

Haupt, Johannes; Lessmann, Stefan

Working Paper

Targeting Cutsomers Under Response-Dependent Costs

IRTG 1792 Discussion Paper, No. 2020-005

Provided in Cooperation with:

Humboldt University Berlin, International Research Training Group 1792 "High Dimensional Nonstationary Time Series"

Suggested Citation: Haupt, Johannes; Lessmann, Stefan (2020) : Targeting Cutsomers Under Response-Dependent Costs, IRTG 1792 Discussion Paper, No. 2020-005, Humboldt-Universität zu Berlin, International Research Training Group 1792 "High Dimensional Nonstationary Time Series", Berlin

This Version is available at:

<https://hdl.handle.net/10419/230811>

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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



Targeting Customers Under Response-Dependent Costs

Johannes Haupt^{*}
Stefan Lessmann^{*}



^{*} Humboldt-Universität zu Berlin, Germany

This research was supported by the Deutsche Forschungsgesellschaft through the International Research Training Group 1792 "High Dimensional Nonstationary Time Series".

<http://irtg1792.hu-berlin.de>
ISSN 2568-5619

TARGETING CUSTOMERS UNDER RESPONSE-DEPENDENT COSTS

PREPRINT

Johannes Haupt

School of Business and Economics
Humboldt-Universität zu Berlin johannes.haupt@hu-berlin.de

Stefan Lessmann

School of Business and Economics
Humboldt-Universität zu Berlin stefan.lessmann@hu-berlin.de

March 25, 2020

ABSTRACT

This study provides a formal analysis of the customer targeting decision problem in settings where the cost for marketing action is stochastic and proposes a framework to efficiently estimate the decision variables for campaign profit optimization. Targeting a customer is profitable if the positive impact of the marketing treatment on the customer and the associated profit to the company is higher than the cost of the treatment. While there is a growing literature on developing causal or uplift models to identify the customers who are impacted most strongly by the marketing action, no research has investigated optimal targeting when the costs of the action are uncertain at the time of the targeting decision. Because marketing incentives are routinely conditioned on a positive response by the customer, e.g. a purchase or contract renewal, stochastic costs are ubiquitous in direct marketing and customer retention campaigns.

This study makes two contributions to the literature, which are evaluated on a coupon targeting campaign in an e-commerce setting. First, the authors formally analyze the targeting decision problem under response-dependent costs. Profit-optimal targeting requires an estimate of the treatment effect on the customer and an estimate of the customer response probability under treatment. The empirical results demonstrate that the consideration of treatment cost substantially increases campaign profit when used for customer targeting in combination with the estimation of the average or customer-level treatment effect. Second, the authors propose a framework to jointly estimate the treatment effect and the response probability combining methods for causal inference with a hurdle mixture model. The proposed causal hurdle model achieves competitive campaign profit while streamlining model building.

The code for the empirical analysis is available on Github.

Keywords Heterogeneous Treatment Effect, Uplift Modeling, Coupon Targeting, Churn/Retention, Campaign Profit

1 Introduction

Data-driven prediction of customer behavior and the automation of campaign targeting are at the core of modern direct marketing [1]. Direct marketing plays a key role in consumer markets with the continuous growth of e-commerce, at 1.8 trillion Euros globally in 2019 [2] as the growth of e-commerce is accompanied by a growth in online and email advertising and in traditional print advertising, e.g. catalog marketing [3]. To make advertising profitable, businesses have shifted away from blanket advertising and select which prospective customers to target. Targeting a customer is profitable if the positive impact of the marketing treatment on the customer and the resulting profit to the company is higher than the cost of the treatment.

Predicting the expected profit requires the estimation of the change in customer behavior if the customer is targeted, known as conditional average treatment effect (CATE) [4]. Estimation of the CATE has been the focus of work under the label of uplift modeling [5] and has received much attention in recent work in statistics [6, 7] and machine learning [8] with the result that heterogeneous response to marketing treatment can be predicted more precisely.

However, profitable targeting must consider the effect of treating a customer in relation to the cost of treatment. Prior research tends to neglect application-specific profit and cost as decision variables and instead assume an external restriction on the number of customers to target [9, 10]. While there exists work that explicitly develops targeting policies that optimize the profit of the marketing campaign [11], these policies are restricted to settings in which the cost of the treatment is known at the time of the targeting decision, e.g. the production and shipping of a catalog.

Many applications in direct marketing include costs that are uncertain at the time of the targeting decision because they are realized only when the customer accepts the marketing offer. These response-dependent costs are present whenever a marketing incentive is conditional on a profitable customer action. Companies use conditional incentives regularly in the form of discounts and the most salient applications have attracted much research, e.g. customer retention [12, 13] or coupon targeting [14, 5]. Because the treatment cost is conditional on the customer action, the uncertainty about the customer action translates into uncertainty about the realization of the cost of the incentive. The targeting decision must then be based on comparing the expected profit to the expected cost of the marketing treatment, which is now uncertain but can be estimated. In addition to the estimation of the CATE, estimation of the expected cost requires a model of the customer decision under treatment. Despite the prevalence of targeted discounts in the industry and the focus of research on customer retention and couponing, the literature has not analyzed the targeting decision problem under response-dependent costs and lacks suitable modeling strategies to estimate both the treatment effect and customer choice efficiently.

This paper makes two contributions to the literature. First, we formally analyze the targeting decision problem under customer response-dependent costs. We show that profit-optimal targeting requires an estimate of the expected change in profit in the form of the treatment effect and estimate of the customer response probability under treatment. Second, we propose a framework to jointly estimate the treatment effect on profit and the absolute response probability. The two proposed models combine methods for causal inference with a hurdle mixture model. We evaluate the effectiveness of our approach on a coupon targeting campaign in an e-commerce setting.

The paper is structured as follows. Section 2 summarizes the existing literature on the estimation of treatment effects in customer targeting and profit-based targeting policies that consider targeting costs. Section 3.1 formally analyzes the targeting decision under response-dependent costs. Section 3.2 introduces hurdle models within causal estimation frameworks as flexible models of treatment effect and customer response. Section 4 introduces the data and experimental design. The results of the experiment are evaluated in Section 5. Section 6 concludes.

2 Literature Review

Customer targeting subsumes research with the goal to identify which customers to target in order to maximize the profit of a marketing campaign. Research on customer targeting has been segmented into work on specific applications such as direct marketing and customer retention management. A starting point of our analysis is that direct marketing and customer churn are characterized by a shared decision problem, whose cost structure has implications for the design of targeting models. Direct marketing and churn management target specific customers with a marketing action through communication channels including website banners, email and print marketing, but differ in the goal of the marketing action. Direct marketing addresses customers to elicit a profitable customer response in the form of a purchase or request for a service. The existing research defines the customer response either as conversion, i.e. if the customer has completed a purchase in the period following the marketing action [15], or as spending, i.e. how much the customer spent following the marketing action [16]. Customer retention management addresses customers to avoid an unfavorable customer action and termination of the customer’s relationship with the company, commonly referred to as customer churn [17]. Positive customer action is defined either as retention, i.e. if the individual remains an active customer, or as customer lifetime value, i.e. the remaining net value of the customer to the company [12]. For our analysis, we refer to the customer action in both settings as customer *response*, which is positive in case a purchase takes place or a customer remains with the company, and to the spending or customer lifetime value as response *value*.

The fundamental decision criterion for customer targeting is the treatment effect due to the marketing action. The treatment effect is the expected change in behavior, measured on response or response value, that is caused by the marketing action. Recent studies on direct marketing and customer retention are careful to stress that the purpose of targeting is to identify the customer with the highest sensitivity to the marketing action [18, 10]. The earlier practice to base targeting decisions on the estimate of response probability favors the targeting of natural responders rather than customers who are impacted by the marketing treatment. [11] show that conversion models may be profitable

in practice when there exists a correlation between customers' natural propensity to respond and their sensitivity to the marketing treatment. As there is no theoretical reason to assume such a correlation, an estimate of the response probability is generally insufficient to determine a profitable targeting policy as we clarify in the formal analysis of the targeting decision problem.

Recent research has therefore focused on the estimation of the treatment effect based on observed customer characteristics, commonly referred to as conditional average treatment effect (CATE). The general applicability of methods for treatment effect estimation has led to developments spread across fields. An comprehensive overview over recent methodology is provided by the following studies and references therein: [4] on *uplift* estimation in information systems, [19] for medical application, [20] and [21] for a more statistical perspective, and [22] for settings with continuous or repeated experiments.

The decision whether to target a customer in the campaign depends on the treatment effect in relation to the cost of the marketing action. While research has focused on the estimation of the treatment effect, insufficient attention has been paid to the cost structures of customer targeting. We distinguish two types of variable costs depending on the type of marketing treatment and communication channel. Applying the marketing action to a customer may entail *targeting-dependent* variable costs that arise whenever the action is taken. In practice, targeting-dependent costs arise for communication with the customer in the form of mail charges or call center fees and for the production of material treatments like catalogs [16].

An important characteristic of targeting-dependent costs is that they arise when the targeting decision is made, independent of its success. This differentiates targeting-dependent from *response-dependent* variable costs, which are incurred only if the customer responds positively after receiving the marketing treatment. Response-dependent costs arise from the design of marketing offers that are conditioned to apply only with a positive customer response. In practice, these offers take the form of free shipping on a future purchase or a discount on an existing service contract [23]. The value of the offer can be fixed, as in the case of coupon codes for free shipping, or relative to the response value, as for discounts on a monthly subscription fee. In both cases, if the customer responds negatively, for example, by terminating the existing contract, then the offer entails no cost for the company.

Beyond variable costs related to the targeting of individual customers, the implementation of a marketing campaign entails *fixed costs* for the design of the marketing action and the development of the targeting policy. While the fixed costs of the campaign are an important strategic consideration, they do not affect the operational targeting decision for individual customers.

The existence of targeting-dependent and response-dependent costs must be taken into account when designing targeting policies to maximize the profit of campaigns in direct marketing and churn. Despite the relevance of targeting costs for the targeting decision, the literature provides little discussion of customer targeting as a policy problem. [24] provide an analytical discussion of profit optimization under exclusively targeting-dependent variable costs. Targeting a customer is then profitable when the incremental value of the marketing action is at least as high as its cost. The assumption of targeting-dependent costs is natural for print advertising and the decision rule is applied by [11] in the setting of catalog marketing, where the targeting cost is incurred by printing and sending a catalog. [23] formulate the campaign profit specific to customer retention campaigns including an estimate of the response value and response-dependent as well as targeting-dependent variable costs. We provide a comprehensive discussion of this formulation, its implicit assumptions and related issues and its relation to our results in Appendix A and summarize our findings here. The churn campaign profit formulation includes a targeting-dependent contact cost and a response-dependent cost of the incentive to the firm in case the offer is accepted, but makes two restrictive assumptions. It implies that treatment effects are strictly positive and assumes a constant probability for customers to accept the offer when treated [25]. Assuming the same response probability for customers who receive the treatment ignores the heterogeneous sensitivity of customers to the treatment and the effect of the treatment on the expected cost, resulting in non-optimal targeting. [26] relax the assumption of a constant response probability and discuss campaign profit from the uplift perspective, but focus on model evaluation rather than model estimation and uphold the assumption of a positive treatment effect. We add to the literature by providing a formal analysis of the general targeting decision problem, which considers variation in treatment effects over customers and guides model estimation under target-dependent and response-dependent costs.

As an alternative to a decision-theoretic approach for expected profit maximization, the literature has suggested the empirical optimization of the targeting policy [27]. A popular approach towards empirical campaign optimization is to determine a threshold for the predicted treatment effect above which customers are targeted. Prior studies heuristically select the threshold that would have targeted the k deciles of the sample with the highest estimated CATE [24, 28, 10, 5]. The optimal proportion of the population to target can be approximated by comparing the group-wise average treatment effect for customers within each decile of the CATE estimates, since a correct ranking of customers by their CATE implies that the average treatment effect in groups with high model estimates must be higher than in groups with low model estimates. The evaluation of the model's ability to rank customers by their expected treatment effect is in line

with industry practice to target a small group of the most profitable prospective customers, but ignores the cost of targeting to determine the size of the campaign. An advantage of the empirical approach is that it remains feasible when the CATE estimates are a biased or badly calibrated estimate of the ITE or when the profit and costs parameters of the campaign are unknown. When there exists heterogeneity in response value or costs, ranking the customer by their expected treatment effect ignores variation in expected profit that is not due to variation in sensitivity to the treatment. Note that response-dependent costs imply variation in expected cost even when the nominal cost of the treatment is constant. Under profit or cost heterogeneity, empirical thresholding of the treatment effect will not result in an optimal targeting policy, as we show in the empirical analysis.

In summary, we find that customer targeting in applications including direct marketing and customer churn requires the consideration of the treatment effect and variable targeting costs. Targeting costs take the form of targeting-dependent costs and response-dependent costs, which are realized if the customer responds positively to the treatment. The next section provides an analysis of the customer targeting problem in settings that include customer-level heterogeneity in variable costs.

3 Methodology

3.1 Optimal Decision Making in Customer Targeting

The customer targeting decision problem is characterized by three components, 1) the value to the marketer conditional on the customer response, 2) the treatment cost conditional on the targeting decision and 3) the treatment cost conditional on the customer response. The existence of response-dependent costs differentiates most retention and coupon campaign settings from the cost setting discussed in previous studies [24], which assumes that all cost components are conditional on the targeting decision, but independent of the customer response.

Let $C_i \in \{0, 1\}$ be a random variable indicating an action by customer i , who is described by a set of observed covariates X_i . We define $C_i = 1$ as an event with a positive impact on business profit, for example, a purchase by the customer for couponing or customer retention in churn modeling. Further, let $V_i \in \mathbb{R}^+$ be the gross profit before targeting costs that is associated with a positive customer action. V_i represents the customer lifetime value in churn prevention or the margin of a purchase in direct marketing and may show substantial variation across customers. For convenience, let $Y_i = C_i \cdot V_i$ be the observed profit of the targeting decision, excluding targeting cost. Note that $Y_i = V_i$ when $C_i = 1$ and $Y_i = 0$ otherwise. The probability of a positive response $p(C = 1|X_i)$ and the expected response value $E[V|X_i]$ are unknown at the time of the marketing decision and need to be estimated given the customer characteristics.

Recall that the variable costs split into two components, the targeting-dependent and response-dependent costs. Let c be a targeting-dependent cost that is constant and independent of the customer characteristics. Targeting-dependent costs can be contact costs, for example, mail charges. Let δ be a response-dependent cost that applies if the customer responds positively after receiving the marketing treatment. The response-dependent cost can be associated with a marketing incentive that is conditioned on a positive customer response, for example, a voucher for free shipping for the current purchase process. The expected response-dependent cost at the time of targeting depends on the probability that the customer will accept the offer. Besides, response-dependent costs may depend on the value of the response. When the marketing treatment is a relative discount, for example, in the form of 10% discount on the current purchase, the nominal discount depends on the completion of the purchase and the purchase amount. The expected offer cost then depends on the probability of a positive customer response and the value of the response. If a customer is not targeted by the campaign then no variable costs occur and $\delta = c = 0$.

Table 1 summarizes decision problems in target marketing by outlining their respective cost structure and anticipates the results of the decision analysis. The decision problems vary in the existence of the treatment- and response-dependent costs, the type of response-dependent incentive and assumptions about the treatment effect on response probability and value. We see that targeting-dependent costs apply to one stream of research with applications in catalog marketing [11] and online banner advertising [29]. The proposed decision framework applies under any combination of variable costs and is crucial whenever there are response-dependent costs. We further differentiate the response-dependent costs into offers with a fixed value, e.g. retention campaigns with a discount upon contract renewal [26], and offers with a value equal to a percentage of the response value, e.g. coupon banners in a webshop [5]. The decision analysis determines the decision variables indicated in the last column. These are the variables required to calculate the expected profit of the targeting decision in the specific setting as a result of our analysis. Note that the set of decision variables may simplify when assuming no treatment effect on the value given conversion for the first and last setting. We will discuss this assumption as a special case below.

Consider an available marketing treatment and let T_i be a variable to indicate if the treatment was applied to customer i . We consider a single treatment and assume $T_i \in \{0, 1\}$, where $T = 1$ indicates that the customer is targeted, the treatment

Table 1: Decision problems in customer targeting and their decision variables

Application Example	Cost			Treatment Effect on		
	Treatment-Depend.	Response-Depend.				
		Fixed	Percentage	Decision	Value	Decision Variables
<i>Advertisement</i>						
Letter and Present ¹	yes	no	no	yes	no	$p(1) - p(0), R$
Online Banner ²	no*	no	no	yes	no	$p(1) - p(0)$
Catalog ³	yes	no	no	yes	yes	$Y(1) - Y(0)$
Online Banner	no*	no	no	yes	yes	$Y(1) - Y(0)$
<i>Discount</i>						
Print Retention Offer ⁴	yes	yes	no	yes	yes	$Y(1) - Y(0), p(1)$
Online Fixed Value	no*	yes	no	yes	yes	$Y(1) - Y(0), p(1)$
Print Discount	yes	no	yes	yes	yes	$Y(1) - Y(0), p(1), R(1)$
Coupon Banner ⁵	no*	no	yes	yes	yes	$Y(1) - Y(0), p(1), R(1)$
Coupon Banner	no*	no	yes	yes	no	$p(1), p(0)$

*We consider online marketing on the company's own website or in the form of email newsletters. Programmatic advertising on third party websites has a complex cost structure due to the underlying auction process.

¹[10] ²[29] ³[11]

⁴[26] ⁵[5]

condition, and $T = 0$ indicates that she is not, the control condition. The following analysis is easily extended to more than one treatment by considering multiple binary comparisons. The treatment is designed to increase the conversion probability of the customer or her value given conversion or both. Following the Neyman-Rubin potential outcome model, we indicate the potential outcomes under treatment using $\cdot(0)$ and $\cdot(1)$. For example, $C_i(1)$ denotes the conversion outcome if customer i is targeted, whereas $C_i(0)$ denotes the conversion outcome if she is not targeted. The individual treatment effect (ITE) on profit is then $\tau_i = Y_i(1) - Y_i(0) = C_i(1)V_i(1) - C_i(0)V_i(0)$. We further distinguish between the ITE on response probability $\tau_i^C = C_i(1) - C_i(0)$ and the ITE on response value $\tau_i^V = V_i(1) - V_i(0)$.

We now begin our analysis of the targeting decision problem. The profit π_i for an individual in the marketing campaign including treatment costs is

$$\pi_i = \begin{cases} C_i(0)V_i(0) & \text{if } T_i = 0 \\ C_i(1)V_i(1) - C_i(1)\delta - c & \text{if } T_i = 1 \end{cases}$$

The general decision problem whether to target a specific customer under response-dependent costs can then be posed as

$$p_i(1)(V_i(1) - \delta) - c > p_i(0) \cdot V_i(0), \quad (1)$$

where we use p_i as a convenient shorthand for $p(C = 1|X = x_i)$. Note how the variable costs affect the campaign profit. The target-dependent costs c are realized before the customer makes any decision and are therefore independent of the customer action. The response-dependent costs δ are realized only when a positive response takes place.

Solving the inequality for the treatment effect yields

$$p_i(1)V_i(1) - p_i(0)V_i(0) > p_i(1)\delta + c \quad (2)$$

The optimal decision naturally depends on the individual treatment effect on the profit on the left side of the equation. However, it also depends on the probability of a positive customer response under treatment as a mitigating factor on the offer cost. Intuitively, the absolute offer costs are a promise from the firm and must be discounted by the chance that the promise will in fact be redeemed by the customer. If the customer does not redeem the offer, then the response-dependent costs are not incurred by the company. The customer targeting decision under response-dependent costs thus differs from the case where $\delta = 0$ because the costs are now stochastic rather than known at the point of the targeting decision.

The optimization of expected profit underlying Eq. 2 implies that, when faced with two customers with an identical CATE, it is more profitable to target the customer who is less likely to respond positively and accept the marketing

offer. The previous practice to target customers with a high response probability after treatment not only disregards the causal effect of the treatment, as previous literature has pointed out [10], but increases the cost of campaigns by targeting customer with high expected response-dependent cost. To clarify the intuition behind this result, consider the treatment of a customer as an investment with probabilistic cost. If the payout of two investments is identical, a rational agent prefers the investment that has lower expected cost. This result suggests that when there is little or no treatment heterogeneity, meaning that the payout of the treatment is identical between customers, it is profitable to target customers with a lower rather than higher probability to respond.

In practical terms, any decision setting with response-dependent costs will require an estimate of the treatment effect $p_i(1)V_i(1) - p_i(0)V_i(0)$ and an estimate of the response probability $p_i(1)$. This result is surprising because previous literature has emphasized uplift models, which provide an estimate of the treatment effect, as a direct replacement of response models, which provide an estimate of the conversion probability. The decision under response-dependent costs requires both a model of the treatment and a model of the conversion probability under treatment.

In application, a positive expected profit may not result in an optimal policy under strategic considerations. Actual targeting campaigns are regularly evaluated by their return on advertising spend (ROAS). The ROAS is defined as the ratio of campaign profit over campaign costs. Note that the same information is sometimes expressed by its inverse as the cost-revenue ratio. The ROAS is a metric of advertising efficiency and as such does not consider campaign size. While it is generally not profit-optimal to maximize efficiency at the cost of targeting fewer customers, a minimum ROAS is often required in practice to satisfy management goals and allocate resources efficiently between marketing channels or campaigns. A side result of our analysis is that the proposed decision rule can be used to set targeting thresholds to reflect a minimum ROAS as

$$\frac{p_i(1)V_i(1) - p_i(0)V_i(0)}{p_i(1) \cdot \delta + c} \geq \text{Target ROAS}$$

We go on to discuss two special cases that arise in digital applications.

First, assume that $c = 0$. In digital marketing settings, there are no variable contact costs if customer communication is digital and automated. In particular, the costs for email targeting and banner campaigns on company's own websites arise in the form of fixed cost into infrastructure, e.g. content management systems and content production. These costs are irrelevant for operational targeting decisions in the short run. The targeting rule is then

$$\begin{aligned} p_i(1)V_i(1) - p_i(0)V_i(0) &> p_i(1) \cdot \delta \\ p_i(1)(V_i(1) - \delta) &> p_i(0)V_i(0) \end{aligned} \quad (3)$$

Assume additionally that offer costs depend linearly on the response value, i.e. $\delta_i = \eta V_i(1)$. The latter assumption corresponds to discount coupons that reduce the checkout amount by a fixed percentage, e.g. 10%, and other forms of dynamic pricing. Percentage discount coupons are frequently used in online marketing as a transparent means to differentiate incentives according to the value of customers and as an incentive that encourages higher spending. The decision rule for discount offers requires an estimate of the expected response value under treatment:

$$\begin{aligned} p_i(1)V_i(1) - p_i(0)V_i(0) &> p_i(1) \cdot \delta_i \\ p_i(1)V_i(1) - p_i(0)V_i(0) &> p_i(1) \cdot \eta \cdot V_i(1) \end{aligned} \quad (4)$$

Second, there exists a special case of the decision problem in Eq. 4 that requires no estimate of the purchase value. Assume that the treatment affects the conversion probability but not the response value, i.e. $V(1) = V(0)$. Then equation 4 reduces to

$$\begin{aligned} (p_i(1) - p_i(0)) \cdot V_i &> p_i(1) \cdot \eta \cdot V_i \\ p_i(1) &> \frac{p_i(0)}{1 - \eta} \end{aligned} \quad (5)$$

Note that under the combined assumptions of a percentage discount with no fixed contact cost and no effect on conversion value, the decision rule becomes independent of the individual purchase value. Intuitively, a negligible communication cost removes the need to make up for the cost of customer targeting. Further making the coupon cost dependent on the response value automatically adjusts the cost to decrease with smaller response values and vice versa. In practice, this setting requires estimation of the purchase probabilities with and without treatment.

The two special cases imply that the cost structure, which is determined by the infrastructure of the campaign and the design of the treatment, can increase or reduce the complexity of the decision problem. In general, when the cost of the treatment is conditioned on additional variables, then the estimation of these variables is relevant for the decision

problem. We can see that percentage discounts on the purchase value introduce an estimate of the purchase value under treatment into the decision (Eq. 4). Similar arguments can be made for more specialized coupon design like a minimum purchase value or a staggered discount increasing with purchase value. The second case shows that specific cost structures may simplify the decision problem. Under the additional assumption of no treatment effect on value, the targeting decision reduces to the estimation of the probabilities of purchase with and without treatment in Eq. 5.

The proposed decision framework is a generalization of marketing decision settings discussed in the literature. Prior research in marketing has considered campaigns with treatment-related but no response-related costs, such as traditional mail catalog marketing [24, 11]. Assuming $\delta = 0$, we can show that the treatment effect on profit Y_i is sufficient for the targeting decision in these cases, which reduces to

$$p_i(1)V_i(1) - p_i(0)V_i(0) > c \quad (6)$$

We see immediately that an estimate of the treatment effect on the profit $p_i(1)V_i(1) - p_i(0)V_i(0) = Y(1) - Y(0)$ is a sufficient decision criterion under the conditions of Eq. 6. If we assume no treatment effect on the value such that $V_i(1) = V_i(0) = V_i$ and assume V_i to be known or modeled independently, we recover

$$p_i(1) - p_i(0) > \frac{c}{V_i}, \quad (7)$$

where the focus lies on the estimation of the treatment effect on the customer response. This recovers the estimation problem addressed by prior research under the label of uplift modeling, although the dependency on the response value is not typically discussed in the literature [4].¹

In summary, the treatment effect is not sufficient for profit-based targeting in settings with response-dependent costs. The additional decision variables required for the targeting decision depend on the cost-structure of the marketing treatment as given in Table 1. For variable costs with a fixed value, the purchase probability under treatment determines the cost as in Eq. 2. For treatment with a value relative to the purchase value, the purchase probability under treatment and the purchase value under treatment jointly determine the effective treatment cost as in Eq. 4. Both cost structures are common in direct marketing. The following sections discuss a model specification to estimate the cost-related decision variables $p(1)$ and $R(1)$ within the model of the treatment effect $Y(1) - Y(0)$.

3.2 Causal Hurdle Models

The targeting decision under response-dependent costs (Eq. 2) requires estimates of the treatment effect $Y(1) - Y(0)$, the response probability under treatment $p_i(1)$ and, for discount coupons, the response value under treatment $V_i(1)$. Alternatively, we can decompose the profit using $Y_i = C_i \cdot V_i$ and estimate the treatment effect as $p_i(1)V_i(1) - p_i(0)V_i(0)$. This formulation makes explicit that the additional decision variables are contained within the treatment effect on profit. The remainder of the study develops a framework to simplify the modeling task based on this observation.

A straight-forward approach to estimate the decision variables is to build several models, where one (causal) model estimates the CATE $\hat{\tau}_i$ and additional models to estimate remaining decision variables under treatment. In the following, we will call this the *distinct modeling* approach. The distinct modeling approach requires one model to estimate $p_i(1)$, a second model to estimate $V_i(1)$ in the case of discount coupons, and one to four models depending on the specific approach to estimate the CATE. In other words, the distinct modeling of each decision variable introduces up to two additional models to the CATE model.

In the following, we propose a framework to avoid additional model complexity and simultaneously estimate the treatment effect on expected profit, the purchase probability and purchase value. The proposed framework exploits the decomposition of the expected profit into the conversion probability $p_i(1)$ and purchase value $V_i(1)$ to collect estimates for $p_i(1)$ and $V_i(1)$ from the treatment effect model. We estimate $V_i(1)$ and $p_i(1)$ jointly within the profit model by modeling the observed profit from customer Y_i as a two-stage hurdle structure.

Hurdle models, known as Tobit II models in econometrics, are mixture models over two distributions, one of which has a point mass at zero, which were previously applied in applications of customer choice [30, 31]. They are convenient to model decisions that involve a binary decision on whether to act, the hurdle, and a conditional decision on the value associated with acting. Hurdle models assume that the occurrence of zeros is entirely driven by a first-stage process, i.e. the second stage value is zero when the first stage decision is a negative response and strictly positive when the first stage decision is a positive response.

¹See [10, fn. 27] for a brief mention of the issue in the case of customer churn.

The hurdle model allows us to decompose the estimation of the profit Y into the estimation of response C and response value R . The probability mass function of the hurdle model is

$$\Pr(Y_i = y|X_i = x) = \begin{cases} (1 - \Pr(C = 1|X_i = x)) \cdot 0 & \text{if } Y_i = 0 \\ \Pr(C = 1|X_i = x) \cdot \Pr(V_i = v|X_i = x) & \text{if } Y_i > 0 \end{cases} \quad (8)$$

where $\Pr(C_i = 1|X = x_i)$ is a model for customer response and $\Pr(V_i = v|X_i = x)$ is a model of response value. If a customer chooses to respond, they decide on their spending behavior in the second stage, which determines the response value to the firm. The profit from a response is zero if a customer chooses not to respond and strictly positive otherwise.

The hurdle model specification has two properties that are relevant in the context of customer choice. First, the separation of the purchase decision and value decision facilitates the estimation and interpretation of each model. In the context of treatment estimation, separating the effect on response probability and response value provides a more nuanced understanding of marketing effectiveness and can be used to improve the treatment. The model structure also accommodates differences in, for example, the relevance of available covariates for each decision step [30]. Second, the models for the prediction of response probability and response value can be estimated separately when the purchase incidence is observed and we assume independent error terms [32, p. 545]. This property will provide additional flexibility when estimating the proposed causal hurdle model in practice.

It remains to integrate the hurdle model into a framework for causal inference. Under the common assumptions of the potential outcome framework, i.e. unconfoundedness, overlap and stable unit treatment value, the CATE can be expressed as the difference between the outcome Y_i conditional on treatment assignment T_i and covariates X_i ,

$$\tau(x) = \mathbb{E}[Y(1) - Y(0)] = \mathbb{E}[Y|X_i = x, T = 1] - \mathbb{E}[Y|X_i = x, T = 0]. \quad (9)$$

We integrate the hurdle model into the standard treatment effect model by modeling profit with the hurdle model

$$\Pr(Y_i = y|T_i = t, X_i = x) = \begin{cases} 1 - \Pr(C = 1|X_i = x, T_i = t) & \text{if } Y_i = 0 \\ \Pr(C = 1|X_i = x, T_i = t) \cdot \Pr(V_i = v|X_i = x, T_i = t) & \text{if } Y_i > 0 \end{cases} \quad (10)$$

where both the conversion probability and the purchase value conditional on conversion depend on the treatment assignment T_i and the covariates X_i .

Following the definition of the hurdle model above and under the assumptions of the potential outcome framework, we specify our causal hurdle model as

$$\begin{aligned} \hat{\tau}(X_i) &= \hat{y}(X_i, 1) - \hat{y}(X_i, 0) \\ &= p(C = 1|X_i, T = 1) \cdot \mathbb{E}[R|X_i, T = 1] - p(C = 1|X_i, T = 0) \cdot \mathbb{E}[R|X_i, T = 0]. \end{aligned} \quad (11)$$

Estimating the treatment effect on response probability and response value separately has an additional advantage if we expect heterogeneity of effect direction and size on customer value and response probability. This is the case if individual customers react differently to the same offer, for example, purchase their basket with higher probability or put additional products into their basket in response to receiving the treatment. Further, we expect the treatment effect on probability and value to be closely connected to the design of the marketing action. Under strong heterogeneity, we expect some customers to react to the marketing treatment by increasing the response value, e.g. putting more products into their basket, while becoming more reluctant to respond at the higher value, e.g. abandon a high-value shopping basket. Explicit estimates of the disentangled treatment effects are then relevant for treatment selection and design.

The formulation in Eq. 10 does not restrict the specific method of causal inference. Figure 1 visualizes the general structure of causal hurdle models and makes the estimation targets explicit. We see that one strategy to estimate all relevant decision variables is to estimate four separate models, i.e. one model each for purchase probability and purchase value times one model each for the treatment and control group. This two-model hurdle model is equivalent to combining the two-model approach for CATE estimation with two hurdle models for which the choice and value components are estimated separately [32, p. 545].

It is possible to simplify the estimation by estimating more than one decision variable jointly. Eq. 9 is the starting point for two approaches to integrating a hurdle model structure into treatment effect models. Figure 1 visualizes the proposed methods to reduce the number of separate models by joint estimation of variables horizontally (solid red), over treatment and control group, or vertically (dashed blue), over purchase probability and value.

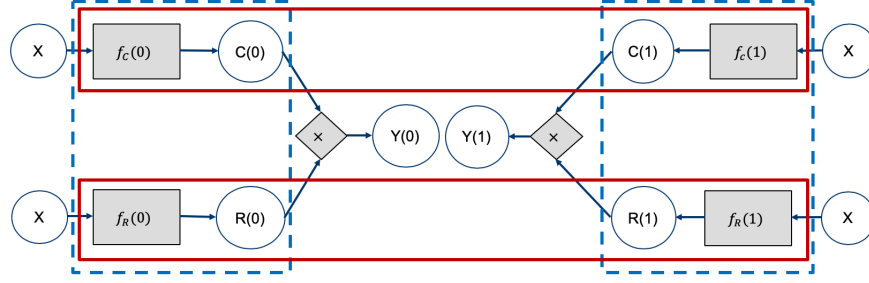


Figure 1: Causal hurdle model structure. Frames indicate the two proposed strategies for estimation in the form of two hurdle models (dashed blue) or two causal single models (solid red)

The *single-model hurdle model* combines the single model approach for causal inference with a two-stage estimation procedure for the hurdle model. A general model for the conditional profit with or without treatment takes the form $y = f(x, t)$ and predicts the return given the covariates X and treatment assignment T . Despite its simplicity, this *single model* approach has been found to provide competitive CATE estimates for sufficiently flexible specifications of $f(\cdot)$ [21]. The single two-stage approach estimates one model for the response probability and one model for the customer value each jointly over the control and treatment group (Figure 1, solid red). Following the single model approach, we include the treatment variable as a covariate into the model. By choosing a flexible parametrization $f(\cdot)$, we can model the conditional average treatment effect through the interaction between T_i and covariates X_i within the model [33].

4 Experimental Design

We evaluate the proposed methodology in an online couponing setting². The decision analysis summarized in Table 1 identifies online couponing as a particularly interesting decision problem because showing a coupon banner to the customer entails no targeting-dependent cost while the coupon value constitutes a substantial response-dependent cost. However, studies on online couponing are scarce in the literature and we are not aware of research considering its cost setting [14, 5].

A German online fashion retailer deploys an automated targeting system that can show website visitors of the online shop an offer for a discount on their purchase. Targeted customers receive a coupon code that provides a discount of € 10 at checkout. The code is made available to the customer through a banner on the webpage, which states the discount offer and displays a coupon code to be entered during the checkout process. The banner is shown repeatedly on subsequent page views within the same session to ensure that the customer is aware of the offer. The discount is subject to common terms and conditions that require a minimum checkout value of € 50 for the coupon to be usable. The operational question of the fashion retailer is to identify the customers whose incremental margin when being targeted is strictly larger than the expected cost due to the coupon. As the theoretical analysis shows, the profit-optimal targeting policy differs from similar target marketing settings discussed in previous studies as the expected cost of the coupon depends on the customer’s purchase probability.

The data contains information on 118,622 anonymized website visitors in the form of 50 variables collected through tracking software and the shop system. Variables include information on the user history, e.g. the number of previous visits, the behavior on the website, e.g. the number of clicks on the website, and the current shopping basket, e.g. the number of items and their total price [34]. 9% of website visitors convert and complete their purchase with a median purchase value of € 75. We remove 1,459 outliers with a substantially higher basket value between € 300 and € 1750 corresponding to the 2.5% percentile of the purchase value distribution.

The data fulfills the assumptions of the potential outcome framework. The unconfoundedness and overlap assumptions are met by design through randomizing treatment assignment, which is common practice in customer targeting applications. The stable unit treatment assumption value requires that social interaction effects between individuals are rare or small in size. This assumption has recently been challenged for the telecommunications industry [35]. The social network that customers form when communicating via the telecommunication network have been found to lead

²The code for the experiment and evaluation is available at <https://github.com/Humboldt-WI/response-dependent-costs>

to substantial positive spillover effects from the targeted customers to their connections. In online shopping, while there is potential for social effects, e.g. sharing information about the availability of coupons on the website, there are no social mechanisms inherent to the purchase process. We therefore assume that any potential social effects are insubstantial for our analysis, but encourage additional research on social effects in couponing applications.

Evaluation of approaches including the estimation of treatment effects is complicated by the fact that the true ITE is unobservable, so comparable studies rely on simulated data [6, 36]. To facilitate the evaluation of the proposed approach through a setting where the treatment effect is known, we conduct an empirical Monte Carlo study combining the observed covariates X and customer spending Y with a simulation of ITE [36]. We simulate the overall treatment effect $\tau(X)$ as a combination of the treatment effect on the conversion probability $\tau_C(X)$ and the purchase value conditional on a purchase $\tau_R(X)$. Each treatment effect is determined by a linear combination of covariates with coefficients drawn randomly following

$$\begin{aligned}\beta_C, \beta_R &\sim \mathcal{N}(0_k, I_k) \\ \tau_C(X) &= X_\tau^\top \beta_C \\ \tau_R(X) &= X_\tau^\top \beta_R\end{aligned}$$

where X_τ is a subset of $k = 11$ selected variables from the full set of variables. Both treatment effects are centered and scaled. To simulate a realistic marketing setting, we scale the ITE distribution to have most of its mass in the range $[0;10]$ and a positive average effect [11, Fig. 12]. For the ITE on response probability, we center the distribution around an ATE of 5 percentage points and truncate the simulated values to the range $[-0.1, 0.15]$. For the ITE on response value, we center the distribution around an ATE of € 1 and truncate the simulated values to the range $[-10, 10]$.

We simulate the potential outcome with and without treatment by flipping the observed outcome label for observations in the treatment group chosen randomly in proportion to their $\tau_C(X_i)$ as in [36]. We do not observe the potential checkout amount for 4680 customers whose outcome we flip from non-converted to converted. We choose not to remove these customers and instead approximate the data generating process of the checkout amount to generate synthetic values. We employ a gradient boosted tree ensemble (GBT) on the customers for which we observed the checkout amount to ensure that the approximating model is sufficiently flexible and predict the unobserved checkout amounts using the tree ensemble. The treatment effect τ_R is then added to the observed, or if unavailable to the synthetic, basket value. The empirical Monte Carlo approach allows us to evaluate our approach against the actual distribution of customers including the real purchase process while controlling the individual treatment effect for evaluation.

We use five-fold cross-validation to compare the causal hurdle model and the distinct modeling approach on the holdout data. Considering the estimation strategies for the treatment effect and the conversion probability results in the eight combinations summarized in Table 2. The proposed approaches to model the treatment effect using hurdle models are grouped as causal hurdle models. They are defined by a two-stage hurdle approach estimating the response probability and the response value given a positive response, with and without treatment. The hurdle specification provides an estimate of the conversion probability under treatment without the need to estimate an additional model. The distinct modeling approaches are defined by the estimation of the treatment effect on profit in one stage. Approaches that estimate the treatment effect on profit require the estimation of a separate model to estimate the conversion probability under treatment. For each approach, we compare a linear specification to a more flexible specification using the GBT. To simplify the analysis for the distinct modeling approaches, we choose the same specification for the models of the treatment effect and the conversion probability.

We consider three approaches for the estimation of the CATE. First, the single-model approach that includes the treatment variable into the model. We test the single-model approach only in combination with the GBT specification because the approach requires a sufficiently flexible model to capture interaction effects between the treatment indicator and covariates. The single-model approach requires the estimation of two models. Under the hurdle model approach, the two models are one single-model including the treatment indicator respectively for the conversion and spending given conversion. Under the distinct modeling approach, the two models are a single-model for the profit and a separate model for conversion under treatment.

Second, the two-model approach that relies on the estimation of separate models for the treatment and control group. For the distinct modeling approach, the two models estimate the expected profit in the treatment and control group, respectively, with a separate model for conversion under treatment. For the hurdle approach, the two models for the treatment and control group are hurdle models that each consist of one model for the conversion and one model for the spending given conversion.

Third, the doubly-robust outcome transformation (DR) due to [37]. The DR approach provides an additional benchmark that has shown strong empirical performance in the econometric literature [20]. Under the DR approach, the

Table 2: Summary of model specifications considered in the experiment

	Architecture			Number
Stages	CATE Model	Conversion Model	Estimator	of Models
<i>Causal Hurdle Models</i>				
hurdle	single-model	-	gbt	2
hurdle	two-model	-	linear	4
hurdle	two-model	-	gbt	4
<i>Distinct Modeling Approaches</i>				
one-stage	single-model	separate	gbt	2
one-stage	two-model	separate	linear	3
one-stage	two-model	separate	gbt	3
one-stage	dr	separate	linear	5
one-stage	dr	separate	gbt	5

treatment effect is estimated using a single model on a transformation of the profit

$$Y_i^{DR} = \mu_1 - \mu_0 + \frac{T_i(Y_i - \mu_1)}{p(T = 1|X_i)} - \frac{(1 - T_i)(Y_i - \mu_0)}{1 - p(T = 1|X_i)}$$

with $\mu_1 = E[Y|X_i = x, T_i = 1]$ and $\mu_0 = E[Y|X_i = x, T_i = 0]$

The expected profit in the treatment and control group, $E[Y|X_i, T_i = 1]$ and $E[Y|X_i, T_i = 0]$, and the probability to receive treatment, $p(T = 1|X_i)$, are estimated by three auxiliary models. For simplicity, we use linear regression to estimate the expected profit and logistic regression to estimate the probability to receive treatment.

5 Empirical Results

Recall that each targeting policy is a combination of an estimate of the treatment effect and an estimate of the treatment cost. The profit generated by the targeting policy depends on the quality of the estimates of the treatment effect and on the quality of the estimates of the conversion probability under treatment. The analysis is therefore structured around the evaluation of the treatment effect estimation and the evaluation of the estimation of the expected individual-specific cost in the proposed hurdle framework. First, we test if the conversion probability estimates are sufficiently informative and economically relevant for profitable targeting. We propose that the expected individual cost is practically relevant and the estimate $p(C|X_i, T = 1)$ is sufficiently precise for profitable targeting. We test their economical relevance through an evaluation of campaign profit. Second, we test if the CATE estimates in the proposed causal hurdle framework are equivalent to CATE estimates under the conventional modeling strategy. The campaign profit under joint model estimation is expected to be at least as high as under the distinct model approach, while being easier to manage in application. Therefore, we evaluate the CATE estimates using statistical metrics on the simulated treatment effect and the evaluation of campaign profit under population-based cost estimates.

Third, we test if the proposed analytical targeting policy has a higher return than empirically optimized policies. The analytical targeting policy the individual treatment effect and response probability that requires a combination of treatment effect estimation with individual-level cost estimation. We test the campaign profit of the policy under the distinct estimation approach and the proposed causal hurdle framework.

The incremental campaign profit must be determined against baseline policies. As a general baseline, we select the sum of profit from individuals in the data when no campaign is run, i.e. no individual is targeted with the marketing treatment. We compare the proposed analytical targeting policy against the alternative *Empirical* policy suggested by our literature review, which determines a targeting threshold that maximizes campaign profit on the training data [38].

5.1 Profit Implications of Individual Cost Estimates

We begin with the prediction of conversion probability under treatment to calculate the expected cost of targeting. Table 3 reports the profit and the fraction of customers treated for campaigns under the proposed targeting policy stated in Eq. 2 (*Analytical*). We evaluate the conversion probability estimates provided under the estimation procedures described by the columns *Conversion Model* and *Estimator*. To calculate the expected cost, the analysis includes the model-based approaches discussed above and, for comparison, the expected conversion rate in the population. The conversion rate assumes a constant conversion probability for all customers, which implies a homogeneous treatment

Table 3: Policy profit for the conversion models evaluated under selected treatment effect estimation methods

Policy	Architecture			Profit	Fraction Treated
	CATE Model	Conversion Model	Estimator		
Baseline	-	-	-	46,236	0.00
Analytical	ATE	Conversion Rate	-	50,830	1.00
Analytical	ATE	Single-Model	GBT	52,931	0.84
Analytical	ATE	Two-Model/Distinct	Linear	51,936	0.76
Analytical	ATE	Two-Model/Distinct	GBT	52,402	0.79
Analytical	Actual	Conversion Rate	-	55,493	0.71
Analytical	Actual	Single-Model	GBT	56,696	0.72
Analytical	Actual	Two-Model/Distinct	Linear	57,361	0.69
Analytical	Actual	Two-Model/Distinct	GBT	57,022	0.69

cost that is often assumed in studies on cost-sensitive learning [39, 40]. The conversion estimates are combined with two estimation procedures of the treatment effect to calculate the campaign profit. The treatment effect for each customer is either estimated to be the average treatment effect over all customers in the training data, denoted as ATE, or presumed to be estimated perfectly, denoted as *Actual*. The ATE policy makes the simplifying assumption that there exists no heterogeneity in treatment effects and is equivalent to the constant acceptance rate of the treatment assumed in prior studies on customer churn [41, 25]. Beyond the comparison to prior research, the ATE policy provides an estimate of the profit implication of the cost-based targeting alone. Presuming perfect estimation of the ITE is unrealistic in practice, since the true treatment effect is unobservable. As a second comparison, the campaign profit under the actual ITE provides an upper bound on campaign profit that would be achievable by estimation of individual-level cost under optimal performance of the treatment effect model.

Table 3 shows that customer-level estimates of the conversion rate provide a more accurate estimate of customer-level costs and translate into higher campaign profit when used in combination with either ATE or CATE estimates. To provide some context for the profit of the models of interest, consider the two simple policies of targeting no or every individual in the population. The Baseline policy, under which no customer is targeted, results in a profit of €46,236. This profit is the result of the natural probability in the customer population to complete a purchase, which we hope to increase with the marketing campaign. Next, consider the average treatment effect of the population and the average conversion rate of the population given treatment. The analytical policy indicates to target all customers given the positive expected average return. The treatment rate of 100% results in a profit of €50,830. The campaign profit defined as the difference between the campaign and no marketing incentive is €4,594.

We now introduce an individual-level targeting policy by estimating the cost of the marketing treatment on the customer level with a response model. Both the two-model and single-model architectures result in a substantial decrease in the fraction of customers treated from 100% to 76%–84%, depending on the estimator. The decrease in treatment ratio is accompanied by an increase in campaign profit between €1,100 and €2,100, again depending on the estimator. This substantial increase of 24–46% in campaign profit compared to universal treatment is the direct result of controlling the expected treatment cost for each customer.

The observed positive impact on profit generalizes to customer-level targeting based on the CATE under treatment effect heterogeneity. A hypothetical targeting policy based on the actual ITE and the average cost results in a campaign profit of €55,493. We again find that campaign profit using customer-level estimates of the treatment cost increase campaign profit by €1,200–€1,900.

Compare now the two-model approach and single model approach with the GBT estimator. The single model GBT results in a campaign profit of €50,830 and €56,696, while the two-model GBT results in a profit of €52,402 and €57,022, for the constant and true treatment estimates respectively. We conclude that the campaign profit under the single-model conversion model is slightly lower than from the campaign profit under the two-model conversion model.

5.2 Profit Implications of Causal Hurdle Models

The analysis was so far restricted to the conversion models and the effect of customer-level cost estimation. Considering the probability of each customer to accept the costly marketing incentive directly results in a substantial profit increase. We therefore conclude that the estimate $p(C|X_i, T = 1)$ is sufficiently precise for profitable targeting and that the expected individual cost is practically relevant for customer targeting. The conclusion applies to campaigns considering heterogeneous treatment effects and population-level estimates of the average treatment effect. In contrast

to prior work [5], our analysis implies that customers with a positive response to treatment can be unprofitable targets due to a high conversion probability after treatment and the associated higher expected treatment cost.

Table 4: Quality of model estimates for the conditional average treatment effect

Architecture			Error	
CATE Model	Stages	Estimator	RMSE	TOL
ATE	-	-	2.75	3387.90
Single-Model	Hurdle	GBT	2.37	3384.91
Two-Model	Hurdle	Linear	5.15	3410.78
Two-Model	Hurdle	GBT	1.94	3381.79
Single-Model	One-Stage	GBT	2.77	3387.49
Two-Model	One-Stage	Linear	4.16	3407.13
Two-Model	One-Stage	GBT	1.94	3381.76
DR	One-Stage	Linear	4.11	3406.10
DR	One-Stage	GBT	2.37	3385.45
Actual	-	-	0.00	3374.99

TOL: Transformed Outcome Loss on the observed outcomes

RMSE: Root Mean Squared Error on the simulated treatment effect.

We now consider the estimation of the CATE for the customer-level prediction of marketing effectiveness. We evaluate the quality of the CATE models using statistical indicators and the resulting profit as part of a targeting policy.

Table 4 shows the root-mean-squared error (RMSE) of the CATE estimates compared to the simulated treatment effect on profit and the transformed outcome loss (TOL) on the observed outcomes. We include the TOL as a feasible metric when the true treatment effect is not simulated and therefore not known [42]. To put the results into context, the ATE estimate provides the baseline obtained by a constant estimator, while the actual ITE in the last row provides the lowest obtainable TOL on the data. Kernel density plots showing the distributional fit of the CATE estimates are available in Figure 2 in the Appendix.

The linear model is consistently outperformed by the GBT and, on average, ranks below the constant treatment effect estimate. The linear model achieves an RMSE of 5.15, 4.16 and 4.11 within the two-model hurdle and the two-model and doubly-robust one-stage architectures, respectively. The noticeably high RMSE for the two-model hurdle architecture is the result of treatment effect estimates with high absolute value for a small number of observations. The good calibration of the linear model may nevertheless ensure its value within a targeting policy.

Under GBT specification, the two-model hurdle architecture compares favorably to the single-model hurdle architecture and models estimating the overall treatment effect directly. The two-model architecture achieves an RMSE of 1.94 for both the hurdle and one-stage model, respectively. The single-model architecture, in comparison, achieves an RMSE of 2.37 and 2.77 for the respective target. The one-stage doubly-robust model with an RMSE of 2.37 performs better than the single-model architecture but worse than the two-model approach. The results suggest that the single-model approach, which requires the least number of models to be estimated, provides worse estimates of the treatment effect than the two-model or DR models. Analysis of the resulting policy profit will clarify if the gap in estimation precision results in a substantial effect on campaign profit in practice.

The campaign profit from customer-level targeting provides an interpretable evaluation of the CATE models. Table 5 reports the campaign profit resulting from each CATE model in combination with a constant targeting cost derived from the population average conversion probability. When applied within a targeting policy, the conclusions drawn from Table 4 are only partially supported.

The linear models are highly profitable when used as part of a targeting policy. With campaign profit of € 54,550, € 54,456 and € 54,459, the linear models are superior to the constant treatment estimate with a profit of € 50,830 despite their higher RMSE. The linear specification is, however, dominated by the GBT specification for all architectures except the single-model.

Within the GBT specification, the single-model approach is substantially less profitable than other architectures. The two-model hurdle model, two-model one-stage model and doubly-robust one-stage model show no substantial difference at a campaign profit of € 55,590, € 55,146 and € 54,629, respectively. Campaign profit is substantially worse for the single-model architecture, with a profit of € 48,840 for the hurdle model and a profit of € 52,795 for the one-stage model. The results confirm that small differences in the precision of the CATE estimates have a practically relevant effect on campaign profit. Despite the hurdle single-model GBT showing a lower RMSE than the ATE baseline in Table 4, it underperforms the baseline of uniform treatment by € 1,990 when applied for targeting. For all

Table 5: Campaign profit for CATE-based targeting under population average cost estimates

Policy	Architecture				Profit	Fraction Treated
	Stages	CATE Model	Estimator	Conversion Model		
Baseline	-	-	-	-	46,236	0.00
Analytical	-	ATE	-	Conversion Rate	50,830	1.00
Analytical	Hurdle	Single-Model	GBT	Conversion Rate	48,840	0.20
Analytical	Hurdle	Two-Model	Linear	Conversion Rate	54,550	0.66
Analytical	Hurdle	Two-Model	GBT	Conversion Rate	55,590	0.70
Analytical	One-Stage	Single-Model	GBT	Conversion Rate	52,795	0.41
Analytical	One-Stage	Two-Model	Linear	Conversion Rate	54,456	0.66
Analytical	One-Stage	Two-Model	GBT	Conversion Rate	55,146	0.72
Analytical	One-Stage	DR	Linear	Conversion Rate	54,459	0.66
Analytical	One-Stage	DR	GBT	Conversion Rate	54,629	0.83
Analytical	-	Actual	-	Conversion Rate	55,493	0.71

other approaches, we observe a substantial increase in campaign profit under the analytical targeting policy relative to uniform targeting in the range of € 3,626–€ 4,760. With regard to the comparison between the hurdle and one-stage approaches, the results suggest that the two-stage hurdle model results in campaign profit equivalent to that of the one-stage approaches.

5.3 Profit Implications of the Proposed Analytical Targeting Policy

The analysis has so far addressed the evaluation of the CATE and conversion estimates separately. We now evaluate the joint impact on campaign profit of the interaction between the proposed treatment and conversion models as part of a targeting policy. Recall that the single- and two-model hurdle models provide an explicit estimate of the conversion probability by design. CATE models that estimate the treatment effect on the profit directly require a separate classification model to predict the conversion rate under treatment.

Table 6 reports the campaign profit under the proposed analytical targeting policy and the empirical thresholding policy introduced in Section 2. Recall that the analytical policy employs the estimated CATE and conversion probability under treatment to calculate the expected profit from targeting the customer using the decision rule proposed in Eq. 2. The empirical policy determines the profit-optimal threshold on the CATE estimates through numeric optimization of the overall campaign profit on the training data. The analytical targeting policy results in a higher campaign profit relative to the baseline for all model architectures and relative to the empirical policy for seven out of eight architectures.

Comparing model architectures, we find that the proposed causal hurdle framework performs at least comparable to the combination of a one-stage treatment effect model with a separate conversion model. All architectures under the analytical and empirical policy increase the campaign profit compared to uniform targeting. Compared to the baseline, the analytical policy increases campaign profit by € 2,051–€ 5,342. The increase in campaign profit by combining estimates of the CATE and expected cost results in a median additional increase of € 1,000 over the treatment-based policy ignoring response-dependent cost reported in Table 5. The campaign profit compared to no targeting lies for the hurdle architectures in the range of € 6,645–€ 9,936 and for the one-stage architectures in the range of € 7,728–€ 9,075. Comparing within the two-model architectures, which predict the treatment effect most precisely, we find no substantial difference at a net campaign profit of around € 9,750 for the hurdle two-model and the one-stage two-model approach. We also find no substantial difference to the DR approach evaluated as a state-of-the-art competitor one-stage benchmark.

However, the single-model architecture performs substantially worse than the two-model approaches within the one-stage and hurdle architectures. This finding is in line with the lower precision of the treatment effect estimates reported in Table 4. We conclude that the proposed two-model hurdle architecture, although not the single-model hurdle architecture, achieves competitive campaign profit to the alternative one-stage, distinct modeling architectures. Despite its disadvantage in estimation, the single-model hurdle architecture improves the effectiveness of the model building process by reducing the number of models that need to be estimated to two compared to the three to five models required by the distinct and two-model hurdle architectures.

Comparing the same model architecture under the analytical and empirical targeting policy, the proposed analytical targeting policy increases campaign profit by an on average € 1000, excluding the single-model approach due its

Table 6: Campaign profit for CATE-based targeting under model-based cost estimation

Policy*	Architecture				Profit	Fraction Treated
	Stages	CATE Model	Conversion Model	Estimator		
Baseline	-	-	-	-	46,236	0.00
Analytical	-	ATE	-	Conversion Rate	50,830	1.00
Analytical	Hurdle	Single-Model	-	GBT	54,665	0.53
Analytical	Hurdle	Two-Model	-	Linear	56,172	0.71
Analytical	Hurdle	Two-Model	-	GBT	56,084	0.71
Analytical	One-Stage	Single-Model	Separate	GBT	52,881	0.49
Analytical	One-Stage	Two-Model	Separate	Linear	56,010	0.66
Analytical	One-Stage	Two-Model	Separate	GBT	55,942	0.68
Analytical	One-Stage	DR	Separate	Linear	56,028	0.66
Analytical	One-Stage	DR	Separate	GBT	55,160	0.75
Empirical	Hurdle	Single-Model	-	GBT	53,964	0.78
Empirical	Hurdle	Two-Model	-	Linear	54,940	0.73
Empirical	Hurdle	Two-Model	-	GBT	55,311	0.70
Empirical	One-Stage	Single-Model	-	GBT	54,546	0.69
Empirical	One-Stage	Two-Model	-	Linear	54,269	0.68
Empirical	One-Stage	Two-Model	-	GBT	55,295	0.70
Empirical	One-Stage	DR	-	Linear	54,481	0.67
Empirical	One-Stage	DR	-	GBT	54,791	0.83

Empirical denotes targeting based on the profit-maximizing threshold on the training data.

weak absolute performance. This result supports the conclusion that the proposed analytical targeting policy increases campaign profit relative to numeric optimization of the decision threshold. Note, however, that the ratio of customer treated by the single-model approaches deviates from the other architectures under the analytical policy, but not under the empirical policy. We interpret these findings as an issue of model calibration for the single-model approach. If either the probability model or the treatment effect model is not well calibrated, expectation calculations will be inaccurate. In this case, empirical thresholding can be an alternative to model recalibration on the level of the policy rather than recalibration of the model estimates. For calibrated models including the GBT when estimated in the two-model architecture, the analytical targeting policy substantially increases policy profit.

We conclude that the proposed analytical decision policy can substantially increase the profitability of targeting models in practice and that the proposed two-model hurdle model architecture is an efficient way to estimate the necessary decision variables in a unified framework.

6 Conclusion

We have presented a general analysis of the customer targeting decision problem under different types of variable costs and proposed a causal hurdle framework to estimate the relevant decision variables efficiently. Our results demonstrate that the consideration of treatment cost substantially increases campaign profit when used for customer targeting independent of whether the treatment effect is considered to vary over customers.

While customer targeting based on expected profit has been used to optimize campaigns, previous analytical frameworks do not include marketing incentives that are conditioned on a profitable customer response, e.g. a retention offer or voucher. We identify these common marketing incentives as a type of stochastic variable cost. Our formal analysis of the targeting decision problem under customer response-dependent costs shows that estimating the expected cost requires an estimate of the customer response conditional on treatment. A central result to the customer targeting literature is that profit-optimal targeting often requires modeling the effect of the marketing treatment and the net customer response under treatment.

In order to estimate the treatment effect and response efficiently, we propose a framework for joint estimation. Our causal hurdle model combines a hurdle model for customer choice with methods for causal inference. The proposed approach is feasible with the single-model and two-model approaches for the estimation of conditional treatment

effects. We find that the causal hurdle model under the two-model specification achieves competitive campaign profit on a coupon targeting campaign in an e-commerce setting, while streamlining model building.

With the increasing relevance of digital marketing and the associated increase in marketing incentives with low targeting-dependent and high response-dependent variable costs, our results are highly relevant for practitioners. We further expect the development of efficient approaches for the estimation of flexible hurdle models and the application of our decision analysis to other applications with stochastic costs as fruitful areas for future research.

References

- [1] David L. Olson and Bongsug Chae. Direct marketing decision support through predictive customer response modeling. *Decision Support Systems*, 54(1):443–451, 2012.
- [2] Statista. eCommerce. Technical report, 2019.
- [3] Statista. Advertising Spending in the Catalog, Mail-order Houses Industry in the United States. Technical report, 2017.
- [4] Floris Devriendt, Darie Moldovan, and Wouter Verbeke. A literature survey and experimental evaluation of the state-of-the-art in uplift modeling: A stepping stone toward the development of prescriptive analytics. *Big Data*, 6(1):13–41, 2018.
- [5] Robin Marco Gubela, Artem Bequé, Fabian Gebert, and Stefan Lessmann. Conversion uplift in e-commerce: A systematic benchmark of modeling strategies. *International Journal of Information Technology & Decision Making*, 18(3):747–791, 2019.
- [6] Susan Athey, Julie Tibshirani, and Stefan Wager. Generalized random forests. *The Annals of Statistics*, 47(2):1148–1178, 2019.
- [7] Scott Powers, Junyang Qian, Kenneth Jung, Alejandro Schuler, Nigam H Shah, Trevor Hastie, and Robert Tibshirani. Some methods for heterogeneous treatment effect estimation in high dimensions. *Statistics in Medicine*, 37(11):1767–1787, 2018.
- [8] Claudia Shi, David M. Blei, and Victor Veitch. Adapting Neural Networks for the Estimation of Treatment Effects. *arXiv preprint*, page arXiv:1906.02120, 2019.
- [9] Robin M. Gubela, Stefan Lessmann, and Szymon Jaroszewicz. Response transformation and profit decomposition for revenue uplift modeling. *European Journal of Operational Research*, 283(2):647–661, 2020.
- [10] Eva Ascarza. Retention futility: Targeting high risk customers might be ineffective. *Journal of Marketing Research*, 55(1), 2018.
- [11] Günter J. Hitsch and Sanjog Misra. Heterogeneous Treatment Effects and Optimal Targeting Policy Evaluation. *SSRN*, 2018.
- [12] Eva Ascarza and Bruce G. S. Hardie. A joint model of usage and churn in contractual settings. *Marketing Science*, 32(4):570–590, 2013.
- [13] Aimée Backiel, Bart Baesens, and Gerda Claeskens. Predicting time-to-churn of prepaid mobile telephone customers using social network analysis. *Journal of the Operational Research Society*, 67(9), 2016.
- [14] Navdeep S. Sahni, Dan Zou, and Pradeep K. Chintagunta. Do targeted discount offers serve as advertising? Evidence from 70 field experiments. *Management Science*, 63(8):2688–2705, 2016.
- [15] Xi (Alan) Zhang, V. Kumar, and Koray Cosguner. Dynamically managing a profitable email marketing program. *Journal of Marketing Research*, 2017.
- [16] Harald Hruschka. Considering endogeneity for optimal catalog allocation in direct marketing. *European Journal of Operational Research*, 206(1):239–247, 2010.
- [17] Sandra Mitrović, Bart Baesens, Wilfried Lemahieu, and Jochen De Weerd. On the operational efficiency of different feature types for telco Churn prediction. *European Journal of Operational Research*, 267(3):1141–1155, 2018.
- [18] Piotr Rzepakowski and Szymon Jaroszewicz. Uplift modeling in direct marketing. *Journal of Telecommunications and Information Technology*, 2:43–50, 2012.
- [19] T. Wendling, K. Jung, A. Callahan, A. Schuler, N. H. Shah, and B. Gallego. Comparing methods for estimation of heterogeneous treatment effects using observational data from health care databases. *Statistics in Medicine*, 37:3309–3324, 2018.

- [20] Michael C. Knaus, Michael Lechner, and Anthony Strittmatter. Machine Learning Estimation of Heterogeneous Causal Effects: Empirical Monte Carlo Evidence. *IZA Discussion Paper*, 12039, 2019.
- [21] Sören R. Künzel, Jasjeet S. Sekhon, Peter J. Bickel, and Bin Yu. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the National Academy of Sciences*, 116(10):4156–4165, 2019.
- [22] Susan Athey and Stefan Wager. Efficient Policy Learning. *arXiv preprint*, page arXiv:1702.02896, 2017.
- [23] Scott A. Neslin, Sunil Gupta, Wagner Kamakura, Junxiang Lu, and Charlotte H. Mason. Defection detection: Measuring and understanding the predictive accuracy of customer churn models. *Journal of Marketing Research*, 43(2):204–211, 2006.
- [24] Behram Hansotia and Brad Rukstales. Incremental value modeling. *Journal of Interactive Marketing*, 16(3):35–46, 2002.
- [25] Aurelie Lemmens and Sunil Gupta. Managing Churn to Maximize Profits. *Social Science Research Network*, page 2964906, 2017.
- [26] Floris Devriendt, Jeroen Berrevoets, and Wouter Verbeke. Why you should stop predicting customer churn and start using uplift models. *Information Sciences*, In Press, 2019.
- [27] Stefan Lessmann, Johannes Haupt, Kristof Coussemment, and Koen W. De Bock. Targeting customers for profit: An ensemble learning framework to support marketing decision making. *Information Sciences*, In Press, 2019.
- [28] Nicholas J. Radcliffe. Using control groups to target on predicted lift: Building and assessing uplift models. *Direct Marketing Analytics Journal*, pages 14–21, 2007.
- [29] Eustache Diemert, Artem Betlei, Christophe Renaudin, and Massih-Reza Amini. A Large Scale Benchmark for Uplift Modeling. In *Proceedings of the AdKDD and TargetAd Workshop, KDD*, London, United Kingdom, 2018. ACM.
- [30] Bas Donkers, Richard Paap, Jedid-Jah Jonker, and Philip Hans Franses. Deriving target selection rules from endogenously selected samples. *Journal of Applied Econometrics*, 21(5):549–562, 2006.
- [31] Merel Van Diepen, Bas Donkers, and Philip Hans Franses. Dynamic and competitive effects of direct mailings: A charitable giving application. *Journal of Marketing Research*, 46(1):120–133, 2009.
- [32] Adrian Colin Cameron and P. K. Trivedi. *Microeconometrics: Methods and Applications*. Cambridge University Press, Cambridge, 2005.
- [33] Jennifer L. Hill. Bayesian Nonparametric Modeling for Causal Inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240, 2011.
- [34] Annika Baumann, Johannes Haupt, Fabian Gebert, and Stefan Lessmann. The price of privacy: An evaluation of the economic value of collecting clickstream data. *Business & Information Systems Engineering*, 61(4):413–431, 2019.
- [35] Eva Ascarza, Peter Ebbes, Oded Netzer, and Matthew Danielson. Beyond the target customer: Social effects of customer relationship management campaigns. *Journal of Marketing Research*, 54(3):347–363, 2017.
- [36] Xinkun Nie and Stefan Wager. Quasi-Oracle Estimation of Heterogeneous Treatment Effects. *arXiv:1712.04912*, 2017.
- [37] James M. Robins and Andrea Rotnitzky. Semiparametric efficiency in multivariate regression models with missing data. *Journal of the American Statistical Association*, 90(429):122–129, 1995.
- [38] Charles X. Ling and Chenghui Li. Data Mining for Direct Marketing: Problems and Solutions. In *Proceedings of the 4th International Conference on Knowledge Discovery and Data Mining*, pages 73–79, Menlo Park, US, 1998. AAAI Press.
- [39] Nicolas Glady, Bart Baesens, and Christophe Croux. Modeling churn using customer lifetime value. *European Journal of Operational Research*, 197(1):402–411, 2009.
- [40] Thomas Verbraken, Cristián Bravo, Richard Weber, and Bart Baesens. Development and application of consumer credit scoring models using profit-based classification measures. *European Journal of Operational Research*, 238(2):505–513, 2014.
- [41] Thomas Verbraken, Wouter Verbeke, and Bart Baesens. A novel profit maximizing metric for measuring classification performance of customer churn prediction models. *IEEE Transactions on Knowledge and Data Engineering*, 25(5):961–973, 2012.
- [42] Susan Athey and Guido W. Imbens. Machine Learning Methods for Estimating Heterogeneous Causal Effects. 2015.

- [43] Wouter Verbeke, Karel Dejaeger, David Martens, Joon Hur, and Bart Baesens. New insights into churn prediction in the telecommunication sector: A profit driven data mining approach. *European Journal of Operational Research*, 218(1):211–229, 2012.

A Relation to Previous Formulations of Churn Campaign Profit

A popular definition of the profit of a customer retention campaign [43, 25, 26] is given by [23]:

$$\Pi = N\alpha [\beta\gamma(V - \delta - c) + \beta(1 - \gamma)(-c) + (1 - \beta)(-\delta - c)] - A$$

with

N : Number of customers

α : Ratio of customers targeted

V : The value of the customer to the company, CLV in their original notation

β : Fraction of (targeted) customers who would churn

γ : Fraction of (targeted) customers who decide to remain when receiving the marketing incentive

δ : The cost of the marketing incentive if it is accepted

c : The cost of contacting the customer with the marketing incentive

A : The fixed cost of running the retention campaign

The number of customers targeted by the campaign and the fixed costs are relevant to calculate the overall campaign profit, but do not affect the targeting decision for a single customer. The profit estimate relevant for customer targeting is thus the part in square brackets:

$$\pi_i = \beta_i\gamma_i(V - \delta) + \beta_i(1 - \gamma_i)(-c) + (1 - \beta_i)(-\delta - c)$$

We will show that this expression is equivalent to the proposed decision policy (Eq. 2) under restrictive assumptions. Using the additive property of the probabilities β_i and $(1 - \beta_i)$ and γ_i and $(1 - \gamma_i)$, we can summarize the terms:

$$\begin{aligned}\pi_i &= \beta_i\gamma_i(V) + \beta_i\gamma_i(-\delta) + (1 - \beta_i)(-\delta) + \beta_i(-c) + (1 - \beta_i)(-c) \\ &= \beta_i\gamma_i V + \beta_i\gamma_i(-\delta) + (1 - \beta_i) - \delta - c \\ &= \beta_i\gamma_i V - \delta(\beta_i\gamma_i + 1 - \beta_i) - c \\ &= \beta_i\gamma_i V - \delta(1 - \beta_i(1 - \gamma_i)) - c\end{aligned}$$

We will target a customer if the profit is positive, i.e.

$$\beta_i\gamma_i V - (1 - \beta_i(1 - \gamma_i))\delta - c > 0 \quad (12)$$

In Eq. 2, we propose the decision rule

$$p_i(1)(V(1) - \delta) - c > p_i(0) \cdot V(0)$$

Assuming that the value of the customer is not influenced by the marketing incentive $V(1) = V(0) = V$ allows us to rearrange the inequality to

$$(p_i(1) - p_i(0))V - p_i(1)\delta - c > 0 \quad (13)$$

Eq. 12 and Eq. 13 are equivalent if the following equalities hold:

$$\begin{aligned}p_i(1) &= (1 - \beta_i(1 - \gamma_i)) \\ p_i(1) - p_i(0) &= \beta_i\gamma_i\end{aligned}$$

In words, we require $p(1)$ to be the complement to the probability for a customer to plan to churn and churn even when offered the treatment. The complementary event is for a customer not to plan to churn or to plan to churn but remain after treatment; or simply, the probability of the customer to stay when given treatment.

We further require $p(1) - p(0)$ to be the probability of a customer to plan to churn and to not churn when offered the treatment. As $\beta_i \cdot \gamma_i \in [0; 1]$, this equality holds under the assumption that the treatment effect is strictly positive, i.e. $p(1) - p(0) \in [0; 1]$. However, we know that the treatment effect on the response probability, $p(1) - p(0)$, is in principle bounded in $[-1, 1]$ and that negative effects are a critical issue in churn campaigns in practice [10]. Under

the previous campaign profit formulation, we see that $\beta_i \gamma_i = 0$ if either β_i or γ_i or both are zero. In words, the campaign has no effect if no customers consider to churn or no customers accept the marketing incentive when offered. This conflicts with the observation that when no customers plan to churn, the campaign may have a net negative effect by priming inattentive customers to churn. Specifically, the shortcoming of the customer profit proposed by [23] is that it implicitly assumes a positive treatment effect by restricting the action space of the customer to $\gamma \in \{\text{Accept treatment, Disregard treatment}\}$.

We conclude that the proposed decision framework is a generalization of [23]’s campaign profit function to cases where a customer may react adversely to the treatment. As an alternative formulation to calculate the overall churn campaign profit, we propose for the general case:

$$\Pi = \sum_{i \in N} \{T_i [(p_i(1) - p_i(0))V_i - p_i(1)\delta - c]\} - A$$

In cases with no or little variation in customer sensitivity to the marketing treatment and a constant customer lifetime value, the churn campaign profit can be simplified to:

$$\Pi = N\alpha [\hat{\tau}_{ATE}V - p(1)\delta - c] - A$$

B Additional Evaluation Results

Table 7 shows the quality of predictions for the conversion probability conditional on treatment. Recall that the single-model hurdle model includes the treatment indicator as a covariate into the model. The two-model hurdle model estimates four separate models, one of which predicts the conversion probability within the treatment group. Note that the default approach, which separates treatment effect estimation and conversion prediction, also requires the estimation of an identical conversion model. This is the redundancy that the proposed causal hurdle framework avoids. We find no substantial difference in the area-under-the-ROC-curve (ROC-AUC) or the Brier score, which indicates model calibration.

Table 7: Quality of model estimates for the prediction of conversion under treatment

Architecture			ROC-AUC	Brier Score
Stages	Specification	Estimator		
Hurdle/Distinct	Two-Model	Linear	0.636	0.103
Hurdle/Distinct	Two-Model	GBT	0.640	0.102
Hurdle Model	Single-Model	GBT	0.636	0.102

Figure 2 depicts the kernel density plot for the treatment estimation approaches and the GBT specification. We combine the out-of-sample estimates for each iteration of the cross-validation procedure to obtain out-of-sample estimates for the full dataset. The dotted line shows the kernel density of the actual ITE.

We observe that no approach fully captures the minor mode of the distribution to the left. The hurdle single-model GBT approach in addition shows a slight shift from the major mode of the distribution that relates to the worse precision reported in Table 4.

The support of the linear model specifications extends beyond the actual range of the simulated treatment effects and beyond the range shown in the figure. For a small set of observations, we observe predicted treatment effects beyond the range $[-100; 100]$ that explain the high statistical error reported in Table 4. For the remaining observations, we observe a reasonable fit to the actual treatment effect distribution. The general fit explains the profitability of the linear specification for the targeting policy as observations with weak support, for which linear extrapolation fails, by definition make up only a minority of cases in the data.

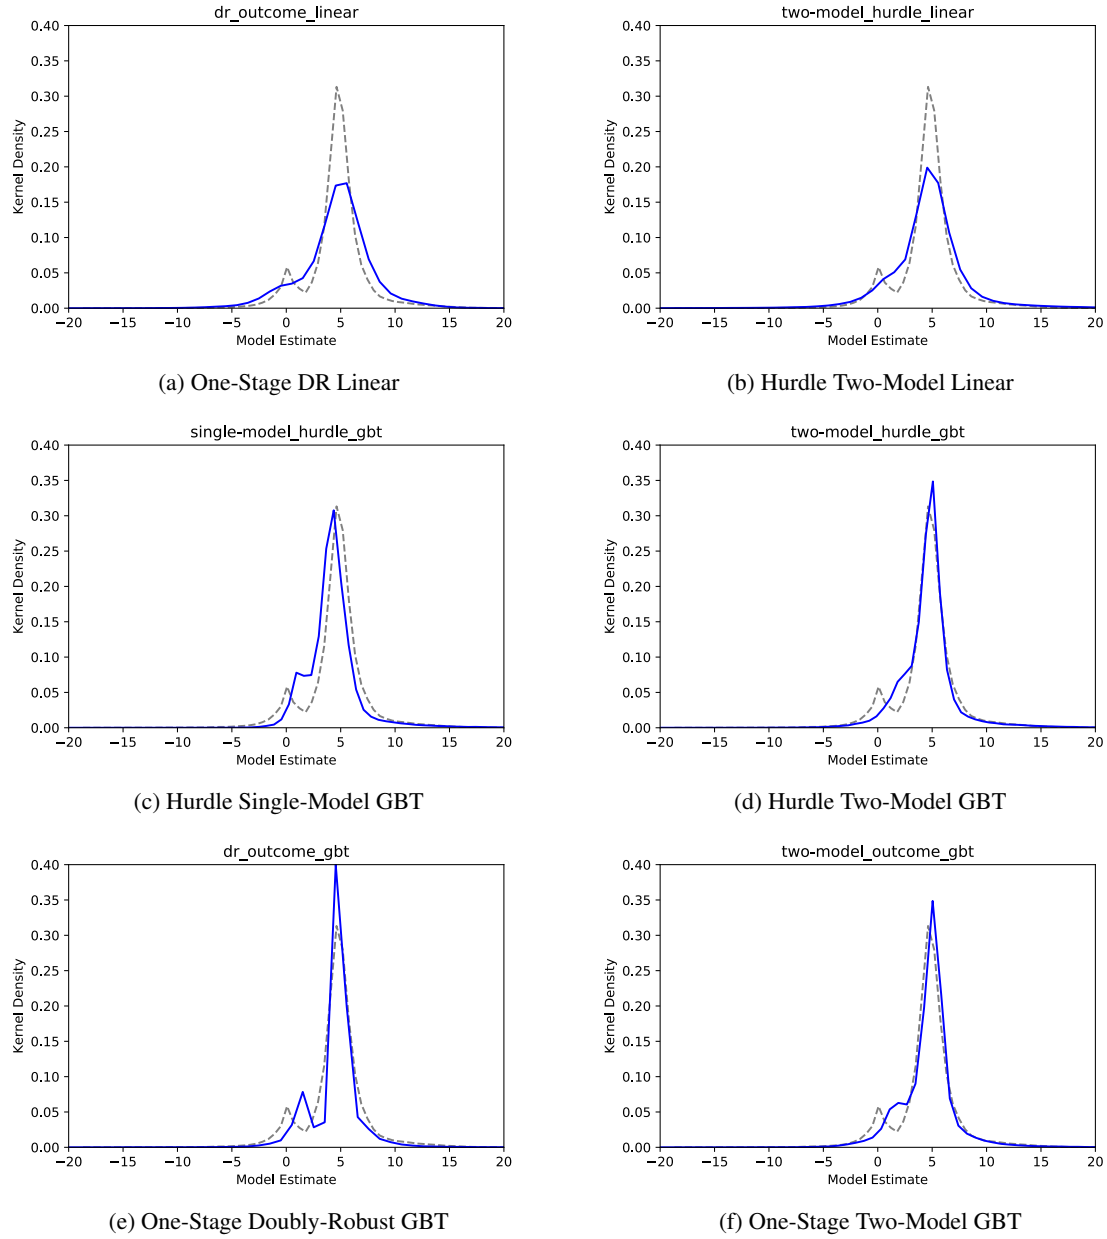


Figure 2: Kernel density plot of the CATE on the outcome as estimated by the hurdle (top rows) and one-stage models (bottom). The dotted line shows the actual individual treatment effect.

IRTG 1792 Discussion Paper Series 2020



For a complete list of Discussion Papers published, please visit
<http://irtg1792.hu-berlin.de>.

- 001 "Estimation and Determinants of Chinese Banks' Total Factor Efficiency: A New Vision Based on Unbalanced Development of Chinese Banks and Their Overall Risk" by Shiyi Chen, Wolfgang K. Härdle, Li Wang, January 2020.
- 002 "Service Data Analytics and Business Intelligence" by Desheng Dang Wu, Wolfgang Karl Härdle, January 2020.
- 003 "Structured climate financing: valuation of CDOs on inhomogeneous asset pools" by Natalie Packham, February 2020.
- 004 "Factorisable Multitask Quantile Regression" by Shih-Kang Chao, Wolfgang K. Härdle, Ming Yuan, February 2020.
- 005 "Targeting Customers Under Response-Dependent Costs" by Johannes Haupt, Stefan Lessmann, March 2020.

IRTG 1792, Spandauer Strasse 1, D-10178 Berlin
<http://irtg1792.hu-berlin.de>

This research was supported by the Deutsche
Forschungsgemeinschaft through the IRTG 1792.