Practical versus Optimal Mechanisms*

Weiran Shen
Institute for Interdisciplinary Information Sciences
Tsinghua University
Beijing, China
emersonswr@gmail.com

Pingzhong Tang
Institute for Interdisciplinary Information Sciences
Tsinghua University
Beijing, China
kenshinping@gmail.com

ABSTRACT

Designing simple mechanisms with desirable revenue guarantees has become a major research agenda in the economics and computation community. However, few mechanisms have been actually applied in industry.

In this paper, we aim to bridge the gap between the "simple versus optimal" theory and practice, and propose a class of parameterized mechanisms, tailored for the sponsored search auction settings. Our mechanisms can balance different objectives by simple parameter tuning, yet at the same time guarantee near optimal revenue in both theoretical and practical senses.

Keywords

auction, revenue maximization, sponsored search, applied mechanism design

1. INTRODUCTION

Designing revenue optimal auctions has been one of the most important themes in economics, ever since Myerson's seminal work [16]. Theories for designing such auctions for the so-called "single parameter" environment have been well-developed [14, 8], and much progress has been made when selling multiple items [23, 24, 7]. Recently, due to interdisciplinary research paradigms such as algorithmic mechanism design [18] and automated mechanism design [6], and its various applications in the sponsored search [9, 25] and similar other domains, it has also become a topic of intensive research at the interface between computer science and economics (e.g. [15, 21, 17]).

Following the vein of algorithmic mechanism design, an important literature, initiated by Hartline and Roughgarden [10], aims to design mechanisms that are simple in their forms (e.g., second price auction with a reserve) and yet guarantee desirable revenue bounds in the worst case. This viewpoint has turned out to be widely adopted in the EC community and has been investigated under a number of extended domains [22, 18, 5, 1, 11, 13, 26, 2, 4]

While a major motivation to design these simple and approximately optimal mechanisms is for the purpose of practicality, unfortunately, to the best of our knowledge, very few of these mechanisms are actually fielded in industry. From an industrial perspective, there are at least the following three concerns when considering implementation of these mechanisms: first of all, all these mechanisms are designed to guarantee the worst-case revenue bounds, while in industry, the evaluation metric tends to be the average-case performance. Secondly, even though most of these papers are able to guarantee constant approximations, say a 2-approximation, of the optimal revenue, they are still not strong enough in the sense that it may indicate that the seller can lose half of the revenue in certain cases. Last but not least, the seller may have other objectives in addition to revenue, which are not guaranteed by these mechanisms. These objectives may change dynamically due to various short-term targets of the company.

In this paper, we aim to bridge the gap between the "simple versus optimal" theory and practice. To address the concerns raised above, we propose the following refined research agenda, targeting specifically the domain of sponsored-search auction design: to design a *parameterized* class of auctions, which

1. has highly desirable worst case guarantees (say, better than 2-approximation) for revenue;
2. gives flexibility to engineers who can freely trade off revenue bounds for other objectives by simply *tuning parameters*;
3. meets industry-level targets via empirical evaluations.

We investigate the sponsored search auction setting, where the seller has several slots for sale and each slot has a certain click-through-rate (CTR). For ease of presentation, a simpler "K identical items" setting is also considered, which is essentially equivalent to the sponsored search setting with the CTR for each slot being 1. In fact, both of the two settings belong to the so-called "single-parameter" setting.

The rest of the paper is organized as follows: section 1.1 summaries our contribution. Section 1.2 lists some existing works in the literature that are related and briefly compares these mechanisms with ours. Section 2 introduces some basic definitions and notations. Section 3 first describes our mechanisms, and then in the two subsections 3.1 and 3.2, analyzes our mechanisms in two different but closely related

*This work was supported in part by the National Natural Science Foundation of China Grant 61303077, 61561146398, a Tsinghua University Initiative Scientific Research Grant and a China Youth 1000-talent program.

Appears in: *Proc. of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017), S. Das, E. Durfee, K. Larson, M. Winikoff (eds.), May 8–12, 2017, São Paulo, Brazil.*

settings. Section 4 describes our simulation setup and results. And the last section 5 concludes the paper.

1.1 Our contribution

With the above agenda in mind, we make the following contributions, for the sponsored search domain:

- We put forward a parameterized class of auctions, which in essence, rank each bidder by a combined (not necessarily linear) function, described by a single parameter α, of its value and Myerson virtual value, allocate the items (CTRs) greedily according to the rank and charge each bidder according to the so-called *payment identity* formula.
- We prove that any auction in the parameterized class is a $(2-\theta)$-approximation of the optimal revenue, where θ is between 0 and 1, as a function of the auction parameter α. Furthermore, given any desired θ between 0 and 1, we give explicitly a mechanism that guarantees a $(2-\theta)$-approximation of the optimal revenue.
- We prove that, as the weight of virtual value increases in the ranking rule, the revenue increases and efficiency drops.
- We empirically evaluate the revenue and efficiency of each auction in this class by simply tuning the parameter α, based on valuations generated from real-bidding data.

1.2 Additional related works

The idea of parameterized auction class has been considered by several existing work in the domain of sponsored search auctions. Lahaie and Pennock [12] consider a class of "squashing" mechanisms. They introduce a parameter α and rank the bidders by $b_i w_i^\alpha$ where w_i is the CTR of bidder i. They find that setting $\alpha < 1$ generally increases the revenue.

Roberts et al. [20] considers the "anchoring" mechanism by ranking the bidders by $(b_i - r)w_i$. They introduce a reserve price parameter r and a reserve score parameter s. Several ranking algorithms, including $b_i w_i / r$, $b_i w_i / s$, $b_i w_i^\alpha / s$ and the "squashing" $b_i w_i^\alpha$ are compared. They show, by simulation, that their "anchoring" mechanism which ranks the bidders by $(b_i - r)w_i$ achieves more revenue and efficiency than other mechanisms.

Bachrach et al. [3] aim towards tradeoffs among different objectives. They use "$\gamma_1$revenue+$\gamma_2$welfare+$\gamma_3$click yield" as their objective function. They show that under the condition that $\gamma_1 + \gamma_2 + \gamma_3 = 1$, their mechanism achieves an γ_1 fraction of the optimal revenue, a γ_2 fraction of the optimal welfare and a γ_3 fraction of the optimal click yield. Their work is similar to ours in the sense that they also consider tradeoffs between efficiency and revenue. In fact, our class of mechanisms includes their mechanism as a special case (up to a constant factor, which does not affect the outcome of the mechanism) by setting $\alpha = \frac{\gamma_1}{\gamma_1+\gamma_2}$ and $P(t) = Q(t) = t + \frac{\gamma_3}{\gamma_1+\gamma_2}$. However, the difference is also very clear: while their goal is to optimize linearly combined objectives and achieve a fraction of the optimal value of these objectives, our goal is to design a class of parameterized mechanisms that are easy for practical use. Also, our mechanisms do not require linear combinations and always achieve an approximation ratio better than 2, which is much stronger than theirs. In addition, the results of the two papers do not imply each other.

In the same spirit, [19] also aims to design parameterized mechanism class in which one can tune worst-case bound by tuning parameters in the facility location domain.

2. PRELIMINARIES

We consider the standard *sponsored search setting*, where there are N bidders competing for several slots, and each bidder aims at one slot. Each bidder i has a private *valuation* $v_i \geq 0$, which is drawn from a publicly known distribution F_i. A *valuation profile* is denoted by $v = (v_1, v_2, \dots, v_N)$. Each bidder i also has a bid b_i that is reported to the seller. Similarly, a *bid profile* is denoted by $b = (b_1, b_2, \dots, b_N)$. We may use v_{-i} and b_{-i} to mean the value profile and bid profile of all bidders except i.

2.1 Mechanisms

A *mechanism* consists of two functions $\mathcal{M} = (x, p)$, where the *allocation rule* x a function $x : \mathbb{R}^N \to [0,1]^N$, which takes as input the bid profile and outputs an N-dimensional vector indicating the quantity of items allocated to each bidder; and the payment rule p is a function $p : \mathbb{R}^N \to \mathbb{R}^N$ that maps the bid profile to an N-dimensional non-negative vector specifying the payment of each bidder. The output of the allocation rule x should be in the set of all feasible allocations $\mathcal{X}$. The bidders are said to have *unit-demand* if each feasible allocation in $\mathcal{X}$ is a binary vector. Further, when the bidders have unit-demand, $\mathcal{X}$ has a set representation $\mathcal{X} = \{X(x) | x \textit{ is feasible}\}$, where $X(x) = \{i | x_i = 1\}$. $\mathcal{X}$ is said to be *downward-closed*, if each subset Y of $X \in \mathcal{X}$ is again a feasible allocation.

2.2 Bayes Nash Equilibria

It is standard to assume that the bidders are *risk neutral* and have *quasi-linear utility*: $u_i(v_i; b_i, b_{-i}) = v_i x_i(b_i, b_{-i}) - p_i(b_i, b_{-i})$. Since the distributions $F = F_1 \times F_2 \times \cdots \times F_N$ are common knowledge but each bidder only knows his own valuation v_i, the bidder's objective is to maximize his utility by reporting a bid $b_i(v_i)$ which is a function of v_i. Thus the utility function for bidder i can be written as $u_i(b_i(v_i), b_{-i}(v_{-i})) = v_i x_i(b_i(v_i), b_{-i}(v_{-i})) - p_i(b_i(v_i), b_{-i}(v_{-i}))$. We say a set of bidding functions $b(v) = (b_1(v_1), \dots, b_N(v_N))$ forms a *Bayes-Nash equilibrium* if for all i, v_i and $b_i'(v_i)$, the following inequality holds:

$$\mathbf{E}_{v_{-i}}\left[u_i(b_i(v_i), b_{-i}(v_{-i}))\right] \geq \mathbf{E}_{v_{-i}}\left[u_i(b_i'(v_i), b_{-i}(v_{-i}))\right]$$

2.3 BIC mechanisms

A mechanism is *Bayesian incentive compatible* (BIC) if $b(v) = v$ is a Bayes-Nash equilibrium. According to the revelation principle [16], it is without loss of generality to consider only BIC mechanisms. So from now on, we do not distinguish between bidders' values and bids. The following lemma by Myerson [16] characterizes the set of BIC mechanisms.

LEMMA 1. *In a single-parameter setting, assume that* $p(0) = 0$, *then a mechanism* $\mathcal{M} = (x, p)$ *is BIC if and only if for each* i, $x_i(v_i)$ *is monotone increasing and* $p_i(v_i)$ *satisfies:*

$$p_i(v_i) = \int_0^{v_i} z \cdot \mathrm{d}x_i(z)$$

The above lemma states that a monotone allocation rule is necessary for a mechanism to be truthful. And given a monotone allocation rule, there exists a unique way to implement a truthful mechanism by setting the payment rule as described in the lemma. Thus, all we need to design is the allocation function.

2.4 Optimal auction

We first define the notation of *virtual valuation*, which follows Myerson [16].

Definition 1 (Virtual Valuation). *Given v_i, the virtual valuation of bidder i is*

$$\varphi_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}$$

where $f_i(v_i)$ is the density function of the distribution $F_i(v_i)$.

Myerson [16] proves that the expected revenue of a mechanism $\mathcal{M}$ is equivalent to the expected virtual valuation. This lemma directly leads to the design of the optimal auction, and is also crucial to our analysis.

Lemma 2 (Myerson's Lemma). *The revenue of any mechanism $\mathcal{M} = (x, p)$ satisfies*

$$REV(\mathcal{M}) = \mathbf{E}_{v \sim F(v)} \left[\sum_{i=1}^{N} x_i(v)\varphi_i(v_i) \right]$$

To maximize revenue, we should set a *reserve price* r_i for each bidder, such that $\varphi_i(r_i) = 0$, and allocate the items to those with the highest virtual valuations.

2.5 Efficiency

Definition 2 (Efficiency). *Given an equilibrium bid profile b, the efficiency of a mechanism $\mathcal{M} = (x, p)$ is defined to be:*

$$EFF(\mathcal{M}) = \mathbf{E}_{v \sim F(v)} \left[\sum_{i=1}^{N} x_i(v) v_i \right]$$

Efficiency is an important metric in practical sponsored search auction design. From the experiences of several major sponsored search teams in China, it is understood that, the number of complaints from advertisers concerning their rankings negatively correlates with the efficiency criteria.

3. GENERALIZED VIRTUAL-EFFICIENT MECHANISMS

In this section, we aim for a tradeoff between optimal revenue and efficiency. Our goal is to provide a spectrum of mechanisms, within which one can easily trade off the two objectives. We call the class of mechanisms *Generalized virtual-efficient mechanisms*, the meaning of which will become clear immediately after its formal definition.

Let $P(\cdot)$ and $Q(\cdot)$ be two increasing functions, and assume that the each bidder's value distribution satisfies the regularity condition [16] (the virtual value function $\varphi(v)$ is increasing with respect to v). We now consider a family of parameterized mechanisms $\mathcal{M}(\alpha, r)$ based on the functions $P(\cdot)$ and $Q(\cdot)$, where α is a real number with $0 \leq \alpha \leq 1$ and r is a reserve price profile. The mechanism $\mathcal{M}(\alpha, r)$ first filters the bidders with the reserve profile r. Then rank the bidders with the function $\alpha P(v) + (1 - \alpha)Q(\varphi(v))$ among those who meet the reserve price conditions, if any, and allocate the items to the agents greedily according to their rankings. Note that $\mathcal{M}(1, r)$ is just the VCG mechanism with reserve profile r, and if $\varphi_i(r_i) = 0$, then $\mathcal{M}(0, r)$ is the optimal auction (denoted by OPT).

3.1 The K identical items setting

In this section, we consider the setting where the seller has K identical items for sale and each bidder only wants exactly one item (unit-demand). This setting is in fact a special case of the sponsored search setting, with all slots have CTRs of 1. We show that in this setting, both the efficiency and revenue change monotonically with respect to the parameter α. In addition, we also provide a tight lower bound for the revenue of our mechanisms.

Theorem 1. *Let $P(\cdot)$ and $Q(\cdot)$ be two increasing functions. Suppose that the seller has K identical items for sale and the bidders have unit-demand. Assume that the distributions F_i that each bidder's valuation is drawn from satisfies the regularity condition. Then $REV(\mathcal{M}(\alpha, r))$ is monotone decreasing with α while $EFF(\mathcal{M}(\alpha, r))$ is monotone increasing with α.*

Proof. For any $0 \leq \alpha_1 < \alpha_2 \leq 1$, let W_1 and W_2 be the set of winners (those with $x_i = 1$) of $\mathcal{M}(\alpha_1, r)$ and $\mathcal{M}(\alpha_2, r)$, respectively. If the number of the bidders who meet the reserve conditions is smaller than or equal to K, then we have $W_1 = W_2$, which contains exactly those who meet the reserve conditions. When the number of the bidders who meet the reserve conditions is greater than K, we only need to consider the case where $W_1 \neq W_2$. In this case, both W_1 and W_2 has K winners. It follows that $|W_1 \backslash W_2| = |W_2 \backslash W_1|$.

Since $|W_1 \setminus W_2| = |W_2 \setminus W_1|$, there exists bijections between the two sets. Let μ be any such bijection. For any $i \in W_1 \setminus W_2$, let $j = \mu(i) \in W_2 \setminus W_1$. Since $\mathcal{M}(\alpha_1, r)$ ranks the bidders by $\alpha_1 P(v) + (1 - \alpha_1)Q(\varphi(v))$ and bidder $j \notin W_1$, we have that bidder i has a higher ranking score than bidder j in mechanism $\mathcal{M}(\alpha_1, r)$:

$$\begin{aligned} &\alpha_1 P(v_i) + (1 - \alpha_1)Q(\varphi_i(v_i)) \\ \geq &\alpha_1 P(v_j) + (1 - \alpha_1)Q(\varphi_j(v_j)) \end{aligned} \tag{1}$$

Similarly, mechanism $\mathcal{M}(\alpha_2, r)$ ranks the bidders by $\alpha_2 P(v) + (1 - \alpha_2)Q(\varphi(v))$, and $i \notin W_2$. It follows that:

$$\begin{aligned} &\alpha_2 P(v_j) + (1 - \alpha_2)Q(\varphi_j(v_j)) \\ \geq &\alpha_2 P(v_i) + (1 - \alpha_2)Q(\varphi_i(v_i)) \end{aligned} \tag{2}$$

Multiplying inequality 1 by α_2, inequality 2 by α_1 and then adding them together yields:

$$\begin{aligned} &\alpha_1\alpha_2 P(v_i) + (1 - \alpha_1)\alpha_2 Q(\varphi_i(v_i)) \\ &\quad + \alpha_1\alpha_2 P(v_j) + \alpha_1(1 - \alpha_2)Q(\varphi_j(v_j)) \\ \geq &\alpha_1\alpha_2 P(v_j) + (1 - \alpha_1)\alpha_2 Q(\varphi_j(v_j)) \\ &\quad + \alpha_1\alpha_2 P(v_i) + \alpha_1(1 - \alpha_2)Q(\varphi_i(v_i)) \end{aligned}$$

With a little rearrangement, we get:

$$(\alpha_2 - \alpha_1)Q(\varphi_i(v_i)) \geq (\alpha_2 - \alpha_1)Q(\varphi_j(v_j))$$

Therefore, $Q(\varphi_i(v_i)) \geq Q(\varphi_j(v_j))$ since $\alpha_2 - \alpha_1 > 0$. It follows that $\varphi_i(v_i) \geq \varphi_j(v_j)$ for $Q(\cdot)$ is an increasing function. Note that the above inequality holds for any $i \in W_1 \setminus W_2$.

Summing over all such i, we have

$$\sum_{i\in W_1\setminus W_2}\varphi_i(v_i)\geq\sum_{i\in W_1\setminus W_2}\varphi_{\mu(i)}(v_{\mu(i)})=\sum_{i\in W_2\setminus W_1}\varphi_i(v_i)$$

Thus

$$\begin{aligned}\sum_{i\in W_1}\varphi_i(v_i)&=\sum_{i\in W_1\cap W_2}\varphi_i(v_i)+\sum_{i\in W_1\setminus W_2}\varphi_i(v_i)\\&\geq\sum_{i\in W_1\cap W_2}\varphi_i(v_i)+\sum_{i\in W_2\setminus W_1}\varphi_i(v_i)\\&=\sum_{i\in W_2}\varphi_i(v_i)\end{aligned}$$

Taking expectation over v yields:

$$\mathbf{E}_{v\sim F(v)}\left[\sum_{i\in W_1}\varphi_i(v_i)\right]\geq\mathbf{E}_{v\sim F(v)}\left[\sum_{i\in W_2}\varphi_i(v_i)\right]$$

which is equivalent to $REV(\mathcal{M}(\alpha_1,r))\geq REV(\mathcal{M}(\alpha_2,r))$ by Myerson's lemma.

Similarly, we can multiply inequality 1 by $1-\alpha_2$ and inequality 2 by $1-\alpha_1$. Adding them together yields:

$$(\alpha_1-\alpha_2)P(v_i)\geq(\alpha_1-\alpha_2)P(v_j)$$

Thus $P(v_i)\leq P(v_j)$ since $\alpha_2-\alpha_1>0$, which implies $v_i\leq v_j$. Similar reasoning gives

$$\sum_{i\in W_1}v_i\leq\sum_{i\in W_2}v_i$$

Taking expectation over v, we have

$$EFF(\mathcal{M}(\alpha_1,r))\leq EFF(\mathcal{M}(\alpha_2,r))$$

which completes the proof. $\square$

Next, we show that if each distribution F_i satisfies the *monotone hazard rate* condition, with appropriate choices of the functions $P(\cdot)$ and $Q(\cdot)$, our mechanisms can achieve approximately optimal revenue.

Definition 3 (Hazard Rate). *Given a probability distribution $F(v)$, the hazard rate with respect to $F(v)$ is defined to be:*

$$h(v)=\frac{f(v)}{1-F(v)}$$

where $f(v)$ is the density function of the distribution $F(v)$.

If the hazard rate $h(v)$ is monotone increasing with respect to v, we say that the corresponding distribution $F(v)$ satisfies the *monotone hazard rate* (MHR) condition. Note that the regularity condition will be automatically satisfied if the hazard rate function is monotone increasing. We first prove a lemma that will be useful for later arguments.

Lemma 3. *Assume v is distributed according to $F(v)$. Let $\varphi(v)$ and $h(v)$ be the corresponding virtual valuation function and hazard rate function. Let $P(\cdot)$ and $Q(\cdot)$ be two functions that are increasing, concave and differentiable, and there exists a constant $c>0$, such that $\forall v, Q'(v)\geq cP'(v)>0$. Then for all $v>0$ and all $0\leq\alpha\leq 1$,*

$$\alpha P(v)+(1-\alpha)Q(\varphi(v))\leq R\left(v-\frac{\theta}{h(v)}\right)$$

where $R(v)=\alpha P(v)+(1-\alpha)Q(v)$, $\theta=\frac{(1-\alpha)c}{\alpha+(1-\alpha)c}$.

Proof. Let $z=v-\frac{\theta}{h(v)}$. Then

$$v=z+\frac{\theta}{h(v)}$$

$$\varphi(v)=v-\frac{1-F(v)}{f(v)}=v-\frac{1}{h(v)}=z-\frac{1-\theta}{h(v)}$$

Since $P(\cdot)$ and $Q(\cdot)$ are concave functions, we have

$$P(v)\leq P(z)+\frac{\theta}{h(v)}P'(z)$$

$$Q(\varphi(v))\leq Q(z)-\frac{1-\theta}{h(v)}Q'(z)$$

Thus

$$\begin{aligned}\alpha P(v)+&(1-\alpha)Q(\varphi(v))\leq\alpha P(z)+(1-\alpha)Q(z)\\&+\frac{\alpha\theta}{h(v)}P'(z)-\frac{(1-\theta)(1-\alpha)}{h(v)}Q'(z)\end{aligned}$$

Since $Q'(z)\geq cP'(z)$ and $0\leq\alpha\leq 1$, $0\leq\theta\leq 1$, $h(x)=\frac{f(x)}{1-F(x)}>0$, we have

$$\begin{aligned}&\frac{\alpha\theta}{h(v)}P'(z)-\frac{(1-\theta)(1-\alpha)}{h(v)}Q'(z)\\\leq&\frac{\alpha\theta}{h(v)}P'(z)-\frac{(1-\theta)(1-\alpha)}{h(v)}cP'(z)\\=&\frac{P'(z)}{h(v)}(\alpha\theta-(1-\theta)(1-\alpha)c)\\=&0\end{aligned}$$

The last equality holds because $\theta=\frac{(1-\alpha)c}{\alpha+(1-\alpha)c}$. Therefore

$$\alpha P(v)+(1-\alpha)Q(\varphi(v))\leq\alpha P(z)+(1-\alpha)Q(z)=R(z)$$

completing the proof. $\square$

Next, we show that under certain technical conditions, our mechanism can achieve near-optimal revenue guarantees with appropriate choices of the functions $P(\cdot)$ and $Q(\cdot)$.

Theorem 2. *Suppose the distribution F_i that each bidder's valuation is drawn from satisfies the monotone hazard rate condition. Let $P(\cdot)$ and $Q(\cdot)$ be two functions that are increasing, concave and differentiable, and there exists a constant $c>0$, such that $\forall v, Q'(v)\geq cP'(v)>0$. Let r^* be the monopoly reserve profile for the bidders, i.e. $\varphi_i(r_i^*)=0$. If either one of the following two conditions is satisfied:*

1. *There are K identical items for sale, and the bidders have unit-demand;*
2. *$P(t)=at$ and $Q(t)=bt$, where a,b are positive constants.*

Then $REV(\mathcal{M}(\alpha,r^*))$ is a $(2-\theta)$-approximation of the optimal mechanism, where $\theta=\frac{(1-\alpha)c}{\alpha+(1-\alpha)c}$.

Remark 1. These approximation ratios are highly desirable. First, since $\theta\geq 0$, all of the approximation ratios are less than 2 (except when $\theta=0$), which guarantees the near-optimality of our mechanism. Second, the approximation ratio depends on the parameter α, which provides more flexibility for practical use. Note that the 2-approximation result by [10] can be immediately obtained from our result, by setting $P(t)=Q(t)=t$ and $\alpha=1$. $\square$

PROOF. Since $\varphi_i(r_i^*) = r_i^* - \frac{1}{h(r_i^*)} = 0$, we have $r_i^* = \frac{1}{h(r_i^*)}$. The MHR condition implies that $\forall v_i > r_i^*$, $\frac{1}{h(r_i^*)} \geq \frac{1}{h(v_i)}$. Thus $\forall v_i > r_i^*$, we have

$$\varphi_i(v_i) + \frac{1-\theta}{h(r_i^*)} \geq \varphi_i(v_i) + \frac{1-\theta}{h(v_i)}$$

which is equivalent to

$$\varphi_i(v_i) + (1-\theta)r_i^* \geq v_i - \frac{\theta}{h(v_i)}$$

It is straightforward that $R(\cdot)$ is an increasing function. Thus $R^{-1}(\cdot)$ exists and is also increasing. So

$$\begin{aligned}\varphi_i(v_i) + (1-\theta)r_i^* &\geq v_i - \frac{\theta}{h(v_i)}\\ &= R^{-1}\left(R\left(v_i - \frac{\theta}{h(v_i)}\right)\right)\\ &\geq R^{-1}\left(\alpha P(v_i) + (1-\alpha)Q(\varphi_i(v_i))\right) \quad (3)\end{aligned}$$

The last inequality holds because of Lemma 3.

Under condition 1, we always allocate the items to the bidders with the highest ranking scores among those who meet the reserve price condition. Let W and W_{OPT} be the set of winners of our mechanism and the optimal mechanism. Then using similar arguments as in the proof of Theorem 1, we have that $|W| = |W_{OPT}|$ and that for any $i \in W \setminus W_{OPT}$ and any $j \in W_{OPT} \setminus W$:

$$\alpha P(v_i) + (1-\alpha)Q(\varphi_i(v_i)) \geq \alpha P(v_j) + (1-\alpha)Q(\varphi_j(v_j))$$

which implies:

$$\begin{aligned}&R^{-1}(\alpha P(v_i) + (1-\alpha)Q(\varphi_i(v_i)))\\ \geq &R^{-1}(\alpha P(v_j) + (1-\alpha)Q(\varphi_j(v_j)))\end{aligned}$$

since $R^{-1}(\cdot)$ is increasing.

Under condition 2,

$$\begin{aligned}R(x) &= \alpha a x + (1-\alpha)bx\\ R^{-1}(x) &= \frac{x}{\alpha a + (1-\alpha)b}\end{aligned}$$

Our mechanism ranks the bidders by $\alpha a v + (1-\alpha)b\varphi(v)$. Thus

$$\begin{aligned}&\sum_{i=1}^{N} \left(\alpha a v_i + (1-\alpha)b\varphi_i(v_i)\right) x_i(v)\\ \geq &\sum_{i=1}^{N} \left(\alpha a v_i + (1-\alpha)b\varphi_i(v_i)\right) x_i^*(v)\end{aligned}$$

where $x(v)$ and $x^*(v)$ is the allocation function of our mechanism and the optimal mechanism, respectively. Equivalently,

$$\begin{aligned}&\sum_{i=1}^{N} \frac{1}{\alpha a + (1-\alpha)b} \left(\alpha a v_i + (1-\alpha)b\varphi_i(v_i)\right) x_i(v)\\ \geq &\sum_{i=1}^{N} \frac{1}{\alpha a + (1-\alpha)b} \left(\alpha a v_i + (1-\alpha)b\varphi_i(v_i)\right) x_i^*(v)\end{aligned}$$

Therefore under both cases, we have:

$$\begin{aligned}&\sum_{i=1}^{N} R^{-1}\left(\alpha P(v_i) + (1-\alpha)Q(\varphi_i(v_i))\right) x_i(v)\\ \geq &\sum_{i=1}^{N} R^{-1}\left(\alpha P(v_i) + (1-\alpha)Q(\varphi_i(v_i))\right) x_i^*(v) \quad (4)\end{aligned}$$

According to Myerson's Lemma, the revenue of our mechanism can be written as the weighted expectation of the virtual valuations:

$$REV(\mathcal{M}(\alpha, r^*)) = \mathbf{E}_{v\sim F(v)}\left[\sum_{i=1}^{N} \varphi_i(v_i)x_i(v)\right] \quad (5)$$

The revenue can also be lower bounded by r^* since the reserve profile is r^*.

$$REV(\mathcal{M}(\alpha, r^*)) \geq \mathbf{E}_{v\sim F(v)}\left[\sum_{i=1}^{N} r_i^* x_i(v)\right] \quad (6)$$

Therefore,

$$\begin{aligned}&(2-\theta)REV(\mathcal{M}(\alpha, r^*))\\ =&\mathbf{E}_{v\sim F(v)}\left[\sum_{i=1}^{N}(\varphi_i(v_i) + (1-\theta)\varphi_i(v_i))x_i(v)\right]\\ \geq &\mathbf{E}_{v\sim F(v)}\left[\sum_{i=1}^{N}(\varphi_i(v_i) + (1-\theta)r_i^*)x_i(v)\right]\\ \geq &\mathbf{E}_{v\sim F(v)}\left[\sum_{i=1}^{N} R^{-1}\left(\alpha P(v_i) + (1-\alpha)Q(\varphi_i(v_i))\right) x_i(v)\right]\\ \geq &\mathbf{E}_{v\sim F(v)}\left[\sum_{i=1}^{N} R^{-1}\left(\alpha P(v_i) + (1-\alpha)Q(\varphi_i(v_i))\right) x_i^*(v)\right]\\ \geq &\mathbf{E}_{v\sim F(v)}\left[\sum_{i=1}^{N} R^{-1}\left(\alpha P(\varphi_i(v_i)) + (1-\alpha)Q(\varphi_i(v_i))\right) x_i^*(v)\right]\\ =&\mathbf{E}_{v\sim F(v)}\left[\sum_{i=1}^{N} \varphi_i(v_i)x_i^*(v)\right]\\ =&REV(OPT)\end{aligned}$$

The first inequality combines equation 5 and 6. The second inequality comes from inequality 3. The third inequality holds because of inequality 4. The fourth inequality comes from the definition of $\varphi(v)$ and the monotonicity of the functions $P(\cdot)$ and $R^{-1}(\cdot)$. And the last equation is a direct application of the Myerson's Lemma. $\square$

REMARK 2. The first condition can actually be extended to the sponsored search environment (see the next section).

Under the second condition, if the set of feasible allocation set $\mathcal{X}$ is downward-closed, our bound for the revenue is actually tight. Consider the following example.

Let $P(t) = Q(t) = t$ and $c = 1$. Suppose there are K items for sale and two groups of bidders EXP and UNI, each containing K bidders, and the bidders inside each group have iid value distributions. Bidders of EXP have a exponential distribution with $F_{EXP}(v) = 1 - e^{-v}$, while bidders of UNI have values that are distributed uniformly in the interval $[\frac{1+\alpha-\varepsilon}{e} - \delta, \frac{1+\alpha-\varepsilon}{e} + \delta]$, where δ, ε are sufficiently small positive numbers. The feasible allocations are those that contain only bidders from the same group. In this case, $\theta = 1 - \alpha$ and our theorem gives an approximation ratio of $1 + \alpha$.

The monopoly reserve prices for bidders in the EXP group and the UNI group are 1 and $\frac{1+\alpha-\varepsilon}{e} - \delta$, respectively. For bidders in EXP, the virtual value $\varphi_{EXP}(v) = v - 1$. While the bidders in UNI have virtual value $\varphi_{UNI}(v) = v - 2\delta(1 - F_{UNI}(v))$, which is highly concentrated at a small neighborhood of $\frac{1+\alpha-\varepsilon}{e}$. Our mechanism uses a ranking score of $v - 1 + \alpha$ and $v - (1-\alpha)2\delta(1 - F_{UNI}(v))$ for the bidders in EXP and UNI. Again, the ranking score of bidders in UNI is highly concentrated around $\frac{1+\alpha-\varepsilon}{e}$. When K is sufficiently large, the number of bidders in EXP that exceeds the reserve price is about $\frac{K}{e}$ and the average value of these bidders is about 2. Thus, both the revenue and the total ranking score of allocating to the group UNI are about $\frac{(1+\alpha-\varepsilon)K}{e}$, while allocating to the group EXP extracts a revenue of only $\frac{K}{e}$ but a total ranking score of $\frac{(1+\alpha)K}{e}$. Thus our mechanism allocates the items to EXP, but the optimal mechanism allocates to UNI. The ratio of revenue is $1 + \alpha - \varepsilon$, which can be arbitrarily close to $1 + \alpha$. $\square$

Note that in Theorem 2, θ always lies in $[0, 1]$ and changes continuously with respect to α. Thus we can achieve any desired approximation ratio in $[1, 2]$ by simply tuning the α, even if the functions $P(\cdot)$ and $Q(\cdot)$ (and thus the constant c) are fixed, which leads to the following immediate corollary:

COROLLARY 1. *Suppose the distribution F_i that each bidder's valuation is drawn from satisfies the monotone hazard rate condition. Let $P(\cdot)$ and $Q(\cdot)$ be two functions that are increasing, concave and differentiable, and there exists a constant $c > 0$, such that $\forall v, Q'(v) \geq cP'(v) > 0$. Let r^* be the monopoly reserve profile for the bidders, i.e. $\varphi_i(r_i^*) = 0$. If either one of the following two conditions is satisfied:*

1. *There are K identical unit of items for sale, and the bidders have unit-demand;*
2. *$P(t) = at$ and $Q(t) = bt$, where a, b are positive constants.*

Then for any $\theta \in [0, 1]$, there exists $\alpha \in [0, 1]$, such that the revenue of the mechanism $\mathcal{M}(\alpha, r^*)$ is a $(2-\theta)$-approximation of that of the optimal mechanism.

3.2 The sponsored search setting

Now we generalize our results to the sponsored search setting. In the standard sponsored search setting, a search engine typically has several slots available for advertisements. These slots have different CTRs and are sold to interested advertisers via auctions. Each keyword corresponds to an auction. When a user enters keyword query, the search engine collects bids from the bidders that are interested in this keyword, and allocate the slots to the winning bidders. If the user clicks on an advertisement, the corresponding advertiser pays according to some payment rules.

Assume that there are K available slots and the j-th slot has a CTR s_j satisfying $s_1 \geq s_2 \geq \cdots \geq s_K \geq 0$. There are N bidders and our mechanism computes for each bidder i a ranking score $R_i(v_i) = \alpha P(v_i) + (1-\alpha)Q(\varphi_i(v_i))$, then allocates the j-th slot to the bidder with the j-th highest score. We show that both the efficiency and the revenue of our mechanism are monotone with respect to α.

Even though the K identical items setting and the sponsored search setting both belong to the "single-parameter" setting, the analysis from the previous subsection cannot be directly applied here since in this setting, different slots have different CTRs, i.e. slots are not identical. We overcome this technical difficulty by decomposing the sponsored search auctions into K sub-auctions with the j-th sub-auction selling j identical slots. Then our results in previous sections still hold for each sub-auction. We aggregate the results for the sub-auctions together to show that they can be extended to this setting.

THEOREM 3. *Let $P(\cdot)$ and $Q(\cdot)$ be two increasing functions. Suppose there are K slots with CTR $s_1 \geq s_2 \geq \cdots \geq s_K \geq 0$ and the distributions F_i that each bidder's valuation is drawn from satisfies the regularity condition. Then $REV(\mathcal{M}(\alpha, r))$ is monotone decreasing with α while $EFF(\mathcal{M}(\alpha, r))$ is monotone increasing with α.*

PROOF. Let $s_{K+1} = 0$ and $x^{(\alpha)}$ be the allocation rule of the mechanism $\mathcal{M}(\alpha, r)$. Note that only pointing out whether a bidder is a winner is not enough in this setting, since the slots have different CTRs. So we let $x_i^{(\alpha)} \in \{s_i | 1 \leq i \leq K+1\}$ to specify the slot that is allocated to bidder i. If bidder i loses in the auction, we say that i wins the $(K+1)$-th slot and $x_i^{(\alpha)} = s_{K+1} = 0$. Thus the efficiency and revenue can be written as:

$$EFF(\mathcal{M}(\alpha, r)) = \mathbf{E}_{v \sim F(v)} \left[\sum_{i=1}^{N} x_i^{(\alpha)} v_i \right]$$

$$REV(\mathcal{M}(\alpha, r)) = \mathbf{E}_{v \sim F(v)} \left[\sum_{i=1}^{N} x_i^{(\alpha)} \varphi_i(v_i) \right]$$

Let $d_j = s_j - s_{j+1} \geq 0, \forall 1 \leq j < K$. We decompose the auction into K sub-auctions with the j-th $(1 \leq j \leq K)$ auction selling the first j slots to N bidders. In the j-th sub-auction, all the j slots for sale have the same CTR of d_j. Thus the j-th sub-auction actually sells j identical items. We apply our mechanism $\mathcal{M}(\alpha, r)$ to these K sub-auctions, i.e. in the j-th sub-auction, we compute the ranking score for each bidder, and allocate the slots to the highest j bidders. Denote the j-th sub-auction by $\mathcal{A}^{(j)}(\alpha, r)$ and let $x_i^{(j,\alpha)} \in \{0, 1\}$ be its allocation rule. Clearly, for any $j > l$, winners of $\mathcal{A}^{(l)}(\alpha, r)$ are also winners of $\mathcal{A}^{(j)}(\alpha, r)$. And if a bidder i wins the l-th $(1 \leq l \leq K)$ slot in the original auction, then i is also among the winners of sub-auctions $j \geq l$. Thus we have that for all i,

$$\begin{aligned} \sum_{j=1}^{K} d_j x_i^{(j,\alpha)} &= \sum_{j=1}^{l-1} d_j \cdot 0 + \sum_{j=l}^{K} d_j \cdot 1 \\ &= \sum_{j=l}^{K} s_j - s_{j+1} \\ &= s_l = x_i^{(\alpha)} \end{aligned} \tag{7}$$

Now consider each sub-auction $\mathcal{A}^{(j)}(\alpha, r)$. For any $0 \leq \alpha_1 < \alpha_2 \leq 1$, according to Theorem 1,

$$\sum_{i=1}^{N} x_i^{(j,\alpha_1)} \varphi_i(v_i) \geq \sum_{i=1}^{N} x_i^{(j,\alpha_2)} \varphi_i(v_i)$$

$$\sum_{i=1}^{N} x_i^{(j,\alpha_1)} v_i \leq \sum_{i=1}^{N} x_i^{(j,\alpha_2)} v_i$$

Multiply the above inequalities by d_j and sum over j, and we obtain:

$$\sum_{j=1}^{K} d_j \left(\sum_{i=1}^{N} x_i^{(j,\alpha_1)} \varphi_i(v_i) \right) \geq \sum_{j=1}^{K} d_j \left(\sum_{i=1}^{N} x_i^{(j,\alpha_2)} \varphi_i(v_i) \right)$$

Switching the order of summation and applying equation 7 gives

$$\sum_{i=1}^{N} \varphi_i(v_i) x_i^{(\alpha_1)} \geq \sum_{i=1}^{N} \varphi_i(v_i) x_i^{(\alpha_2)}$$

Taking expectation yields $REV(\mathcal{M}(\alpha_1, r)) \geq REV(\mathcal{M}(\alpha_2, r))$. Similarly

$$\sum_{i=1}^{N} v_i x_i^{(\alpha_1)} \leq \sum_{i=1}^{N} v_i x_i^{(\alpha_2)}$$

It follows that $EFF(\mathcal{M}(\alpha_1, r)) \leq EFF(\mathcal{M}(\alpha_2, r))$, which completes the proof. □

THEOREM 4. *Let $P(\cdot)$ and $Q(\cdot)$ be two functions that are increasing, concave and differentiable, and there exists a constant $c > 0$ such that $Q'(t) \geq cP'(t) > 0, \forall t$. Suppose there are K slots with CTRs $s_1 \geq s_2 \geq \cdots \geq s_K \geq 0$ and the distributions F_i that each bidder's valuation is drawn from satisfies the monotone hazard rate condition. Let r^* be the monopoly reserve profile for the bidders. Then $REV(\mathcal{M}(\alpha, r^*))$ is a $(2-\theta)$-approximation of the optimal mechanism, where $\theta = \frac{(1-\alpha)c}{\alpha+(1-\alpha)c}$.*

PROOF. We also decompose the original auction into K sub-auctions and follow the notations defined in the proof of Theorem 3. Let $x^{(OPT)}$ be the allocation function of the original auction and $x^{(j,OPT)} \in \{0,1\}(1 \leq j \leq K)$ be the allocation function of the corresponding j-th sub-auction. It is clear that for each bidder i

$$x_i^{(OPT)} = \sum_{j=1}^{K} d_j x_i^{(j,OPT)}$$

And from Theorem 2, we have that for each sub-auction j:

$$(2-\theta) \sum_{i=1}^{N} \varphi_i(v_i) x_i^{(j,\alpha)} \geq \sum_{i=1}^{N} \varphi_i(v_i) x_i^{(j,OPT)}$$

Multiply the above inequality by d_j and sum over all sub-auctions j and we have

$$(2-\theta) \sum_{j=1}^{K} d_j \left(\sum_{i=1}^{N} \varphi_i(v_i) x_i^{(j,\alpha)} \right) \geq \sum_{j=1}^{K} d_j \left(\sum_{i=1}^{N} \varphi_i(v_i) x_i^{(j,OPT)} \right)$$

Applying the fact $x_i^{(\alpha)} = \sum_{j=1}^{K} d_j x_i^{(j,\alpha)}$ gives

$$(2-\theta) \sum_{i=1}^{N} \varphi_i(v_i) x_i^{(\alpha)} \geq \sum_{i=1}^{N} \varphi_i(v_i) x_i^{(OPT)}$$

Taking expectation over v yields

$$(2-\theta) REV(\mathcal{M}(\alpha, r^*)) \geq REV(OPT)$$

completing the proof. □

4. EXPERIMENTS

In Theorem 1, we show that in the "K identical items" setting, both the efficiency and revenue of our mechanism are monotone with respect to the parameter α.

To verify our results, we first consider a relatively simple case. We assume that there are 3 identical items for sale and 10 interested bidders. The value of bidder i is uniformly distributed in the interval $[0, u_i]$, where u_i is again uniformly distributed in $[1, 2]$. No reserve prices are set for all bidders. We set the functions $P(t) = Q(t) = t$ and $c = 1$, so the only parameter in the simulation is α. The numbers in the figures are normalized since the absolute value is not important.

Figure 1 shows that the efficiency changes almost linearly with respect to α within a large range of values. The revenue, however, changes slowly when α is small and rapidly when α is large. Therefore we can set an appropriate α value to achieve a great efficiency gain but only suffer from a slight revenue loss.

We also evaluate our mechanisms in the sponsored search setting. We use real data from one of the major search engines in Chinese. We first select a keyword with over 600 bidders. There are 3 slots for sale, with the first slots having a CTR of about 0.1. We extract two weeks' data for the keyword. Each bidder's valuations for the keyword is fitted to a lognormal distribution[1], and assume that these bidders have independent distributions. Then we run the auction 1000 times. Inside each auction we draw a sample bid for each bidder according to their respective distributions, and allocate the 3 slots based on our ranking algorithm. We use the average efficiency and revenue of the 1000 auctions as the per-impression efficiency and revenue. We still choose $P(t) = Q(t) = t$ and $c = 1$ for this simulation but use the monopoly reserve price for each bidder. Figure 2 shows similar trends as in the simple case.

Next we evaluate our mechanism on the most profitable 100 keywords. Each keyword has nearly 500 interested bidders on average and the most popular has over 1500 bidders. We adopt the same simulation method as described above for each keyword. After computing the per-impression efficiency and revenue of these 100 keywords, we multiply them by their respective occurrence frequency and then add them together to compute the total efficiency and revenue.

[1]Note that the lognormal distribution does not satisfy the regularity condition. Nevertheless, the simulation is complementary to our theoretical results, showing that our mechanism works under this setting as well.

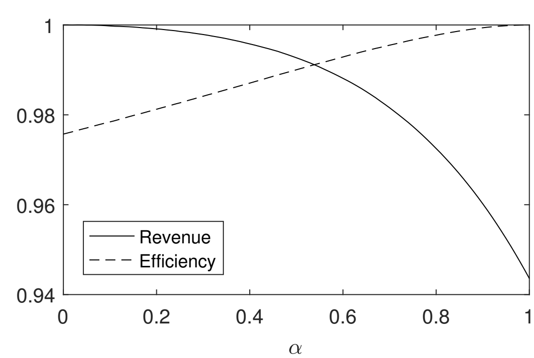

Figure 1: Efficiency and revenue of the simple case

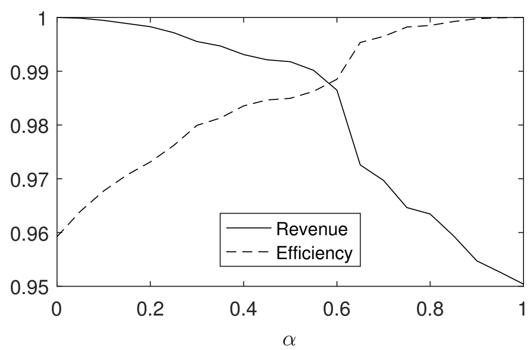

Figure 2: Efficiency and revenue of the keyword with over 600 bidders

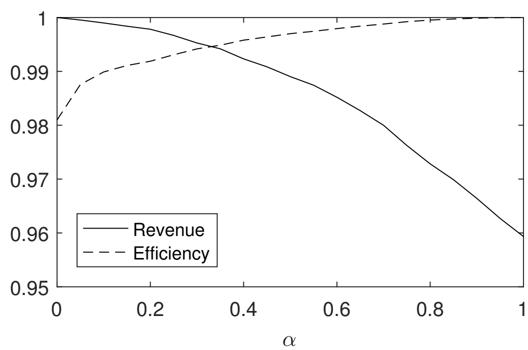

Figure 3: Efficiency and revenue of the sponsored search setting

Figure 3 shows that the efficiency is sensitive with respect to α (changes quickly) when α is small, which may cause some inconvenience for parameter tuning if the objective is to guarantee a certain amount of efficiency. However, this problem could be solved by simply changing the function $Q(t) = t$ to $Q(t) = \frac{0.5}{1.2}t^{1.2} + \frac{0.5}{0.8}t^{0.8}$. It is clear that we still have $Q'(t) \geq 1, \forall t > 0$. This ranking algorithm causes the efficiency to change almost linearly (shown in Figure 4), making it easier to tune the parameter α.

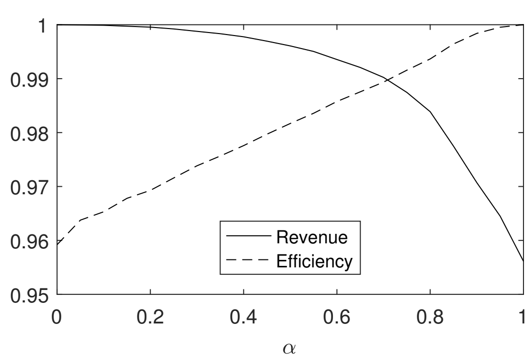

Figure 4: Efficiency and revenue of the sponsored search setting with a different $Q(\cdot)$ function

5. CONCLUSION AND FUTURE WORK

In this paper, we propose a parameterized class of mechanisms that are tailored specifically for the sponsored search setting. We show that under certain conditions, our mechanisms guarantee a highly desirable revenue bounds, while in the meantime provide great flexibility to engineers who can freely trade off between revenue and efficiency. We also show that our mechanisms perform well in realistic settings by simulations on real search engine data. A better way to thoroughly test our mechanisms is through a large-scale field experiment. We are currently implementing our framework with Baidu (the major search engine in China) sponsored search team.

REFERENCES

[1] S. Alaei, H. Fu, N. Haghpanah, and J. Hartline. The simple economics of approximately optimal auctions. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, pages 628–637. IEEE, 2013.

[2] S. Alaei, J. Hartline, R. Niazadeh, E. Pountourakis, and Y. Yuan. Optimal auctions vs. anonymous pricing. In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on*, pages 1446–1463. IEEE, 2015.

[3] Y. Bachrach, S. Ceppi, I. A. Kash, P. Key, and D. Kurokawa. Optimising trade-offs among stakeholders in ad auctions. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 75–92. ACM, 2014.

[4] Y. Cai, N. R. Devanur, and S. M. Weinberg. A duality based unified approach to bayesian mechanism design. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC)*, 2016.

[5] S. Chawla, H. Fu, and A. Karlin. Approximate revenue maximization in interdependent value settings. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 277–294. ACM, 2014.

[6] V. Conitzer and T. Sandholm. Automated mechanism design for a self-interested designer. In *Proceedings of the 4th ACM conference on Electronic commerce*, pages 232–233. ACM, 2003.

[7] C. Daskalakis, A. Deckelbaum, and C. Tzamos. Strong duality for a multiple-good monopolist. *Econometrica*, to appear.

[8] C. Deng and S. Pekec. Money for nothing: exploiting negative externalities. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 361–370. ACM, 2011.

[9] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *The American economic review*, 97(1):242–259, 2007.

[10] J. D. Hartline and T. Roughgarden. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM conference on Electronic commerce*, pages 225–234. ACM, 2009.

[11] Z. Huang, Y. Mansour, and T. Roughgarden. Making the most of your samples. In *Proceedings of the Sixteenth ACM Conference on Economics and*

Computation, pages 45–60. ACM, 2015.

[12] S. Lahaie and D. M. Pennock. Revenue analysis of a family of ranking rules for keyword auctions. In *Proceedings of the 8th ACM conference on Electronic commerce*, pages 50–56. ACM, 2007.

[13] X. Li and A. C.-C. Yao. On revenue maximization for selling multiple independently distributed items. *Proceedings of the National Academy of Sciences*, 110(28):11232–11237, 2013.

[14] E. Maskin, J. Riley, and F. Hahn. Optimal multi-unit auctions. *The economics of missing markets, information, and games*, 1989.

[15] V. Mirrokni, R. P. Leme, P. Tang, and S. Zuo. Dynamic auctions with bank accounts. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 2016.

[16] R. B. Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

[17] D. Nekipelov, V. Syrgkanis, and E. Tardos. Econometrics for learning agents. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 1–18. ACM, 2015.

[18] N. Nisan and A. Ronen. Algorithmic mechanism design. In *Proceedings of the thirty-first annual ACM symposium on Theory of computing*, pages 129–140. ACM, 1999.

[19] A. D. Procaccia, D. Wajc, and H. Zhang. Approximation-variance tradeoffs in mechanism design. working paper, 2016.

[20] B. Roberts, D. Gunawardena, I. A. Kash, and P. Key. Ranking and tradeoffs in sponsored search auctions. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, pages 751–766. ACM, 2013.

[21] P. Tang and T. Sandholm. Optimal auctions for spiteful bidders. In *AAAI*, 2012.

[22] P. Tang and T. W. Sandholm. Approximating optimal combinatorial auctions for complements using restricted welfare maximization. AAAI, 2011.

[23] P. Tang and Z. Wang. Optimal auctions for negatively correlated items. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 103–120. ACM, 2016.

[24] P. Tang and Z. Wang. Optimal mechanisms with simple menus. *Journal of Mathematical Economics*, 2017.

[25] H. R. Varian. Position auctions. *international Journal of industrial Organization*, 25(6):1163–1178, 2007.

[26] A. C.-C. Yao. An n-to-1 bidder reduction for multi-item auctions and its applications. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 92–109. SIAM, 2015.