Free Lunch! Retrospective Uplift Modeling for Dynamic Promotions Recommendation within ROI Constraints

Dmitri Goldenberg*

Javier Albert*

dima.goldenberg@booking.com

javier.albert@booking.com

Booking.com, Tel Aviv, Israel

Lucas Bernardi[†]

Pablo Estevez

lucas.bernardi@booking.com

pablo.estevez@booking.com

Booking.com, Amsterdam, Netherlands

ABSTRACT

Promotions and discounts have become key components of modern e-commerce platforms. For online travel platforms (OTPs), popular promotions include room upgrades, free meals and transportation services. By offering these promotions, customers can get more value for their money, while both the OTP and its travel partners may grow their loyal customer base. However, the promotions usually incur a cost that, if uncontrolled, can become unsustainable. Consequently, for a promotion to be viable, its associated costs must be balanced by incremental revenue within set financial constraints. Personalized treatment assignment can be used to satisfy such constraints. This paper introduces a novel uplift modeling technique, relying on the Knapsack Problem formulation, that dynamically optimizes the incremental treatment's outcome subject to the required Return on Investment (ROI) constraints. The technique leverages *Retrospective Estimation*, a modeling approach that relies solely on data from positive outcome examples. The method also addresses training data bias, long term effects, and seasonality challenges via online-dynamic calibration. This approach was tested via offline experiments and online randomized controlled trials at Booking.com - a leading OTP with millions of customers worldwide, resulting in a significant increase in the target outcome while staying within the required financial constraints and outperforming other approaches.

CCS CONCEPTS

• **Computing methodologies** → **Learning from implicit feedback**; • **Information systems** → **Personalization**.

KEYWORDS

Uplift Modeling, Machine Learning, Causal Inference, Benefits Recommendation, Implicit Feedback

*Both authors contributed equally to this research.

[†]Corresponding author.

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

RecSys '20, September 22–26, 2020, Virtual Event, Brazil

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-7583-2/20/09.

<https://doi.org/10.1145/3383313.3412215>

ACM Reference Format:

Dmitri Goldenberg, Javier Albert, Lucas Bernardi, and Pablo Estevez. 2020. Free Lunch! Retrospective Uplift Modeling for Dynamic Promotions Recommendation within ROI Constraints. In *Fourteenth ACM Conference on Recommender Systems (RecSys '20)*, September 22–26, 2020, Virtual Event, Brazil. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3383313.3412215>

1 INTRODUCTION

Modern online travel platforms (OTPs) offer a vast variety of products and services to travellers, such as flights, accommodations and ground transportation. Recently, offering a promotion such as a room upgrade, free meals, or discounted car rental has become a popular approach for offering customers more value for their money when completing a purchase. A promotion is expected to have a positive Average Treatment Effect (ATE) on outcomes like customer base growth or customer satisfaction, but can be viable only under set Return on Investment (ROI) constraints. ROI is a common measure in the field of marketing campaigns [19], which represents the ratio of profit to investment as follows:

$$ROI = \frac{\Delta Revenue - \Delta Investment}{\Delta Investment} \quad (1)$$

In order to ensure that a promotional campaign is sustainable over time, the target ROI is often restricted with a lower bound, such as $ROI \geq 0$. This creates a trade-off between the increment in target outcome and the cost of offering the promotion. A good example of such a promotion could be offering free lunch as an incentive for booking an accommodation; the promotional campaign will be sustainable only if the incremental value in accommodation booking rate can cover the extra cost of providing the free lunch. By offering the promotion to a subset of customers, a company can provide greater value to each customer while maintaining financial constraints that keep the promotional campaign sustainable over time [22]. One popular approach for efficiently assigning promotions to a subset of customers is targeted personalization [17]. In this approach, uplift models [5] are commonly used to estimate the expected causal effect of the treatment on the outcome in order to distinguish between the *voluntary buyers* and the *persuadables*, who would only purchase as a response to the incentive [15]. This can be achieved by estimating the Conditional Average Treatment Effect (CATE) of the promotion. $CATE_Y(x)$ is defined as the conditional increment in probability of completing the purchase $Pr(Y = 1)$ caused by the incentive, given the customer characteristics x .

CATE estimation can be conducted via various meta-learning techniques [8], such as the two-models approach [21], transformed outcome [10], uplift trees [23], and other meta-learners [14]. Some

previous work aimed to address allocation of promotions by estimating the marginal effect [26], incorporating net value optimization within the meta-learner [27] and modeling the optimization as a min-flow problem [18]. Such estimations might be unreliable [4, 16], however, due to high variance in the training dataset, which is heavily dependant on a majority of non-transactional visits having no intent to purchase. In addition, offline solutions might be biased towards historical data and require dynamic calibration [28] according to long-term changes and seasonality trends, which are particularly common in the travel industry [2] and in promotional campaigns [12]. Dynamic adaptive strategy adjustments is a common practice in cases of budget constraints and were found effective in ads management [17], influence maximization [7, 24], marketing budget allocation [1, 15, 26] and discount personalization [6].

In this paper we present a novel technique based on *Retrospective Estimation*, a modeling approach that relies solely on data from positive outcome examples. This is extended by a method to dynamically maximize the target outcome of the treatment within ROI constraints. Our main contributions are: (1) Optimization framework setup for uplift within ROI constraints; (2) *Retrospective Estimation* technique, relying only on data from positive examples; (3) Dynamic model calibration for online applications; and (4) Large scale experimental validation. The study proceeds as follows: The second section formalizes the problem we address and the constraints imposed by business and product requirements. The third section covers the suggested method for uplift modeling under ROI constraints, retrospective estimation technique, and dynamic calibration of the method. The fourth section covers validation of the method (Randomized Controlled Experiment), including details on data gathering, experimental setup and results.

2 PROBLEM FORMULATION

In this work, we focus on maximizing the overall number of customers completing a purchase, while deciding whether or not to offer the promotion in order to comply with the global constraint of $ROI \geq 0$. In addition to Y (a random variable representing the completion of a purchase), we consider the random variables R and C , representing the monetary revenue associated with the purchase and the cost associated with the promotion, respectively. We assume that there are no revenue and no cost without a purchase.

We adopt the Potential Outcomes framework [9], which allows us to express the causal effects of the promotions on all three different outcomes: for a customer i , $Y_i(1)$ represents the potential purchase if i is given the promotion ($T = 1$), while $R_i(1)$ and $C_i(1)$ represent the potential revenue and costs. Likewise, $Y_i(0)$, $R_i(0)$, and $C_i(0)$ represent the potential outcomes if the promotion is not given. We can thus define the conditional average treatment effect on all these variables, for a given customer with pre-promotion covariates x , as follows: $CATE_Y(x) = E(Y_i(1) - Y_i(0)|X_i = x)$, $CATE_R(x) = E(R_i(1) - R_i(0)|X_i = x)$ and $CATE_C(x) = E(C_i(1) - C_i(0)|X_i = x)$. We also define the profit outcome $\Pi_i = R_i - C_i$ and the loss as $\mathcal{L} = -\Pi$. Their associated $CATE_\Pi(x) = CATE_R(x) - CATE_C(x) = -CATE_{\mathcal{L}}(x)$ are simply the average increment in profit (or loss). Our goal is to find a function F , that given x , decides if the customer is offered the promotion or not, while maximizing the total incremental number of transactions

under incremental ROI constraints. Therefore, our problem is to learn F , given a sample of customers U , where some received the promotion and others did not. Formally, our objective is to learn F from tuples $(x_i, T_i, Y_i, R_i, C_i)$, $0 < i < n$, where $T_i = 1$ if the customer received the promotion and $T_i = 0$ otherwise.

2.1 Solution Framework

We model F with a scoring function g and a threshold θ , that outputs its decision as $\mathbb{1}[g(x) \geq \theta]$. Function g needs to represent the increment in probability of booking as well as the associated revenue and costs. By relying of existing data, we can reassign the promotions in U by solving the following optimization problem:

$$\begin{aligned} & \text{Maximize} \quad \sum_{i \in U} z_i \cdot CATE_Y(x_i) \\ & \text{subject to:} \quad \sum_{i \in U} z_i \cdot CATE_{\mathcal{L}}(x_i) \leq 0 \end{aligned} \quad (2)$$

Here, $z_i \in \{0, 1\}$ is the assignment variable indicating if customer i is offered the promotion or not. The target function is maximizing the total treatment effect, while the constraint specifies the condition of incremental total loss being non-positive (equivalent to positive ROI), by comparing the incremental costs and revenues potentially incurred by each incentivized customer. This sets a *Binary Knapsack Problem* with negative utilities and weights. We derive that $CATE_Y \leq 0 \implies CATE_{\mathcal{L}} \geq 0$, since the revenue can only come from extra purchases. Applying the transformation described in [25] results in the residual problem of a standard binary Knapsack Problem for customers with $CATE_Y > 0$ and $CATE_{\mathcal{L}} > 0$. It has a new constraint constant given by $\sum_j CATE_{\mathcal{L}}^j$ for customers where $CATE_Y^j > 0$ and $CATE_{\mathcal{L}}^j \leq 0$.

According to the greedy algorithm for the *Fractional Knapsack Problem*, we can approximate the solution by sorting customers U by $utility_i/weight_i$ descending. In our problem this ratio maps to $CATE_Y^i/CATE_{\mathcal{L}}^i$. All customers are treated in that order, until the constraint is no longer satisfied. Denoting j as the last treated customer, we can set $g(x) = CATE_Y(x)/CATE_{\mathcal{L}}(x)$ and the threshold $\theta^* = CATE_Y^j/CATE_{\mathcal{L}}^j$. Assuming the sample U is a representative of future customers, we can use this rule to decide who to target with our promotion. This framework reduces the problem to learning $\frac{CATE_Y(x)}{CATE_{\mathcal{L}}(x)}$; $\text{sign}[CATE_Y(x)]$; and $\text{sign}[CATE_{\mathcal{L}}(x)]$. The following sections explore different approaches.

3 METHOD FOR UPLIFT MODELING UNDER ROI CONSTRAINTS

We performed a comparison of two common uplift modeling techniques and introduced two novel modeling approaches:

- **Two-Models** [21]: a difference between two individual estimations of the effect under each of the treatments.
- **Transformed outcome** [10]: direct estimation of $CATE_Y(x)$ with a single model.
- **Fractional Approximation**: greedy sorting score relying on *Two-Models* output - $CATE_Y(x)/CATE_{\mathcal{L}}(x)$.
- **Retrospective Estimation** (3.1): direct estimation of greedy sorting score, relying only on positive examples.

All methods rely on data from an online randomized controlled experiment [11, 13], which allows us to write the potential outcomes $\mathbb{E}[Y_i(1)|X_i = x]$ as $\mathbb{E}[Y_i|T_i = 1, X_i = x]$ (and their analogous for Revenue, Cost and Profit). These expectations are in turn estimated using machine learning. All methods are trained with underlying *Xgboost* models [3] on data of >100 millions of interactions on Booking.com website, lasting over several weeks. The model performance was tuned with time-based cross validation and evaluated as described in section 3.3. With $\hat{C}(x) = \mathbb{E}[C|x, T = 1, Y = 1]$ and $\hat{R}_i(x) = \mathbb{E}[R|x, T = i, Y = 1]$, the full greedy criterion $\frac{CATE_Y(x)}{CATE_L(x)}$ can be written in terms of expectations as follows:

$$\frac{Pr(Y = 1|x, T = 1) - Pr(Y = 1|x, T = 0)}{Pr(Y = 1|x, T = 1)[\hat{C}(x) - \hat{R}_1(x)] + Pr(Y = 1|x, T = 0)\hat{R}_0(x)} \quad (3)$$

Fractional approximation technique relies on direct estimations of all the probabilities and expectations from formula 3. The **Two-Models** and **Transformed Outcome** methods that are estimating only the numerator of Eq. 3, are used as benchmark solutions to the unconstrained problem.

3.1 Retrospective Estimation

This technique relies on the following response distribution factorization:

$$Pr(Y = 1|x, T = 1) = \frac{Pr(T = 1|x, Y = 1)Pr(Y = 1|x)}{Pr(T = 1|x)} \quad (4)$$

The denominator is the treatment propensity, which we can set to 0.5. This decomposition does not allow us to estimate $CATE_Y(x)$ directly, but instead provides an expression for the ratio between $Pr(Y = 1|x, T = 1)$ and $Pr(Y = 1|x, T = 0)$:

$$\frac{Pr(Y = 1|x, T = 1)}{Pr(Y = 1|x, T = 0)} = \frac{Pr(T = 1|x, Y = 1)}{1 - Pr(T = 1|x, Y = 1)} = \frac{S(x)}{1 - S(x)} \quad (5)$$

Here $S(x) = Pr(T = 1|x, Y = 1)$ is the probability of a positive example to result from the treatment $T = 1$. This allows us to write the full greedy criterion in terms of $S(x)$:

$$\frac{CATE_Y(x)}{CATE_L(x)} = \frac{2S(x) - 1}{S(x)[\hat{C}(x) - \hat{R}_1(x)] + [1 - S(x)]\hat{R}_0(x)} \quad (6)$$

$$\text{sign}[CATE_Y(x)] = \text{sign}[S(x) - 0.5] \quad (7)$$

$$\text{sign}[CATE_L(x)] = \text{sign}\left[\frac{\hat{R}_1(x)}{2\hat{R}_1(x) - \hat{C}(x)} - S(x)\right] \quad (8)$$

Estimating $S(x)$, $\hat{R}_1(x)$, $\hat{R}_0(x)$ and $\hat{C}(x)$ requires data only from customers where $Y = 1$, making this approach robust to the noise in the general set of visitors (including scrapes and crawlers - a significant portion of traffic who rarely purchase) and less prone to class imbalance issues. It also requires a smaller transactional dataset, which is often easier to collect. Furthermore, assuming that the costs and revenues per booking are all independent of x and that the revenue per booking is independent of treatment T ($\hat{R}_i(x) = \hat{R}$ and $\hat{C}(x) = \hat{C}$) we can represent the ratio by equation 9:

$$\frac{2S(x) - 1}{S(x)[\hat{C} - \hat{R}] + [1 - S(x)]\hat{R}} \quad (9)$$

We can derive positive values of $CATE_L(x) > 0$ if $S(x) < \frac{\hat{R}}{2\hat{R} - \hat{C}}$ and $CATE_Y(x) > 0$ if $S(x) > 0.5$. This means that we can solve the greedy problem just by modeling $S(x) = Pr(T = 1|x, Y = 1)$.

3.2 Dynamic calibration strategy

Fitting a function g and a threshold θ offline provides a solution to the offline samples assignment. However, during an online campaign, the general traffic distribution might differ and therefore the optimal threshold θ^* can shift accordingly. We propose a simple curve fitting technique to address this problem. As seen in figure 2, we assume a non-linear monotonic relation between the portion of the exposed population Q_θ and the resulted ROI. We suggest to model this relation as an exponential curve, by learning the parameters a, b and c such that: $\widehat{ROI}(Q) = ae^{bQ} + c$ using the Levenberg-Marquardt least-squares curve fitting algorithm [20] to select a new θ^* such that $\widehat{ROI}(Q_{\theta^*}) = 0$. Using this parametric curve fitting method allows incorporating historical and up-to-date data in ROI estimation, by refitting the curve every selected time-period.

3.3 Evaluating an Uplift Model: Qini and Qini-ROI curves

The methods are compared with *Qini curves* [21] that represent the potential cumulative treatment effect achieved by targeting a selected top ranked population based on the model predictions (see figure 1). The *Qini Coefficient* (also called *AUUC* - Area Under the Uplift Curve) [5] is used to quantify and evaluate the overall quality of an uplift model.

Following the same ranking procedure, we propose the *Qini-ROI curve*, representing the expected ROI that can be achieved by targeting a selected portion of the top ranked population (see figure 2). According to the described optimization problem (2.1), $ROI \geq 0$ is a binding constraint to the treatment assignment, therefore the optimal model threshold is at the operating point where the treatment effect is the highest and also $ROI \geq 0$.

Thus, the models' performance is evaluated by *Maximal ATE at ROI ≥ 0* metric. The *AUUC* and *Maximal targeted population at ROI ≥ 0* metrics are used to evaluate the class-separation and the population coverage of the methods.

4 EXPERIMENTAL STUDY

4.1 Randomized controlled trial of indiscriminated treatment assignment

In order to test the suggested method, we conducted a series of offline and online experiments. As a preceding step to our method, we conducted a randomized controlled trial to assess the potential impact of a given promotion. The experiment took the form of an online A/B test, in which the treatment group was offered a promotion and the control group was not. The experiment was conducted on Booking.com website and lasted for several weeks. In our experiment, all of the customers in the treatment group received an offer of a free benefit as an incentive for booking an accommodation. The individual outcome metrics - completing a booking, promotion cost, and incremental revenue - were aggregated and compared between the control (no promotion) and treatment (free benefit) groups, producing an estimation of the *ATE* and *ROI* of the promotion treatment. The experiment resulted in a conclusively positive treatment effect ($ATE > 0$), but with an $ROI < 0$ value, which is insufficient in order to keep offering the promotion.

Figure 1: Qini curves comparison

4.2 Uplift methods training and offline evaluation

The data gathered during the A/B test was used as training instances for uplift modeling. The dataset consisted of >100 million website visit instances, associated with customer's current and historical search queries and transactions. Each instance was attributed with associated transaction, revenue, and cost, if applicable. We evaluated the models' performance on a validation set according to the metrics described in 3.3. Figure 1 presents the comparison of the models' uplift performance; Figure 2 presents the comparison of the models' ROI profile; key values are summarized in table 1. The stars at each curve represent the best operating points (in terms of ATE) for each method, s.t. $ROI \geq 0$.

The *Transformed Outcome* method achieves the highest AUUC score (0.912), however, its $ROI = 0$ intersection is only at 2.4% of population, achieving 38.6% of the treatment effect relative to giving the promotion to all the users. The *Two-Models* method has the highest potential treatment effect (109%) in an unconstrained setup, but it performs poorly in key metrics and doesn't have any operating point with $ROI \geq 0$. The *Fractional Approximation* method provides the highest population coverage (38%) at $ROI \geq 0$ and achieves 65.8% of the full treatment effect. Recalling the objective function of the optimisation problem, the *Retrospective Estimation* method provides the best solution. It leads with a great advantage on maximal treatment effect achieving 79.7% of the possible uplift and is a runner-up in both AUUC (0.805) and Population at $ROI = 0$ (31.4%) metrics. It is worth highlighting that the *Retrospective Estimation* performance relies solely on positive examples data.

Table 1: Uplift methods' performance comparison

Method	AUUC	Max. Population at $ROI=0$	Max. ATE at $ROI=0$
Two-Models	0.638	0%	0%
Transformed Outcome	0.912	2.4%	38.6%
Fractional Approximation	0.721	38%	65.8%
Retrospective Estimation	0.806	31.4%	79.7%

Figure 2: Qini-ROI curves comparison

4.3 Randomized controlled trial with personalized treatment assignment

We recreated the randomized controlled trial setup as an online validation of our technique 4.1. For this iteration, visitor traffic was split into four randomly-assigned treatment groups:

- **A - Control Group** - no promotion offering
- **B - Undiscriminated Treatment** - promotion is offered to everyone
- **C - Personalized Treatment** - promotion is offered to a population with $Score(x) \geq \theta$
- **D - Dynamically Personalized Treatment** - promotion is offered to a population with $Score(x) \geq \theta_d$

We used the *Retrospective Estimation* technique for the underlying personalized treatment assignment method in groups **C** and **D**, due to its superior performance on the target metric. Model score threshold (θ) was selected based on offline evaluation of the model, such that 31.4% of the population are expected to receive the promotion. We re-calibrated the threshold for the *Dynamically Personalized Treatment* (θ_d) after each predefined time-period. A randomized controlled trial (over 8 update periods) was conducted under the same promotion conditions as in the first experiment (an offer of a free benefit as an incentive for an accommodation booking on the website). The goal of the experiment was to evaluate the incremental ATE and the ROI of the treatments groups compared to control group **A**. Figure 3 presents the measured cumulative ROI (dashed lines) and the relative ATE (solid lines) of the treatments, where the *diamond* markers ($\diamond$) represent the dynamic calibration points during the experiment. Treatment **B** resulted in ROI values below 0, similarly to the previous iteration of the experiment. Treatment **C** (Personalized group) provided consistently high ROI values ($ROI = 0.36$ at the last time period) and its total cumulative uplift is 66% of that achieved by fully-treated group **B**. Treatment **D** starts with a high $ROI = 0.3$ in the first period and converges to $ROI = 0.24$ at the end of the experiment. Dynamic adjustment of the threshold in group **D** allowed for exposing the treatment to 47% of the population, achieving on average 74% of the uplift achieved by treatment **B**, including 80% during the last period.

Figure 3: Experimental results (achieved ROI and portion of ATE) of personalized promotions assignment

5 CONCLUSIONS

In this work, we introduced an optimization framework to solve the personalized promotion assignment by relying on the Knapsack Problem and presented a novel method to estimate the required sorting score. Classical uplift models excelled in AUUC metric, but under-performed in ROI-related measures as they do not account for the cost factor.

Our novel *Retrospective Estimation* technique relies solely on data from positive outcome examples, which makes its training easier and more scalable in big-data environments, especially for cases where only the outcomes' data is available. The method achieved 80% of the possible uplift, by targeting only 31% of the population and meeting the $ROI \geq 0$ constraint during offline evaluation. It outperforms the direct *Fractional Estimation* model, potentially thanks to robustness to noise in the training data, which opens an opportunity for further research on new applications.

We suggested a dynamic calibration technique, which allows adjusting the model threshold and exposing a different portion of the population to the treatment, by relying on online performance. We evaluated our suggested techniques through a massive online randomized controlled trial and showed that a personalized treatment assignment can turn an under-performing promotions campaign into a campaign with a viable ROI and significant uplift ($ATE = 66\%$ of the possible effect). Applying the dynamic threshold calibration to online data improved the uplift of the treatment by up to 20%, while staying within the constraints, and provided a long term solution for a rapidly changing environment.

Our contributions lay the foundation for future research to apply the *Retrospective Estimation* technique to other uplift modeling and personalization applications. This sets the groundwork for new optimization approaches and provides solid empirical evidence of the benefits of online personalized promotion allocation and dynamic calibration. It creates more value to the customers and increases the loyal customer base for the vendors and the platform.

ACKNOWLEDGMENTS

We would like to thank the entire Value Team for their contribution to the project and to Adam Horowitz, Tamir Lahav, Guy Tsybe, Pavel Levin, Irene Teinemaa and Nitzan Mekel-Bobrov for their support and insightful discussions.

REFERENCES

- [1] Ashwinkumar Badanidiyuru, Robert Kleinberg, and Yaron Singer. 2012. Learning on a budget: posted price mechanisms for online procurement. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. 128–145.
- [2] Lucas Bernardi, Themistoklis Mavridis, and Pablo Estevez. 2019. 150 Successful Machine Learning Models: 6 Lessons Learned at Booking. com. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1743–1751.
- [3] Tianqi Chen and Carlos Guestrin. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*. 785–794.
- [4] Alex Deng and Victor Hu. 2015. Diluted treatment effect estimation for trigger analysis in online controlled experiments. In *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining*. 349–358.
- [5] Floris Devriendt, Darie Moldovan, and Wouter Verbeke. 2018. 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 (2018), 13–41.
- [6] Marc Fischer, Sönke Albers, Nils Wagner, and Monika Frie. 2011. Dynamic Marketing Budget Allocation across Countries, Products, and Marketing Activities. *Marketing Science*, forthcoming. (2011).
- [7] Dmitri Goldenberg, Alon Sela, and Erez Shmueli. 2018. Timing matters: Influence maximization in social networks through scheduled seeding. *IEEE Transactions on Computational Social Systems* 5, 3 (2018), 621–638.
- [8] Pierre Gutierrez and Jean-Yves Grardy. 2017. Causal inference and uplift modelling: A review of the literature. In *International Conference on Predictive Applications and APIs*. 1–13.
- [9] Guido W Imbens and Donald B Rubin. 2010. Rubin causal model. In *Microeconometrics*. Springer, 229–241.
- [10] Maciej Jaskowski and Szymon Jaroszewicz. 2012. Uplift modeling for clinical trial data. In *ICML Workshop on Clinical Data Analysis*.
- [11] Raphael Lopez Kaufman, Jegar Pitchforth, and Lukas Vermeer. 2017. Democratizing online controlled experiments at Booking. com. *arXiv preprint arXiv:1710.08217* (2017).
- [12] Hyeon Min Kim and Thomas Kramer. 2006. "Pay 80%" versus "get 20% off": The effect of novel discount presentation on consumers' deal perceptions. *Marketing Letters* 17, 4 (2006), 311–321.
- [13] Ron Kohavi, Randal M Henne, and Dan Sommerfield. 2007. Practical guide to controlled experiments on the web: listen to your customers not to the hippo. In *Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining*. 959–967.
- [14] Sren R Knzl, 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.
- [15] Yi-Ting Lai, Ke Wang, Daymond Ling, Hua Shi, and Jason Zhang. 2006. Direct marketing when there are voluntary buyers. In *Sixth International Conference on Data Mining (ICDM'06)*. IEEE, 922–927.
- [16] Randall A Lewis and Justin M Rao. 2015. The unfavorable economics of measuring the returns to advertising. *The Quarterly Journal of Economics* 130, 4 (2015), 1941–1973.
- [17] Ying-Chun Lin, Chi-Hsuan Huang, Chu-Cheng Hsieh, Yu-Chen Shu, and Kun-Ta Chuang. 2017. Monetary discount strategies for real-time promotion campaign. In *Proceedings of the 26th International Conference on World Wide Web*. 1123–1132.
- [18] Rahul Makhijani, Shreya Chakrabarti, Dale Struble, and Yi Liu. 2019. LORE: a large-scale offer recommendation engine with eligibility and capacity constraints. In *Proceedings of the 13th ACM Conference on Recommender Systems*. 160–168.

- [19] Saurabh Mishra and Madhvendra Mishra. 2011. Current industry approaches towards Marketing ROI an Empirical study. *The European Journal of Business Management* 3, 6 (2011).
- [20] Jorge J Moré. 1978. The Levenberg-Marquardt algorithm: implementation and theory. In *Numerical analysis*. Springer, 105–116.
- [21] Nicholas J Radcliffe. 2007. Using control groups to target on predicted lift: Building and assessing uplift models. *Direct Marketing Analytics Journal* 1 (2007), 1421.
- [22] Thomas Reutterer, Andreas Mild, Martin Natter, and Alfred Taudes. 2006. A dynamic segmentation approach for targeting and customizing direct marketing campaigns. *Journal of interactive Marketing* 20, 3–4 (2006), 43–57.
- [23] Piotr Rzepakowski and Szymon Jaroszewicz. 2012. Decision trees for uplift modeling with single and multiple treatments. *Knowledge and Information Systems* 32, 2 (2012), 303–327.
- [24] Alon Sela, Dmitri Goldenberg, Irad Ben-Gal, and Erez Shmueli. 2018. Active viral marketing: Incorporating continuous active seeding efforts into the diffusion model. *Expert Systems with Applications* 107 (2018), 45–60.
- [25] Paolo Toth and Silvano Martello. 1990. *Knapsack problems: Algorithms and computer implementations*. Wiley, 14–15 pages.
- [26] Kui Zhao, Junhao Hua, Ling Yan, Qi Zhang, Huan Xu, and Cheng Yang. 2019. A Unified Framework for Marketing Budget Allocation. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1820–1830.
- [27] 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.
- [28] Yunhong Zhou, Deeparnab Chakrabarty, and Rajan Lukose. 2008. Budget constrained bidding in keyword auctions and online knapsack problems. In *International Workshop on Internet and Network Economics*. Springer, 566–576.