PACING WITH ROI: BUDGET ALLOCATION IN SPONSORED SEARCH

Mohamed A. Abd-Elmagid*
Department of ECE
Virginia Tech
Blacksburg, VA 24061, USA
{maelaziz}@vt.edu

Djordje Gligorijevic †
Meta AI
Menlo Park, California, USA
{gligorijevic}@meta.com

Arnob Ghosh
Department of ECE
New Jersey Institute of Technology
Newark, NJ 07102, USA
{arnob.ghosh}@njit.edu

Andrew Perrault
Department of CSE
The Ohio State University
Columbus, OH 43210, USA
{perrault.17}@osu.edu

Xinxin Shu
eBay
San Jose, California, USA
{xshu}@ebay.com

Ness B. Shroff
Department of ECE and CSE
The Ohio State University
Columbus, OH 43210, USA
{shroff.11}@osu.edu

Abraham Bagherjeiran
eBay
San Jose, California, USA
{abagherjeiran}@ebay.com

ABSTRACT

Sponsored search is a core revenue engine for online marketplaces, where ad slots are sold via auctions and advertisers bid on keywords subject to daily budgets. The central challenge is budget pacing—allocating spending over time to balance immediate opportunities against potentially better future ones. Existing pacing methods are largely heuristic, offering no Return-on-Investment (ROI) guarantees for advertisers and often ignoring minimum spending requirements from the platform's side. We propose new theoretically-motivated budget pacing algorithms that account for the budget and ROI constraints of the advertisers as well as minimum spending constraints of the platform. Evaluated in eBay's sponsored search environment, our algorithms show that the form of the ROI constraint materially shapes the tradeoff between the advertiser's utilities (e.g., impressions, clicks, cost per click) and the platform's revenue, and they consistently outperform widely used PID-based pacing heuristics, and a state-of-the-art budget pacing approach.

1 INTRODUCTION

Sponsored search advertising has become an essential tool for online marketplaces like eBay, enabling sellers to connect with new customers. Through these programs, sellers can boost the visibility and sales of their products by bidding for sponsored placements within search results or on product pages. In particular, sellers define budgets and bids, which may be set at the keyword, product, or campaign level depending on the ad format. The platform then uses this information to create automated agents that place bids on behalf of advertisers in ad auctions through which the

*This work was done while the author was with the Department of ECE at The Ohio State University.
†This work was done while the author was with eBay.

impression opportunities/slots are being sold. Allocating small budgets or high maximum bids can cause advertisers to quickly deplete their budgets and miss out bidding on valuable opportunities. Moreover, from the platform's perspective, ensuring a consistent and appropriate spending pattern for its sellers is quite important for maintaining their long-term trust and partnership Karande et al. (2013); Lee et al. (2013). This has motivated the need to develop efficient budget pacing aware automated agents, which is the main focus of this paper.

Related Work. Prior research on budget pacing can be grouped into two categories. One adopts a pragmatic, empirically motivated perspective; the other develops theoretical frameworks that provide rigorous performance guarantees, albeit under stronger assumptions. Specifically, prior empirical studies proposed several budget pacing solutions from various angles, such as bipartite graph allocation methodologies Mehta et al. (2007), control theory Karlsson & Zhang (2013); Zhang et al. (2016), simultaneously optimizing business metrics such as click-through-rate (CTR) Karande et al. (2013); Lee et al. (2013), or other empirical approaches Agarwal et al. (2014); Zhang et al. (2014). All of these empirical solutions were implemented via one of the following two budget pacing approaches: i) throttling, where the ad campaigns are excluded from joining ad auctions probabilistically, and ii) bid shading, where the advertiser's expenditure is controlled via reducing their posed bids. Further, these prior empirical studies did not provide formal theoretical guarantees for the performance of their approaches.

The prior theoretical studies approached the problem of budget pacing using tools such as game theory, mechanism design and learning theory, with many of them establishing game-theoretic equilibria of budgeted auctions Balseiro et al. (2017); Charles et al. (2013); Ciocan & Iyer (2021); Conitzer et al. (2022) or devised budget pacing strategies that satisfy either performance guarantees for aggregate/individual utilities or certain incentive properties Balseiro et al. (2015); Balseiro & Gur (2019); Chen et al. (2024b); Gaitonde et al. (2022). These prior theoretical studies proposed methods that can also be considered as either throttling or bid shading. Further, they obtained theoretical guarantees under stringent assumptions that fail to hold in a noisy real-world environment, or relied on auction designs that are rarely used in practice Balseiro et al. (2019); Goel et al. (2020). These stringent assumptions include: i) the assumption of having a stationary environment where the maximum bid of each advertiser is drawn from a fixed distribution, and ii) the consideration of a non-asymptotic time horizon model. Note that in practice, the advertisers provide maximum bids in an arbitrary fashion (i.e., the environment is non-stationary), and their budgets are reset daily or weekly (i.e., a non-asymptotic time horizon is considered).

To bridge the gap between the empirical and theoretical viewpoints, it is necessary to develop theoretically-inspired budget pacing approaches that suit the empirical considerations of sponsored search programs. Towards this objective, a very recent study Chen et al. (2024a) investigated the performance of an optimization-based bid shading approach within a real-world sponsored search environment. However, the bidding strategy in Chen et al. (2024a) was developed while only accounting for advertisers' budget constraints. In many practical scenarios, satisfying some form of an ROI constraint in needed from the advertiser's perspective Feng et al. (2023); Lucier et al. (2024). In particular, the ROI constraint requires the ratio of the advertiser's total value to its total payment to be at least some specified target. Enforcing an ROI target improves the advertiser's retention and long-run spending, mitigates adverse selection (especially for small budgets), and prevents pacing policies from buying low-quality impressions merely to exhaust budgets. Motivated by these limitations, this paper presents a unified online optimization-based framework which develops budget pacing algorithms for the practical setting of having an ROI constraint along with the budget constraints. Specifically, we consider several forms of practical interest for the ROI constraint and answer the following open question:

How can the choice of the form of the ROI constraint impact the platform's business objectives and the advertisers' utilities?

Contributions. This paper considers a sponsored search advertising program in which search ad placements are sold to a set of advertisers/campaigns via second-price multi-slot auctions. For this model, we study the problem of maximizing the expected utility of each campaign subject to its budget, minimum spending, and ROI constraints. Specifically, we consider three different forms of the ROI constraint that provide guarantees for each campaign's achievable utilities, including its expected gain, expected number of clicks, and number of impressions (as will be discussed in detail in Section 2). For each form of the ROI constraint, we consider a Lagrangian-based approach to analytically derive the bid value of each campaign as a function of the dual variables associated with different constraints. The derived expressions of the bid value are then used to develop online optimization-based budget pacing algorithms under different forms of the ROI constraint considered in this paper. The efficacy of our proposed algorithms in real-word environments is validated by quantifying their performances within eBay's sponsored search environment. We conduct extensive experiments to answer a number of practical research questions. In particular, our results demonstrate that the structure of the ROI constraint can significantly impact the tradeoff between the advertisers' utilities and the platform's revenue. Our results also show that the platform can flexibly control this tradeoff through a minimum spending constraint. Further, our experiments demonstrate that the performance of our developed algorithms significantly outperform that of PID controllers, which are widely used in the industry as heuristics Zhang et al. (2016) for budget pacing.

Connection to mechanism design and strategic decision making. Sponsored search is a *repeated mechanism* in which the platform runs many multi-slot auctions while advertisers face intertemporal constraints (budgets, ROI targets). In practice, platforms deploy *automated bidding agents* that translate these high-level goals—and platform objectives such as minimum-spend commitments—into per-auction bids, thereby shaping effective strategies, allocations, payments, and long-run participation.

We design such agents via an optimization-based framework that jointly enforces budget, minimum-spend, and ROI constraints. ROI can be viewed as a *viability* constraint: persistent ROI shortfalls lead to churn and adverse selection. Our resulting dual-based bid-shading rules are interpretable as shadow prices and adapt aggressiveness to budget scarcity and ROI slack. Finally, we show that the *ROI semantics* (utility-, click-, or impression-based) are a first-order design choice, inducing different bidding structures and shifting the advertiser-outcome–platform-revenue trade-off—a central concern in mechanism design and strategic decision making.

2 Problem Statement and Online Optimization-based Budget Pacing Algorithm

In a sponsored search program, there are a number of search activities (impressions) $i \in \mathcal{I}$ being initiated at any given time $t \in [T]$. Whenever a user initiates a search, each ad campaign of the sponsored search program, denoted by $k \in [K]$, is considered for retrieval to an internal ad auction that determines the allocation of the sponsored placements for the search result page. For each impression i, campaign k determines a maximum bid it can post, denoted by $v_{k,i} \in [0, \bar{v}]$. Retrieval is then performed based on the campaigns' targeting strategies, which can include manual and auto targeting as well as their extensions. If campaign k's targeting strategy does not align with impression i, we simply let $v_{k,i} = 0$. Here, the maximum $v_{k,i}$ can be viewed as a proxy for the ad campaign's valuation of the ad slot. In addition to its maximum bid for each impression, the ad campaign also provides the platform with (i) its total budget, denoted by B_k, which is the total monetary value to be spent throughout the horizon; and (ii) its target spending curve, captured by $\rho_k \in \Delta_T$, where $\rho_{k,t}$ is the percentage of budget that campaign k wishes to allocate up to round t. The target spending curve determines the rate at which the campaign wishes to spend its budget over time. Some common examples include the uniform spending curve, the traffic curve, the CTR curve, etc.

Second-price multi-slot ads auction. For each search activity/impression i, the sponsored search program conducts a second-price multi-slot ads auction to determine the allocation of slots. External competition, including other ad programs or ad deals, can also enter this auction and thus impact the campaigns' strategies. Each campaign in the sponsored search program should submit a bid $b_{k,i} \in [0, v_{k,i}]$, bounded by its maximum bid value. If $b_{k,i} = 0$, campaign k opts out of this auction. Having collected the bids from all campaigns, the platform computes an ad expected value for each campaign: $r_{k,i} = b_{k,i} p_{k,i}$, which is determined using campaign k's bid value $b_{k,i}$ and its response probability $p_{k,i}$ (i.e., click-through rate for cost-per-click goal type). The campaigns are then ranked based on their ad expected values. If impression i offers N sponsored slots, the campaigns with the top N ad expected values would be allocated a slot in descending order. Each winning campaign is charged a clearing price $c_{k,i}$ if a user response is recorded (i.e., click). Here, the clearing price is computed based on the ad expected value of the campaign that immediately follows, defined as $d_{k,i} = \max_{i': r_{k,i'} < r_{k,i}} (r_{k,i'})$. The clearing price is subsequently determined as $c_{k,i} = d_{k,i} / p_{k,i}$, which is inherently less or equal to the bid $b_{k,i}$. For ease of presentation, we consider single-slot auctions when formulating the campaign's optimization problem and pacing algorithm. However, the same optimization framework would readily extend to a multi-slot setup, as discussed by Gaitonde et al. (2022). In our real-world experiments in Section 3, we also consider multiple ad slots.

Problem statement. Campaign k needs to determine $x_{k,i} \in \{0,1\}$, i.e., whether it wishes to win impression i. Note that if campaign k is not budget constrained, it would always wish to win the impression. However, with the budget constraint in mind, campaign k needs to focus on winning impressions that yield the highest utilities. The expected utility that campaign k receives from impression i is

$$x_{k,i}\left(v_{k,i} - c_{k,i}\right) p_{k,i} = x_{k,i}\left(v_{k,i} p_{k,i} - d_{k,i}\right), \tag{1}$$

where $v_{k,i} - c_{k,i}$ captures the utility gained by campaign k if it is shown in search i and receives a user response. Note that the advertiser only pays if its ad is clicked which occurs with probability $p_{k,i}$. The equality in (1) follows from the definition of the clearing price ($c_{k,i} = d_{k,i}/p_{k,i}$). Similarly, we can write the expected spending of campaign k on impression i as follows:

$$x_{k,i} c_{k,i} p_{k,i} = x_{k,i} d_{k,i}, \tag{2}$$

From (1) and (2), a campaign k that seeks to maximize its total expected utility throughout the horizon over $\{x_{k,i}\}_{i \in \mathcal{I}_t}, t \in [T]$, subject to its budget, minimum spending, and ROI constraints can thus solve the following optimization problem:

Objective. Each campaign k maximizes its expected gain: $\sum_{t=1}^{T} \sum_{i \in \mathcal{I}_t} x_{k,i}\left(v_{k,i} p_{k,i} - d_{k,i}\right)$, subject to the following constraints:

Budget constraint. $\sum_{t=1}^{T} \sum_{i \in \mathcal{I}_t} x_{k,i} d_{k,i} \leq B_k$.

Minimum spending constraint. $\sum_{t=1}^{T} \sum_{i \in \mathcal{I}_t} x_{k,i} d_{k,i} \geq \alpha_k B_k$.

ROI constraint.

Note that $\alpha_k \in [0,1]$ is the minimum percentage of budget that the platform would like campaign k to at least spend throughout the horizon. Also, our focus in this paper is on the following three different forms of practical interest for the ROI constraint:

- In the first form, the ROI constraint indicates that the expected total payment of the advertiser is less than or equal to a multiple of the expected utility it receives from winning auctions for impressions/items of interest. In that case, the ROI constraint can be expressed as: $\sum_{t=1}^{T} \sum_{i \in \mathcal{I}_t} x_{k,i} d_{k,i} \leq w_k \sum_{t=1}^{T} \sum_{i \in \mathcal{I}_t} x_{k,i}\left(v_{k,i} p_{k,i} - d_{k,i}\right)$.
- In the second form, the ROI constraint indicates that the expected total payment of the advertiser is less than or equal to a multiple of the expected number of clicks for items associated with the auctions it wins. In that case, the ROI constraint can be expressed as: $\sum_{t=1}^{T} \sum_{i \in \mathcal{I}_t} x_{k,i} d_{k,i} \leq w_k \sum_{t=1}^{T} \sum_{i \in \mathcal{I}_t} x_{k,i} p_{k,i}$.
- In the third form, the ROI constraint indicates that the expected total payment of the advertiser is less than or equal to a multiple of the number of impressions it receives (or equivalently, the number of auctions it wins). In that case, the ROI constraint can be expressed as: $\sum_{t=1}^{T} \sum_{i \in \mathcal{I}_t} x_{k,i} d_{k,i} \leq w_k \sum_{t=1}^{T} \sum_{i \in \mathcal{I}_t} x_{k,i}$.

We consider a Lagrangian-based approach to analytically derive the bid value (as a function of the dual variables associated with different constraints) for each campaign under the three different forms of the ROI constraint considered in this paper. Note that the bidding strategies in Chen et al. (2024a) hold as special cases of those derived in this paper. The detailed analysis is provided in the Appendix (A.2-A.4), and the algorithm is detailed in Algorithm 1. To improve readability, we also provide here the derived bid value under the first form of the ROI constraint as an example:

$$b_{k,i} = \frac{\left(1 + \eta_k\right) v_{k,i}}{1 + \eta_k + \left(\frac{\eta_k}{w_k} + \mu_k - \gamma_k\right)^{+}}, \tag{3}$$

where μ_k, γ_k and η_k are the dual variables associated with the budget, minimum spending, and ROI constraints, respectively, and $(\cdot)^{+} = \max(\cdot, 0)$. Further, the dual variables are updated in an online manner using a subgradient method as follows:

$$\mu_{k,t+1} = \left[\mu_{k,t} - \epsilon_{k,t}\left(\rho_{k,t} B_k - \bar{z}_{k,t}\right)\right]^{+}, \tag{4}$$

Algorithm 1 Generalized Adaptive Pacing Algorithm.

Input: $B_k, \alpha_k, w_k, \rho_k, \epsilon_{k,t}, \epsilon'_{k,t}, \epsilon''_{k,t}$
Initialize dual variables $\mu_{k,t}, \gamma_{k,t}, \eta_{k,t}$
for $t \in [T]$ **do**
  Whenever campaign k joins the auction for impression $i \in \mathcal{I}_t$, with maximum bid $v_{k,i}$
    (a) Post bid $b_{k,i}$ as in (3)
    (b) Realized spending $\bar{z}_{k,i} = d_{k,i}/p_{k,i}$ if campaign k wins the auction and gets clicked; $\bar{z}_{k,i} = 0$ otherwise
  Compute total realized spending $\bar{z}_{k,t} = \sum_{i \in \mathcal{I}_t} \bar{z}_{k,i}$ and $\bar{v}_{k,t}$
  Update dual variables as in (4)-(6)
end for

$$\gamma_{k,t+1} = \left[\gamma_{k,t} - \epsilon'_{k,t}\left(\bar{z}_{k,t} - \alpha_k \rho_{k,t} B_k\right)\right]^+, \tag{5}$$

$$\eta_{k,t+1} = \left[\eta_{k,t} - \epsilon''_{k,t}\left(\bar{v}_{k,t} - \bar{w}_k \bar{z}_{k,t}\right)\right]^+, \tag{6}$$

where $\epsilon_{k,t}, \epsilon'_{k,t}$ and $\epsilon''_{k,t}$ are the step sizes for updating the dual variables, and $\bar{w}_k = 1 + \frac{1}{w_k}$. In addition, $\bar{z}_{k,t}$ and $\bar{v}_{k,t}$ can be respectively expressed as

$$\bar{z}_{k,t} = \sum_{i \in \mathcal{I}_t} d_{k,i} \mathbf{1}\{(\eta_k + 1) v_{k,i} p_{k,i} \geq d_{k,i}(1 + \mu_k - \gamma_k + \bar{w}_k \eta_k)\},$$

$$\bar{v}_{k,t} = \sum_{i \in \mathcal{I}_t} v_{k,i} p_{k,i} \mathbf{1}\{(\eta_k + 1) v_{k,i} p_{k,i} \geq d_{k,i}(1 + \mu_k - \gamma_k + \bar{w}_k \eta_k)\}.$$

Business metrics. The budget pacing algorithms are evaluated in the next section based on a number of key system-level performance metrics. Specifically, we consider the following metrics: the number of impressions (Imps) won by the campaigns in our sponsored program; the number of clicks (Clks); total ad revenue (Rev); average click-through rate (CTR); average cost-per-click (CPC), a seller-oriented metric that evaluates their cost; surface rate (SR), which is the percentage of sponsored slots won by campaigns in the sponsored search program (recall that other ad programs/deals can also join the auctions and compete for visibility). Finally, another central metric we assess is the pacing error (PE), defined as: $\mathrm{PE} = \frac{1}{T}\sum_1^T \frac{|pS_t - pT_t|}{pT_t}$, with pS_t and pT_t as the fraction of total spend and traffic at t-th round. This metric measures smoothness of system-level spend over a day using the traffic curve as a reference. Note that a lower pacing error is more ideal.

3 Practical Research Questions and Experimental Results

In this section, we provide experimental results evaluating the performance of our proposed algorithms in practice and making comparisons with other commonly adopted budget pacing approaches. We leverage eBay's sponsored search test bed, and evaluate the performance of our algorithms using key business metrics mentioned in Section 2. For completeness, we elaborate on eBay's sponsored search test bed in Appendix A.1. In the following experiments, we address a number of research questions of practical importance to online platforms. Note that all of the results are expressed as percentage variations relative to the metrics observed under no-pacing (that is, each campaign k always posts maximum bid $v_{i,t}$ whenever it joins an auction for impression i, until its budget gets depleted). We let each round t span one minute and the entire time horizon is set to be $T = 1440$, which is the duration of one day. The default target spending curve used by our algorithms is the traffic curve, where $\rho_{k,t} = |\mathcal{I}_t| / \sum_{t=1}^{1440} |\mathcal{I}_t|$ represents the fraction of traffic during the t-th minute.

3.1 How does the form of the ROI constraint impact the business metrics?

We start our experiments by showing the impact of the choice of the ROI constraint on the business metrics. Towards this objective, first we compare our algorithm under budget + ROI constraints (i.e., no minimum spending constraints) to the adaptive pacing baseline Balseiro & Gur (2019), which enforces only the budget constraint. Note that when we only have budget and ROI constraints, our

Table 1: The key business and pacing performance metrics comparing AdapPac to adaptive pacing with different ROI constraints ($\epsilon_{k,t} = \epsilon''_{k,t} = 0.01$), under varying hyperparameter w_k. Results are shown as percentage variations relative to no-pacing.

	Imps	Clks	Rev	CTR	CPC	SR	PE
AdapPac	6.12%	5.45%	-5.26%	-0.47%	-10.06%	6.00%	-18.90%
Utility, $w_k = 1$	4.57%	4.59%	-10.74%	0.17%	-14.53%	4.60%	-2.03%
Utility, $w_k = 2$	4.92%	4.91%	-7.05%	0.18%	-11.32%	4.88%	-7.45%
Utility, $w_k = 10$	5.41%	5.21%	-4.10%	-0.02%	-8.83%	5.36%	-13.17%
Utility, $w_k = 50$	5.31%	5.18%	-3.62%	0.05%	-8.39%	5.28%	-14.71%
Utility, $w_k = 100$	5.41%	5.30%	-3.53%	0.05%	-8.41%	5.34%	-14.79%
ExpClks, $w_k = 1$	0.16%	-0.57%	-19.77%	-0.60%	-19.20%	1.90%	4.05%
ExpClks, $w_k = 2$	1.85%	1.37%	-18.88%	-0.34%	-19.77%	3.22%	2.95%
ExpClks, $w_k = 10$	5.71%	5.71%	-15.92%	0.10%	-20.24%	5.80%	0.55%
ExpClks, $w_k = 50$	6.23%	5.88%	-10.27%	-0.11%	-15.01%	5.93%	-5.37%
ExpClks, $w_k = 100$	5.78%	5.52%	-6.76%	-0.06%	-11.45%	5.76%	-10.99%
Imps, $w_k = 1$	6.34%	5.64%	-12.29%	-0.46%	-16.67%	6.13%	-1.76%
Imps, $w_k = 2$	6.18%	5.45%	-8.94%	-0.52%	-13.40%	5.94%	-7.13%
Imps, $w_k = 10$	**6.11%**	**5.54%**	**-3.79%**	**-0.34%**	**-8.77%**	**6.01%**	**-17.17%**
Imps, $w_k = 50$	6.23%	5.52%	-4.71%	-0.47%	-9.61%	6.04%	-18.51%
Imps, $w_k = 100$	6.23%	5.52%	-5.21%	-0.47%	-10.07%	6.07%	-18.04%

algorithm is referred to as the adaptive pacing with ROI constraint algorithm and the bid of campaign k on impression i ($b_{k,i}$) can be obtained by setting γ_k to zero in (3), (9) or (14) depending on the form of the ROI constraint considered. Table 1 compares the performance of our adaptive pacing with ROI constraint algorithm with that of the adaptive pacing (AdapPac) algorithm Balseiro & Gur (2019). The results in Table 1 are obtained under different forms of ROI constraint considered in this paper (where the abbreviations Utility, ExpClks and Imps respectively refer to the first, second and third forms of the ROI constraint).

Comparison with no-pacing. Compared to no-pacing, it can be seen from Table 1 that both adaptive pacing and our proposed algorithm under different forms of the ROI constraint notably improve several business metrics such as impressions, clicks, and surface rate. This, in turn, indicates that campaigns are utilizing their budgets more effectively and appearing more frequently in search results, compared to no-pacing. The pacing error and CPC also decrease significantly, further enhancing advertiser utilities. We observe a slight decrease in total ad revenue precisely because campaigns post bids less than their maximum bids, which reduces the clearing price (as seen in CPC) under the second-price auction mechanism.

Our proposed algorithm vs. Adaptive pacing. Table 1 shows that the choice of the form of the ROI constraint along with the value of w_k plays a key role in improving business metrics (compared to AdapPac). For the Utility and ExpClks variants, larger w_k generally increases impressions, clicks, and revenue. Under the Imps variant, the most impressions/clicks occur at $w_k = 1$ while revenue peaks at $w_k = 10$. The ExpClks form delivers especially strong gains in impressions, clicks, and surface rate—particularly at $w_k \in \{50, 100\}$. The Utility form yields the largest lift in total ad revenue. Finally, within Imps, choosing $w_k = 10$ raises revenue without reducing impressions or clicks versus AdapPac.

Tradeoff between the utilities campaigns receive and the revenue collected by the platform. Table1 shows a clear trade-off between advertiser utility and platform revenue. Impressions/clicks are maximized under ExpClks at $w_k = 50$, but this comes with a substantial drop in total ad revenue. Conversely, Utility at $w_k = 100$ yields the highest revenue while reducing impressions/clicks. A moderate setting under Imps (e.g., $w_k = 10$) strikes a practical balance—raising revenue without sacrificing impressions/clicks relative to AdapPac. Hence, the platform can choose the ROI formulation to match business goals. For example, if revenue loss is capped at 4%, Imps with $w_k = 10$ delivers the largest gains in impressions/clicks, CPC, surface rate, and pacing error among the feasible options in Table 1.

What are the ideal ranges of values for the step sizes $\epsilon_{k,t}$ and $\epsilon''_{k,t}$? In Table 2, we evaluate the performance of our adaptive pacing with ROI constraint algorithm under different values of the step sizes $\epsilon_{k,t}$ and $\epsilon''_{k,t}$ (used to update the dual variables μ_k and η_k associated with the budget and ROI

Table 2: Performance of adaptive pacing with ROI constraint-Utility ($w_k = 10$ and $T = 120$) under varying step sizes $\epsilon_{k,t}$ and $\epsilon''_{k,t}$.

$\epsilon_{k,t}$	$\epsilon''_{k,t}$	Imps	Clks	Rev	CTR	CPC	SR	PE
0.01	**0.01**	1.32%	0.49%	-9.96%	-0.85%	-10.77%	2.03%	11.56%
0.01	**0.1**	0.97%	0.77%	-5.21%	-0.30%	-6.09%	1.15%	4.58%
0.01	**0.5**	0.58%	0.13%	-3.69%	-0.49%	-3.95%	0.52%	-2.71%
0.01	5	0.36%	-0.18%	-3.27%	-0.55%	-3.23%	0.23%	0.53%
0.01	10	0.30%	-0.31%	-3.23%	-0.63%	-3.06%	0.12%	0.68%
0.1	0.01	-6.00%	-8.37%	-19.21%	-2.56%	-12.28%	-4.94%	31.46%
0.1	0.1	-1.10%	-1.64%	-13.70%	-0.41%	-12.60%	-0.35%	17.93%
0.1	0.5	-0.76%	-0.97%	-9.52%	-0.13%	-8.97%	-0.07%	12.28%
0.1	5	-1.05%	-1.00%	-5.81%	0.14%	-4.88%	-1.06%	-4.22%
0.1	10	-1.09%	-1.25%	-5.70%	-0.09%	-4.49%	-1.07%	-2.55%
0.5	0.01	-19.05%	-22.71%	-27.70%	-5.02%	-7.60%	-18.27%	80.09%
0.5	0.1	-9.38%	-11.55%	-20.41%	-2.45%	-10.62%	-8.56%	36.60%
0.5	0.5	-5.23%	-6.84%	-16.59%	-1.60%	-11.02%	-4.87%	27.78%
0.5	5	-3.34%	-4.20%	-10.09%	-0.80%	-6.37%	-3.38%	13.82%
0.5	10	-3.23%	-3.74%	-8.67%	-0.45%	-5.35%	-3.36%	10.26%
5	0.01	-31.13%	-35.21%	-36.10%	-6.81%	-2.60%	-32.26%	134.21%
5	0.1	-26.33%	-30.29%	-32.17%	-5.96%	-3.97%	-26.61%	100.89%
5	0.5	-18.24%	-22.56%	-26.87%	-5.67%	-6.18%	-18.27%	72.29%
5	5	-10.48%	-12.21%	-20.64%	-1.88%	-10.30%	-11.33%	42.44%
5	10	-8.45%	-10.24%	-17.45%	-1.99%	-8.66%	-8.67%	36.76%
10	0.01	-32.44%	-36.74%	-37.42%	-7.29%	-2.23%	-34.13%	143.40%
10	0.1	-29.47%	-33.52%	-34.53%	-6.45%	-2.65%	-30.38%	115.34%
10	0.5	-22.55%	-26.60%	-29.74%	-5.73%	-5.41%	-22.70%	80.41%
10	5	-11.56%	-13.14%	-21.92%	-1.61%	-10.58%	-12.46%	32.68%
10	10	-10.49%	-11.96%	-20.48%	-1.48%	-10.31%	-11.44%	28.74%

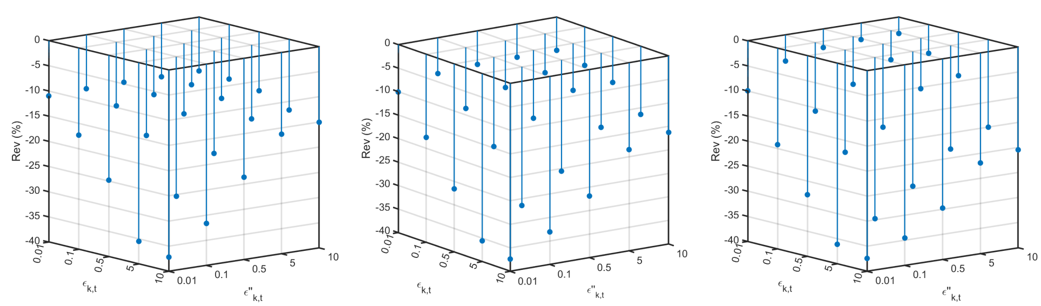

(a) (b) (c)

Figure 1: Impact of varing step sizes $\epsilon_{k,t}$ and $\epsilon''_{k,t}$ on the revenue for the adaptive pacing with ROI constraint-Utility : (a) $w_k = 1$, (b) $w_k = 2$, and (c) $w_k = 10$.

constraints, respectively) when considering the Utility form of the ROI constraint. Two patterns emerge. First, we observe that increasing $\epsilon''_{k,t}$ for a fixed value of $\epsilon_{k,t}$ leads to an improvement in the revenue. Second, $\epsilon_{k,t} = 0.01$ achieves the highest numbers of impressions and clicks. In order to achieve positive percentages of impressions and clicks (compared to no-pacing) while not much dropping the revenue, one can conclude from Table 2 that $\epsilon_{k,t} = 0.01$ and $\epsilon''_{k,t} \in [0.01, 0.5]$ are the best choices for tuning the step sizes. We obtained similar insights under the Imps and ExpClks forms of the ROI constraint (results omitted due to space limitations). This justifies our choice for the values of the step sizes used in all the experiments conducted in this paper. For better visualization, Fig. 1 shows how varying the step sizes affects platform revenue across different values of w_k.

Table 3: The key business and pacing performance metrics comparing AdapPacWithPenalty to generalized adaptive pacing with different ROI constraints ($\epsilon_{k,t} = \epsilon'_{k,t} = \epsilon''_{k,t} = 0.01$), under varying hyperparameter w_k.

	Imps	Clks	Rev	CTR	CPC	SR	PE
AdapPacWithPenalty	5.56%	5.05%	-3.28%	-0.31%	-7.98%	5.60%	-20.32%
Utility, $w_k = 1$	4.70%	4.90%	-7.78%	0.35%	-12.02%	4.65%	-7.47%
Utility, $w_k = 2$	4.69%	5.01%	-3.70%	0.45%	-8.29%	4.75%	-11.88%
Utility, $w_k = 10$	5.00%	4.90%	-2.05%	0.07%	-6.71%	4.90%	-14.62%
Utility, $w_k = 50$	4.97%	4.88%	-2.03%	0.05%	-6.65%	4.93%	-14.82%
Utility, $w_k = 100$	4.97%	4.87%	-2.00%	0.04%	-6.63%	4.91%	-15.06%
ExpClks, $w_k = 1$	2.57%	2.48%	-14.79%	0.03%	-16.68%	3.48%	-6.91%
ExpClks, $w_k = 2$	3.99%	3.92%	-14.08%	0.12%	-17.12%	4.59%	-7.97%
ExpClks, $w_k = 10$	6.30%	6.19%	-11.59%	0.06%	-16.50%	6.15%	-7.95%
ExpClks, $w_k = 50$	5.56%	5.34%	-6.03%	-0.01%	-10.66%	5.52%	-12.15%
ExpClks, $w_k = 100$	5.30%	5.15%	-3.53%	0.03%	-8.15%	5.23%	-14.93%
Imps, $w_k = 1$	5.88%	5.39%	-7.94%	-0.27%	-12.33%	5.67%	-9.66%
Imps, $w_k = 2$	5.48%	5.02%	-5.14%	-0.21%	-9.45%	5.35%	-12.79%
Imps, $w_k = 10$	5.40%	4.86%	-2.22%	-0.28%	-6.77%	5.45%	-16.30%
Imps, $w_k = 50$	5.62%	5.11%	-2.70%	-0.32%	-7.52%	5.59%	-19.33%
Imps, $w_k = 100$	5.66%	5.10%	-3.01%	-0.35%	-7.76%	5.59%	-19.75%

3.2 How does the minimum spending constraint impact the business metrics?

Our objective now is to quantify how the minimum spending constraint impacts the tradeoff between the utilities campaigns receive and the revenue collected by the platform (as described in Subsection 3.1). Towards this objective, under different forms of the ROI constraint, we evaluate the performance of our proposed algorithm after adding the minimum spending constraint in Table 3. Note that when considering the minimum spending constraint along with the budget and ROI constraints, we refer to our algorithm as the generalized adaptive pacing algorithm. Although the platform can set the minimum spending percentage α_k based on its business need, α_k is considered in our experiments to be the percentage of the budget spent by campaign k under no-pacing for fair comparison purposes.

Controlling the tradeoff between the advertisers' utilities and the platform's revenue. By comparing each row in Table 1 with its corresponding one in Table 3 (i.e., under the same form of the ROI constraint and value of w_k used), one can notice that the revenue is notably increased in Table 3. This improvement in revenue is also associated with a slight reduction in the number of impressions/clicks. This indicates that setting α_k to be the percentage of the budget spent by campaign k under no-pacing is a good choice to balance the advertisers' utilities and the platform's revenue. Indeed, the platform can arbitrarily control this tradeoff for each campaign through the choice of α_k. Further, by comparing our proposed algorithms under different forms of the ROI constraints with the AdapPacWithPenalty algorithm (where only the budget and minimum spending constraints are considered), one can conclude that by appropriately choosing w_k, the incorporation of the ROI constraint can maintain the same achievable number of impressions and clicks under the AdapPacWithPenalty algorithm as well as increase the revenue. For example, under the Imps ROI form, choosing $w_k \in \{50, 100\}$ exhibits this behavior.

Impact of the target spending curve on business metrics. Table 4 evaluates our algorithms under different target spending profiles $\boldsymbol{\rho}_k$ which impacts the dual variables (cf.(4)): (i) *Traffic curve*, with $\rho_{k,t}$ proportional to the share of historical traffic in minute t; (ii) *Uniform curve*, with $\rho_{k,t} = 1/1440$ (uniform daily spend); (iii) *CTR curve*, with $\rho_{k,t}$ proportional to the historical CTR profile. Across all ROI formulations, the traffic curve delivers the largest gains in impressions and clicks and while ensuring the lowest pacing error (Table 4). This shows that our proposed approach improves performance across different target spending curves.

3.3 How does our proposed algorithms measure up against commonly adopted methods?

After demonstrating the efficacy of our theoretically-motivated budget pacing algorithms with ROI constraints in the sponsored search environment, it is obvious to wonder about how they measure

Table 4: Performance of generalized adaptive pacing ($\epsilon_{k,t} = 0.01, \epsilon'_{k,t} = \epsilon''_{k,t} = 0.1, w_k = 10$) under different target spending curves.

	Imps	Clks	Rev	CTR	CPC	SR	PE
Utility, Traffic	**2.03%**	**2.17%**	**-0.50%**	0.11%	**-2.71%**	**2.29%**	**-4.73%**
Utility, Uniform	1.51%	1.62%	-0.58%	**0.12%**	-2.24%	1.77%	-1.88%
Utility, CTR	1.53%	1.61%	-0.55%	**0.12%**	-2.18%	1.78%	-1.38%
ExpClks, Traffic	**5.56%**	**5.89%**	-13.21%	**0.45%**	**-17.86%**	**5.55%**	**-8.20%**
ExpClks, Uniform	5.29%	5.46%	-12.89%	0.30%	-17.28%	5.46%	1.38%
ExpClks, CTR	5.29%	5.47%	**-12.77%**	0.28%	-17.20%	5.42%	2.17%
Imps, Traffic	**3.58%**	**3.33%**	-0.81%	-0.10%	**-4.17%**	**3.99%**	**-9.97%**
Imps, Uniform	2.92%	2.80%	**-0.75%**	**0.01%**	-3.54%	3.68%	-5.62%
Imps, CTR	2.95%	2.75%	-0.76%	-0.10%	-3.44%	3.71%	-4.94%

against other commonly adopted budget pacing methods. In order to address this, we focus here on comparing the performance of our algorithms with that of the PID controllers (which are widely used in the industry as heuristics for budget pacing Zhang et al. (2016)).

In order to easily understand how PID controllers can be used to handle our optimization problem with both budget and ROI constraints, it is useful to first elaborate on the interesting connection between the PID controller and the adaptive pacing algorithm Balseiro & Gur (2019). Recall that the adaptive pacing algorithm was proposed in Balseiro & Gur (2019) to solve a special case of the optimization problem considered in this paper, where the budget constraint was only considered. In particular, when the budget constraint is only considered in the optimization problem, the bid value $b_{k,i}$ in (3) reduces to: $b_{k,i} = \frac{v_{k,i}}{1+\mu_k}$. Further, the dual variable associated with the budget constraint μ_k can be updated in an online manner using a subgradient method as: $\mu_{k,t+1} = [\mu_{k,t} - \epsilon_{k,t}\left(\rho_{k,t} B_k - \bar{z}_{k,t}\right)]^+$, where $\bar{z}_{k,t}$ can be expressed as: $\bar{z}_{k,t} = \sum_{i \in \mathcal{I}_t} d_{k,i} \mathbf{1}\{v_{k,i} p_{k,i} \geq d_{k,i}(1+\mu_k)\}$. On the other hand, if a dual-based PID controller is used to update the dual multiplier μ_k, the update would be expressed as: $\mu_{k,t+1} = \left(\mu_{k,t} - \lambda_{\mathrm{P}} g_t - \lambda_{\mathrm{I}} \sum_{n=0}^{t-1} g_{t-n} - \lambda_{\mathrm{D}}\left(g_t - g_{t-1}\right)\right)^+$, where $\lambda_{\mathrm{P}}, \lambda_{\mathrm{I}}, \lambda_{\mathrm{D}} > 0$ are the step sizes associated with the proportional (P),integral (I), and derivative (D) terms, respectively. Further, the error term g_t is equal to the subgradient of the Lagrangian at time t, i.e., $g_t = \rho_{k,t} B_k - \bar{z}_{k,t}$. One can clearly see that the adaptive pacing algorithm is indeed a special case of the PID controller (also known as a P controller) where $\lambda_{\mathrm{I}} = \lambda_{\mathrm{D}} = 0$. The PID controller additionally incorporates past momentum and optimism through the integral and derivative terms, respectively.

Now, we compare the performance of our theoretically-motivated algorithms with that of PID controllers when both the budget and ROI constraints are considered. To handle this case using PID controllers, we use two dual-based PID controllers to control the two dual variables μ_k and η_k associated with the budget and ROI constraints, respectively. By inspecting the update rules for μ_k and η_k in (4) and (6) after setting γ_k to zero (i.e., when our adaptive pacing with ROI constraint-utility algorithm is considered), one can realize that these update rules can be achieved with the use of two P controllers (which is a special case of two dual-based PID controllers). This establishes the connection between our proposed algorithm (when both the budget and ROI constraints are considered) and two dual-based PID controllers, and shows a similar analogy to the connection between the adaptive pacing algorithm and one dual-based PID controller when the budget constraint is only considered. In practice, the PID controllers are often designed to update control variables (i.e., the dual variables μ_k and η_k in our case), and the control variables then directly impact the bid factor defined using some function $\mathcal{F}$. That is, for $i \in \mathcal{I}_t$, the campaign k posts bid of the following form $b_{k,i} = v_{k,i} \mathcal{F}(\mu_{k,t}, \eta_{k,t})$. Note from (3) that under the first form of the ROI constraint, the bid value for our adaptive pacing with ROI constraint algorithm (i.e., after setting γ_k to zero in (3)) can be expressed as: $b_{k,i} = v_{k,i}\left(1+c\right)^{-1}$, where $c = \frac{1}{1+\eta_k}\left[\frac{\eta_k}{\mu_k} + \mu_k\right]^+$. Inspired by that, when the PID controllers are used while considering the first form of the ROI constraint, $\mathcal{F}(\mu_{k,t}, \eta_{k,t})$ should be a decreasing function of c. Following the same procedure, one can derive the expression of c under the other forms of the ROI constraint.

In Table 5, we consider the following different forms of $\mathcal{F}(\mu_{k,t}, \eta_{k,t})$ Chen et al. (2024a): i) Linear: $\mathcal{F}(\mu_{k,t}, \eta_{k,t}) = -\frac{1}{5}c + 1$, ii) Exponential: $\mathcal{F}(\mu_{k,t}, \eta_{k,t}) = \exp(-c)$, and iii) AdapPacWROI:

Table 5: The key business and pacing performance metrics comparing PID controllers with different forms of $\mathcal{F}(\mu_{k,t}, \eta_{k,t})$ to adaptive pacing with different ROI constraints ($w_k = 10$).

	Imps	Clks	Rev	CTR	CPC	SR	PE
Utility	**6.14%**	**5.41%**	-5.28%	-0.51%	**-10.04%**	**5.99%**	**-18.75%**
PID-Utility, AdapPacWROI	1.01%	1.04%	-0.74%	**0.08%**	-1.82%	1.11%	-1.75%
PID-Utility, Exponential	1.17%	1.18%	-0.98%	0.04%	-2.17%	1.34%	-2.00%
PID-Utility, Linear	0.32%	0.38%	**-0.30%**	**0.08%**	-0.74%	0.43%	-0.08%
ExpClks	**5.71%**	**5.71%**	-15.92%	**0.10%**	**-20.24%**	**5.80%**	0.55%
PID-ExpClks, AdapPacWROI	1.54%	1.59%	-1.86%	0.08%	-3.37%	1.77%	-2.52%
PID-ExpClks, Exponential	1.02%	0.95%	-2.48%	-0.06%	-3.38%	1.07%	**-3.51%**
PID-ExpClks, Linear	0.53%	0.60%	**-1.39%**	**0.10%**	-1.98%	0.56%	-1.82%
Imps	**6.12%**	**5.52%**	-3.72%	-0.40%	**-8.71%**	**6.03%**	**-17.15%**
PID-Imps, AdapPacWROI	1.08%	1.08%	**-0.60%**	**0.03%**	-1.69%	1.29%	-1.22%
PID-Imps, Exponential	1.24%	1.06%	-1.58%	-0.14%	-2.62%	1.33%	-2.90%
PID-Imps, Linear	0.31%	0.27%	-0.99%	-0.01%	-1.33%	0.24%	-1.54%

$\mathcal{F}(\mu_{k,t}, \eta_{k,t}) = \frac{1}{1+c}$. Note that the abbreviation AdapPacWROI refers to the case where the bid value derived in our proposed algorithm (when considering both budget and ROI constraints) is used in designing the two dual-based PID controllers. We observe from Table 5 that under each form of the ROI constraint considered in this paper, the performance of our proposed algorithm significantly outperforms that of the dual-based PID controllers in terms of increasing the numbers of impressions, clicks and surface rate as well as reducing the pacing error and cost-per-click for advertisers. This indeed demonstrates the efficacy of implementing our proposed algorithms in real-world environments, compared to widely used heuristics for budget pacing such as PID controllers.

4 Conclusion

This paper developed new online optimization-based budget pacing algorithms that account for the budget and ROI constraints of the advertisers as well as the minimum spending constraints of the platform. In particular, the ROI constraint was constructed using three different forms of practical interest. For each form of the ROI constraint considered in this paper, a Lagrangian-based approach was utilized to analytically derive the bid value of each campaign as a function of the dual variables associated with different constraints. We demonstrated the efficacy of our developed algorithms by investigating their performances within the eBay's sponsored search environment.

Our extensive experiments provided several system design insights as well as answers to a number of practical research questions. For instance, our results demonstrated that the chosen form of the ROI constraint plays a key role in the achievable tradeoff between the advertisers' and the platform's interests. They also showed that the platform could control this tradeoff through the minimum spending constraint. Furthermore, our experiments showed the impact of different system design parameters, including the ROI constraint multiplier, step sizes for updating the dual variables, and spending patterns, on different business metrics. Finally, our results quantified the significant improvements in the business metrics associated with the implementation of our developed algorithms, compared to PID controllers which are widely used in industry as heuristics for budget pacing.

Our focus in this paper was on optimizing budget pacing for each advertiser within a single sponsored search program. A promising direction of future work is to incorporate the cross-channel marketing effects into our optimization problem, where the budget pacing for each advertiser can be optimized across multiple advertising channels.

Acknowledgments

This work was supported in part by eBay eRUPT program and the NSF AI Institute (AI-EDGE) 2112471.

References

Deepak Agarwal, Souvik Ghosh, Kai Wei, and Siyu You. Budget pacing for targeted online advertisements at linkedin. In *Proc. of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 1613–1619, 2014.

Santiago Balseiro, Anthony Kim, Mohammad Mahdian, and Vahab Mirrokni. Budget management strategies in repeated auctions. In *Proc. of the 26th International Conference on World Wide Web*, pp. 15–23, 2017.

Santiago R Balseiro and Yonatan Gur. Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Management Science*, 65(9):3952–3968, 2019.

Santiago R Balseiro, Omar Besbes, and Gabriel Y Weintraub. Repeated auctions with budgets in ad exchanges: Approximations and design. *Management Science*, 61(4):864–884, 2015.

Santiago R Balseiro, Omar Besbes, and Gabriel Y Weintraub. Dynamic mechanism design with budget-constrained buyers under limited commitment. *Operations Research*, 67(3):711–730, 2019.

Denis Charles, Deeparnab Chakrabarty, Max Chickering, Nikhil R Devanur, and Lei Wang. Budget smoothing for internet ad auctions: a game theoretic approach. In *Proc. of the fourteenth ACM conference on Electronic commerce*, pp. 163–180, 2013.

Qinyi Chen, Phuong Ha Nguyen, and Djordje Gligorijevic. Optimization-based budget pacing in ebay sponsored search. In *Companion Proc. of the ACM Web Conference 2024*, pp. 328–337, 2024a.

Zhaohua Chen, Chang Wang, Qian Wang, Yuqi Pan, Zhuming Shi, Zheng Cai, Yukun Ren, Zhihua Zhu, and Xiaotie Deng. Dynamic budget throttling in repeated second-price auctions. In *Proc. of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 9598–9606, 2024b.

Dragos Florin Ciocan and Krishnamurthy Iyer. Tractable equilibria in sponsored search with endogenous budgets. *Operations research*, 69(1):227–244, 2021.

Vincent Conitzer, Christian Kroer, Eric Sodomka, and Nicolas E Stier-Moses. Multiplicative pacing equilibria in auction markets. *Operations Research*, 70(2):963–989, 2022.

Zhe Feng, Swati Padmanabhan, and Di Wang. Online bidding algorithms for return-on-spend constrained advertisers. In *Proc. of the ACM Web Conference 2023*, pp. 3550–3560, 2023.

Jason Gaitonde, Yingkai Li, Bar Light, Brendan Lucier, and Aleksandrs Slivkins. Budget pacing in repeated auctions: Regret and efficiency without convergence. *arXiv preprint arXiv:2205.08674*, 2022.

Gagan Goel, Vahab Mirrokni, and Renato Paes Leme. Clinching auctions with online supply. *Games and Economic Behavior*, 123:342–358, 2020.

Chinmay Karande, Aranyak Mehta, and Ramakrishnan Srikant. Optimizing budget constrained spend in search advertising. In *Proc. of the sixth ACM international conference on Web search and data mining*, pp. 697–706, 2013.

Niklas Karlsson and Jianlong Zhang. Applications of feedback control in online advertising. In *2013 American control conference*, pp. 6008–6013, 2013.

Kuang-Chih Lee, Ali Jalali, and Ali Dasdan. Real time bid optimization with smooth budget delivery in online advertising. In *Proc. of the seventh international workshop on data mining for online advertising*, pp. 1–9, 2013.

Brendan Lucier, Sarath Pattathil, Aleksandrs Slivkins, and Mengxiao Zhang. Autobidders with budget and roi constraints: Efficiency, regret, and pacing dynamics. In *The Thirty Seventh Annual Conference on Learning Theory*, pp. 3642–3643. PMLR, 2024.

Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. Adwords and generalized online matching. *Journal of the ACM (JACM)*, 54(5):22–es, 2007.

Phuong Ha Nguyen, Djordje Gligorijevic, Arnab Borah, Gajanan Adalinge, and Abraham Bagherjeiran. Practical budget pacing algorithms and simulation test bed for ebay marketplace sponsored search. In *AdKDD@ KDD*, 2023.

Weinan Zhang, Shuai Yuan, and Jun Wang. Optimal real-time bidding for display advertising. In *Proc. of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 1077–1086, 2014.

Weinan Zhang, Yifei Rong, Jun Wang, Tianchi Zhu, and Xiaofan Wang. Feedback control of real-time display advertising. In *Proc. of the Ninth ACM International Conference on Web Search and Data Mining*, pp. 407–416, 2016.

A Appendix

A.1 eBay Sponsored Search Test Bed

The sponsored search test bed is a simulator constructed from historical user search logs that reproduces the full product environment over a one-day horizon, corresponding to the typical budget duration of sponsored search campaigns. To ensure scalability, the entire simulation period is partitioned into small time intervals, within which each search request and ad auction is modeled independently. Although real-time budget signal estimation is theoretically optimal, its implementation can hinder research efficiency and may destabilize production systems due to its substantial computational overhead. Accordingly, near real-time budget updates with a one-minute resolution are employed as a practical trade-off between computational cost and modeling accuracy.

Targeting-based retrieval (a core component of real-world sponsored search environments that requires sorting by relevance and other critical metrics) is typically complex and computationally intensive. To facilitate this process, a distinct targeting set is constructed for each search query using logged recall sets. The bid associated with an item within a campaign depends on its targeting strategy, particularly keyword targeting, which may overlap across ad groups. As a result, the expected value of an advertisement is computed by prioritizing the ad group with the highest bid, consistent with the behavior of the production system. User response probabilities are generated from available features, while the user responses that result in spend are simulated using a counterfactual modeling approach.

The quality of the test bed was previously evaluated in Nguyen et al. (2023) through comparisons with results obtained from naive simulation baselines. While discrepancies may arise due to the approximations inherent in the test bed, the reported results demonstrate that the key components of the sponsored search environment, particularly response probability generation, lead to substantial improvements across multiple performance metrics.

A.2 First Form of the ROI Constraint

To solve the above optimization problem in a real-time fashion, we consider a Lagrangian-based approach. In particular, the Lagrangian can be expressed as

$$\mathcal{L}\ \left(\{x_{k,i}\}_{i\in\mathcal{I}_t, t\in[T]}, \mu_k, \gamma_k, \eta_k\right) = \sum_{t=1}^{T}\sum_{i\in\mathcal{I}_t} x_{k,i}\left(v_{k,i}p_{k,i} - d_{k,i}\right) + \mu_k\left(B_k - \sum_{t=1}^{T}\sum_{i\in\mathcal{I}_t} x_{k,i}d_{k,i}\right)$$
$$+ \gamma_k\left(\sum_{t=1}^{T}\sum_{i\in\mathcal{I}_t} x_{k,i}d_{k,i} - \alpha_k B_k\right) + \eta_k\left(\sum_{t=1}^{T}\sum_{i\in\mathcal{I}_t} x_{k,i}\left(v_{k,i}p_{k,i} - \bar{w}_k d_{k,i}\right)\right), \tag{7}$$

where μ_k, γ_k and η_k are the dual variables associated with the budget, minimum spending, and ROI constraints, respectively, and $\bar{w}_k = 1 + \frac{1}{w_k}$. Therefore, the dual function $\mathcal{G}(\mu_k, \gamma_k, \eta_k)$ can be obtained as $\mathcal{G}(\mu_k, \gamma_k, \eta_k) = \max_{\mathbf{x}_k} \mathcal{L}(\mathbf{x}_k, \mu_k, \gamma_k, \eta_k)$, where $\mathbf{x}_k$ is a vector that contains the variables $\{x_{k,i}\}_{i\in\mathcal{I}_t}, t \in [T]$. Further, the dual problem will be: $\min_{\mu_k,\gamma_k,\eta_k} \mathcal{G}(\mu_k, \gamma_k, \eta_k)$. To obtain the dual function by maximizing the Lagrangian in (7) over $\{x_{k,i}\}_{i\in\mathcal{I}_t}, t \in [T]$, it is useful to note that (7) can be rewritten as : $\mathcal{L}(\mathbf{x}_k, \mu_k, \gamma_k, \eta_k) =$

$$\sum_{t=1}^{T} \left[\rho_{k,t} B_k \left(\mu_k - \gamma_k \alpha_k\right) + \sum_{i \in \mathcal{I}_t} x_{k,i} \big[(\eta_k + 1) v_{k,i} p_{k,i} - d_{k,i}(1 + \mu_k - \gamma_k + \bar{w}_k \eta_k) \big] \right].$$

Thus, the optimal $\mathbf{x}_k$ that maximizes the Lagrangian is such that $x_{k,i} = \mathbf{1}\{(\eta_k + 1)v_{k,i}p_{k,i} \geq d_{k,i}(1 + \mu_k - \gamma_k + \bar{w}_k \eta_k)\}$, where $\mathbf{1}\{\cdot\}$ is the indicator function. That is, the campaign k wants to win all auctions such that $(\eta_k + 1)v_{k,i}p_{k,i} \geq d_{k,i}(1 + \mu_k - \gamma_k + \bar{w}_k \eta_k)$. This suggests setting the bid of campaign k on impression i as:

$$b_{k,i} = \frac{(1 + \eta_k)\, v_{k,i}}{1 + \mu_k - \gamma_k + \bar{w}_k \eta_k} \overset{(a)}{\rightarrow} \frac{(1 + \eta_k)\, v_{k,i}}{1 + \eta_k + \left(\frac{\eta_k}{w_k} + \mu_k - \gamma_k \right)^+},$$

where step (a) ensures that $0 \leq b_{k,i} \leq v_{k,i}$, and $(\cdot)^+ = \max(\cdot, 0)$. After substituting the optimal $\mathbf{x}_k$ in the Lagrangian, we can update the dual variables in an online manner using a subgradient method as follows:

$$\mu_{k,t+1} = \left[\mu_{k,t} - \epsilon_{k,t} \left(\rho_{k,t} B_k - \bar{z}_{k,t}\right)\right]^+ ,$$

$$\gamma_{k,t+1} = \left[\gamma_{k,t} - \epsilon'_{k,t} \left(\bar{z}_{k,t} - \alpha_k \rho_{k,t} B_k\right)\right]^+ ,$$

$$\eta_{k,t+1} = \left[\eta_{k,t} - \epsilon''_{k,t} \left(\bar{v}_{k,t} - \bar{w}_k \bar{z}_{k,t}\right)\right]^+ ,$$

where $\epsilon_{k,t}, \epsilon'_{k,t}$ and $\epsilon''_{k,t}$ are the step sizes for updating the dual variables. In addition, $\bar{z}_{k,t}$ and $\bar{v}_{k,t}$ can be respectively expressed as

$$\bar{z}_{k,t} = \sum_{i \in \mathcal{I}_t} d_{k,i} \mathbf{1}\{(\eta_k + 1)v_{k,i}p_{k,i} \geq d_{k,i}(1 + \mu_k - \gamma_k + \bar{w}_k \eta_k)\},$$

$$\bar{v}_{k,t} = \sum_{i \in \mathcal{I}_t} v_{k,i} p_{k,i} \mathbf{1}\{(\eta_k + 1)v_{k,i}p_{k,i} \geq d_{k,i}(1 + \mu_k - \gamma_k + \bar{w}_k \eta_k)\}.$$

A.3 Second Form of the ROI Constraint

The Lagrangian here can be expressed as

$$\mathcal{L} \left(\{x_{k,i}\}_{i \in \mathcal{I}_t, t \in [T]}, \mu_k, \gamma_k, \eta_k \right) = \sum_{t=1}^{T} \Bigg[\rho_{k,t} B_k \left(\mu_k - \gamma_k \alpha_k\right) + \sum_{i \in \mathcal{I}_t} x_{k,i} \big[p_{k,i}(v_{k,i} + \eta_k w_k) - d_{k,i}(1 + \mu_k - \gamma_k + \eta_k) \big] \Bigg]. \tag{8}$$

Thus, the optimal $\mathbf{x}_k$ that maximizes the Lagrangian is such that $x_{k,i} = \mathbf{1}\{(v_{k,i} + \eta_k w_k)p_{k,i} \geq d_{k,i}(1 + \mu_k - \gamma_k + \eta_k)\}$. That is, the campaign k wants to win all auctions such that $(v_{k,i} + \eta_k w_k)p_{k,i} \geq d_{k,i}(1 + \mu_k - \gamma_k + \eta_k)$. This suggests setting the bid of campaign k on impression i as:

$$b_{k,i} = \frac{v_{k,i} + \eta_k w_k}{1 + \mu_k - \gamma_k + \eta_k} \overset{(a)}{\rightarrow} \frac{v_{k,i} + \eta_k w_k}{1 + (\mu_k - \gamma_k + \eta_k)^{++}}, \tag{9}$$

where step (a) ensures that $0 \leq b_{k,i} \leq v_{k,i}$, and $(\mu_k - \gamma_k + \eta_k)^{++} = \max(\mu_k - \gamma_k + \eta_k, \frac{\eta_k w_k}{v_{k,i}})$. After substituting the optimal $\mathbf{x}_k$ in the Lagrangian, we can update the dual variables in an online manner using a subgradient method as follows:

$$\mu_{k,t+1} = \left[\mu_{k,t} - \epsilon_{k,t} \left(\rho_{k,t} B_k - \bar{z}_{k,t}\right)\right]^+ , \tag{10}$$

$$\gamma_{k,t+1} = \left[\gamma_{k,t} - \epsilon'_{k,t} \left(\bar{z}_{k,t} - \alpha_k \rho_{k,t} B_k\right)\right]^+ , \tag{11}$$

$$\eta_{k,t+1} = \left[\eta_{k,t} - \epsilon''_{k,t} \left(\bar{p}_{k,t} - \bar{z}_{k,t}\right)\right]^+ , \tag{12}$$

where $\bar{p}_{k,t}$ can be expressed as

$$\bar{p}_{k,t} = \sum_{i \in \mathcal{I}_t} w_k p_{k,i} \mathbf{1}\{(v_{k,i} + \eta_k w_k)p_{k,i} \geq d_{k,i}(1 + \mu_k - \gamma_k + \eta_k)\}.$$

A.4 THIRD FORM OF THE ROI CONSTRAINT

The Lagrangian here can be expressed as

$$\mathcal{L}\ \left(\{x_{k,i}\}_{i\in\mathcal{I}_t, t\in[T]}, \mu_k, \gamma_k, \eta_k\right) = \sum_{t=1}^{T}\Bigg[\rho_{k,t}B_k\left(\mu_k - \gamma_k\alpha_k\right) + \sum_{i\in\mathcal{I}_t} x_{k,i}\big[p_{k,i}v_{k,i} + \eta_k w_k - d_{k,i}(1+\mu_k-\gamma_k+\eta_k)\big]\Bigg]. \quad (13)$$

Thus, the optimal $\mathbf{x}_k$ that maximizes the Lagrangian is such that $x_{k,i} = \mathbf{1}\{p_{k,i}v_{k,i} + \eta_k w_k \geq d_{k,i}(1+\mu_k-\gamma_k+\eta_k)\}$. That is, the campaign k wants to win all auctions such that $p_{k,i}v_{k,i}+\eta_k w_k \geq d_{k,i}(1+\mu_k-\gamma_k+\eta_k)$. This suggests setting the bid of campaign k on impression i as:

$$b_{k,i} = \frac{p_{k,i}v_{k,i} + \eta_k w_k}{p_{k,i}(1+\mu_k-\gamma_k+\eta_k)} \overset{\text{(a)}}{\to} \frac{p_{k,i}v_{k,i} + \eta_k w_k}{p_{k,i}[1+(\mu_k-\gamma_k+\eta_k)^{+*}]}, \quad (14)$$

where step (a) ensures that $0 \leq b_{k,i} \leq v_{k,i}$, and $(\mu_k-\gamma_k+\eta_k)^{+*} = \max(\mu_k-\gamma_k+\eta_k, \frac{\eta_k w_k}{p_{k,i}v_{k,i}})$. After substituting the optimal $\mathbf{x}_k$ in the Lagrangian, we can update the dual variables in an online manner using a subgradient method as follows:

$$\mu_{k,t+1} = \left[\mu_{k,t} - \epsilon_{k,t}\left(\rho_{k,t}B_k - \bar{z}_{k,t}\right)\right]^{+}, \quad (15)$$

$$\gamma_{k,t+1} = \left[\gamma_{k,t} - \epsilon'_{k,t}\left(\bar{z}_{k,t} - \alpha_k\rho_{k,t}B_k\right)\right]^{+}, \quad (16)$$

$$\eta_{k,t+1} = \left[\eta_{k,t} - \epsilon''_{k,t}\left(\bar{w}_{k,t} - \bar{z}_{k,t}\right)\right]^{+}, \quad (17)$$

where $\bar{w}_{k,t}$ can be expressed as

$$\bar{w}_{k,t} = \sum_{i\in\mathcal{I}_t} w_k\mathbf{1}\{p_{k,i}v_{k,i} + \eta_k w_k \geq d_{k,i}(1+\mu_k-\gamma_k+\eta_k)\}.$$