Causal Machine Learning for Promotions: Industry Evidence and Applications

Muxi Xu
Muxi.Xu@doordash.com
DoorDash Inc.
New York City, New York, USA

Kun Hu
Kun.Hu@doordash.com
DoorDash Inc.
New York City, New York, USA

Sudeep Das
sudeep.das2@doordash.com
DoorDash Inc.
San Francisco, California, USA

Bruce Wang
b.wang@doordash.com
DoorDash Inc.
New York City, New York, USA

ABSTRACT

Promotional campaigns can grow a user base but often lead to overspending when not tailored to individuals. Using DoorDash data, we show that causal modeling improves efficiency by identifying users who are most responsive to deeper discounts and selecting cost-effective offers. We further personalize discount amounts to maximize incremental impact within a fixed promotional budget. Offline and online tests demonstrate significant gains in user adoption and cost efficiency.

CCS CONCEPTS

• **Applied computing → Marketing**; • **Information systems → Online shopping**; • **Mathematics of computing → Causal networks**; • **Computing methodologies → Causal reasoning and diagnostics**.

KEYWORDS

Causal Inference, Recommendation System, Promotion Strategy

ACM Reference Format:
Muxi Xu, Kun Hu, Sudeep Das, and Bruce Wang. 2025. Causal Machine Learning for Promotions: Industry Evidence and Applications. In *Proceedings of KDD 2025 Workshop - Causal Inference and Machine Learning in Practice (KDD 2025).* ACM, New York, NY, USA, 6 pages. https://doi.org/XXXXXXX.XXXXXXX

1 INTRODUCTION

Promotional campaigns, whether percentage discounts (e.g., "X % off orders over $ Y") or perks such as free delivery, remain one of the most powerful levers for accelerating growth in e-commerce. By lowering transaction costs or adding perceived value, promotions entice first-time buyers, nudge fence-sitters to convert, and reactivate dormant consumers. The downside is cost: a blanket 30% discount shown to millions of users can erode margins faster than it drives incremental orders. Although many platforms attempt to mitigate this with simple business rules ("show a coupon only to new users" or "cap the daily promo budget"), such heuristics do not account for the large, user-specific variation in promo sensitivity and therefore leave both revenue and consumer experience on the table.

DoorDash uses promotions to encourage consumers to try new services or order more frequently. Figure 1 illustrates an example of a discount offered to DoorDash consumers. In practice, DoorDash may run hundreds of active promotions every day.

We hypothesize that consumers exhibit varying degrees of sensitivity or "appetite" for these promotions. By personalizing promotional offers, both in terms of *who* receives them and *how much* to offer, we can deliver a more satisfying consumer experience while improving cost efficiency for the business. This motivates us to use causal inference frameworks such as [8] and [9] , which aim to estimate the true incremental effect of an intervention. Specifically, we use causal machine learning to estimate the true incremental effect of a given promotion on each user, and couple it with an optimization layer that selects the best action under budget constraints. A core innovation of our work is to provide guidance on *when to model promotions as discrete or continuous treatment variables*, and how to account for the budget constraints. We show that this modeling choice depends on the campaign mechanics:

- When the treatment is a fixed-form offer (e.g., "free delivery" or "30% off for X days"), the key decision is targeting: determining *who* should receive the promotion. In these cases, it is appropriate to treat the treatment variable as discrete.
- When a campaign offers many tunable incentive levels (e.g., varying percent- or dollar-off) that meaningfully affect outcomes, it is best to model the treatment as continuous. This allows the model to capture price sensitivity and optimize discount depth.

Our framework is validated across two real-world case studies: one focusing on personalized targeting for fixed campaigns, and the other on personalized discount depth. These examples illustrate how different problem formulations motivate different causal learning strategies.

We summarize the key contributions of this paper as follows:

- **Modeling Framework for Discrete vs. Continuous Treatments** We provide practical criteria and empirical evidence

KDD 2025, August 3–7, 2025, Toronto, ON, Canada

ACM ISBN 978-1-4503-XXXX-X/2018/06
https://doi.org/XXXXXXX.XXXXXXX

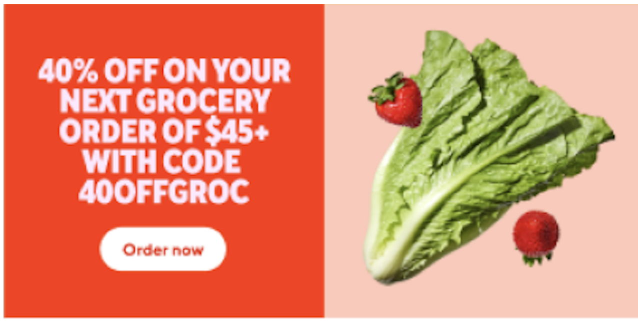

Figure 1: Example of a discount-based promotion shown to DoorDash consumers.

to guide the choice of treatment formulation depending on the structure of the promotion.

- **Real-World Validation of Causal Models** We compare *S*-, *T*-, and *DML*-learners offline and validate the selected model in a large-scale online A/B test, showing measurable gains in order rate and cost efficiency improvement.
- **End-to-End Optimization** We couple the causal estimates with a cost-aware optimizer that assigns the best promotion to each user, subject to global budget constraints.
- **Operational Learnings** We address implementation challenges such as repeated promo re-entry, campaign re-targeting, and offline-to-online consistency, and demonstrate how the system scales across campaigns and verticals.

2 METHODOLOGY

We divide our approach into two key stages: first, we estimate a consumer's response to a given promotion, and second, we use these estimates to design promotions that maximize the desired business outcome.

2.1 Causal response estimate

Our first objective is to estimate how promotions affect consumer behavior, specifically each user's incremental purchase probability. Framing promotions as treatments, we apply a causal inference approach to estimate their effect on outcomes to reduce confounding and isolate true impact.

2.1.1 Modeling Targets. In practice, modeling promotion effectiveness is more complex than standard binary treatment setups. Real-world promotions vary not only in whether they are shown (binary), but also in type (e.g., order vs. delivery discount) and intensity (e.g., 10%, 20%, 40% discounts). To accommodate this, we classify modeling targets into two broad categories:

- **Discrete (Binary / Multi-arm) Promotions.** In this setup, each user receives either no promotion or one of several mutually exclusive promotional variants (e.g., *no promo*, *10% off*, *non-monetary reward*). This structure naturally lends itself to binary or multi-class treatment modeling. Formally, each user i has a set of potential outcomes $Y_i(t)$ for treatment $t \in \{0, 1, \ldots, K\}$, where $t = 0$ denotes no promotion and $t = k$ indicates the k-th promotion variant. The vector x_i denotes the observed features for user i. The Conditional Average Treatment Effect (CATE) in the binary case is:

$$\tau(x_i) = \mathbb{E}\left[Y_i(1) - Y_i(0) \mid x_i\right]$$

For multi-arm settings, we generalize this to arm-specific CATE:

$$\tau_k(x_i) = \mathbb{E}\left[Y_i(k) - Y_i(0) \mid x_i\right], \quad k = 1, \ldots, K$$

- **Continuous Promotions.** Some campaigns define discount levels as a continuous or ordinal variable (e.g., 0%, 10%, 20%, 30%). In these settings, the treatment variable T is numeric and represents varying *depths* of a single promotion type. Although it is possible to model such treatments as discrete categories, this approach discards the inherent structure of the action space and ignores potential smoothness or monotonicity across treatment levels. Instead, by modeling T as a continuous variable, we can use regression-based uplift models or estimate treatment-response curves $\tau(t, x)$ that characterize how the expected outcome evolves with treatment intensity:

$$\tau(t, x_i) = \frac{\partial}{\partial t}\mathbb{E}\left[Y_i \mid T = t, x_i\right]$$

In practice, however, treatment effects of this general form are not identifiable without further assumptions. To enable tractable estimation, we often impose structure on $\tau(t, x)$, for instance, by assuming linearity in t by defining $\tau(t, x) = \tau(x)t$. Such assumptions make it feasible to apply frameworks like Double Machine Learning. Beyond linear specifications, recent work has introduced more flexible approaches. For example, Zhang et al. [12] use Radial Basis Functions (RBFs) to encode the action space, enabling smoother generalization across discount levels while preserving model interpretability.

In our work, we support both setups. We treat discrete variants with multi-arm treatment learners and use ordinal encodings or bin-wise uplift models to approximate continuous treatment effects.

2.1.2 Causal Estimation Methods. To estimate $\tau_k(x)$ or $\tau(t, x)$, we evaluated several general causal inference approaches. These methods follow the frameworks introduced by Künzel et al. [6] and Chernozhukov et al. [4]. Each method is trained on internal promotion logs and benchmarked based on uplift evaluation metrics.

S-Learner. A single model $f_S(X, T)$ is trained on pooled data. The treatment indicator T is passed as an input feature. Individual uplift is computed as:

$$\hat{\tau}(X) = f_S(X, 1) - f_S(X, 0)$$

T-Learner. Two separate models are trained—$f_1(X)$ on treated users and $f_0(X)$ on controls. The treatment effect is estimated as:

$$\hat{\tau}(X) = f_1(X) - f_0(X)$$

Double Machine Learning (DML). A two-stage approach that controls for confounding by residualizing both outcome and treatment:

$$\mu_y(X) \approx \mathbb{E}[Y \mid X], \quad \mu_t(X) \approx \mathbb{P}(T = 1 \mid X)$$
$$\widetilde{Y} = Y - \mu_y(X), \quad \widetilde{T} = T - \mu_t(X)$$
$$\hat{\tau}(X) = \text{Regress } \widetilde{Y} \text{ on } \widetilde{T}$$

DML generalizes to both binary and multi-arm treatments and is particularly robust under non-randomized exposure.

2.2 Promotion Optimization

Once we estimate the user-level uplift for each promotion arm, we aim to assign the optimal promotion to each user in order to maximize total incremental adoption while accounting for the cost of delivery and business constraints such as budget or eligibility.

2.2.1 Optimization Objective. We define a reward function that combines uplift and cost:

$$\text{Reward}(i, k) = \tau_k(x_i) - \lambda \cdot c_k$$

where:

$$\begin{aligned} x_i &: \text{Feature vector for user } i \\ \tau_k(x_i) &: \text{Estimated uplift for user } i \text{ under promotion } k \\ c_k &: \text{Expected cost of applying promotion } k \\ \lambda &: \text{Trade-off parameter between uplift and cost} \end{aligned}$$

We then select the optimal promotion using:

$$\pi^*(x_i) = \arg\max_{k \in \mathcal{K}} \text{Reward}(i, k)$$

where $\mathcal{K}$ includes all available promotion arms, including $k = 0$ for no promotion.

2.2.2 Constraints and Deployment Strategy. In production, the optimization must satisfy several constraints, including a GLOBAL BUDGET constraint ($\sum_i c_{\pi(i)} \le \mathcal{B}$), QUOTA LIMITS on certain promotion arms, and ELIGIBILITY RULES that restrict offers to specific users, geographies, or merchant categories.

2.2.3 Continuous Discount Handling. When discount levels vary continuously (e.g., 10%, 20%, 30%), we treat promotion depth as a real-valued variable t and optimize:

$$t_i^* = \arg\max_{t \in [t_{\min}, t_{\max}]} \left[\tau(t, x_i) - \lambda \cdot c(t)\right]$$

where $\tau(t, x_i)$ is the estimated uplift and $c(t)$ the expected cost. When t takes a few discrete values with distinct effects, the problem can be transformed to a discrete treatment case.

2.2.4 Final Policy Deployment. The final policy $\pi(x_i)$ is calculated offline and applied offline in real time, subject to budget and other constraints. It is periodically updated as new data or goals emerge.

The details of how we solve the constrained optimization problems are presented in Sections 3.1 and 3.2. In general, our solution draws inspiration from budget-aware optimization frameworks developed in the context of e-Commerce and advertising. For example, Tang and Wang [10] propose algorithms to maximize promotional return under budget constraints using marginal ROI ranking, while Bottou et al. [3] frame ad assignment as a counterfactual policy optimization problem using offline data.

3 EMPIRICAL STUDY

Uplift modeling techniques in marketing have a long history, with the early work of Lo [7] proposing the true lift model as an alternative to traditional response modeling for campaign targeting. In this section, we demonstrate how our causal modeling framework is applied in real-world marketing scenarios at DoorDash. We present two case studies that reflect common promotion challenges: (1) promotion (re)engagement for non-restaurant trial promotions, and (2) personalized discount depth for restaurant consumers. Each case illustrates how we select modeling targets, estimate uplift using causal learners, and optimize promotions under real-world constraints. We evaluated both offline metrics (e.g., Qini score, uplift@k) and online A/B tests to validate the end-to-end effectiveness of our approach.

3.1 Personalized Discount Targeting for Non-Restaurant Delivery

To grow the user base and improve retention, we launched a set of (re)engagement promotions to encourage consumers to try non-restaurant deliveries at DoorDash (e.g., grocery, household items, or personal care) by offering relevant promotions shortly after their restaurant order. Historically, we tested several predefined strategies:

- **Cross-vertical promo** — a fixed percent-off incentive usable across all Non-restaurant categories
- **Vertical-specific promo** — the same discount structure, but restricted to certain verticals such as grocery or retail
- **Delivery fee discount** — waived delivery fees for eligible non-restaurant orders

All offers could be single- or multi-use (subject to redemption limits) and were valid within a restricted time window. These campaigns aim to determine whether encouraging exploration of other non-restaurant use cases drives stronger incremental adoption and higher order frequency. Crucially, the promotion content and incentive levels were defined in advance; our task was to identify which users should receive each promotion. Thus, the core modeling challenge was one of **targeting**—deciding who should be exposed to each pre-specified offer—not one of optimizing promotion type or amount.

3.1.1 Causal Modeling Targets. We frame this as a treatment effect estimation problem under a fixed promotion policy. For each promotion variant, we model the incremental effect of showing the offer versus withholding it. Although the campaign included multiple discount tiers (e.g., varying across consumer segments), we observed that:

(i) Adoption rates were broadly similar across different discount levels in historical data
(ii) The available data was too sparse to support fine-grained modeling

As a result, we consolidate offers into two broad categories, *dollar off* and *no delivery fee*—and model each as a discrete treatment arm. This simplifies the task to estimating a multi-arm Conditional Average Treatment Effect (CATE) per user:

$$\tau_k(x) = \mathbb{E}[Y_i(k) - Y_i(0) \mid x_i], \quad k \in \{\text{dollar off, no delivery fee}\}$$

3.1.2 Causal Modeling Approach. As discussed in Section 3.1.2, we evaluated several causal learners, including the S-Learner, T-Learner, and Double Machine Learning (DML). For the T-Learner, we tested both the standard variant and a Category Transformer

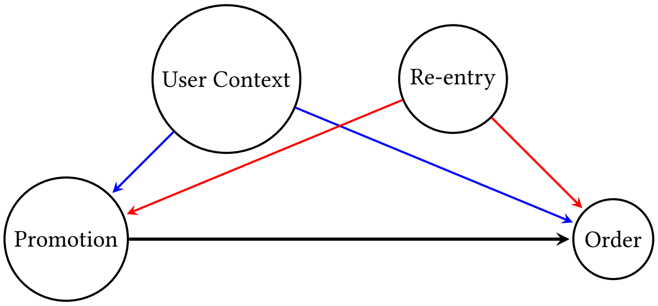

Figure 2: Campaign Model Causal Diagram

adaptation, as detailed in Section 3.1.4. The final model was trained on historical campaign logs and configured to estimate multi-arm treatment effects.

3.1.3 Training Data and Segmentation Strategy. We train our models on internal impression-level logs, where each record captures a user's exposure to a specific promotion. Campaigns are typically launched uniformly to eligible users, with a small randomized hold-out group as control. The data includes features on user behavior, order history, vertical engagement, and promo response. We find that model performance is highly sensitive to how training data is constructed, particularly by user lifecycle stage segmentation and exposure history treatments.

Segmented Training by Consumer Type. To encourage generalization, we initially pool data across many historical category-level campaigns. However, we find that indiscriminately combining users across different lifecycle stages (e.g., new vs. existing) can degrade out-of-sample uplift prediction, even when stage indicators are included. Behaviorally heterogeneous data often introduces confounding that overwhelms model capacity. Instead, we segment training by consumer type such as *new users*, based on order recency and frequency. Models trained on more homogeneous cohorts yield more accurate treatment effect estimates and more stable uplift rankings.

Handling Re-entry Events. Repeated exposures pose another modeling challenge: users may re-enter the same campaign after a cooldown period (e.g., 28 days), creating correlated impressions that bias estimated effects. We evaluate three strategies:

- **First-exposure only:** Retain only first impression per user campaign, assuming the first touch has the strongest effect.
- **Per-entry exposure:** Treat each exposure as a separate sample, and include re-entry context features (e.g., number of prior exposures).
- **Sequential modeling:** Treat reentries as a multi-step process using sequence models or reinforcement learning.

For this promotion campaign, we adopt the **per-entry** strategy to retain more training data while accounting for re-entry dynamics through derived features.

The final causal model structure is visualized in the causal diagram (figure 2).

Table 1: Offline Model Performance (Qini, uplift@30)

Model	OutSample		InSample	
	Qini	Uplift@30	Qini	Uplift@30
S-Learner	0.49%	0.77%	31.81%	1.91%
T-Learner	0.09%	0.12%	1.13%	0.18%
Cat Transformer	1.18%	0.19%	8.31%	0.62%
DML	8.03%	1.11%	13.73%	0.77%

3.1.4 Causal Model Offline Result. We evaluated four causal modeling approaches using historical campaign data, measuring both in-sample and out-of-sample performance via k-fold cross-validation. Table 1 reports results across two key metrics: Qini Score, which assesses uplift ranking versus random targeting, and Uplift@30, which captures incremental gain in the top 30% of scored users. Double Machine Learning (DML) achieved the strongest out-of-sample results, with a Qini score of 8.03% and Uplift@30 of 1.11%, outperforming all baselines. The S-Learner exhibited strong in-sample metrics but failed to generalize, indicating overfitting. The T-Learner underperformed consistently. We also tested a Category Transformer adapted from our production ranker. While it achieved high in-sample Qini (8.31%), its generalization lagged behind DML, suggesting campaign-specific overfitting.

Given its balance of accuracy and robustness, we adopt DML as our default uplift modeling approach in production.

3.1.5 Promotion Strategy and Online Evaluation. We evaluate the model within a live promotion campaign. The current setup uses uniform targeting, with a single promotion type available per consumer segment. To introduce personalization, we adopt a top percentile targeting strategy: we identify the optimal uplift threshold using out-of-sample validation data and apply this cutoff in the online experiment. The full procedure is detailed in Algorithm 1.

Algorithm 1 Uplift-Based Percentile Targeting

Require: Estimated uplifts $\tau_k(x_i)$ for all users i and promotions k; validation set with observed outcomes

```
for all users i do
    Set reward: r_i = τ(x_i)
end for
Sort users by r_i descending
for all thresholds P ∈ {0.01, ..., 1.0} do
    Compute total observed uplift G_P on top-P validation users
end for
Let P* = arg max_P G_P
for all users i in deployment set do
    if rank(i) ≤ P* then
        Assign top-scoring promotion k*(x_i)
    else
        Assign no promotion
    end if
end for
```

We evaluate the system via an online A/B test comparing a **control group**, which receives uniform promotions, to a **treatment**

group, which receives targeted promotions only if their predicted uplift exceeds a percentile threshold. Assigned offers are optimized per user. Performance is measured by **overall order lift** and **cost efficiency** (cost index).

Table 2: Online A/B test results: personalized vs. uniform targeting

Group	Order Rate Lift	Cost Index
Control – no discount	–	–
Treatment – uniform targeting	2.43%	100
Treatment – causal targeting	2.44%	55

While both strategies achieved similar topline lift, the personalized approach reduced cost per incremental order by nearly half. This underscores the value of causal targeting: even when the promotion itself is fixed, choosing *who* to show it to can substantially improve efficiency.

3.1.6 Discussion. While this case demonstrates that improved targeting can significantly reduce cost with minimal impact on order rate lift, the broader marketing literature suggests that such benefits are not universal. In some contexts, targeting can backfire due to adverse consumer reactions. For instance, Goldfarb and Tucker [5] find that highly targeted but intrusive ads may reduce effectiveness, while Ascarza [1] shows that targeting high-risk customers for retention may be inefficient when those customers are unlikely to respond. These examples highlight that the effectiveness of targeting depends not only on predictive accuracy, but also on consumer perception and behavioral heterogeneity.

3.2 Case Study: Discount Depth Optimization for Restaurant Consumers

3.2.1 Context. Building on the previous case of fixed-form promotion targeting, we now address a more granular problem: optimizing the *discount depth* (e.g., 0%, 10%, 20%, 30%) for each consumer, where 0% represents no discount. This formulation generalizes targeting as it subsumes the decision of whether to offer a promotion. Unlike discrete treatment modeling, we treat the discount as a continuous variable and aim to personalize it to maximize a cost-adjusted reward function (Section 2.2.1).

3.2.2 Data. We use a one-month holdout dataset that includes both control and treated groups:

- **Control group:** Consumers who did not receive any discount.
- **Treated group:** Consumers exposed to varying discount levels.

Feature engineering mirrors Case 1 (Section 3.1) and includes prior order volume, vertical mix, promo redemption history, and merchant engagement. Each row is labeled with a binary outcome indicating whether the user placed an order following exposure.

Handling the Treatment Variable. There are two key complications associated with the treatment variable that motivate us to model it as a *continuous* rather than a *discrete* variable.

First, the observed treatment (discount) may vary across impressions and is not always recorded with full fidelity. For instance, the intended discount level is denoted by $\tilde{T}$; however, due to constraints such as discount caps or eligibility rules (which may vary by merchant or consumer), the actual discount experienced by the user, T, may differ from $\tilde{T}$—that is, $T \neq \tilde{T}$. This introduces a form of **treatment non-compliance**, where the assigned offer differs from the offer actually received. Such mismatches can bias causal effect estimation, particularly when the deviation is systematically related to user or merchant features. To mitigate this, we log the effective treatment T for each impression and apply Double Machine Learning (DML) to obtain debiased estimates (see later sections for details).

Second, the support of the treatment variable in the training data differs from that of the actual policy. For example, historical promotions may have offered only discrete values such as \$2 or \$4 off, while the production policy aims to select from a broader, finer-grained range, for instance, \$0.50 to \$5.00 in \$0.50 increments. Modeling the discount level as a continuous variable allows generalization beyond discrete training values and supports smooth interpolation across the action space.

3.2.3 Modeling Approach. In this setting, the treatment variable is a scalar: the offered discount amount $t \in [0, t_{\max}]$. Rather than comparing between discrete arms, we aim to estimate a user-level treatment response function:

$$\tau(t, x) = \mathbb{E}[Y \mid T = t, X = x] - \mathbb{E}[Y \mid T = 0, X = x]$$

To simplify estimation, we assume that $\tau(t, x)$ is linear in t: i.e., $\tau(t, x) = \tau(x) \cdot t$. This linearity assumption is empirically supported based on our data: when plotting average treatment effects across discount depths, the relationship appears approximately linear within the action space. Note that Zhang et al. (2024) used a different approach to encode continuous treatments without assuming linearity, by applying radial basis function (RBF) embeddings over discount depth to enable smooth generalization across similar actions while preserving non-linear relationships [11].

To address the non-random nature of treatment assignment, we apply Double Machine Learning (DML):

- Fit a **treatment model** $\mu_t(x)$ to predict the discount level T assigned to each user x.
- Fit an **outcome model** $\mu_y(x)$ to predict the adoption probability Y.
- Compute residuals: $\widetilde{T} = T - \mu_t(x)$ and $\widetilde{Y} = Y - \mu_y(x)$.
- Estimate the partial effect of $\widetilde{T}$ on $\widetilde{Y}$ via regression, yielding a debiased estimate of $\partial Y / \partial T$.

This approach allows us to estimate user-specific marginal response curves and infer the optimal discount t_i^* for each user.

3.2.4 Constrained Optimization. We frame the final assignment as a constrained optimization problem. For each user i, the goal is to choose the optimal discount t to maximize:

$$\text{Reward}(i, t) = \tau(x_i) \cdot t - \lambda \cdot c_i(t)$$

subject to the global budget constraint:

$$\sum_i c_i(t) \leq \mathcal{B}$$

While this problem is tractable offline (when the full user population is known in advance), it is less practical in online settings, where user traffic and budget evolve in real time. Moreover, the budget $\mathcal{B}$ is often a soft rather than hard constraint. In practice, we care more about maintaining the cost-to-reward ratio below a predefined efficiency threshold.

To this end, we reformulate the constraint as:

$$\frac{\sum_i c_i(t)}{\sum_i \text{Reward}(i,t)} \leq \theta$$

where θ is a predetermined threshold reflecting acceptable cost efficiency.

At optimality, the marginal cost-to-reward ratio should be equalized across all users receiving a promotion. This yields an implementable rule for assignment:

$$t_i^* = \max\left\{ t \leq t_{\max} : \frac{c_i(t)}{\text{Reward}(i,t)} \leq \theta \right\}$$

3.2.5 Evaluation. We evaluate the impact of our personalized discount assignment via a long-term online A/B test. In this experiment, one of the treatment groups received varying levels of discounts, and the other a fixed amount (as the baseline), while the control group received no discount.

Table 3 summarizes the key results. The personalized strategy achieved s higher order rate lift while significantly reducing overall cost by more than 70%.

Table 3: Online A/B test results: personalized vs. uniform discount amount

Group	Order Rate Lift	Cost Index
Control – no discount	–	–
Baseline – uniform amount	2.21%	100
Treatment – optimized amount	2.56%	28

3.2.6 Discussion. This case illustrates the need for robust modeling when personalization extends beyond binary treatment design. By treating discount depth as a continuous policy variable, we enable fine-grained optimization while accounting for endogenous treatment assignment. Though more complex than multi-arm setups, this method generalizes naturally to pricing, subsidy, or incentive calibration problems at scale.

4 CONCLUSION

We present a causal machine learning framework for optimizing promotions in large-scale, multi-vertical marketplaces. By combining double-debiased treatment effect estimation (DML) with cost-aware optimization, we enable personalized assignment of both promotion type and incentive depth, balancing impact with budget efficiency.

A core contribution of this work is to clarify when promotions should be modeled as discrete versus continuous treatments. Through two case studies, we show that the appropriate modeling choice depends on campaign structure: for fixed-form promotions where the offer is predefined (especially when the discount range is narrow or does not meaningfully affect outcome metrics), targeting decisions are best handled with discrete treatment models; for tunable incentives such as percent-off or fee reductions, a continuous treatment framework enables more precise optimization based on price sensitivity.

Online A/B tests confirm that personalized promotion achieves comparable order lift at significantly lower cost than uniform assignment, validating the modeling approach and its real-world utility. Beyond model performance, we address real-world deployment challenges, including re-exposure effects, eligibility filtering, and budget constraints, which are critical to realizing uplift gains at scale.

Looking ahead, we see several opportunities to extend this work, particularly in the direction of sequential decision-making for lifecycle promotions, adaptive experimentation pipelines, and integration with reinforcement learning. Our framework also aligns closely with recent advances in contextual bandits and online policy optimization under budget constraints. For example, Bastani and Bayati [2] develop algorithms for high-dimensional decision-making with theoretical guarantees on regret, while Zhou et al. [13] propose a budget-aware contextual bandit approach for personalized discounting in e-commerce. These threads suggest promising directions for scaling causal personalization to more dynamic and constrained environments. We hope this work contributes to the growing literature on causal inference for promotion systems and inspires further research into scalable, cost-effective personalization strategies.

ACKNOWLEDGMENTS

We thank Charlie Furrer for his work on an enhanced version of the continuous treatment CATE model introduced in Case Study 2 (Section 3.2).

REFERENCES

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

[2] Hamsa Bastani and Mohsen Bayati. 2020. Online decision-making with high-dimensional covariates. *Operations Research* 68, 1 (2020), 276–294.

[3] Léon Bottou, Jonas Peters, Joaquin Quiñonero-Candela, Denis X Charles, David M Chickering, Elon Portugaly, Dipankar Ray, Patrice Simard, and Edward Snelson. 2013. Counterfactual reasoning and learning systems: The example of computational advertising. *Journal of Machine Learning Research* 14 (2013), 3207–3260.

[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. *The Econometrics Journal* 21, 1 (2018), C1–C68.

[5] Avi Goldfarb and Catherine E Tucker. 2011. Online display advertising: Targeting and obtrusiveness. *Marketing Science* 30, 3 (2011), 389–404.

[6] Sören R Künzel, Jasjeet S Sekhon, Peter J Bickel, and Bin Yu. 2019. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the National Academy of Sciences* 116, 10 (2019), 4156–4165.

[7] Victor S Y Lo. 2002. The True Lift Model: A Novel Data Mining Approach to Response Modeling in Database Marketing. *SIGKDD Explorations* 4, 2 (2002), 78–86.

[8] Judea Pearl. 2009. *Causality: Models, Reasoning and Inference.* Cambridge University Press.

[9] Donald B Rubin. 1974. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology* 66, 5 (1974), 688–701.

[10] Lei Tang and Yiling Wang. 2017. Optimizing targeted promotions under budget constraints in e-commerce. *INFORMS Journal on Computing* 29, 1 (2017), 43–56.

[11] Jason Shuo Zhang, Benjamin Howson, Panayiota Savva, and Eleanor Loh. 2024. DISCO: An End-to-End Bandit Framework for Personalised Discount Allocation. arXiv:2406.06433 [cs.LG] https://arxiv.org/abs/2406.06433

[12] Shuo Zhang, Benjamin Howson, Panayiota Savva, and Eleanor Loh. 2024. Causal Machine Learning for Promotion Design. *Journal of Marketing Science* (2024).

[13] Yixuan Zhou, Yunzhe Wu, Anuj Gupta, and Lihong Li. 2023. Budget-Aware Discount Personalization in E-Commerce Using Contextual Bandits. In *Advances in Neural Information Processing Systems (NeurIPS).*