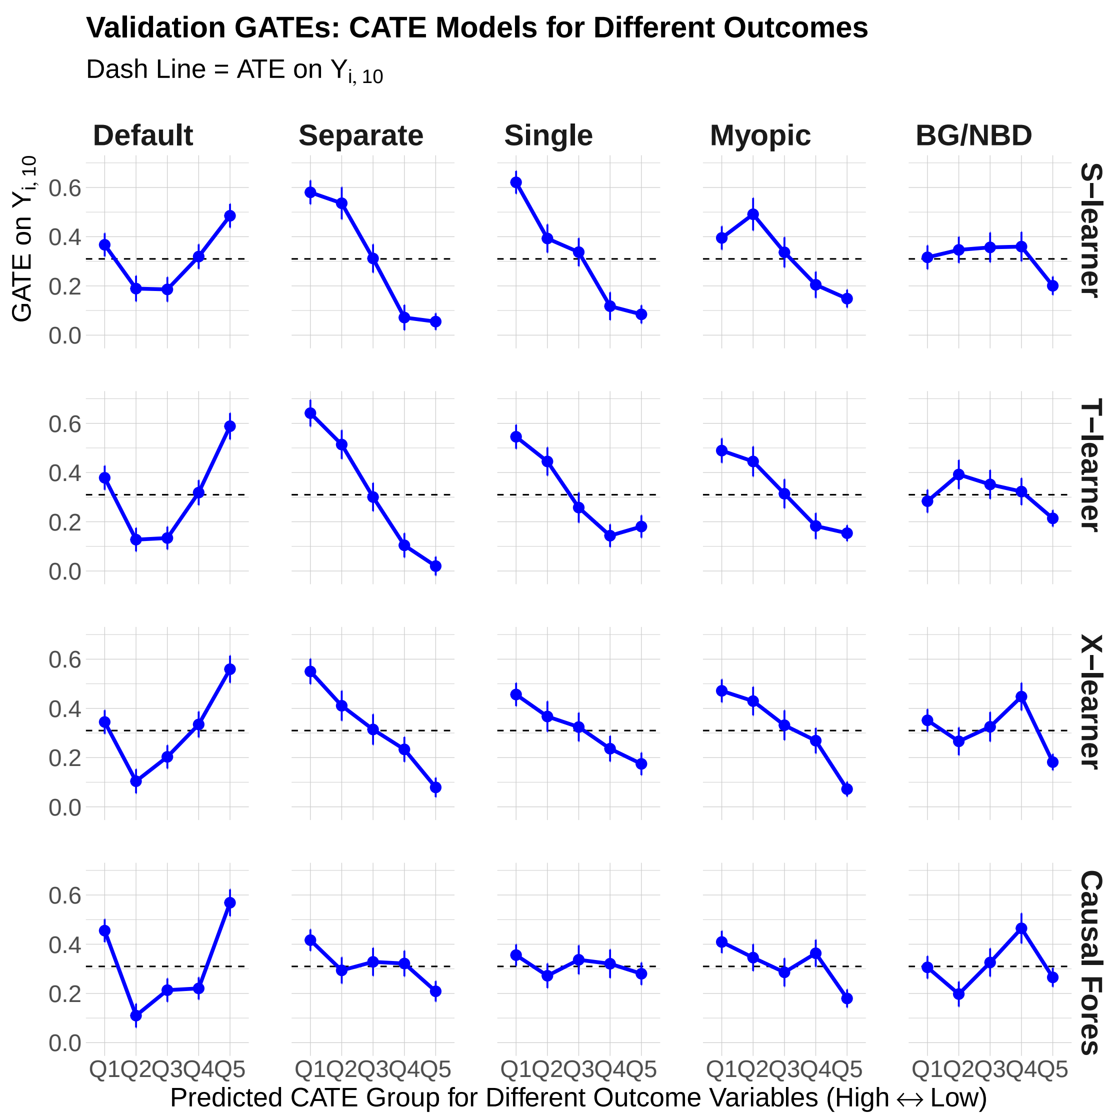


*Note.* Each point represents the mean of bootstrap results on the validation customers together with the two standard error interval (the errorbar). Groups $\mathcal{Q}_1^{\widehat{\tau}_R}, \cdots, \mathcal{Q}_5^{\widehat{\tau}_R}$ are defined by decreasing order of treatment effect, as predicted by the CATE model for different outcome variables. GATEs are computed from $Y_{i,10}$.

### C.8. How Many Periods Should the Focal Firm Use?

Figure App-6 compares the expected profit improvement from targeting based on surrogate indices constructed using different periods of outcomes. The result suggests that using one-period outcome in surrogate models gives the highest profit, and targeting based on short-term signals consistently outperforms or is as good as targeting based on the actual long-term outcome.

**Figure App-6  Expected Profit Improvement: Comparison of Different Periods Used for Surrogate Models.**

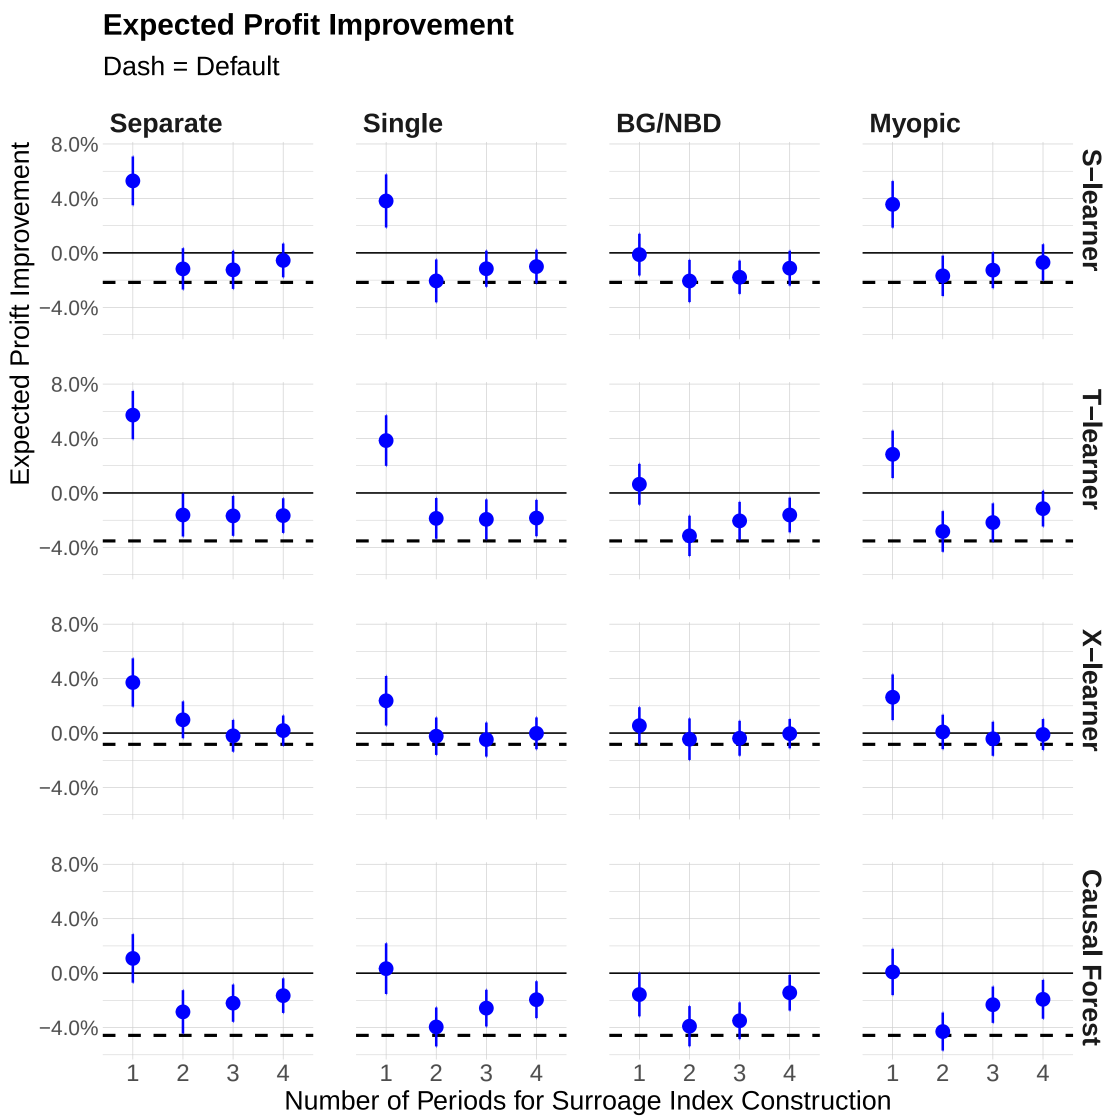


*Note.* Each point represents the mean of bootstrap results on the validation customers together with the two standard error interval (the errorbar). The dashed line reports the expected profit improvement when the targeting policy is based on predicted CATEs on $Y_{i,10}$.

## References

Athey S, Tibshirani J, Wager S (2019) Generalized random forests. *The Annals of Statistics* 47(2):1148–1178.

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

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

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

Sverdrup E, Kanodia A, Zhou Z, Athey S, Wager S (2020) policytree: Policy learning via doubly robust empirical welfare maximization over trees. *Journal of Open Source Software* 5(50):2232.

Wager S, Athey S (2018) Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association* 113(523):1228–1242.