

Autobidding Auctions in the Presence of User Costs

Yuan Deng
Google Research
USA
dengyuan@google.com

Jieming Mao
Google Research
USA
maojm@google.com

Vahab Mirrokni
Google Research
USA
mirrokni@google.com

Hanrui Zhang
Carnegie Mellon University
USA
hanruiz1@cs.cmu.edu

Song Zuo
Google Research
USA
szuo@google.com

ABSTRACT

We study autobidding ad auctions with user costs, where each bidder is value-maximizing subject to a return-over-investment (ROI) constraint, and the seller aims to maximize the social welfare taking into consideration the user's cost of viewing an ad. We show that in the worst case, the approximation ratio of social welfare by running the vanilla VCG auctions with user costs could be as bad as 0. To improve the performance of VCG, we propose a new variant of VCG based on properly chosen cost multipliers, and prove that there exist auction-dependent and bidder-dependent cost multipliers that guarantee approximation ratios of $1/2$ and $1/4$ respectively in terms of the social welfare. We validate our theory via experiments that demonstrate significant improvement on real-world auction data.

CCS CONCEPTS

• **Theory of computation** → **Algorithmic game theory and mechanism design.**

KEYWORDS

autobidding, mechanism design, online advertising

ACM Reference Format:

Yuan Deng, Jieming Mao, Vahab Mirrokni, Hanrui Zhang, and Song Zuo. 2023. Autobidding Auctions in the Presence of User Costs. In *Proceedings of The Web Conference (WWW'23)*. ACM, New York, NY, USA, 9 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

1 INTRODUCTION

As 2022, over 90% of online advertising transact through automation technologies [27]. As one key application of automation in online advertising, auto-bidding significantly simplifies the tasks of advertisers by delegating the bidding to automated agents who submit real-time bids on behalf of the advertisers in each ad auction to optimize with pre-specified high-level objectives and constraints. One popular and representative example of auto-bidding strategies

is to maximize the total advertiser value subject to a target return-on-investment (ROI) constraint declared by the advertiser. This type of auto-bidders is referred as *value maximizers* by a recent line of research [5, 13] as these auto-bidders aim to maximize the total advertiser value from ad auctions, such as the (weighted) number of clicks or conversions, while guaranteeing the ROI no less than the advertiser-specified threshold.

The differences in behavior models between the classic (quasi-linear) *utility maximizers* (who optimize for value minus payment) and value maximizers results in different conclusions for many basic problems in auction theory. As one demonstrating example, Vickrey-Clarke-Groves (VCG) auction is known to achieve the optimal social welfare when bidders are utility maximizers. However, Aggarwal et al. [2] showed that in the worst case, the welfare of VCG auction can be as bad as $1/2$ of the optimal social welfare when bidders are value maximizers. Similar distinctions have been discovered by recent works revisiting the effectiveness of various existing mechanisms in the presence of value maximizers [5, 12, 24, 25].

User Costs. User costs to the auctioneer do not receive much attention in the literature on auction design for online advertising with utility maximizers, despite of the prevalence of the externalities on user experience of displaying ads to users [1]. One reason is that when bidders are utility maximizers, many commonly studied auctions and their properties directly generalize from models without user costs to models with user costs. For example, VCG auction by definition directly applies to the model with user costs and continues to achieve the optimal social welfare. Abrams and Schwarz [1] generalize General-Second-Price (GSP) auction to accommodate the cost from user experience impact in online advertising by treating the original bids minus the costs as the effective bids from advertisers. They further prove that GSP preserves the property that there exists an equilibrium achieving the optimal social welfare after the generalization.

Challenges from User Cost in Autobidding. However, such nice generalization of concepts and results for user costs no longer works when the bidders are value maximizers. As we will show later in Section 4.2, VCG auction can have 0 approximation to the optimal social welfare in the worst case when bidders are value maximizers. One intuitive explanation of the worsened approximation guarantee of VCG auction in this case lies behind the so-called *uniform-bidding* strategy widely adopted by auto-bidding agents [8, 17]. Under uniform-bidding strategies, each bidder bids its value of each auction times the *bid multiplier* that is constant across all

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WWW'23, April 30 – May 4, 2023, Austin, TX

© 2023 Association for Computing Machinery.
ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00
<https://doi.org/XXXXXXXX.XXXXXXX>

auctions for this bidder. From the perspective of value maximizers, the best response is always achievable by uniform-bidding with a properly optimized bid multiplier under VCG auction with/without user costs [2, 13]. In contrast, from the perspective of the auctioneer, bidders participating in the same auction could have different bid multipliers, meaning that their base values are scaled up by different magnitudes. Thus, the ranking of their bids can be very different from the efficient ranking, leading to an inefficient allocation. The inefficiency becomes worse off in the presence of user costs, since the cost associated to each bidder is not scaled by the corresponding bid multiplier in VCG auction. For bidders with bid multipliers much larger than 1, the effect of their costs in ranking is then vanished, resulting in additional allocation inefficiency.

Based on the above observation, one naïve and minimal fix to VCG auction is to artificially scale user costs by a multiplier and hope that VCG auction can retain a good approximation guarantee with a properly chosen *cost multiplier*. Unfortunately, using one global cost multiplier is not enough as we will show later in Section 4.2, the approximation guarantee for VCG auction with a global cost multiplier is still 0 in the worst case.

Our Results. In this paper, we present how the constant approximation of social welfare can be retained with one step further on top of the above idea, namely, VCG auction with cost multipliers. In particular, we prove that there exists a set of cost multipliers corresponding to each auction (i.e., auction-dependent but bidder-independent cost multipliers) such that the welfare in any equilibrium among value maximizers under VCG auction with those cost multipliers is at least 1/2 of the optimal social welfare (Section 4.3). In addition, we further show that there exists a set of cost multipliers corresponding to each bidder (i.e., bidder-dependent but auction-independent cost multipliers) such that the welfare performance is at least 1/4 of the optimal social welfare (Section 4.4).

As a complement to our theoretical results, in Section 5, we empirically evaluate the performance of several “showcase” candidates following the intuition of the construction in our proofs using semi-synthetic auction data derived from real auction data of a major online advertising platform. We note that these proposed “showcase” candidates do not aim for practical use but mainly serve as proof-of-concept examples for the class of auctions with provable approximation guarantees.

1.1 Related Work

Since the seminal work of Vickrey–Clarke–Groves (VCG) auctions [10, 21, 31] and Myerson’s auction [28], auction theory has found numerous practical applications. Examples includes combinatorial auctions for reallocating radio frequencies [11], generalized second-price auctions (GSP) and dynamic auctions for online advertising [3, 15, 26, 30]. Most of the work in auction theory assume bidders are utility maximizers.

More recently, there’s a line of research focusing on value maximizing bidders motivated by auto-bidding (e.g. target CPA bidders and target ROAS bidders) in online advertising. Aggarwal et al. [2] characterizes the optimal bidding strategies in truthful auction for value maximizing bidders and shows that in the worst case, the welfare of VCG auction can be as bad as 1/2 of the optimal social welfare. Balseiro et al. [7], Deng et al. [12, 13] show how boosts and

reserves can be used to improve the welfare efficiency in auctions with value maximizing bidders. Balseiro et al. [5, 6] characterize the revenue-optimal single-stage auctions with value maximizers with or without budget constraints under various information structure. In addition to the above, there are other related work studying slightly different bidding models, e.g. Babaiouff et al. [4], Goel et al. [18, 19], Golrezaei et al. [20] study utility maximizers with ROI constraints, and Fadaei and Bichler [16], Wilkens et al. [32, 33] study value maximizing bidders with per-auction target constraints.

As far as we know, Abrams and Schwarz [1] is the first paper to study auction design with user costs for utility maximizing bidders, and Li et al. [23] extends the study to include position-specific information. User costs are motivated by studies on “banner blindness” [9, 14, 22] which show how user experience affects future user behavior with ads.

2 PRELIMINARIES

Ad auctions with user costs. Following prior work on autobidding ad auctions [2, 7, 13, 25], we consider the following multi-auction model. There are n bidders $[n] = \{1, 2, \dots, n\}$ and m auctions $[m] = \{1, 2, \dots, m\}$. We generally use i to index bidders, and j to index auctions. In each auction j , each bidder i has a *value* $v_{i,j}$, as well as a *user cost* $c_{i,j}$. The user cost is suffered by the user who views the ad provided by bidder i in auction j when i wins. As such, it does not affect the behavior of the bidder directly, but counts negatively towards the social welfare. We assume the value $v_{i,j}$ is private, but the user cost $c_{i,j}$ is public and can be directly observed by the auction mechanism. In each auction j , each bidder i submits a *bid* $b_{i,j}$, and all bidders’ bids $\{b_{i',j}\}_{i'}$ and user costs $\{c_{i',j}\}_{i'}$ together determine each i ’s *allocation* $x_{i,j}$ (i.e., $x_{i,j} = 1$ if i wins in j , and 0 otherwise) and *payment* $p_{i,j}$, in the way prescribed by the auction mechanism (we will discuss auction mechanisms momentarily). We omit the dependence of $x_{i,j}$ and $p_{i,j}$ on $\{b_{i',j}\}_{i'}$ when it is clear from the context.

ROI-constrained value maximizers. We assume bidders are ROI-constrained value maximizers. That is, each bidder i maximizes the total value i receives in all auctions, subject to the constraint that the ratio between this total value and the total payment i makes in all auctions is at least some target quantity. Following prior work [7, 12, 13], without loss of generality, we assume this target ratio is 1 for all bidders. Formally, each bidder i solves the following optimization problem to decide how to bid in all auctions:

$$\begin{aligned} \max \quad & \sum_j x_{i,j} \cdot v_{i,j} \\ \text{subject to} \quad & \sum_j (x_{i,j} \cdot v_{i,j} - p_{i,j}) \geq 0. \end{aligned}$$

We always assume there is a way of bidding such that the ROI constraint is satisfied (this is true for all quasilinear truthful auction mechanisms).

Quasilinear truthfulness and uniform bidding. In this paper, we focus on auction mechanisms that are quasilinear truthful. An auction mechanism is quasilinear truthful, if each bidder would be motivated to bid their true value if they were to maximize their quasilinear utility. In other words, for each i and j , fixing any $v_{i,j}$

and $\{b_{i',j}\}_{i' \neq i}$, the following holds for any possible bid b :

$$\begin{aligned} & x_{i,j}(b_{i,j} = v_{i,j}, \{b_{i',j}\}_{i' \neq i}) \cdot v_{i,j} - p_{i,j}(b_{i,j} = v_{i,j}, \{b_{i',j}\}_{i' \neq i}) \\ & \geq x_{i,j}(b_{i,j} = b, \{b_{i',j}\}_{i' \neq i}) \cdot v_{i,j} - p_{i,j}(b_{i,j} = b, \{b_{i',j}\}_{i' \neq i}). \end{aligned}$$

One example of quasilinear truthful mechanisms is the *second-price auction*: in each auction j , the bidder i with the highest bid wins, and the payment i makes is the second highest bid. In addition, quasilinear truthful mechanisms enjoy the following nice property: without loss of generality, the optimal bidding strategy of an ROI-constrained value maximizer is always *uniform bidding* [2]. That is, for each bidder i , there exists a bid multiplier $\theta_i \geq 1$, such that fixing other bidders' bids, the optimal (i.e., value-maximizing subject to the ROI constraint) bidding strategy of i is to bid $b_{i,j} = \theta_i \cdot v_{i,j}$ in each auction j . Note that the bid multiplier is never smaller than 1, which in particular means no bidder should ever bid below their true value. In the rest of the paper, we only consider such bidding strategies, and sometimes use the corresponding bid multiplier to represent a bidding strategy.

Equilibria and the price of anarchy. To measure the efficiency of an auction mechanism, we consider the behavior of bidders in equilibrium under that mechanism. We say the bidding strategies $\{\theta_i\}_i$ form an equilibrium, if each bidder i 's bidding strategy θ_i is a best response to all other bidders' strategies $\theta_{-i} = \{\theta_{i'}\}_{i' \neq i}$. Moreover, we use the price of anarchy (PoA) to quantify the efficiency of a mechanism. The PoA of a mechanism is the ratio between the worst-case welfare in equilibrium under that mechanism and the optimal welfare. Formally, fixing a mechanism,

$$\text{PoA} = \inf_{\substack{n, m, \{v_{i,j}\}, \{c_{i,j}\} \\ \{\theta_i\}_i \text{ form an equilibrium}}} \frac{\sum_{i,j} x_{i,j} \cdot (v_{i,j} - c_{i,j})}{\sum_j \max \{0, \max_i (v_{i,j} - c_{i,j})\}}.$$

There are 2 things worth noting in the above definition:

- If every bidder would contribute negatively to the social welfare if they were to win in a particular auction j (i.e., $\max_i (v_{i,j} - c_{i,j}) < 0$), then auction j contributes 0 to the optimal welfare.
- In order for the PoA to be well defined, we always assume the optimal welfare is positive.

3 CHALLENGES POSED BY USER COSTS

To gain some intuition, we first briefly review existing approaches to autobidding auctions through quasilinear truthful mechanisms, particularly the second-price auction, and discuss why it fails in the presence of user costs. In a second-price auction, the bidder with the highest bid wins, and the payment of the winner is the second highest bid. It is known that the second-price auction has a PoA of $1/2$ [2] when (1) bidders are ROI-constrained value maximizers and (2) there is no user cost. Variants of the second-price auction has been analyzed in other cost-free settings, sometimes leading to improved efficiency guarantees (e.g., in settings with machine-learned advice [7]). Nonetheless, in the core of these analyses lies the following conceptually simple argument:

THEOREM 3.1 ([2, 7, 13]). *Suppose there is no user cost. The second-price auction achieves a PoA of $1/2$.*

PROOF. Suppose there is no user cost. For each auction j , consider the “rightful winner” $\text{rw}(j)$ who has the highest value in j , i.e., $\text{rw}(j) = \text{argmax}_i v_{i,j}$ (ties are broken arbitrarily). Since the second-price auction is quasilinear truthful, in equilibrium, all bidders, including the rightful winner, must bid at least their own value, i.e., $b_{i,j} \geq v_{i,j}$ for all i and j . Therefore, in each auction j , one of the following two cases must happen: (1) the rightful winner $\text{rw}(j)$ wins, and collects value $v_{i,j}$ from auction j , or (2) the rightful winner $\text{rw}(j)$ does not win, in which case the second highest bid is at least $v_{\text{rw}(j),j} = \max_i v_{i,j}$, and therefore the winner must pay at least this amount in auction j .

Summing over all auctions, the above reasoning implies that the sum of the following two quantities must be at least the optimal welfare: (1) the total value in all auctions collected by the respective rightful winners, and (2) the total payment made in all auctions by bidders other than the respective rightful winners. The first quantity clearly lower bounds the welfare in equilibrium. In fact, the second quantity is also a lower bound of the welfare: recall that the ROI constraints require that the total value each bidder receives should be at least the total payment that bidder makes, and summing over all bidders, one can see the second quantity also lower bounds the welfare in equilibrium. Now since the sum of the two quantities is at least the optimal welfare, at least one of the two must be no smaller than half of the optimal welfare, which means the welfare in equilibrium is no smaller than half of the optimal welfare. This gives a lower bound of $1/2$ on the PoA of the second-price auction. \square

In words, the above argument lower bounds the total value collected in equilibrium and the total payment made in equilibrium, respectively, and use the larger one between the two as a lower bound of the welfare in equilibrium. Given this argument, it might appear that the most natural mechanism in the presence of user costs is the second price auction with “cost-adjusted bids”. That is, bidders are sorted by the following adjusted bid: $b_{i,j} - c_{i,j}$ (negative adjusted bids are discarded); the bidder with the highest adjusted bid wins, and pays the minimum bid to win (by Myerson’s characterization [28], this is the unique payment rule that guarantees quasilinear truthfulness).¹

Now let us try to generalize the argument above to this new mechanism. For each auction j , the rightful winner $\text{rw}(j)$ is the one with the largest cost-adjusted value, i.e., $\text{rw}(j) = \text{argmax}_i (v_{i,j} - c_{i,j})$. Again there are two cases: (1) if the rightful winner $\text{rw}(j)$ wins in auction j , then the welfare collected by $\text{rw}(j)$ is $v_{\text{rw}(j),j} - c_{\text{rw}(j),j}$; (2) if the rightful winner does not win, then the second highest cost-adjusted bid must be at least $\max_i (v_{i,j} - c_{i,j})$, so payment made by the winner (say i^*) is at least $c_{i^*,j} + \max_i (v_{i,j} - c_{i,j})$.

Now one might be tempted to apply the following (faulty) reasoning and argue the PoA is again at least $1/2$: in case (1), the contribution of auction j to the total value collected minus user cost incurred is $v_{\text{rw}(j),j} - c_{\text{rw}(j),j} = \max_i (v_{i,j} - c_{i,j})$, which is the same as the contribution of j to the optimal welfare. Similarly, in case (2), the contribution of j to the total payment made minus user

¹To make this concrete, suppose there are 3 bidders in some auction j . The bids and costs are: $b_{1,j} = 5$, $b_{2,j} = 3$, $b_{3,j} = 4$, $c_{1,j} = 1$, $c_{2,j} = 2$, and $c_{3,j} = 1$. Then the cost-adjusted bids are $b_{1,j} - c_{1,j} = 4$, $b_{2,j} - c_{2,j} = 1$, and $b_{3,j} - c_{3,j} = 3$, so bidder 1 wins. Bidder 3 has the second highest cost-adjusted bid, so the minimum bid that bidder 1 has to make in order to win is $b_{3,j} - c_{3,j} + c_{1,j} = 3 + 1 = 4$, so bidder 1 pays $p_{1,j} = 4$.

cost incurred is $\max_i(v_{i,j} - c_{i,j})$, which is again the same as the contribution of j to the optimal welfare. Now if we sum over all j , then we know that the sum of (1) the total value minus user cost and (2) the total payment minus user cost is at least the optimal welfare, so the larger one is at least half of the optimal welfare, which gives a PoA lower bound of $1/2$.

The problem of this reasoning is in the last step: when summing over j , the user cost is only counted once, so what we actually get is: the sum of the total value and the total payment, minus the total user cost incurred, is at least the optimal welfare. This unfortunately does not give us any nontrivial bound on the welfare in equilibrium.

In fact, we will show in Section 4.2 that the second-price auction with cost-adjusted bids has a PoA of 0. Indeed, to handle user costs we need fundamentally new ideas and techniques.

4 OUR MECHANISMS

We present two quasilinear truthful mechanisms with provable efficiency guarantees in the presence of user costs. The first mechanism is anonymous, i.e., the outcome that any bidder receives is independent of that bidder's identity (or index). The second mechanism is auction-oblivious, i.e., it works in the same way across all auctions, and may treat different bidders in different ways.

4.1 Cost Multipliers

We first introduce the key idea behind both of our mechanisms, the use of cost multipliers. Instead of adjusting each bidder's bid by subtracting the user cost as is (as in the second-price auction with cost-adjusted bids discussed in the previous section), we apply a multiplier to the cost before performing the subtraction. We illustrate the power of this idea in the single-bidder setting, which also serves as a warm-up for our subsequent discussion.

Cost multipliers in the single-bidder setting. Suppose $n = 1$. For each auction j , let v_j , c_j , and b_j be the only bidder's value, user cost, and bid in j , respectively. Let S be the set of auctions where the bidder's value is no smaller than the respective cost, i.e., $S = \{j \mid v_j \geq c_j\}$. Then, the optimal welfare is simply $\sum_{j \in S} (v_j - c_j)$.

Let $\alpha \geq 1$ be the unique number such that $\sum_{j \in S} v_j = \alpha \cdot \sum_{j \in S} c_j$. Consider the following mechanism: in each auction j , the only bidder wins if the bid $b_j \geq \alpha \cdot c_j$, in which case the bidder pays the minimum bid to win, i.e., $p_j = \alpha \cdot c_j$; otherwise, no bidder wins and no payment happens. Observe that the bidder's optimal bidding strategy in response to this mechanism is to set their bid multiplier to exactly α , such that $b_j = \alpha \cdot v_j$ in each j . This is because when the bid multiplier is α , the bidder wins in all auctions in S and no other auctions, so the total value collected by the bidder is $\sum_{j \in S} v_j$. On the other hand, in each $j \in S$, the payment made by the bidder is $p_j = \alpha \cdot c_j$, so by the choice of α , the total payment is $\sum_{j \in S} \alpha \cdot c_j = \sum_{j \in S} v_j$, which is the same as the total value. In other words, when the bid multiplier is α , the bidder's ROI constraint is binding. This means there is no way to win in more auctions without violating the ROI constraint, because the marginal ROI of winning in any other auction is strictly smaller than $1/\alpha \leq 1$. So, the only equilibrium is fully efficient, and the PoA of this mechanism in the single-bidder setting is 1.

4.2 Mechanisms with Global Cost Multipliers

However, it turns out that it is not enough to use a *global* cost multiplier. In particular, mechanisms with global cost multipliers can get arbitrarily worse PoA.

THEOREM 4.1. *For any $\delta \in (0, 1/3)$, there exists an instance in which any mechanism with a global cost multiplier gets PoA $\leq 3\delta$.*

PROOF. Consider the instance with $n = \lfloor 1/\delta \rfloor$ and $m = 2n$. Bidder i only has non-zero values and costs in auction $2i - 1$ and $2i$. The values and costs of bidder i are specified as the follow table:

	Auction $2i - 1$	Auction $2i$
Values $(v_{i,2i-1}, v_{i,2i})$	δ	$1 + 1/\delta^i$
Costs $(c_{i,2i-1}, c_{i,2i})$	$1 - \delta$	$1/\delta^i$

It is easy to check that the optimal welfare can be achieved by letting bidder i only wins auction $2i$, i.e.

$$\text{OPT} = \sum_j \max \left\{ 0, \max_i (v_{i,j} - c_{i,j}) \right\} = n.$$

Notice that

$$\frac{v_{i,2i}}{c_{i,2i}} = 1 + \delta^i,$$

and

$$\frac{v_{i,2i-1} + v_{i,2i}}{c_{i,2i-1} + c_{i,2i}} = 1 + \frac{2\delta}{1 - \delta + \delta^{-i}} > 1 + \frac{2\delta}{2\delta^{-i}} = 1 + \delta^{i+1}.$$

For any mechanism with a global cost multiplier α , if $\alpha > 1 + \delta^i$, bidder i will win nothing and contribute 0 to the welfare. And if $\alpha \leq 1 + \delta^{i+1}$, bidder i will win both auction $2i - 1$ and auction $2i$, and the overall welfare is $\delta + 1 + 1/\delta^i - (1 - \delta + 1/\delta^i) = 2\delta$.

Therefore, for each bidder $i \in [n]$, its contribution to welfare is at most 2δ if $\alpha \notin (1 + \delta^{i+1}, 1 + \delta^i]$, and at most 1 otherwise. Note that these intervals $(1 + \delta^{i+1}, 1 + \delta^i]$ are disjoint, and thus the overall welfare of a mechanism with any global multiplier α is at most

$$2\delta \cdot (n - 1) + 1 < (2\delta + 1/n) \cdot n \leq 3\delta n = 3\delta \cdot \text{OPT}. \quad \square$$

Such an impossibility result implies that it is necessary to incorporate additional information into the design of cost multipliers to be effective. In Section 4.3 and Section 4.4, we show the existence of effective auction-dependent cost multipliers and bidder-dependent cost multipliers, respectively. The constructions of these cost multipliers in the existence proofs rely on knowing additional auction-specific and/or bidder-specific value information, which may not be accessible in a real system. One could consider applying data-driven techniques to construct cost multipliers in practice, which is outside of the scope of this paper.

4.3 Auction-Dependent Cost Multipliers

In this section, we show that there exist auction-dependent (and bidder-independent) cost multipliers with a PoA of $1/2$.

The mechanism. In each auction j , if $\max_i (v_{i,j} - c_{i,j}) < 0$, then no bidder wins (we let $\text{rw}(j) = 0$ in such cases). Otherwise, let $\text{rw}(j) = \text{argmax}_i (v_{i,j} - c_{i,j})$ (breaking ties arbitrarily), and let $\alpha_j = (v_{\text{rw}(j),j} - c_{\text{rw}(j),j}) / (2c_{\text{rw}(j),j})$. Sort all bidders by the following score: $b_{i,j} - (1 + \alpha_j) \cdot c_{i,j}$. Let $i_1(j)$ and $i_2(j)$ be the bidders with the highest and second highest scores respectively (again breaking

ties arbitrarily). $i_1(j)$ wins (i.e., $x_{i_1(j),j} = 1$) if their score is at least 0, and pays the minimum bid to win, i.e.,

$$p_{i_1(j),j} = \max\{b_{i_2(j),j} - (1 + \alpha_j) \cdot c_{i_2(j),j}, 0\} + (1 + \alpha_j) \cdot c_{i_1(j),j}.$$

If $i_1(j)$'s score is smaller than 0, then no bidder wins.

One may check that the above mechanism is in fact quasilinear truthful. There is one ambiguity in the mechanism: for an auction j , if $c_{rw(j),j} = 0$, then the mechanism sets $\alpha_j = \infty$, which may introduce undefined behavior. Although this corner case is unlikely to be a problem in practice, for consistency in theory, we use the following convention: if $\alpha_j = \infty$, then for any bidder i , $\alpha_j \cdot c_{i,j} = \infty$ if $c_{i,j} > 0$, and $\alpha_j \cdot c_{i,j} = \frac{1}{2}v_{rw(j),j}$ if $c_{i,j} = 0$. One may check that with this convention, our analysis correctly handles infinite multipliers.

Analysis of the mechanism. The efficiency of the mechanism is captured by the following theorem.

THEOREM 4.2. *The mechanism with auction-dependent cost multipliers has a PoA of 1/2.*

PROOF. First observe that by Myerson's characterization [28], the mechanism is quasilinear truthful, and so each bidder i 's bid in each auction j is at least their value, i.e., $b_{i,j} \geq v_{i,j}$. Using this fact, we will lower bound the difference between the total payment made in all auctions and the total cost incurred in all auctions in equilibrium. That is, we lower bound the following quantity:

$$\sum_{i,j} (p_{i,j} - x_{i,j} \cdot c_{i,j}).$$

Due to the ROI constraints, this is a lower bound of the welfare in equilibrium. Fix any equilibrium. We compare the contribution of each auction to the above difference against the contribution of the same auction to the optimal welfare. For each auction j , consider the following 3 cases:

- $\max_i (v_{i,j} - c_{i,j}) < 0$ and no bidder wins. In this case, the payment made and the cost incurred in j are both 0. On the other hand, the contribution of j to the optimal welfare is also 0. So the contribution of j to the difference is (trivially) at least 1/2 of the contribution of j to the optimal welfare.
- $rw(j)$ wins in auction j , i.e., $i_1(j) = rw(j)$. In this case, by the choice of α_j , the payment that $rw(j)$ makes is at least

$$p_{rw(j),j} \geq (1 + \alpha_j) \cdot c_{rw(j),j} = c_{rw(j),j} + \frac{1}{2}(v_{rw(j),j} - c_{rw(j),j}).$$

The cost incurred in j is $c_{rw(j),j}$, so the contribution of j to the difference is at least $\frac{1}{2}(v_{rw(j),j} - c_{rw(j),j})$. On the other hand, the contribution of j to the optimal welfare is $v_{rw(j),j} - c_{rw(j),j}$, and the ratio between the contribution to the difference and that to the optimal welfare is at least 1/2.

- A bidder other than the rightful winner wins. In this case, we know the second highest score is at least the score of the rightful winner $rw(j)$. Since the rightful winner bids at least their value, and by the choice of α , we have

$$\begin{aligned} b_{i_2(j),j} - (1 + \alpha_j) \cdot c_{i_2(j),j} &\geq b_{rw(j),j} - (1 + \alpha_j) \cdot c_{rw(j),j} \\ &\geq v_{rw(j),j} - (1 + \alpha_j) \cdot c_{rw(j),j} \\ &= \frac{1}{2}(v_{rw(j),j} - c_{rw(j),j}). \end{aligned}$$

Then the payment of the winner is at least

$$\begin{aligned} p_{i_1(j),j} &\geq b_{i_2(j),j} - (1 + \alpha_j) \cdot c_{i_2(j),j} + (1 + \alpha_j) \cdot c_{i_1(j),j} \\ &\geq \frac{1}{2}(v_{rw(j),j} - c_{rw(j),j}) + (1 + \alpha_j) \cdot c_{i_1(j),j}. \end{aligned}$$

The cost incurred in j is $c_{i_1(j),j}$, so the contribution of j to the difference is at least

$$\begin{aligned} p_{i_1(j),j} - c_{i_1(j),j} &\geq \frac{1}{2}(v_{rw(j),j} - c_{rw(j),j}) + \alpha_j \\ &\geq \frac{1}{2}(v_{rw(j),j} - c_{rw(j),j}). \end{aligned}$$

In other words, the contribution of j to the difference is at least 1/2 of the contribution of j to the optimal welfare.

Summarizing the 3 cases, we see that the contribution of each auction j to the difference between the total payment and the total cost is at least 1/2 of the contribution of the same auction to the optimal welfare. As a result, the difference between the total payment and the total cost is at least 1/2 of the optimal welfare. By the ROI constraints of the bidders, the former quantity is a lower bound of the welfare in equilibrium, which gives a lower bound of 1/2 on the PoA. \square

4.4 Bidder-Dependent Cost Multipliers

We now proceed to show the existence of bidder-dependent (and auction-independent) cost multipliers with a PoA of 1/4.

The mechanism. For each bidder i , let S_i be the set of auctions where i is the rightful winner (as defined in the previous subsection), i.e., $S_i = \{j \mid rw(j) = i\}$. Note that by definition, for each i and $j \in S_i$, $v_{i,j} - c_{i,j} \geq 0$. For each bidder i , let $\alpha_i \geq 0$ be the unique number such that

$$\sum_{j \in S_i} v_{i,j} = (1 + 2\alpha_i) \cdot \sum_{j \in S_i} c_{i,j}.$$

In each auction j , first discard any bidder i where $b_{i,j} - (1 + \alpha_i) \cdot c_{i,j} < 0$. If all bidders are discarded, then no bidder wins in auction j . Otherwise, let each bidder i 's score be the cost-adjusted bid $b_{i,j} - c_{i,j}$ and sort all remaining bidders by their scores. The bidder among the remaining ones with the highest score, $i_1(j)$, wins, and pays the minimum bid to win. That is, if $i_1(j)$ is the only remaining bidder, then $i_1(j)$ pays

$$p_{i_1(j),j} = (1 + \alpha_{i_1(j)}) \cdot c_{i_1(j),j}.$$

Otherwise, let $i_2(j)$ be the bidder among the remaining ones with the second highest score. $i_1(j)$ pays

$$p_{i_1(j),j} = \max\{(1 + \alpha_{i_1(j)}) \cdot c_{i_1(j),j}, b_{i_2(j),j} - c_{i_2(j),j} + c_{i_1(j),j}\}.$$

Again, one can check the above mechanism is in fact quasilinear truthful. We also note that this mechanism, unlike our anonymous mechanism, uses different quantities to prescreen and sort bidders. In particular, the quantity used for prescreening is no larger than the score used for sorting. This is necessary for the mechanism to provide the desired efficiency guarantee.

Analysis of the mechanism. The efficiency of the mechanism is captured by the following theorem.

THEOREM 4.3. *The mechanism with bidder-dependent cost multipliers has a PoA of 1/4.*

The proof of the above theorem is fairly involved. We first introduce some useful notions and properties. For each bidder i , let i 's core auctions T_i be the subset of S_i where i 's value is at least $(1 + \alpha_i)$ times i 's user cost, i.e.,

$$T_i = \{j \in S_i \mid v_{i,j} \geq (1 + \alpha_i) \cdot c_{i,j}\}.$$

We will use the following fact regarding T_i :

LEMMA 4.4. *For each bidder i ,*

$$\sum_{j \in T_i} (v_{i,j} - c_{i,j}) \geq \frac{1}{2} \sum_{j \in S_i} (v_{i,j} - c_{i,j}).$$

PROOF. For each i ,

$$\begin{aligned} \sum_{j \in T_i} (v_{i,j} - c_{i,j}) &= \sum_{j \in S_i} (v_{i,j} - c_{i,j}) - \sum_{j \in S_i \setminus T_i} (v_{i,j} - c_{i,j}) \\ &= 2\alpha_i \cdot \sum_{j \in S_i} c_{i,j} - \sum_{j \in S_i \setminus T_i} (v_{i,j} - c_{i,j}) \\ &\hspace{15em} \text{(choice of } \alpha_i) \\ &\geq 2\alpha_i \cdot \sum_{j \in S_i} c_{i,j} - \sum_{j \in S_i \setminus T_i} \alpha_i \cdot c_{i,j} \\ &\hspace{15em} \text{(definition of } T_i) \\ &\geq \alpha_i \cdot \sum_{j \in S_i} c_{i,j} \\ &= \frac{1}{2} \sum_{j \in S_i} (v_{i,j} - c_{i,j}). \hspace{5em} \text{(choice of } \alpha_i) \end{aligned}$$

This establishes the claim. \square

Fix an equilibrium. We partition all bidders into *aggressive* ones B_{agg} and *conservative* ones B_{con} , based on how large or small a bidder's bid multiplier is. We say a bidder i is aggressive ($i \in B_{\text{agg}}$), if i 's bid multiplier θ_i is at least α_i ; i is conservative otherwise:

$$B_{\text{agg}} = \{i \mid \theta_i \geq \alpha_i\}, \quad B_{\text{con}} = \{i \mid \theta_i < \alpha_i\}.$$

The plan is to consider two quantities separately, which each lower bound the welfare in equilibrium. The first quantity, say A , is the total value collected, minus the total cost imposed, by all conservative bidders in their respective core auctions. Formally,

$$A = \sum_{i \in B_{\text{con}}, j \in T_i} x_{i,j} \cdot (v_{i,j} - c_{i,j}).$$

The second quantity, say B , is the total payment made by all bidders, minus the total cost imposed, in auctions (1) where an aggressive bidder is the rightful winner, or (2) which are among the core auctions of a conservative bidder and that bidder does not win. Formally,

$$B = \sum_{i \in B_{\text{agg}}, j \in S_i} \sum_{i' \neq i} (p_{i',j} - x_{i',j} \cdot c_{i',j}) + \sum_{i \in B_{\text{con}}, j \in T_i} \sum_{i' \neq i} (p_{i',j} - x_{i',j} \cdot c_{i',j}).$$

The following lemma states that each of the two quantities is a lower bound of the welfare in equilibrium.

LEMMA 4.5.

$$\max\{A, B\} \leq \sum_{i,j} x_{i,j} \cdot (v_{i,j} - c_{i,j}).$$

PROOF. Consider A first. We first decompose the welfare into the contribution of each bidder.

$$\sum_{i,j} x_{i,j} \cdot (v_{i,j} - c_{i,j}) = \sum_i \sum_{j: x_{i,j}=1} (v_{i,j} - c_{i,j}).$$

Observe that each bidder i 's contribution is nonnegative. This is because for each i ,

$$\begin{aligned} \sum_{j: x_{i,j}=1} (v_{i,j} - c_{i,j}) &\geq \sum_{j: x_{i,j}=1} (p_{i,j} - c_{i,j}) \hspace{2em} (i\text{'s ROI constraint}) \\ &\geq \sum_{j: x_{i,j}=1} ((1 + \alpha_i) \cdot c_{i,j} - c_{i,j}) \hspace{2em} \text{(payment rule)} \\ &= \alpha_i \cdot \sum_{j: x_{i,j}=1} c_{i,j} \geq 0. \hspace{5em} (\alpha_i \geq 0) \end{aligned}$$

So we have

$$\sum_{i,j} x_{i,j} \cdot (v_{i,j} - c_{i,j}) \geq \sum_{i \in B_{\text{con}}} \sum_{j: x_{i,j}=1} (v_{i,j} - c_{i,j}).$$

For each conservative bidder $i \in B_{\text{con}}$, further observe that the contribution of each auction j to the right hand side of the above inequality is nonnegative. This is because in order for i to win in an auction j , i 's bid must at least satisfy

$$b_{i,j} \geq (1 + \alpha_i) \cdot c_{i,j}.$$

On the other hand, since i is conservative, we know

$$b_{i,j} = \theta_i \cdot v_{i,j} \leq \alpha_i \cdot v_{i,j}.$$

So in any auction j , $x_{i,j} = 1$ only if $v_{i,j} \geq c_{i,j}$. Given this, we can further relax the above inequality and get

$$\sum_{i,j} x_{i,j} \cdot (v_{i,j} - c_{i,j}) \geq \sum_{i \in B_{\text{con}}} \sum_{j \in T_i: x_{i,j}=1} (v_{i,j} - c_{i,j}) = A.$$

Now consider B . We have

$$\begin{aligned} &\sum_{i,j} x_{i,j} \cdot (v_{i,j} - c_{i,j}) \\ &= \sum_i \sum_{j: x_{i,j}=1} (v_{i,j} - c_{i,j}) \\ &\geq \sum_i \sum_{j: x_{i,j}=1} (p_{i,j} - c_{i,j}) \hspace{2em} \text{(each } i\text{'s ROI constraint)} \\ &\geq \sum_{i \in B_{\text{agg}}, j \in S_i} \sum_{i' \neq i} (p_{i',j} - x_{i',j} \cdot c_{i',j}) \\ &\quad + \sum_{i \in B_{\text{con}}, j \in T_i} \sum_{i' \neq i} (p_{i',j} - x_{i',j} \cdot c_{i',j}) \\ &= B, \end{aligned}$$

where the last inequality is because by the payment rule, for each i and j ,

$$p_{i,j} - x_{i,j} \cdot c_{i,j} \geq x_{i,j} \cdot ((1 + \alpha_i) \cdot c_{i,j} - c_{i,j}) \geq 0.$$

So discarding the contribution of some auctions and some bidders cannot make the sum larger. This finishes the proof. \square

Now we are ready to prove Theorem 4.3.

PROOF OF THEOREM 4.3. Given Lemma 4.5, the welfare in equilibrium can be bounded in the following way:

$$\sum_{i,j} x_{i,j} \cdot (v_{i,j} - c_{i,j}) \geq \max\{A, B\} \geq \frac{1}{2}(A + B).$$

So we only need to compare $A + B$ against the optimal welfare. Below we argue the former is at least $1/2$ of the latter, which gives a lower bound of $1/4$ on the PoA.

First observe that the optimal welfare can be written as

$$\sum_j \max_i \{\max(v_{i,j} - c_{i,j}), 0\} = \sum_i \sum_{j \in S_i} (v_{i,j} - c_{i,j}).$$

Below we consider each bidder i separately, and show that the contribution of auctions in S_i to $A + B$ is at least half of the contribution of the same auctions to the optimal welfare, i.e., $\sum_{j \in S_i} (v_{i,j} - c_{i,j})$. First consider each $i \in B_{\text{agg}}$. In this case, each $j \in S_i$ has a winner. This is because the bid $b_{i,j}$ of i in j satisfies $b_{i,j} = \theta_i \cdot v_{i,j} \geq (1 + \alpha_i) \cdot v_{i,j} \geq (1 + \alpha_i) \cdot c_{i,j}$, so at least one bidder (i.e., i) passes the prescreening stage of the mechanism. Also, observe that each $j \in B_{\text{agg}}$ contributes nothing to A . So we only need to argue that the total contribution of all auctions in S_i to B is at least $1/2$ of the contribution of the same auctions to the optimal welfare. In fact, if i wins in auction j , then the payment i makes in j is at least $(1 + \alpha_i) \cdot c_{i,j}$, so the contribution to B is at least $\alpha_i \cdot c_{i,j}$. Summing over $j \in S_i$ and by the choice of α_i , the total contribution to B (and therefore to $A + B$) is at least

$$\sum_{j \in S_i} \alpha_i \cdot c_{i,j} = \frac{1}{2} \sum_{j \in S_i} (v_{i,j} - c_{i,j}),$$

which is $1/2$ of the contribution to the optimal welfare.

Now consider each $i \in B_{\text{con}}$. For such a bidder i , only auctions in T_i may contribute to $A + B$. Observe that each $j \in T_i$ has a winner. This is because the bid $b_{i,j}$ is at least $v_{i,j}$, which by the choice of T_i satisfies $v_{i,j} \geq (1 + \alpha_i) \cdot c_{i,j}$. So $b_{i,j} \geq (1 + \alpha_i) \cdot c_{i,j}$, and at least one bidder (i.e., i) passes the prescreening stage of the mechanism. Below we argue that the contribution of each $j \in T_i$ to $A + B$ is at least the contribution of the same auction to the optimal welfare. If i wins in j , then j contributes to A , and the contribution is precisely $v_{i,j} - c_{i,j}$, which is the same as the contribution of j to the optimal welfare. If i does not win in j , then the score of $i_2(j)$ is at least the score of i , i.e.,

$$b_{i_2(j),j} - c_{i_2(j),j} \geq b_{i,j} - c_{i,j} \geq v_{i,j} - c_{i,j}.$$

This means the payment made by the winner, $i_1(j)$ must satisfy

$$p_{i_1(j),j} \geq b_{i_2(j),j} - c_{i_2(j),j} + c_{i_1(j),j} \geq v_{i,j} - c_{i,j} + c_{i_1(j),j}.$$

So the contribution of j to B is at least $v_{i,j} - c_{i,j}$, which is the contribution of j to the optimal welfare. Summing over $j \in T_i$, we see that the total contribution is at least

$$\sum_{j \in T_i} (v_{i,j} - c_{i,j}),$$

which by Lemma 4.4 is at least

$$\frac{1}{2} \sum_{j \in S_i} (v_{i,j} - c_{i,j}),$$

which is the contribution of auctions in S_i to the optimal welfare.

Now we can conclude the proof by observing

$$\begin{aligned} \sum_{i,j} x_{i,j} \cdot (v_{i,j} - c_{i,j}) &\geq \max\{A, B\} && \text{(Lemma 4.5)} \\ &\geq \frac{1}{2}(A + B) \\ &\geq \frac{1}{4} \sum_i \sum_{j \in S_i} (v_{i,j} - c_{i,j}) \\ &= \frac{1}{4} \sum_j \max_i \{\max(v_{i,j} - c_{i,j}), 0\}. \end{aligned}$$

This gives a lower bound of $1/4$ on the PoA of the mechanism. \square

5 EXPERIMENTS

We present our experimental results using a semi-synthetic dataset derived from real auction data of a major search engine. It is worth highlighting that rather than estimating the efficiency potential of an auction system implemented in practice, the main goal of our experiments is to validate our theoretical findings. There could be many practical aspects in an auction system that would not be considered in theory.

We simulate second-price auctions with bids from auto-bidders only. To exclude practical noises from the real system, we generate artificial target constraints instead of using the real ones. In line with the recent literature [6, 13], we simulate the auto-bidder response by gradient descent on their bid multipliers in the log space until convergence [29].

5.1 Experiment Setup

To evaluate the performance of a new mechanism, we first obtain a starting point by pre-training the uniform bid multipliers for 25 iterations to converge to an equilibrium in an environment with second-price auctions of cost multiplier 1. We will use this starting point as the baseline for our performance measurement. After obtaining a starting point, we simulate auctions with different mechanisms for another 25 iterations with auto-bidder response to get the final equilibria for performance measurement of the mechanisms.

We evaluate the performance of four mechanisms with different ways of applying cost multipliers. We note that these proposed “showcase” candidates mainly serve as proof-of-concept examples for the class of auctions with provable approximation guarantees and they are not meant for practical use as their constructions rely on knowing additional value information, which may not be accessible in a real system.

Among these four mechanism, one mechanism applies auction-dependent cost multipliers proposed in Section 4.3, denoted by AUCTION-DEPENDENT; and the other mechanism applies bidder-dependent cost multipliers proposed in Section 4.4, denoted by BIDDER-DEPENDENT.

Cost Multipliers based on Auction Competitiveness. In addition to these two mechanisms, we also examine the performance of two related mechanisms leveraging *auction competitiveness*, partially inspired by BIDDER-DEPENDENT. Instead of looking at the ratio between value and cost, we measure the ratio between bid and

payment. In particular, COMP_i satisfies

$$\sum_j b_{i,j} \cdot x_{i,j} = \frac{1}{\text{COMP}_i} \sum_j p_{i,j}.$$

Here, COMP_i measures the auction competitiveness for bidder i in the auctions. When COMP_i is large, i.e., the ratio between bidder i 's total bid and total payment is small, it implies that the auctions bidder i participates in are competitive as the total payment is close to bidder i 's total bid. Similarly, when COMP_i is small, it implies that the auctions bidder i participates in are not very competitive as the total payment is far away from bidder i 's total bid.

Intuitively, we would like to set a higher cost multiplier for bidders with a smaller auction competitiveness. In the bidder-dependent variant, denoted by BIDDER-DEPENDENT-COMP, we use $\alpha_i = 1/\text{COMP}_i$ as the cost multiplier for bidder i ; while in the auction-dependent variant, denoted by AUCTION-DEPENDENT-COMP, for each auction j , we set $\alpha_j = 1/\text{COMP}_{rw(j)}$ to be the auction competitiveness of the rightful winner for auction j .

5.2 Experimental Results

All metrics we report are relative to the gap between the starting point and the optimal solution. More precisely, let $\kappa_{baseline}$ be the welfare of the starting point, κ_m be the welfare under mechanism m , and OPT be the optimal welfare (revenue). Then the reported percentage is computed by $(\kappa_m - \kappa_{baseline})/(\text{OPT} - \kappa_{baseline})$.

As shown in table 1, we can see that all our mechanisms get better welfare than the baseline. We only report welfare here because revenue would be mostly equal to welfare in equilibrium.

Mechanism	Welfare
AUCTION-DEPENDENT	14%
BIDDER-DEPENDENT	9%
AUCTION-DEPENDENT-COMP	29%
BIDDER-DEPENDENT-COMP	20%

Table 1: Welfare lift of different mechanisms

Observe that for mechanisms (AUCTION-DEPENDENT, BIDDER-DEPENDENT) with theoretical guarantees, both of them improve upon the baseline. Interestingly, the other two mechanisms based on auction competitiveness have even better performance in our experiments. It would be an interesting open problem to better understand the mechanisms based on auction competitiveness.

6 CONCLUSION

In this paper, we study the impact of user costs in autobidding auctions. We show that running vanilla VCG may result in poor performance. To improve the performance of VCG, we propose new variants with cost multipliers and provably show that they can guarantee constant approximations in terms of social welfare. We also validate our results via empirical experiments. For future works, as our paper examines one possibility to improve VCG by cost multipliers, it is interesting to explore other variants of VCG that may improve the performance. It is also very interesting to consider bridging the gap between theory and practice by designing practical mechanisms inspired by our theoretical results.

REFERENCES

- [1] Zoë Abrams and Michael Schwarz. 2007. Ad Auction Design and User Experience. In *Proceedings of the 3rd International Conference on Internet and Network Economics* (San Diego, CA, USA) (WINE'07). Springer-Verlag, Berlin, Heidelberg, 529–534.
- [2] Gagan Aggarwal, Ashwinkumar Badanidiyuru, and Aranyak Mehta. 2019. Autobidding with Constraints. In *International Conference on Web and Internet Economics*. Springer, 17–30.
- [3] Gagan Aggarwal, Ashish Goel, and Rajeev Motwani. 2006. Truthful auctions for pricing search keywords. In *Proceedings of the 7th ACM Conference on Electronic Commerce*. 1–7.
- [4] Moshe Babaioff, Richard Cole, Jason Hartline, Nicole Immorlica, and Brendan Lucier. 2021. Non-quasi-linear Agents in Quasi-linear Mechanisms. In *Proceedings of 12th Innovations in Theoretical Computer Science*.
- [5] Santiago Balseiro, Yuan Deng, Jieming Mao, Vahab Mirrokni, and Song Zuo. 2021. The Landscape of Auto-Bidding Auctions: Value versus Utility Maximization. In *Proceedings of the 22nd ACM Conference on Economics and Computation*.
- [6] Santiago Balseiro, Yuan Deng, Jieming Mao, Vahab Mirrokni, and Song Zuo. 2022. Optimal Mechanisms for Value Maximizers with Budget Constraints via Target Clipping. In *Proceedings of the 23rd ACM Conference on Economics and Computation*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4081457
- [7] Santiago R. Balseiro, Yuan Deng, Jieming Mao, Vahab S. Mirrokni, and Song Zuo. 2021. Robust Auction Design in the Auto-bidding World. In *Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021*. 17777–17788.
- [8] MohammadHossein Bateni, Jon Feldman, Vahab Mirrokni, and Sam Chiu-wai Wong. 2014. Multiplicative bidding in online advertising. In *Proceedings of the fifteenth ACM conference on Economics and computation*. 715–732.
- [9] Jan Panero Benway and David M. Lane. 1998. Banner Blindness: Web Searchers Often Miss "Obvious" Links.
- [10] Edward H Clarke. 1971. Multipart pricing of public goods. *Public choice* 11, 1 (1971), 17–33.
- [11] Peter Cramton, Yoav Shoham, and Richard Steinberg. 2006. *Combinatorial Auctions*. MIT Press.
- [12] Yuan Deng, Jieming Mao, Vahab Mirrokni, Hanrui Zhang, and Song Zuo. 2022. Efficiency of the First-Price Auction in the Autobidding World. *arXiv preprint arXiv:2208.10650* (2022).
- [13] Yuan Deng, Jieming Mao, Vahab Mirrokni, and Song Zuo. 2021. Towards Efficient Auctions in an Auto-bidding World. In *Proceedings of The Web Conference 2021*.
- [14] Xavier Drèze and François-Xavier Husherr. 2003. Internet advertising: Is anybody watching? *Journal of Interactive Marketing* 17, 4 (2003), 8–23. <https://doi.org/10.1002/dir.10063> arXiv:<https://doi.org/10.1002/dir.10063>
- [15] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review* 97, 1 (2007), 242–259.
- [16] Salman Fadaei and Martin Bichler. 2016. Truthfulness and approximation with value-maximizing bidders. In *International Symposium on Algorithmic Game Theory*. Springer, 235–246.
- [17] Jon Feldman, Shanmugavelayutham Muthukrishnan, Martin Pal, and Cliff Stein. 2007. Budget optimization in search-based advertising auctions. In *Proceedings of the 8th ACM conference on Electronic commerce*. 40–49.
- [18] Gagan Goel, Vahab Mirrokni, and Renato Paes Leme. 2019. Pareto efficient auctions with interest rates. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 1989–1995.
- [19] Gagan Goel, Vahab Mirrokni, and Renato Paes Leme. 2014. Clinching auctions beyond hard budget constraints. In *Proceedings of the fifteenth ACM conference on Economics and computation*. 167–184.
- [20] Negin Golrezaei, Ilan Lobel, and Renato Paes Leme. 2021. Auction design for ROI-constrained buyers. *Proceedings of The Web Conference 2021* (2021).
- [21] Theodore Groves. 1973. Incentives in teams. *Econometrica: Journal of the Econometric Society* (1973), 617–631.
- [22] Changhoan Cho and Hongsik John Cheon. 2004. Why do people avoid advertising on the Internet. *Journal of Advertising* (2004), 89–97.
- [23] Jun Li, De Liu, and Shulin Liu. 2013. Optimal Keyword Auctions for Optimal User Experiences. *Decis. Support Syst.* 56, C (dec 2013), 450–461.
- [24] Christopher Liaw, Aranyak Mehta, and Andres Perloth. 2022. Efficiency of non-truthful auctions under auto-bidding. <https://doi.org/10.48550/ARXIV.2207.03630>
- [25] Aranyak Sitanshu Mehta. 2022. Auction design in an auto-bidding setting: Randomization improves efficiency beyond VCG. In *Proceedings of the ACM Web Conference 2022*. 173–181.
- [26] Vahab Mirrokni, Renato Paes Leme, Pingzhong Tang, and Song Zuo. 2020. Non-Clairvoyant Dynamic Mechanism Design. *Econometrica* 88, 5 (2020), 1939–1963.
- [27] Evelyn Mitchell. 2022. Programmatic Advertising Explainer. <https://www.insiderintelligence.com/content/programmatic-advertiser-explainer>. Accessed: Oct 1, 2022.
- [28] Roger B Myerson. 1981. Optimal auction design. *Mathematics of Operations Research* 6, 1 (1981), 58–73.

- [29] Yuri Nesterov. 2013. *Introductory lectures on convex optimization: A basic course*. Vol. 87. Springer Science & Business Media.
- [30] Hal R Varian. 2007. Position auctions. *International Journal of Industrial Organization* 25, 6 (2007), 1163–1178.
- [31] William Vickrey. 1961. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance* 16, 1 (1961), 8–37.
- [32] Christopher A Wilkens, Ruggiero Cavallo, and Rad Niazadeh. 2017. GSP: the cinderella of mechanism design. In *Proceedings of the 26th International Conference on World Wide Web*. 25–32.
- [33] Christopher A Wilkens, Ruggiero Cavallo, Rad Niazadeh, and Samuel Taggart. 2016. Mechanism design for value maximizers. *arXiv preprint arXiv:1607.04362* (2016).