## Segment Discovery: Enhancing E-commerce Targeting

QIQI LI\*, Amazon, USA ROOPALI SINGH, Amazon, USA CHARIN POLPANUMAS, Amazon, Japan TANNER FIEZ, Amazon, USA NAMITA KUMAR, Amazon, USA SHREYA CHAKRABARTI, Amazon, USA

Modern e-commerce services frequently target customers with incentives or interventions to engage them in their products such as games, shopping, video streaming, etc. This customer engagement increases acquisition of more customers and retention of existing ones, leading to more business for the company while improving customer experience. Often, customers are either randomly targeted or targeted based on the propensity of desirable behavior. However, such policies can be suboptimal as they do not target the set of customers who would benefit the most from the intervention and they may also not take account of any constraints. In this paper, we propose a policy framework based on uplift modeling and constrained optimization that identifies customers to target for a use-case specific intervention so as to maximize the value to the business, while taking account of any given constraints. We demonstrate improvement over state-of-the-art targeting approaches using two large-scale experimental studies and a production implementation.

#### **ACM Reference Format:**

Qiqi Li, Roopali Singh, Charin Polpanumas, Tanner Fiez, Namita Kumar, and Shreya Chakrabarti. 2024. Segment Discovery: Enhancing E-commerce Targeting. In *Proceedings of Recommender Systems Conference, CONSEQUENCES Workshop. (RecSys'24 CONSEQUENCES Workshop)*. ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/nnnnnn.nnnnnnn

## 1 INTRODUCTION

Promotions and discounts have become key components of modern e-commerce services. Popular promotions include discounts, bundled offers, free services, etc. By offering these promotions, companies aim to increase revenue and customer base, while also improving customer experience. However, such promotions usually incur a cost and can become unsustainable without any guardrails in place. A popular approach is to target customers with high or low propensity for desired behavior. For example, a retail company is likely to target customers who are at risk of leaving if they want to retain its customers by offering certain incentives. However, previous studies show that this strategy is ineffective and could be detrimental towards the company objectives [2, 6, 7]. Moreover, additional analysis needs to be done for the choice of propensity score threshold for targeting (e.g., target anyone whose propensity to leave is higher than 0.8), because the wrong threshold may lead to sub-optimal outcomes [2].

Each customer responds differently to the same promotion. This motivates the development of our personalized targeting approach that showcases these promotions to a set of customers that will have a positive response due to the

\*Also with Univeristy of California, Santa Barbara

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

treatment offered, while controlling for any business constraints. The contributions of this work are as follows: 1. We present a structured two-step methodology for customizing customer targeting strategies in the e-commerce industry. Our framework considers specific business constraints during the allocation of treatments to customers, aiming to optimize desired business outcomes. 2. We illustrate the practical application of our framework by devising targeting strategies for three distinct business scenarios, each characterized by unique outcome objectives and operational constraints. 3. We demonstrate the superiority of our approach over commonly used targeting methods based on propensity thresholds in e-commerce settings across three key business applications using offline policy estimation techniques. 4. We validate the efficacy of our proposed targeting policies against the current baseline approach using a large-scale online A/B test.

**Related Work**. Previously, a similar approach has been used in context of rehabilitation program design [17] while some companies have directly incorporated cost into the outcome variable of the causal uplift models [18] and performed constrained optimization for a specific business use-case [1, 8]. However, our approach offers a more generalized framework that is applicable to a wide variety of customer targeting problems.

## 2 METHODOLOGY

The goal is to identify a set of customers to target with a treatment, while optimizing for a metric taking account of any constraints. Consider a set of customers  $C = \{1, 2, ..., N\}$  and a set of treatments  $\mathcal{T} = \{t_0, t_1, ..., t_K\}$  to be applied to this customer set. Suppose we have some experimental data, where we observe  $\{X_i, T_i, Y_i\}$  for each customer i = 1, 2, ..., N. Here,  $X_i$  is a set of covariates associated with customer i and  $Y_i$  is the observed outcome of interest when customer i receives treatment  $T_i \in \mathcal{T}$ . We aim to find an optimal targeting policy  $\Pi$ , that assigns a treatment to each customer given certain attributes X, while optimizing for Y. The proposed policy is evaluated using metrics discussed in Section 2.2 and specific business metrics for each use case.

#### 2.1 Solution Framework

Given the problem setup, the goal is to find a set of customers to target with a specific treatment such that it optimizes the outcome, given certain constraints. We propose a two-stage approach: (1) estimate the impact of a treatment on the outcome for each customer, (2) find an optimal set of customers such that the estimated impact is maximized, given a constraint.

Stage 1: Uplift Model. Given the potential outcomes framework [11], we define  $Y_i(t_k)$  as the potential outcome corresponding to customer *i*, if they were to receive treatment  $t_k \in \mathcal{T}$ . We are interested in estimating conditional average treatment effect (CATE) of treatment  $t_k$ . Assuming  $t_0$  as the control treatment level, CATE of treatment  $t_k$  for an individual *i* can be defined as

$$\tau_k(x_i) = \mathbb{E}[Y_i(t_k) \mid X_i = x_i] - \mathbb{E}[Y_i(t_0) \mid X_i = x_i]$$

A key assumption for identifying CATE is unconfoundedness (strong ignorability) which means that conditional on covariates, the potential outcomes are independent of the treatment assignment. When it holds, current literature offers a wide range of models for causal estimation such as meta-learners (e.g., S, T and X-learners) and forest-based estimators (e.g., forest doubly robust Learner[10], causal forest [16] and causal forest double machine learning estimator [3]). Manuscript submitted to ACM

Stage 2: Constrained Optimization. Next step is to find an optimal targeting policy,  $\Pi$ , which assigns treatments to customers such that it maximizes the uplift measured in stage 1. A policy  $\Pi$  can be defined as a matrix of order  $N \times K$  as  $\Pi = [\pi_{ik}]_{N \times K+1}$ , where

$$\pi_{ik} = \begin{cases} 1, \text{ if treatment } t_k \text{ is assigned to customer } i \\ 0, \text{ if treatment } t_k \text{ is not assigned to customer } i \end{cases}$$

such that  $\sum_k \pi_{ik} = 1$ . The optimal policy  $\Pi$  can be found such that,

$$\max_{\Pi} \sum_{i=1}^{N} \sum_{k=1}^{K} \pi_{ik} w_i \hat{\tau}_k(x_i), \text{ s.t. } \pi_{ik} \in \{0,1\}, \sum_{k} \pi_{ik} = 1, g(\Pi) \le c,$$

where  $\hat{\tau}_k(x_i)$  is the estimate of  $\tau_k(x_i)$ ,  $w_i$  is a predefined weight for each customer, g(.) is a convex function of the policy  $\Pi$  and c is a predefined constant. For example, there can be budget constraint on the number of customers assigned to each treatment i.e.,  $g_k(\Pi) = \sum_i \pi_{ik} \leq c_k$  for a treatment k.

#### 2.2 Evaluation Metrics

*Uplift Model Evaluation.* A common evaluation metric for uplift models is the area under the cumulative uplift curve [9, 12]. We rank the customers by their predicted uplift from the corresponding causal model. The cumulative uplift for a customer ranked  $r^{th}$  by the predicted uplift, can be defined as:

$$\text{Cumulative Uplift}_{r} = \left(\frac{\sum_{i=1}^{r} Y_{i} \mathbb{1}\{\pi_{i0} = 0 \& T_{i} \neq t_{0}\}}{\sum_{i=1}^{r} \mathbb{1}\{\pi_{i0} = 0 \& T_{i} \neq t_{0}\}} - \frac{\sum_{i=1}^{r} Y_{i} \mathbb{1}\{\pi_{i0} = 1 \& T_{i} = t_{0}\}}{\sum_{i=1}^{r} \mathbb{1}\{\pi_{i0} = 1 \& T_{i} = t_{0}\}}\right) * \frac{r}{|C|}$$

where  $Y_i$  is the observed outcome,  $t_0$  is the control treatment and  $\pi_{i0} = 1$ , if customer *i* is assigned treatment  $t_0$  by policy  $\Pi$  else 0. We can compute the area under the curve (AUC), where the larger AUC value indicates a better model.

*Offline Policy Evaluation.* The observed data comes from randomized experiments (logging policy) that maps customers to treatments. Assuming that the logging policy is customer independent, we compare our proposed policy,  $\Pi$ , with different targeting policies using the following counterfactual policy evaluation metrics:

(1) Inverse Propensity Score [13]: 
$$\frac{1}{N} \left( \sum_{k=0}^{K} \frac{\sum_{i=1}^{N} Z_i \mathbb{1}\{\pi_{ik} = 1 \& T_i = t_k\}}{p_k} \right)$$
(2) Self-Normalized IPS [14]: 
$$\frac{\sum_{k=0}^{K} \frac{\sum_{i=1}^{N} Z_i \mathbb{1}\{\pi_{ik} = 1 \& T_i = t_k\}}{\sum_{k=0}^{K} \frac{\sum_{i=1}^{N} \mathbb{1}\{\pi_{ik} = 1 \& T_i = t_k\}}{p_k}}$$

where  $Z_i$  can be the outcome of interest (e.g., revenue or completion rate) and  $p_k = 1/K, k = 0, 1, ..., K$ .

Depending upon the application, the outcome of interest,  $Z_i$ , might be different. Typically, it is the same as the objective for uplift models i.e.,  $Z_i = Y_i$  but sometimes, we may be interested in evaluating more than one business objective. For example, when targeting customers with discount offers, we would want a targeting policy that increases net revenue while maintaining customer experience and in this case,  $Z_i$  can reflect revenue or customer experience. We discuss specific business metrics when we study different applications of this problem.



Fig. 1. "True" uplift based on simple difference across different retention scores.

### **3 APPLICATIONS**

We leverage data from large-scale experiments that support different business purposes to test the proposed methodology. We assess the potential impact of given treatments via offline policy estimates. The experiments took the form of an online randomized A/B test, in which the treatment group was given the said treatment and the control group was not. In all experiments, we have binary treatment levels i.e.,  $\{t_0, t_1\}$  or  $\{0, 1\}$  for simplicity and estimate the CATE for each customer  $i \in C$  as  $\tau(x_i) = \mathbb{E}[Y_i(1) | X_i = x_i] - \mathbb{E}[Y_i(0) | X_i = x_i]$  using uplift models mentioned in Section 2.1. All policies are compared based on this randomized experiment data via offline policy evaluation presented in Sec 2.2. We discuss the specific details and outcome of each experiment below.

### 3.1 Increase Customer Retention

In this scenario, we aim to identify the customers to target with an intervention that improves customer retention. Consider the business case of a company offering paid services aiming to enhance customer engagement and loyalty. The company defines churn as the absence of any service payment within the last month. To mitigate churn, the company intends to identify at-risk customers and engage them proactively with messages promoting additional services. One approach is to predict each customer's likelihood to continue using the services, termed as their retention score. Customers with lower retention scores (specifically below 0.391) are selected for targeted interventions based on this prediction model. While the company has meticulously selected the 0.391 threshold through historical data analysis, relying solely on this retention score threshold for targeting interventions may not be optimal.

To explore customer reaction to such interventions, we ran a randomized experiment where customers in the treatment group received a messaging template encouraging them to learn more about the services offered by the company. We then clustered the customers into 100 buckets of equal size in an incremental manner going from minimum value of retention scores to maximum and calculated the "true" uplift,  $\bar{\tau}^b$ , for each bucket b as  $\bar{\tau}^b = \frac{\sum_{i \in b} Y_i \mathbb{1}\{T_i=1\}}{\sum_{i \in b} \mathbb{1}\{T_i=1\}} - \frac{\sum_{i \in b} Y_i \mathbb{1}\{T_i=0\}}{\sum_{i \in b} \mathbb{1}\{T_i=0\}}$ . Figure 1 shows this "true" uplift based on simple difference in means along the retention score percentiles. We observe, customers with higher retention scores (> 0.6) have close to 0 uplift, meaning that they are not affected by the targeting message and will probably continue to stay nonetheless. However, the customers with lower scores tend to have negative uplift, which means that the targeting message is doing more harm than good. This shows that the targeting policy based simply on retention score threshold is not optimal for business.

To define optimal targeting segment, we built uplift models using the data from the randomized experiment data. The outcome Y is binary i.e., 1 if the customer does not churn else 0. Intuitively, covariate X for uplift model can be the retention score, as it directly captures the likelihood of the desired outcome. However, sometimes a single variable may not be able to capture the variability in the customers. So, we also explore using the top customer features that were used to build the retention model as the set of covariates, X. Since higher uplift means more treatment impact, we aim to target customers with the higher uplift, keeping the proportion of customers to be targeted as low. In this case, the constraint can be to restrict the proportion of customers to be targeted by c, if there is a budget constraint. This would essentially mean targeting the top  $c \times N$  customers with highest uplift. However, we do not specify c explicitly but rather look for proportion of customers with positive uplift, that can be targeted by different policies. Thus, we do not need a constraint optimization in this case. The proportion of customers we target based on a policy will be referred to as the targeting proportion hereafter. The end goal of targeting customers is to increase retention so we choose the policy with maximum IPS and SNIPS corresponding to  $Z_i = Y_i$ . Based on uplift model evaluation, we propose a policy based on Causal Forest trained on customer features (see Table A.4 in Appendix). Just by targeting 6.5% of the population, our proposed policy gets a relative lift of 2.35% in retention rate when compared to the old policy that targets based on retention score, a lift of 3.905% when compared to not targeting anyone and a lift of 2.79% in retention rate when compared to targeting everyone (see Table 1).

Table 1. Relative lift in metrics by the proposed policy over baseline policies for customer retention.

	Lift in Retention Rate			
Baseline Policy	IPS	SNIPS		
Retention score < 0.391	2.35%	0.647%		
Targeting no one	3.905%	1.302%		
Targeting everyone	2.79%	0.215%		

#### 3.2 Maximize Event Revenue

In this business scenario, the objective is to boost customer participation in shopping events and promotional activities. Often times e-commerce websites run deal events, where customers are invited to make purchase over some spend threshold to get certain rewards back on their order. We ran a marketing experiment at a commercial scale, where we offered  $P_1$  or  $P_2$  rewards back  $(P_1 < P_2)$  if the customers purchased goods over a fixed amount within a time frame. They randomly assigned  $P_1$  or  $P_2$  rewards back offer to customers, which wasn't optimal given business goals. Ultimately, the goal was to minimize a population-level efficiency metric defined as, e%iS =  $\frac{Average Reward Expense}{Average Sale - Baseline Average Sale}$ , where average sale is attributed to the purchases that happened during the days of the event and baseline average sale refers to the sales had the campaign not run. Since the company had to let everyone participate in the promotion, this baseline average sale was counter-factually calculated. While the goal is to target customer with a suitable offer that improves customer experience and increases revenue, we do not want to deteriorate sales by more than 1% compared to assigning maximum reward (i.e.,  $P_2$ ) for everyone. The customers have a better experience when they participate in these campaigns and are able to complete the offer by purchasing goods worth the defined amount within the allocated time frame and receive the rewards.

Here, control denotes giving  $P_1$  and treatment denotes giving  $P_2$  rewards. Let Sales<sub>i</sub> be during the deal event and Rewards<sub>i</sub> be rewards for customer  $i \in C$ . In order to maximize the overall revenue while controlling the cost, we define Manuscript submitted to ACM

	Lift	in Revenue	Lift in C	Completion Rate	
Baseline Policy	IPS	SNIPS	IPS	SNIPS	e%iS
Completion score < 0.15	0.317%	0.254%	1.02%	1.02%	-11.879%
Only $P_1$	0.993%	0.934%	2.062%	2.062%	-55.368%
Only $P_2$	0.724%	0.665%	-1.980%	-1.980%	-11.879%

Table 2. Relative lift in metrics by the proposed policy over baseline policies for revenue maximization.

outcome *Y* as  $Y_i$  = Sales<sub>*i*</sub> – Rewards<sub>*i*</sub>. Similar to the application in section 3.1, we train uplift models based on (1) the propensity of a customer to complete the offer (referred as completion score hereafter), and (2) customer features. Assuming each customer is valued the same i.e.,  $w_i \equiv 1 \forall i \in C$ , the optimal policy  $\Pi$  can be found such that,

$$\max_{\pi} \sum_{i=1}^{N} \pi_{ik} \hat{\tau}(x_i), \text{ subject to } \pi_{ik} \in \{0,1\} , \ \sum_{k=0}^{1} \pi_{ik} = 1 , 1 - \frac{\text{Sales}_{\Pi}}{\text{Sales}_{\Pi^*}} \le 0.01$$

where  $\operatorname{Sales}_{\Pi} = \frac{\sum_{k=0}^{1} \sum_{i=1}^{N} Sale_{i} \mathbbm{1}\{\pi_{ik}=1 \& T_{i}=k\}}{\sum_{i=1}^{1} \sum_{i=1}^{N} \sum_{i=1}^{N} \mathbbm{1}\{\pi_{ik}=1 \& T_{i}=k\}}$  denotes the average expected sale under the targeting policy  $\Pi$  and  $\operatorname{Sales}_{\Pi}^{*} = \frac{\sum_{i=1}^{N} Sale_{i} \mathbbm{1}\{T_{i}=1\}}{\sum_{i=1}^{N} \mathbbm{1}\{T_{i}=1\}}$  denotes the average expected sale when giving all customers the maximum-reward treatment  $(P_{2})$ . The constraint here restricts the deterioration of sales by more than 1% compared to assigning maximum reward for everyone. Since the population is large, the optimization is time consuming. To make it scalable, we create equal-sized groups of customers based on the predicted uplift  $\hat{\tau}$  and then optimize over these groups, treating each group as an individual. Consequently, we assign a treatment to a group i.e., everyone in a group receives the same treatment.

We considered a targeting strategy where we can offer higher rewards ( $P_2$ ) to the customers with low propensity to complete the offer so that they're more motivated to shop. Since goal of this campaign is to increase revenue, minimize e%iS and improve customer experience, we compare our proposed policy with the aforementioned old policy based on these metrics. Based on offline evaluation on the randomized experiment data, we recommend the policy associated with Causal Forest DML trained on completion score (see Table A.5 in Appendix). Our recommended policy outperforms other policies in terms of revenue and e%iS, while having a better completion rate than the baseline targeting policy associated with completion scores threshold, see Table 2.

### 4 ONLINE EXPERIMENT: OPTIMAL OFFER FOR REVENUE MAXIMIZATION

Sometimes, customers are offered rewards when they make purchases over a certain amount in shopping events. The objective is to allocate offer optimally to maximize the total revenue while maintaining customer experience. In these shopping events, two thresholds for the purchase amount, say  $S_1$  and  $S_2$ , were randomly assigned to customers where  $S_2 > S_1$ . An intuitive hypothesis was that lowering the spend threshold will incentivize customers to complete the offer and spend more on average. However, previous campaign data showed that the treatment with  $S_1$  spend threshold had 0.30% lower average revenue as compared to  $S_2$  spend threshold, while more customers completed the offer with  $S_1$ . This result motivated us to make optimal offer allocations to customers that maximizes the total revenue from the event, while maintaining customer experience.

Here, T = 0 denotes giving spend threshold of  $S_2$  and T = 1 refers to giving spend threshold of  $S_1$ . Our outcome is Y defined as  $Y_i = \text{Sales}_i - \text{Rewards}_i$  and the covariate X can be the completion score or the customer features. We aim to target everyone with a positive uplift and hence, there is no constraint optimization in this case. Based on offline evaluation, we propose a policy using previous campaign data based on Causal Forest trained on customer features (see Manuscript submitted to ACM

0.022%

0.214%

2.225%

	]	Lift in Revenue	Lift in Completion Rate	
Baseline Policy	IPS	SNIPS	IPS	SNIPS

-0.204%

0.008%

2.015%

Table 3. Relative lift in metrics by the proposed policy when compared with different baseline policies for assigning spend thresholds to customers to receive rewards.

-3.017%

2.506%

-12.494%

-3.287%

2.221%

-12.737%

Table A.6 in Appendix) because it outperforms all baseline policies in terms of revenue (see Table 3). The proposed
policy only assigns the lower spend threshold, $S_1$ , to 12% of the population while the policy based on completion score
assigns it to 39% of the population. Note that giving customers $S_1$ purchase threshold, is more costly than offering $S_2$
purchase threshold. Our policy has better completion rate than if we were to offer only $S_2$ threshold and has better
revenue than if we were to offer only $S_1$ threshold.

We then conducted a large-scale commercial experiment to assess the effectiveness of our proposed policy. We compared the performance of our recommended targeting policy with a strong static baseline policy of offering  $S_2$  spend threshold to all customers based on past campaign (see Table A.7 in Appendix). In this randomized experiment, customers in control (C) group received  $S_2$  threshold only, while the treatment (T) group customers were offered either  $S_1$  or  $S_2$  spend threshold based on our proposed policy. Our proposed policy recommended  $S_1$  spend threshold to about 11% of the customers in the treatment group. The treatment group achieved statistically significant lift in revenue (0.36%; p = 0.040) and completion rate (5.49%; p = 0.000).

## 5 CONCLUSION

Completion score < 0.5

S<sub>2</sub> Only

 $S_1$  Only

This work provides a principled approach to customer targeting scenarios in e-commerce industry. Based on uplift modeling and constrained optimization, we propose a generalized targeting policy framework for allocating treatments to customers such that the outcome of interest is optimized while taking account of any given business constraints. We demonstrate improvement over baseline targeting approaches in two business applications using offline evaluations and validate our proposed policy via a large-scale online experiment. This framework enables businesses to personalize their marketing campaigns thereby improving value to the business and customer experience.

In all applications, we had binary treatment levels. However, in some cases, we may have infinite continuous treatment levels. For example, one may want to offer customers x% points, somewhere between say, 5% to 15%. There are a few estimation methods, such as double/debiased machine learning [4], that support estimation of continuous treatment effect. With such models, this framework can be naturally extended to the case of continuous treatment levels. Our framework is also applicable to observational studies, given the unconfoundedness assumption, that is, we will need to measure and control for all potential confounders. When the unobserved confounding is a concern, other methods such as proximal causal inference[15] and sensitivity analysis [5, 19] can be used to make reliable causal conclusions.

### REFERENCES

- Javier Albert and Dmitri Goldenberg. 2022. E-commerce promotions personalization via online multiple-choice knapsack with uplift modeling. In Proceedings of the 31st ACM International Conference on Information & Knowledge Management. 2863–2872.
- [2] Eva Ascarza. 2018. Retention futility: Targeting high-risk customers might be ineffective. Journal of Marketing Research 55, 1 (2018), 80-98.
- [3] Susan Athey, Julie Tibshirani, and Stefan Wager. 2019. Generalized random forests. (2019).
- [4] Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. 2018. Double/debiased machine learning for treatment and structural parameters.
- [5] Victor Chernozhukov, Carlos Cinelli, Whitney Newey, Amit Sharma, and Vasilis Syrgkanis. 2022. Long story short: Omitted variable bias in causal machine learning. Technical Report. National Bureau of Economic Research.
- [6] Floris Devriendt, Jeroen Berrevoets, and Wouter Verbeke. 2021. Why you should stop predicting customer churn and start using uplift models. Information Sciences 548 (2021), 497–515.
- [7] Shuyang Du, James Lee, and Farzin Ghaffarizadeh. 2019. Improve User Retention with Causal Learning. In The 2019 ACM SIGKDD Workshop on Causal Discovery. PMLR, 34–49.
- [8] Dmitri Goldenberg, Javier Albert, Lucas Bernardi, and Pablo Estevez. 2020. Free lunch! retrospective uplift modeling for dynamic promotions recommendation within roi constraints. In Proceedings of the 14th ACM Conference on Recommender Systems. 486–491.
- [9] Pierre Gutierrez and Jean-Yves Gérardy. 2017. Causal inference and uplift modelling: A review of the literature. In International conference on predictive applications and APIs. PMLR, 1–13.
- [10] Miruna Oprescu, Vasilis Syrgkanis, and Zhiwei Steven Wu. 2019. Orthogonal random forest for causal inference. In International Conference on Machine Learning. PMLR, 4932–4941.
- [11] Donald B Rubin. 2005. Causal Inference Using Potential Outcomes. J. Amer. Statist. Assoc. 100, 469 (2005), 322-331.
- [12] Piotr Rzepakowski and Szymon Jaroszewicz. 2012. Decision trees for uplift modeling with single and multiple treatments. *Knowledge and Information Systems* 32 (2012), 303–327.
- [13] Adith Swaminathan and Thorsten Joachims. 2015. Counterfactual risk minimization: Learning from logged bandit feedback. In International Conference on Machine Learning. PMLR, 814–823.
- [14] Adith Swaminathan and Thorsten Joachims. 2015. The self-normalized estimator for counterfactual learning. advances in neural information processing systems 28 (2015).
- [15] Eric J Tchetgen Tchetgen, Andrew Ying, Yifan Cui, Xu Shi, and Wang Miao. 2020. An introduction to proximal causal learning. arXiv preprint arXiv:2009.10982 (2020).
- [16] Stefan Wager and Susan Athey. 2018. Estimation and inference of heterogeneous treatment effects using random forests. J. Amer. Statist. Assoc. 113, 523 (2018), 1228–1242.
- [17] Amulya Yadav, Roopali Singh, Nikolas Siapoutis, Anamika Barman-Adhikari, and Yu Liang. 2021. Optimal and non-discriminative rehabilitation program design for opioid addiction among homeless youth. In Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence. 4389–4395.
- [18] Zhenyu Zhao and Totte Harinen. 2019. Uplift modeling for multiple treatments with cost optimization. In 2019 IEEE International Conference on Data Science and Advanced Analytics (DSAA). IEEE, 422–431.
- [19] Jiajing Zheng, Jiaxi Wu, Alexander D'Amour, and Alexander Franks. 2023. Sensitivity to Unobserved Confounding in Studies with Factor-structured Outcomes. J. Amer. Statist. Assoc. (2023), 1–12.

## A APPENDIX

# A.1 Offline Application Results

			<b>Retention Rate</b>		
Policy	Targeting proportion	AUC	IPS	SNIPS	
	Based on reten	tion scores			
FDR Learner	4.22%	0.00231	93.2%	93.2%	
Causal Forest	6.29%	0.00226	93.7%	93.2%	
Causal Forest DML	34.69%	0.00200	92.4%	93.0%	
Based on customer features					
FDR Learner	0.48%	0.00344	93.3%	93.2%	
<b>Causal Forest</b>	6.47%	0.00355	95.8%	93.4%	
Causal Forest DML	30.57%	0.00335	92.6%	93.1%	

Table A.4. The targeting proportion and offline policy evaluation metrics for the different policies for member retention.

Table A.5. The targeting proportion, AUC and offline policy evaluation metrics for the different policies for the deal with different points back offer.

		Revenue			
Policy	Targeting proportion	AUC	IPS	SNIPS	e%iS
	Based on	completion s	cores		
FDR Learner	8.83%	2.371	1402.91	1402.42	88.35%
Causal Forest	94.48%	0.938	1411.67	1410.52	48.41%
Causal Forest DML	67.47%	4.373	1417.27	1416.44	35.46%
	Based on	customer fea	tures		
FDR Learner	97.85%	-0.023	1412.02	1412.17	45.20%
Causal Forest	10.91%	4.670	1408.70	1409.03	52.22%
Causal Forest DML	99.49%	-0.673	1412.35	1412.01	45.30%

Table A.6.	The targeting propo	rtion, AUC and offlin	e policy evaluation	n metrics for the	different policie	s for the deal v	with different
spend three	esholds.						

			Reve	enue	
Policy	Targeting proportion	AUC	IPS	SNIPS	
Based on completion scores					
FDR Learner	34.49%	5.748	4,646.53	4,643.30	
Causal Forest	28.75%	29.455	4,686.02	4,678.98	
Causal Forest DML	33.85%	10.144	4,647.33	4,646.82	
Based on customer features					
FDR Learner	12.79%	21.788	4,680.94	4,662.42	
<b>Causal Forest</b>	11.93%	35.804	4,698.25	4,688.58	
Causal Forest DML	11.70%	22.566	4,676.99	4,646.82	

## A.2 Online Experiment Results

Table A.7. Past experiment results comparing static spend threshold of  $S_2$  (C) vs  $S_1$  (T).

Treatment	Average Revenue	Completion Rate
<i>S</i> <sub>2</sub> (C)	1,477.98	0.89%
<i>S</i> <sub>1</sub> (T)	1,474.32	1.05%
Lift (T-C)	-0.25%	18.18%
p-value (T-C)	0.098	0.000

We deep-dive further into how our proposed policy might have performed if we were to assign every one, whether in control or treatment groups, the policy recommended threshold of  $S_1$  or  $S_2$  (see Table A.8). The completion rate of customers who are recommended  $S_2$  threshold, is not very different (2.75% uplift) for control and treatment groups. However, the completion rate of customers who are recommended  $S_1$  threshold, is very different (18.37% uplift) for control and treatment groups. This could be attributed to the fact that these customers were much more likely to complete the offer when given the  $S_1$  spend threshold.

Table A.8. Difference in completion rates for customers who the model recommended  $S_1$  vs  $S_2$  spend threshold in control and treatment groups

Threshold	Group	# Customers Triggered	Completion Rate	Completion Rate Uplift
<i>S</i> <sub>1</sub>	С	1,051,297	0.98%	-
$S_1$	Т	527,663	1.16%	18.37%
$S_2$	С	8,197,866	0.73%	-
$S_2$	Т	4,103,641	0.75%	2.74%