Budget Management Strategies in Repeated Auctions*

Santiago Balseiro
Duke University
srb43@duke.edu

Anthony Kim†
Stanford University
tonyekim@stanford.edu

Mohammad Mahdian
Google Research
mahdian@alum.mit.edu

Vahab Mirrokni
Google Research
mirrokni@theory.lcs.mit.edu

ABSTRACT

In online advertising, advertisers purchase ad placements by participating in a long sequence of repeated auctions. One of the most important features advertising platforms often provide, and advertisers often use, is budget management, which allows advertisers to control their cumulative expenditures. Advertisers typically declare the maximum daily amount they are willing to pay, and the platform adjusts allocations and payments to guarantee that cumulative expenditures do not exceed budgets. There are multiple ways to achieve this goal, and each one, when applied to all budget-constrained advertisers simultaneously, steers the system toward a different equilibrium. While previous research focused on online stochastic optimization techniques or game-theoretic equilibria of such settings, our goal in this paper is to compare the "system equilibria" of a range of budget management strategies in terms of the seller's profit and buyers' utility. In particular, we consider six different budget management strategies including probabilistic throttling, thresholding, bid shading, reserve pricing, and multiplicative boosting. We show these methods admit a system equilibrium in a rather general setting, and prove dominance relations between them in a simplified setting. Our study sheds light on the impact of budget management strategies on the tradeoff between the seller's profit and buyers' utility. Finally, we also empirically compare the system equilibria of these strategies using real ad auction data in sponsored search and randomly generated bids. The empirical study confirms our theoretical findings about the relative performances of budget management strategies.

*We thank Omar Besbes and Gabriel Weintraub for helpful discussions that motivated this research. The full paper version is available at `http://ssrn.com/abstract=2858261`.

†Part of this work was done when the author was an intern at Google, Inc.

WWW 2017, April 3–7, 2017, Perth, Australia.
ACM 978-1-4503-4913-0/17/04.
http://dx.doi.org/10.1145/3038912.3052682

Keywords

Online advertising; ad auctions; budget management

1. INTRODUCTION

Search engines and online ad exchanges run billions of auctions every day to sell advertising opportunities. A common feature such platforms often offer to their advertisers is budget management: advertisers can specify the maximum daily amount they are willing to spend, and the platform guarantees that they will never have to pay more than the budget. There are various ways to enforce such budget constraints. For example, the platform can 1) simply stop serving an advertiser as soon as his budget is exhausted; 2) probabilistically throttle the advertiser throughout the day at a rate that ensures the advertiser exhausts his budget close to the end of the day; or 3) achieve the same rate of spend by shading the bids of the advertiser. These methods are different in terms of the values generated for the advertisers and for the platform.

The goal of this paper is to compare a range of budget management mechanisms in the context of repeated auctions for online advertising applications. While previous research focused on online stochastic optimization techniques [4, 7, 19] or game theoretic equilibria of such settings [1, 5, 14], we investigate the, so called, system equilibria achieved under different mechanisms when the budgets of all advertisers are handled through a single mechanism and compare in terms of the seller's profit and buyers' utility. The system equilibria is not a game-theoretic notion of equilibria, as we are not considering the strategic interactions of budget-constrained advertisers bidding to maximize their utilities, but rather the market outcome when budgets are managed by a particular mechanism implemented by the platform and buyers report their values truthfully.

In the standard models of utility maximization, it is often in the advertisers' best interest not to reveal their budget constraints and instead submit bids that maximize their utilities subject to their budgets (in other words, directly manage budgets themselves instead of allowing the platform to manage their budgets). In practice, however, advertisers do not follow such game-theoretic equilibrium strategies, and a significant percentage of advertisers use the platform's budget management mechanism. This may be due to a variety of reasons such as the simplicity of relying on the platform's budget management mechanism or the advertisers' complex organizational structure and incentives. While understanding the inconsistency of budget-constrained advertisers' behaviors with the utility maximization models is an interesting (and relevant) research question, it is not the main focus of the present paper. In this paper, we analyze and compare the system equilibria induced by various budget management mechanisms, both theoretically and empirically, assuming the platform implements these mechanisms on behalf of the advertisers and advertisers truthfully report their values.

Our Contributions.

In this paper, we present a new framework for studying and comparing budget management strategies, and study a variety of budget planning algorithms. We first introduce a model for studying "system equilibria" of budget management mechanisms in the context of repeated auctions (Section 2). We then present, in Section 3, a number of budget management mechanisms that can be easily deployed in the context of repeated auctions: probabilistic throttling, thresholding, bid shading, reserve pricing, and multiplicative boosting. Some of these mechanisms have been previously proposed in the literature, and others are natural mechanisms or mechanisms that interpolate between known mechanisms. Our first theoretical result shows the existence of system equilibria for all these budget management mechanisms. As our main theoretical result, we present a rigorous analysis of the performances of system equilibria in a simplified model for each of these mechanisms, i.e., in a symmetric model where all advertisers have the same budget and all bids are drawn i.i.d. from the same distribution (Section 4). In particular, we formally prove dominance relationships among a wide range of budget management mechanisms in terms of the seller's profit and buyers' utility. Finally, we also empirically compare the system equilibria of these strategies using real ad auction data, as well as randomly generated bids (Section 5). The empirical study confirms our theoretical findings about the relative performances of these budget management strategies in the simplified model. In particular, we show that while bid shading provides higher buyer utility, mechanisms such as thresholding and throttling provide higher profit.

Related Work.

Online budget management strategies have been extensively studied in problems of online stochastic matching [8, 10, 15, 18], and online budgeted allocation [4, 7, 19] for repeated first-price auctions. As a result, the insights from these studies have been applied mainly in the context of display ads in which pricing is negotiated ahead of time and no major pricing happens during the allocation period [8, 9]. A number of recent papers [1, 5, 6, 14] have addressed the above issue, but they study different equilibrium concepts, and do not discuss the majority of the budget management strategies discussed in this paper. Moreover, we provide either the first or a different theoretical analysis (on a new model) for the three budget management strategies that are commonly studied in these papers and ours. In particular, [1] and [13] propose approximate equilibrium concepts to study the dynamic interactions among budget-constrained advertisers in repeated second-price auctions. They show that bid shading is an approximately optimal best-response for utility-maximizing advertisers in large markets. In addition, variants of throttling and thresholding mechanisms have been studied in [5, 14]. While these two papers show the existence of a variant of fixed-point equilibrium for throttling and thresholding functions, we prove the existence of equilibria in a different and more general model. More notably, these two papers do not provide a theoretical analysis of the profit or utility obtained by these mechanisms, and to the best of our knowledge, we are the first to provide a theoretical analysis of the profit and utility properties of the throttling and thresholding mechanisms.

Another line of work that addresses the budget constraints in online repeated auctions is that of optimal auction design in the presence of budgets [2, 11, 12]. These papers, however, study more complex auction formats that are not used in repeated ad auctions in practice, to the best of our knowledge. In particular, [11, 12] do not consider the case of repeated second-price auctions which is the most commonly used auction format for sponsored search and ad exchange applications.

2. MODEL

We consider a repeated auction setting in which a seller allocates an item to n buyers in each round. The seller has T items for sale over the time horizon with a fixed opportunity cost $c > 0$ per item. Buyer i has a budget B_i that limits his cumulative expenditure over the T rounds. Motivated by the fact that a buyer participates in a very large number of auctions, we assume that the budget constraint needs to be satisfied only in expectation. The values of buyer i for each item are drawn independently from a cumulative distribution function (CDF) F_i with probability density function (PDF) f_i. Values are assumed to be drawn independently across buyers. We denote by $\mathcal{X}_i$ the support of the distribution of values of buyer i, and by $\mathcal{X} = \mathcal{X}_1 \times \ldots \mathcal{X}_n$ the product space of supports for all buyers. We consider value distributions with PDFs f_i that have contiguous (possibly, unbounded) support containing 0 and bounded expected value, and CDFs F_i that are absolutely continuous and satisfy $\lim_{x\to\infty} x\bar{F}_i(x) = 0$. The technical condition that $\lim_{x\to\infty} x\bar{F}_i(x) = 0$ is mild and satisfied by most general distributions. For example, by Markov's inequality, it is sufficient that $\mathbb{E}\left[x_i^{1+\epsilon}\right] < \infty$ for some $\epsilon > 0$.

When appropriate, we assume distributions with additional properties such as regularity and increasing generalized failure rates. A distribution is *regular* if the virtual value function $\psi_i(x) := x - \frac{\bar{F}_i(x)}{f_i(x)}$ is increasing in x, where $\bar{F}_i(x) := 1 - F_i(x)$. A distribution has an *increasing generalized failure rate (GFR)* if the generalized failure rate $\xi_i(x) := \frac{x f_i(x)}{\bar{F}_i(x)}$ is an increasing function. Note that any distribution with an increasing GFR is also regular. These assumptions are common in the pricing and auction theory literature, and many distributions satisfy these conditions (see, e.g., [20] and [17]).[1]

We consider different budget management mechanisms designed to help buyers spend exactly or close to their budgets in expectation. We restrict attention to stationary (or stateless) mechanisms, which have the practically appealing property that budgets are depleted at a constant rate, or smoothly over time. These mechanisms are derived with some modifications from the standard second-price auction with a reserve price. In addition, we consider Bayesian optimal mechanisms that satisfy the buyers' budget constraints in the ex-ante sense. As will be shown in Section 3, the mechanisms can be implemented with n parameters, one parameter for each buyer, except for one optimal mechanism. We assume buyers bid truthfully and are only interested in their total expenditures fully meeting their budget constraints, and the seller implements the mechanisms on their behalf. The assumption essentially decouples budget management and bidding, and facilitates both theoretical and empirical studies of budget management strategies.

More formally, we denote by $\mathcal{S}_i = [\underline{s}_i, \bar{s}_i]$ the platform parameter space for buyer i and by $\mathcal{S} = \mathcal{S}_1 \times \ldots \times \mathcal{S}_n$ the product space of parameter spaces for all buyers. A *budget management mechanism* consists of a pair (Q, M), where $Q : \mathcal{X} \times \mathcal{S} \to \Delta^n$ is an allocation rule from the space of bids and parameters to the probability simplex $\Delta^n = \{q \in \mathbb{R}^n : q_i \geq 0, \sum_{i=1}^n q_i \leq 1\}$, and $M : \mathcal{X} \times \mathcal{S} \to \mathbb{R}^n$ is a payment rule. That is, when buyers bid $x \in \mathcal{X}$ and the platform parameters are $s \in \mathcal{S}$, the function $Q_i(x; s)$ gives the probability that the item is assigned to buyer i, while the function $M_i(x; s)$ gives the payment buyer i must make to the seller.

Given a mechanism (Q, M), we let $G_i : \mathcal{S} \to \mathbb{R}$ be the expected expenditure and $V_i : \mathcal{S} \to \mathbb{R}$ the gross value obtained by buyer i over the T rounds as a function of the platform parameters. These functions are given by $G_i(s) = T\mathbb{E}_x [M_i(x; s)]$ and $V_i(s) = T\mathbb{E}_x [x_i Q_i(x; s)]$, respectively, where $\mathbb{E}_x [-]$ denotes the expectation with respect to the buyers' values and we used that values are identically distributed across auctions.

Depending on the parameters, the seller can implement mechanisms that lead to different outcomes. For the budget management problem, we are interested in those outcomes, that we call system equilibria, in which the platform determines, for each buyer, the parameter that maximizes his expected utility subject to his budget constraint. While a system equilibrium is budget feasible, i.e., $G_i(s) \leq B_i$ for all buyer i, it also satisfies that a buyer has no incentive to unilaterally change his parameter, if given the chance. We formally define system equilibria as follows:

Definition 1 (System Equilibrium). *Given a budget management mechanism (Q, M), a vector of platform parameters $s^* = (s_i^*)_{i=1}^n \in \mathcal{S}$ is a* system equilibrium *if for each buyer i:*

$$s_i^* \in \arg\max_{s_i \in \mathcal{S}_i} \; V_i(s_i, s_{-i}^*) - G_i(s_i, s_{-i}^*) \qquad (1)$$
$$\text{s.t. } \; G_i(s_i, s_{-i}^*) \leq B_i \,.$$

Suppose that the expected utility $V_i(s) - G_i(s)$ and expected expenditure $G_i(s)$ of each buyer is non-increasing in s_i for all s_{-i}. Then, a system equilibrium $s^* \in \mathcal{S}$ can be alternatively characterized as a solution of the following non-linear complementarity problem:

$$G_i(s^*) \leq B_i \quad \perp \quad s_i^* \geq \underline{s}_i, \, \forall i \in [n] \,,$$

where $\perp$ indicates a complementarity condition between the expenditure and parameter, i.e., at least one inequality should be met with equality. Intuitively, this result states that, in a system equilibrium, if a buyer has excess budget, then the parameter is set to $s_i = \underline{s}_i$ and his expenditure is maximized. Otherwise, the parameter is chosen so that the buyer depletes his budget in expectation.

We compare the system equilibria of budget management mechanisms with respect to two objectives: the profit of the seller (P) and the utility of buyers (U). Denote by $I : \mathcal{S} \to \mathbb{R}$ the expected number of items sold over the horizon, which is given by $I(s) = T\sum_{i=1}^n \mathbb{E}_x [Q_i(x; s)]$. Then, these objectives are given by $P(s) = \sum_{i=1}^n G_i(s) - cI(s)$ and $U(s) = \sum_{i=1}^n (V_i(s) - G_i(s))$.

3. MECHANISMS

We present the budget management mechanisms to be considered in this paper. The mechanisms are derived with simple modifications from the standard second-price auction with the reserve price equal to the opportunity cost. In particular, each mechanism has a parametrization that leads to the same outcome as the standard second-price auction with a reserve. For the definitions of mechanisms' allocation and payment rules, see Table 1. Note that we break ties arbitrarily.

The mechanisms are designed to control the buyers' expenditures by a combination of more strict allocation rules and/or discounted payments implemented via reserve prices, reduced competition among buyers, and bid shading factors. All the mechanisms are individually rational and have monotone allocation rules, i.e., a buyer's higher bid leads to a higher probability of allocation when other bids are fixed, but are not necessarily incentive compatible when considered with corresponding payment rules.

The throttling mechanism controls expenditures by excluding a buyer independently and at random with a fixed probability, and then running a second-price auction with reserve c. This may result in the buyer missing good opportunities, for example, from not participating when his value is high. The thresholding mechanism allows the buyer to participate in an auction when his bid is above a fixed threshold, guaranteeing that the buyer does not miss the items he values the most. The reserve pricing mechanism has a similar allocation rule to the thresholding mechanism: each buyer has a personalized reserve price and can only win when his bid is above the reserve. However, in the reserve pricing mechanism, the winner is

[1] For instance, the uniform, exponential, triangular, truncated normal, gamma, Weibull, and log-normal distributions have increasing GFR's.

Mechanism (Abbrev.)/Its Parameters	Allocation/Payment Rules
Bid Shading (S) $(\mu_i \geq 0, \forall i)$	Each buyer i has a parameter $\mu_i \geq 0$ Buyer i wins if $\frac{x_i}{1+\mu_i} \geq \max_{j \neq i} \frac{x_j}{1+\mu_j} \vee c$ pays $\max_{j \neq i} \frac{x_j}{1+\mu_j} \vee c$
Multiplicative Boosting (MB) $(\delta_i \geq 1, \forall i)$	Each buyer i has a parameter $\delta_i \geq 1$ Buyer i wins if $x_i \geq \delta_i \cdot (\max_{j \neq i} x_j \vee c)$ pays $\max_{j \neq i} x_j \vee c$
Reserve Pricing (R) $(r_i \geq c, \forall i)$	Each buyer i has a parameter $r_i \geq c$ Buyer i wins if $x_i \geq \max_{\substack{j \neq i \\ x_j \geq r_j}} x_j \vee r_i$ pays $\max_{\substack{j \neq i \\ x_j \geq r_j}} x_j \vee r_i$
Thresholding (T) $(\tau_i \geq c, \forall i)$	Each buyer i has a parameter $\tau_i \geq c$ Buyer i wins if $x_i \geq \max_{\substack{j \neq i \\ x_j \geq \tau_j}} x_j \vee \tau_i$ pays $\max_{\substack{j \neq i \\ x_j \geq \tau_j}} x_j \vee c$
Throttling (TO) $(\theta_i \in [0,1], \forall i)$	Each buyer i has a parameter θ_i and $I_i = 1$ with prob. $1 - \theta_i$ Buyer i wins if $I_i = 1$ and $x_i \geq \max_{\substack{j \neq i \\ I_j = 1}} x_j \vee c$ pays $\max_{\substack{j \neq i \\ I_j = 1}} x_j \vee c$
Profit Optimal (PO) $(\gamma_i \in [0,1], \forall i)$	Each buyer i has a virtual value function ψ_i and parameter $\gamma_i \in [0,1]$ Buyer i wins if $(1-\gamma_i)\psi_i(x_i) \geq c$ and $(1-\gamma_i)\psi_i(x_i) \geq \max_{j \neq i}(1-\gamma_j)\psi_j(x_j)$ pays $\inf\{z : (1-\gamma_i)\psi_i(z) \geq c$ and $(1-\gamma_i)\psi_i(z) \geq (1-\gamma_j)\psi_j(x_j), \forall j \neq i\}$
Utility Optimal (UO) $(\mu \geq \gamma_i \geq 0, \forall i)$	There is a system parameter $\mu \geq 0$ Each buyer i has a virtual value function ψ_i and parameter $\gamma_i \in [0,\mu]$ Buyer i wins if $x_i + (\mu - \gamma_i)\psi_i(x_i) \geq (1+\mu)c$ and $x_i + (\mu-\gamma_i)\psi_i(x_i) \geq \max_{j \neq i}\{x_j + (\mu - \gamma_j)\psi_j(x_j)\}$ pays $\inf\{z : z + (\mu-\gamma_i)\psi_i(z) \geq (1+\mu)c$ and $z + (\mu-\gamma_i)\psi_i(z) \geq x_j + (\mu-\gamma_j)\psi_j(x_j), \forall j \neq i\}$

Table 1: Definitions of budget management mechanisms with their abbreviations and parameters ($\vee$ denotes the max operator). Buyers bid $x_1, \ldots, x_n$.

charged the maximum of the second-highest bid and his reserve price, leading to higher payments.

The bid shading mechanism allows buyers to participate in all auctions with bids shaded by a constant multiplicative factor. This guarantees that each buyer wins the items with the highest bang for the buck (ratio of utility to payment). [1] showed that bid shading is an approximately optimal best-response for buyers in large markets with large number of auctions and players: strategic buyers should optimally shade bids to account for the option value for future good opportunities. Multiplicative boosting works similarly to bid shading in that buyers win the items with the highest bang for the buck. However, in multiplicative boosting, instead of shading bids, the allocation rule is directly modified, leading to higher payments.

As benchmarks, we consider Bayesian optimal mechanisms with respect to the objectives that satisfy the buyers' budget constraints in the ex-ante sense. Note these mechanisms are incentive compatible. Assuming the buyers' value distributions are regular, we can derive the Bayesian optimal mechanisms using the well-known Myerson Lemma (see, e.g., Chapter 5 in Krishna [16]). Due to the space constraints, we defer the derivation details to full version of the paper [3].

We assume the seller implements the mechanisms on behalf of the buyers, and buyers bid their values truthfully. Note that the mechanisms can be implemented in terms of parameters with one parameter per buyer (with the exception of the utility optimal mechanism which has an additional parameter). The parameters provide a mean to control buyers' expected expenditures and would required minimal system overhead in practical applications.

3.1 Existence of System Equilibria

In the general asymmetric setting where buyers can have different value distributions and budgets, we show the existence of system equilibria. To prove the existence of system equilibria, we use the following result based on mathematical fixed-point theorems. The proof of this result uses a standard application of the Maximum Theorem and Kakutani Fixed-Point Theorem.

PROPOSITION 1. *A budget management mechanism admits a system equilibrium if for each buyer i:*

1. *The parameter space* $\mathcal{S}_i = [\underline{s}_i, \bar{s}_i]$ *satisfies* $0 \leq \underline{s}_i \leq \bar{s}_i < \infty$.

2. *The expected utility $U_i(s) = V_i(s) - G_i(s)$ is jointly continuous and quasi-concave in s_i for all $s_{-i} \in \mathcal{S}_{-i}$.*

3. *The expected expenditure $G_i(s)$ is jointly continuous and strictly decreasing in s_i for all $s_{-i} \in \mathcal{S}_{-i}$.*

4. *There exists some $s_i^0 \in \mathcal{S}_i$ such that $G_i\left(s_i^0, s_{-i}\right) \leq B_i$ for all $s_{-i} \in \mathcal{S}_{-i}$.*

We present the main result of the section with a proof sketch. All complete proofs are deferred to the full version of the paper [3] because of space considerations.

Theorem 2. *In the asymmetric setting:*

1. *Bid shading (S), multiplicative boosting (MB), thresholding (T) and throttling (TO) admit a system equilibrium.*

2. *Reserve pricing (R) admits a system equilibrium when the ranges of reserve prices are suitably restricted, and bid distributions have strictly increasing GFRs.*

We prove the result by showing that the budget management mechanisms considered satisfy the assumptions in the statement of Proposition 1 when the spaces of feasible parameters $\mathcal{S}$ are suitably chosen. In the case of reserve pricing mechanism, we need to choose parameter spaces over which the expected expenditure functions are strictly decreasing.

Proof Sketch of Theorem 2. For a sketch, we show that the bid shading mechanism has a system equilibrium by choosing the parameter space $\mathcal{S}_i$ such that Conditions (2)-(4) of Proposition 1 hold. Since the buyers' value distributions are independent and have absolutely continuous cumulative distribution functions F_i, the expected utilities $U_i(s)$ and expected expenditures $G_i(s)$ are jointly continuous.

For Condition (2), we note that increasing buyer i's parameter in the bid shading mechanism leads to lower utility, since it results in fewer realized allocations while the payments stay the same in each allocation, when other buyers' parameters are fixed. Then, the buyer's utility function is monotonically decreasing and, hence, quasi-concave.

For Conditions (3)-(4), we see that the expenditure function for buyer i, G_i, strictly decreases as the buyer i's parameter is increased. If the support of f_i is bounded, i.e., $\sup \mathcal{X}_i < \infty$, we determine the parameter space in terms of the support such that G_i strictly decreases to 0 over $\mathcal{S}_i$. Let $d_i = \max_{j \neq i} x_j/(1+\mu_j)$ be the highest competing bid faced by buyer i when competitors shade their values by $1+\mu_j$, which is distributed as $H_i(x; \mu_{-i}) = \mathbb{P}\{d_i \leq x\} = \prod_{j \neq i} F_j((1+\mu_j)x)$ because values are independent. The mechanism will allocate to buyer i whenever $x_i/(1+\mu_i) \geq d_i \vee c$ and charge $d_i \vee c$ in the case of winning. Thus, the expected expenditure of buyer i is

$$\begin{aligned} G_i(\mu_i, \mu_{-i}) = {} & c\bar{F}_i((1+\mu_i)c)H_i(c; \mu_{-i}) \\ & + \int_c^\infty x\bar{F}_i((1+\mu_i)x) \, \mathrm{d}H_i(x; \mu_{-i}) \, . \end{aligned}$$

Taking derivatives, we obtain

$$\begin{aligned} \frac{\partial G_i}{\partial \mu_i} = {} & -c^2 f_i((1+\mu_i)c)H_i(c; \mu_{-i}) \\ & - \int_c^\infty x^2 f_i((1+\mu_i)x) \, \mathrm{d}H_i(x; \mu_{-i}) \, , \end{aligned}$$

and $\frac{\partial G_i}{\partial \mu_i} < 0$ if $(1+\mu_i)c$ lies within the support. It suffices to take $\mathcal{S}_i = [0, \sup \mathcal{X}_i/c - 1]$.

If the support is unbounded, we note that the expected expenditure of buyer i can be rewritten and upper bounded as follows

$$\begin{aligned} G_i(\mu_i, \mu_{-i}) &= \mathbb{E}\left[(c \vee d_i)\mathbf{1}\{x_i/(1+\mu_i) \geq c \vee d_i\}\right] \\ &\leq \mathbb{E}\left[c\mathbf{1}\{x_i/(1+\mu_i) \geq c\}\right] \\ &\quad + \mathbb{E}\left[d_i\mathbf{1}\{x_i/(1+\mu_i) \geq c\}\right] \\ &= (c + \mathbb{E}[d_i])\bar{F}_i((1+\mu_i)c) \\ &\leq (c + \textstyle\sum_{j \neq i} \mathbb{E}[x_j])\bar{F}_i((1+\mu_i)c) \, , \end{aligned}$$

where the last inequality follows because $d_i \leq \sum_{j \neq i} \frac{x_j}{1+\mu_j} \leq \sum_{j \neq i} x_j$. The upper bound goes to 0 as μ_i goes to infinity and the expected expenditure approaches 0 in the limit. We take $\bar{s}_i$ to be sufficiently large such that each budget constraint can be met independently. □

4. DOMINANCE RELATIONS IN THE SYMMETRIC SETTING

Budget management mechanisms can lead to different tradeoffs between objective values. In the symmetric setting when all buyers are of the same type, we show dominance relations between mechanisms, that is, we show that some mechanisms lead to system equilibria with higher objective values regardless of the budget level and distribution of values.

We assume all buyers have the same distribution of values $F(\cdot)$ and budget B, and restrict attention to symmetric equilibria in which the platforms uses the same parameter for all buyers. Let $G : [\underline{s}, \bar{s}] \to \mathbb{R}$ be the expected expenditure of one buyer when the same parameter is used for all buyers. Similarly, we let $U : [\underline{s}, \bar{s}] \to \mathbb{R}$ be the expected utility of one buyer. The next result shows that, in symmetric settings, the mechanisms in consideration admit a unique symmetric equilibrium, and this equilibrium satisfies a simple non-linear complementarity condition.

Proposition 3. *In the symmetric setting, mechanisms (S), (MB), (R), (T) and (TO) admit a unique symmetric system equilibrium under the assumptions of Theorem 2. Moreover, the symmetric system equilibrium $s^* \in [\underline{s}, \bar{s}]$ is the unique solution of:*

$$G(s^*) \leq B \quad \perp \quad s^* \geq \underline{s} \, .$$

Equilibrium uniqueness facilitates comparison between budget management mechanisms. We say mechanism (Q^π, M^π) *dominates* mechanism $(Q^{\pi'}, M^{\pi'})$ in terms of some objective if for each possible budget level, the system equilibrium of (Q^π, M^π) leads to a higher objective value than the system equilibrium of $(Q^{\pi'}, M^{\pi'})$. On the other hand, we say that a *non-dominance* relation holds

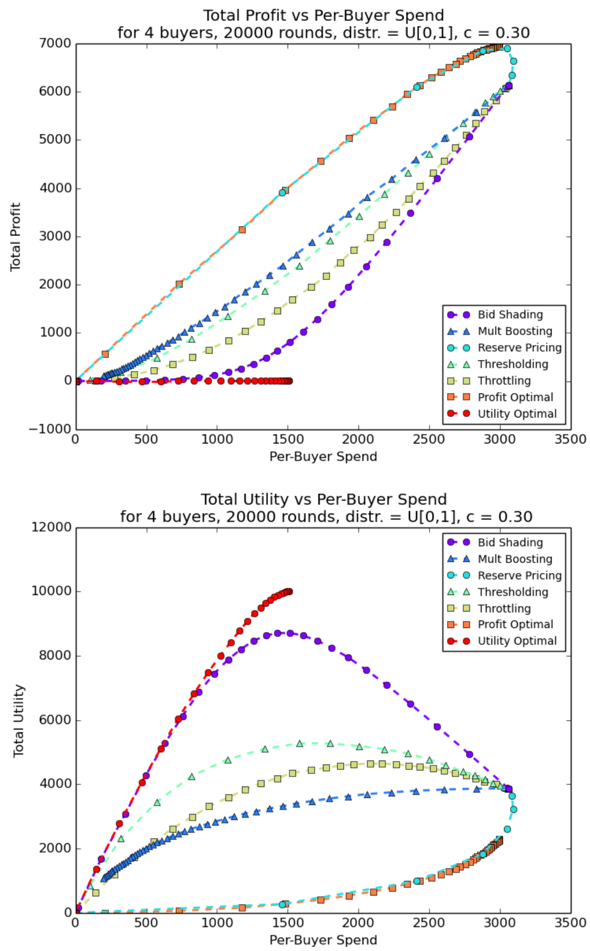

Figure 1: Performance curves of budget management mechanisms showing dominance and non-dominance relations in a symmetric setting.

when the dominance relation between mechanisms is different, hence, indeterminate, depending on the problem setting such as the number of buyers, opportunity cost, distribution, and budget. We use $\geq$ to denote dominance relations, as in $\pi \geq \pi'$, and $\|$ for non-dominance relations.

In order to pictorially compare budget mechanisms we introduce the concept of *performance curves*, which relies on delineating the expected expenditure and expected objective value achievable by a mechanism. Fix a mechanism (Q^π, M^π) with parameter set $\mathcal{S}^\pi$, and denote by $G^\pi : \mathcal{S}^\pi \to \mathbb{R}$ the expected expenditure function and by $O^\pi : \mathcal{S}^\pi \to \mathbb{R}$ some objective function. The performance curve is given by the set of points in $\mathbb{R}^2$ in $\{(G^\pi(s), O^\pi(s))\}_{s\in\mathcal{S}^\pi}$. When the expected expenditure functions of mechanisms π and π' are monotone, we have a dominance relation $\pi \geq \pi'$ if the performance curve of π lies above the performance curve of π'. Similarly, a non-dominance relation is equivalent to the corresponding curves crossing for a certain setting, or exhibiting different dominance relations depending on the setting. For different objectives considered in this paper, we refer to the performance curves with respect to these objectives as profit and utility curves. See Figure 1 for the performance curves in a symmetric setting with 4 buyers and uniformly distributed values.

Let $\bar{B}$ be the unconstrained expected expenditure of one buyer under the second-price auction with reserve c without a budget constraint. Then, the mechanisms can realize any budget level in the range $[0, \bar{B}]$ and achieve some objective values since the expected per-buyer expenditure approaches 0 for large enough values of the platform parameter. For Bayesian optimal mechanisms, the parametrization does not necessarily cover the whole range and we consider their extended parametrization, i.e., if an optimal mechanism already achieves the maximal objective value feasible at a low budget level, the same objective value is achieved at all higher budget levels.

In the next section we prove dominance relations for the most general class of distributions considered in this paper. For a few relations, we prove them for some restricted classes of distributions but conjecture that they hold for the general distributions. We use $\geq_{IF}$ to denote a dominance relation restricted to distributions with an increasing GFR, and $\geq_U$ to denote one restricted to uniform distributions.

4.1 Seller's Profit

The following theorem compares the seller's profit across different budget management mechanisms, covering all possible pairwise comparisons between mechanisms:

THEOREM 4. *The following dominance and non-dominance relations hold for the seller's profit:*

1. *Profit Optimal $=$ Reserve Pricing $\geq$ Thresholding $\geq$ Throttling $\geq_{IF}$ Bid Shading $\geq$ Utility Optimal*
2. *Reserve Pricing $\geq_U$ Mult. Boosting $\geq_U$ Bid Shading*
3. *Mult. Boosting $\|$ Thresholding; Mult. Boosting $\|$ Throttling*

In order to show a dominance relation, we either prove it directly or show an implication in terms of the expected expenditure of one buyer (G) and the expected number of impressions sold (I). Recall that the profit of the seller is given by $P = nG - cI$. When budgets are binding, the seller is extracting the maximum possible revenue and, because items are costly, a budget management mechanism leads to higher profits if it depletes budgets while selling fewer impressions in expectation.

For instance, we show that reserve pricing dominates thresholding by showing that $I^{\mathrm{T}}(\tau) < I^{\mathrm{R}}(r) \Rightarrow G^{\mathrm{T}}(\tau) < G^{\mathrm{R}}(r)$. This implies that at every equilibrium where budgets are depleted, i.e., $G^{\mathrm{T}}(\tau^*) = G^{\mathrm{R}}(r^*) = B$, then the expected number of impressions sold satisfy $I^{\mathrm{T}}(\tau^*) \geq I^{\mathrm{R}}(r^*)$, and thus $P^{\mathrm{R}}(r^*) = nB - cI^{\mathrm{R}}(r^*) \geq nB - cI^{\mathrm{T}}(\tau^*) = P^{\mathrm{T}}(\tau^*)$. It then follows that the reserve pricing mechanism dominates the thresholding mechanism in terms of the profit objective. For non-dominance relations, we provide instances where performance curves of the mechanisms cross. We prove the dominance relation between the reserve pricing and thresholding mechanisms in detail below. The proofs of other dominance relations are technical and deferred to the full version of the paper [3].

Proof Sketch of Theorem 4. For a sketch, we show the dominance relation between the reserve pricing and thresholding mechanisms. Reserve pricing mechanism $(Q^{\mathrm{R}}, M^{\mathrm{R}})$ dominates thresholding mechanism $(Q^{\mathrm{T}}, M^{\mathrm{T}})$ if, for every budget level B, we have $P^{\mathrm{R}}(r^*) \geq P^{\mathrm{T}}(\tau^*)$, where r^* and τ^* are the unique system equilibrium parameters of the mechanisms. For each mechanism π, let $\bar{G}^\pi = \max_{s \in \mathcal{S}^\pi} G^\pi(s) = G^\pi(\underline{s}^\pi)$ be the maximum expenditure and $\bar{I}^\pi = \max_{s \in \mathcal{S}^\pi} I^\pi(s) = I^\pi(\underline{s}^\pi)$ be the maximum number of items sold. Note $\bar{G}^{\mathrm{R}} \geq \bar{G}^{\mathrm{T}}$ and it is possible that $\bar{G}^{\mathrm{R}} > \bar{G}^{\mathrm{T}}$ when $\underline{s}^{\mathrm{R}} = \max\{\psi^{-1}(0), c\} > c$. We assume such case and prove the dominance relation; the equality case follows similarly. For simplicity, we drop the multiplicative factor T that appears in both G and I.

For budget levels $B < \bar{G}^{\mathrm{T}}$, we first show that $I^{\mathrm{T}}(\tau) < I^{\mathrm{R}}(r) \Rightarrow G^{\mathrm{T}}(\tau) < G^{\mathrm{R}}(r)$. Since $1 - F^n(\tau) = I^{\mathrm{T}}(\tau) < I^{\mathrm{R}}(r) = 1 - F^n(r)$, we obtain that $\tau > r$. Furthermore, G^{T} is strictly decreasing in τ, and it follows that

$$\begin{aligned} G^{\mathrm{T}}(\tau) &< G^{\mathrm{T}}(r) \\ &= cF^{n-1}(r)\bar{F}(r) + \int_r^\infty \bar{F}(x) x dF^{n-1}(x) \\ &\leq rF^{n-1}(r)\bar{F}(r) + \int_r^\infty \bar{F}(x) x dF^{n-1}(x) \\ &= G^{\mathrm{R}}(r) \ . \end{aligned}$$

Then, at every equilibrium where budgets are depleted, i.e., $G^{\mathrm{T}}(\tau^*) = G^{\mathrm{R}}(r^*) = B$, the expected number of impressions sold satisfy $I^{\mathrm{T}}(\tau^*) \geq I^{\mathrm{R}}(r^*)$. Thus, $P^{\mathrm{R}}(r^*) = nB - cI^{\mathrm{R}}(r^*) \geq nB - cI^{\mathrm{T}}(\tau^*) = P^{\mathrm{T}}(\tau^*)$.

If $B > \bar{G}^{\mathrm{R}}$, the corresponding system equilibria are: $r^* = \underline{s}^{\mathrm{R}} = \psi^{-1}(0)$ and $\tau^* = \underline{s}^{\mathrm{T}} = c$. Note G^{R} increases as r increases from c to $\psi^{-1}(0)$ while I^{R} decreases over the same range. Hence, $P^{\mathrm{R}}(r^*) > P^{\mathrm{R}}(c) = P^{\mathrm{T}}(c) = P^{\mathrm{T}}(\tau^*)$.

If $\bar{G}^{\mathrm{T}} < B < \bar{G}^{\mathrm{R}}$, the thresholding mechanism's equilibrium parameter is $\tau^* = \underline{s}^{\mathrm{T}}$ and the reserve pricing mechanism's parameter r^* is such that the budget is depleted, i.e., $G^{\mathrm{R}}(r^*) = B$ where $r^* \geq \psi^{-1}(0)$. Note there exists $\tilde{s}^{\mathrm{R}} \in [c, \psi^{-1}(0)]$ such that $G^{\mathrm{R}}(\tilde{s}^{\mathrm{R}}) = B$ and $P^{\mathrm{R}}(\tilde{s}^{\mathrm{R}}) > P^{\mathrm{T}}(\tau^*)$. As I^{R} strictly increases, $P^{\mathrm{R}}(r^*) > P^{\mathrm{R}}(\tilde{s}^{\mathrm{R}}) > P^{\mathrm{T}}(\tau^*)$. □

We next provide some intuition for the results. In the symmetric setting, the profit optimal mechanism can be implemented via a second-price auction with a single reserve price. Thus, the profit optimal and reserve pricing mechanisms coincide. Recall that reserve pricing and thresholding mechanisms implement a similar allocation rule, with the difference that in reserve pricing the payments are higher (at the same parameter levels) because the winner is charged the maximum of the second-highest bid and the reserve price. Since expenditure is decreasing with the threshold, at equilibrium, the threshold needs to be set lower than the reserve price in order to deplete budgets. Lower thresholds lead to more items allocated, and thus lower seller profits. The throttling mechanism naturally leads to lower profits than thresholding, because, in throttling, buyers might miss good opportunities in which their values, and thus payments, are high.

Budget management mechanisms such as thresholding and throttling control expenditures by excluding buyers, while bid shading does so by reducing bids. For example, in throttling there are fewer competitors per auction with their original bids, while in bid shading all buyers compete in each auction with lower bids. It turns out that, from the seller's perspective, throttling is more effective in controlling expenditures, and budgets can be depleted with fewer items sold, leading to higher seller profits. Recall that multiplicative boosting works similarly to bid shading in that buyers win the items with the highest bang for the buck. While bid shading impacts both allocations and payments, in multiplicative boosting, only the allocation rule is modified and the payment rule coincides with that of a second-price auction with the original bids. This naturally leads to higher payments, and thus higher seller profits compared to bid shading.

4.2 Buyers' Utility

The following theorem compares the buyers' utility across different budget management mechanisms, covering nearly all possible pairwise comparisons between mechanisms:

Theorem 5. *The following dominance and non-dominance relations hold for buyers' utility:*

1. *Utility Optimal $\geq$ Bid Shading $\geq$ Thresholding $\geq$ Reserve Pricing = Profit Optimal*
2. *Bid Shading $\geq_U$ Mult. Boosting $\geq_U$ Reserve Pricing*
3. *Throttling $\|$ Thresholding; Throttling $\|$ Mult. Boosting*

We use the same approach to prove dominance and non-dominance relations as in the case of the profit objective. Recall that the buyers' utility is given by $U = nV - nG$, where V is the expected gross value obtained by a buyer. The analyses are more involved than in the case of profit objective, since we have to work with the integral expressions for V as opposed to closed-form expressions for I. We defer the proofs to the full version of the paper [3].

Not surprisingly, most dominance relationships in terms of buyers' utility and seller's profit are reversed, since the seller obtains higher profits –mostly– at the buyers' expense. Roughly speaking, a budget management mechanism obtains higher profits, at the same expenditure level, by selling fewer impressions. Fewer impressions could lead to lower gross value for the buyers, and thus lower utility. This reasoning is incomplete because it fails to take into account the actual value of the items sold. The proof of Theorem 5 makes this analysis precise by showing, for example, that at the same expenditure level, bid shading leads to higher gross value than thresholding.

Theorem 5 shows that bid shading leads to higher buyer utility than most budget management mechanism (with the exception of the utility optimal mechanism). Bid shading is known to be an optimal best-response and a strategic buyer is better off unilaterally deviating to bid shading whenever possible. However, this does not necessarily imply that a buyer is better off when all buyers simultaneously shade bids at a system equilibrium. In shading, competitors' bids are lower, but, compared to thresholding, there are more competitors per auction, which might result in buyers paying more for the same

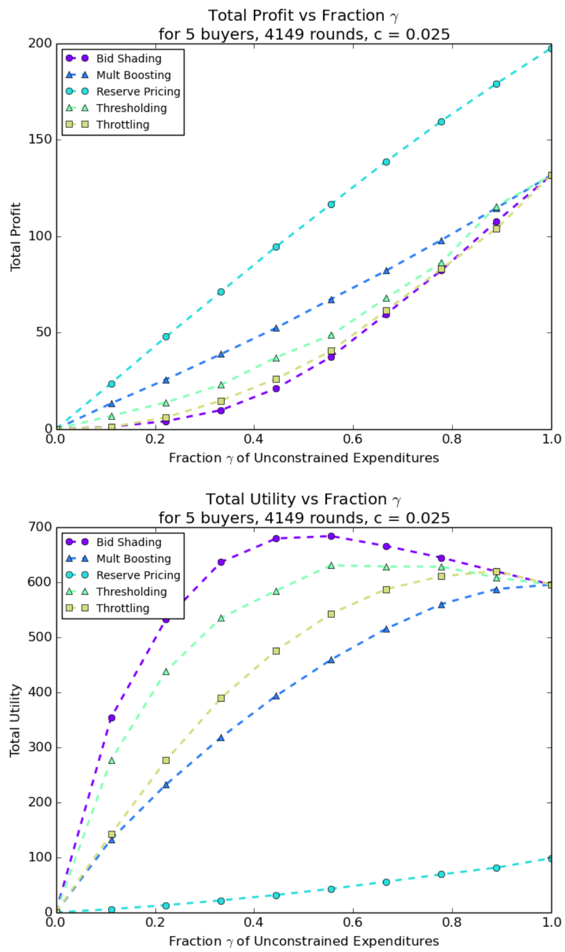

Figure 2: Empirical results for a real bid data in the general asymmetric setting from Section 5.1.

inventory. Interestingly, for buyers, lower bids are more beneficial than fewer competitors, and bid shading leads to the highest buyer utility.

5. EMPIRICAL STUDY

In this section, we design a simple iterative algorithm for computing system equilibria of budget management mechanisms in the general asymmetric setting. Using anonymized, rescaled real bid data, we show the algorithm computes system equilibria within a small margin of error in the budget constraints and that the objective values achieved at these equilibria agree with our theoretical results in Section 4.

Given a mechanism, the Best-Response Algorithm provided in Algorithm 1 iteratively updates the mechanism's parameters until their convergence. In each iteration (Line 2), we set buyers' parameters equal to their best-response parameters were they to unilaterally choose them. The best-response parameter computation is equivalent to solving the nonlinear complementarity problem:

$$G_i(x, s_{-i}) \leq B_i \perp x \geq \underline{s}_i\,, \quad \forall i.$$

Regarding the termination condition, we say that the parameters have converged if the total magnitude of their per-iteration updates becomes smaller than a preset threshold in terms of the ℓ_2 distance.

Note that we update the parameters s_i in place, in some arbitrary order fixed at initialization, and a buyer's best-response parameter can depend on another's in the current iteration. While we do not provide a convergence guarantee for Algorithm 1, in our experiments we obtain system equilibria with small margin of error in the budget constraints. Additionally, the algorithm converges to the same equilibrium even if we change the order of buyers or the initial parameters.

Algorithm 1 Best-Response Algorithm

1: $s_i^0 = s_i, \forall i;\ k = 1$
2: **repeat**
3:     **for** each buyer $i = 1, \ldots, n$ **do**
4:         **if** $G_i(\underline{s}_i, s_{-i}) < B_i$ **then**
5:             $s_i = \underline{s}_i$
6:         **else**
7:             $s_i = \min\{x : G_i(x, s_{-i}) \leq B_i\}$
8:         **end if**
9:             $s^k = s;\ \Delta = \|s^k - s^{k-1}\|_2^2;\ k = k + 1$
10:     **end for**
11: **until** $\Delta < \epsilon$

5.1 Setup

Using real advertiser bid data, we generated anonymized data with bids rescaled uniformly. The range of rescaled bid values was $[0, 2]$. In each auction, a different subset of advertisers may participate, and a particular advertiser may participate in a small fraction of all available auctions.

We took the top 5 advertisers by number of submitted bids and focused on those auctions, 4149 of them, where all these top advertisers participated. For the opportunity cost of impressions, we chose an arbitrary fixed number $c = 0.025$. We let the advertisers' base budgets $(B_1, \ldots, B_5)$ be the unconstrained expenditures under the standard second-price auction with reserve price c. Recall, the budget management mechanisms reduce to the second-price auction with reserve price c, by default, when no budget enforcement measures are in effect.

To generate the performance curves, we computed system equilibria of the mechanisms using Algorithm 1 and the corresponding objective values achieved under the budget constraints $(\gamma B_1, \ldots, \gamma B_5)$ for various fractions $\gamma \in [0, 1]$; that is, advertiser i's budget would be γB_i. In this setup, the performance curves were one-dimensional functions in the fraction γ.

In the full version of the paper [3], we report experiments with other setups. Our findings are similar and thus not reported here. For example, we looked at 5000 randomly selected auctions where at least one of the top 5 advertisers participated and completed the bid data by assuming that buyers not present bid zero. We additionally considered the case of varying opportunity costs and experiments with synthetically generated data based on parametrically fitted distributions. In one such experiment, the opportunity costs were computed by taking the maximum bids from non-top-5 advertisers.

5.2 Results

We present the empirical performance curves of the mechanisms in Figure 2. The results agree with our theoretical findings in Section 4. More specifically, they do not violate the dominance relations theoretically proved in the symmetric setting. As expected, reserve pricing performed best in terms of the profit objective, but worst in terms of the utility objective. On the other hand, bid shading was the best with respect to the utility objective, but the worst with respect to the profit objective.

6. CONCLUSION

In this paper, we compare the system equilibria of different budget management mechanisms provided by advertising platforms to control the expenditures of advertisers. Budget management mechanisms control expenditures by reducing bids (bid shading), modifying the allocation (multiplicative boosting), excluding buyers (throttling and thresholding), or imposing reserve prices (reserve pricing). Our study sheds light on the impact of these budget management strategies on the tradeoff between the seller's profit and buyers' utility. Via a combination of theoretical analyses and empirical studies with real bid data, we show that from the seller's perspective, imposing reserve prices and excluding buyers are more effective in controlling expenditures, and budgets can be depleted with fewer items sold, leading to higher seller profits. From the buyers' perspective, lower bids are more beneficial than fewer competitors, and bid shading leads to the highest buyer utility. We hope that the insights herein developed guide the design and operation of budget management mechanisms in the future.

References

[1] Balseiro, S. R., O. Besbes, and G. Y. Weintraub (2015). Repeated auctions with budgets in ad exchanges: Approximations and design. *Management Science 61*(4), 864–884.

[2] Balseiro, S. R., O. Besbes, and G. Y. Weintraub (2016). Dynamic mechanism design with budget constrained buyers under limited commitment. Working Paper.

[3] Balseiro, S. R., A. Kim, M. Mahdian, and V. Mirrokni. Budget management strategies in repeated auctions. (2016, October 24) Available at SSRN: `https://ssrn.com/abstract=2858261`.

[4] Buchbinder, N., K. Jain, and J. Naor (2007). Online primal-dual algorithms for maximizing ad-auctions revenue. In *ESA*, pp. 253–264. Springer.

[5] Charles, D. X., D. Chakrabarty, M. Chickering, N. R. Devanur, and L. Wang (2013). Budget smoothing for internet ad auctions: a game theoretic approach. In *ACM Conference on Electronic Commerce, EC '13, Philadelphia, PA, USA, June 16-20, 2013*, pp. 163–180.

[6] Ciocan, D. F. and K. Iyer (2017). Truthful equilibrium for sponsored search with endogenous budgets. Working paper.

[7] Devanur, N. and T. Hayes (2009). The adwords problem: Online keyword matching with budgeted bidders under random permutations. In *EC*, pp. 71–78.

[8] Feldman, J., M. Henzinger, N. Korula, V. S. Mirrokni, and C. Stein (2010). Online stochastic packing applied to display ad allocation. In *ESA*, pp. 182–194. Springer.

[9] Feldman, J., N. Korula, V. Mirrokni, S. Muthukrishnan, and M. Pal (2009). Online ad assignment with free disposal. In *WINE*.

[10] Feldman, J., A. Mehta, V. S. Mirrokni, and S. Muthukrishnan (2009). Online stochastic matching: Beating 1-1/e. In *50th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2009, October 25-27, 2009, Atlanta, Georgia, USA*, pp. 117–126.

[11] Goel, G., V. S. Mirrokni, and R. P. Leme (2013). Clinching auction with online supply. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013*, pp. 605–619.

[12] Goel, G., V. S. Mirrokni, and R. P. Leme (2015). Polyhedral clinching auctions and the adwords polytope. *J. ACM 62*(3), 18.

[13] Gummadi, R., P. Key, and A. Proutiere (2012). Optimal bidding strategies and equilibria in dynamic auctions with budget constraints. Working paper.

[14] Karande, C., A. Mehta, and R. Srikant (2013). Optimizing budget constrained spend in search advertising. In *Sixth ACM International Conference on Web Search and Data Mining, WSDM 2013, Rome, Italy, February 4-8, 2013*, pp. 697–706.

[15] Karande, C., A. Mehta, and P. Tripathi (2011). Online bipartite matching with unknown distributions. In *STOC*, pp. 587–596.

[16] Krishna, V. (2010). *Auction Theory* (Second Edition ed.). San Diego: Academic Press.

[17] Lariviere, M. A. (2006, May). A note on probability distributions with increasing generalized failure rates. *Operations Research 54*, 602–604.

[18] Mahdian, M. and Q. Yan (2011). Online bipartite matching with random arrivals: A strongly factor revealing LP approach. In *STOC*, pp. 597–606.

[19] Mehta, A., A. Saberi, U. Vazirani, and V. Vazirani (2007). Adwords and generalized online matching. *J. ACM 54*(5), 22.

[20] Myerson, R. (1981). Optimal auction design. *Mathematics of Operations Research 6*(1), 58–73.