



Auction Design in an Auto-bidding Setting: Randomization Improves Efficiency Beyond VCG

Aranyak Mehta
 Google Research
 Mountain View, CA, USA
 aranyak@google.com

ABSTRACT

Auto-bidding is an area of increasing importance in the domain of online advertising. We study the problem of designing auctions in an auto-bidding setting with the goal of maximizing welfare at system equilibrium. Previous results showed that the price of anarchy (PoA) under VCG is 2 and also that this is tight even with two bidders. This raises an interesting question as to whether VCG yields the best efficiency in this setting, or whether the PoA can be improved upon. We present a prior-free randomized auction in which the PoA is approx. 1.896 for the case of two bidders, proving that one can achieve an efficiency strictly better than that under VCG in this setting. We also provide a stark impossibility result for the problem in general as the number of bidders increases – we show that no (randomized) anonymous truthful auction can have a PoA strictly better than 2 asymptotically as the number of bidders per query increases. While it was shown in previous work that one can improve on the PoA of 2 if the auction is allowed to use the bidder’s values for the queries in addition to the bidder’s bids, we note that our randomized auction is prior-free and does not use such additional information; our impossibility result also applies to auctions without additional value information.

CCS CONCEPTS

• Applied computing → Economics.

KEYWORDS

auction design, auto-bidding, equilibrium

ACM Reference Format:

Aranyak Mehta. 2022. Auction Design in an Auto-bidding Setting: Randomization Improves Efficiency Beyond VCG. In *Proceedings of the ACM Web Conference 2022 (WWW '22)*, April 25–29, 2022, Virtual Event, Lyon, France. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3485447.3512062>

1 INTRODUCTION

Auto-bidding is an area of increasing importance in the online advertising ecosystem (see, e.g., [9, 13]), with significant innovation and adoption in recent years. With auto-bidding, each advertiser states its goals and constraints to an auto-bidding agent, which then converts those into per-auction bids. In a setting where all

advertisers use auto-bidding, there one agent per advertiser, bidding optimally with respect to the other agents. We are then interested in the properties of the system equilibrium.

A prototypical auto-bidding setting is that of target-cost-per-acquisition (tCPA) in which the advertiser aims to maximize their conversions (sales after an ad-click) subject to an ROI constraint that the average cost-per-conversion is no bigger than a given target. Target-return-on-ad-spend (tROAS) generalizes tCPA to take into account the value of a conversion as well. Recent research captures the theoretical problems in this setting. In [1], the authors introduced the problem of auto-bidding under advertiser goals and constraints and presented two results: First, they presented a bidding formula to bid optimally into truthful auctions. Secondly, they proved a price of anarchy result: If all advertisers adopt such optimal auto-bidding to bid into a (per-query) VCG auction, then the welfare at equilibrium is at least a half of the value obtained an an optimal allocation. They also presented an instance of the problem (with two bidders) in which there exists an equilibrium with welfare equal to a half of the optimal, showing that their analysis is tight. Thus, they showed that the price of anarchy (PoA) under VCG is 2. Further research [6] improves on the PoA of 2 by allowing the auction access to the values of the advertisers. For example, they show that by adding a boost to each bidder’s bid equal to c times the value for the ad, and then running a second-price auction yields a PoA of $\frac{2+c}{1+c}$.

In this paper, we investigate the design of auctions to improve on the PoA in the setting introduced in [1]. As opposed to the direction taken in [6], we consider the setting where the auction does not have access to any other information (e.g., the values) besides the bid. This restriction has practical motivation: the auction may not have access to the values of each impression, e.g., the bidding system which converts the values, objectives, and constraints into bids may be a separate system, or could even be owned by an third party different from the auction platform. Also, as opposed to the Bayesian setting, we do not assume that the auction has access to the distribution of the values or targets, or indeed if there is such an underlying distribution – i.e., our auction is *prior-free*. We ask the question: *Is it possible to improve over VCG in the prior-free setting without any additional information besides the bids?*

1.1 Results

We first consider the case of two bidders, and introduce a simple prior-free randomized truthful auction $\text{RAND}(\alpha, p)$ (parameterized by a constant $\alpha \geq 1$, and probability $p \in (0, 1/2]$) and show (**Theorem 1**) that with a choice of $p = \frac{2}{5}$ and $\alpha \in [1, \frac{1+\sqrt{85}}{6}]$, the auction has a price of anarchy no worse than $\frac{2\alpha+6+\frac{2}{\alpha}}{2\alpha+1+\frac{2}{\alpha}}$. Choosing



This work is licensed under a Creative Commons Attribution International 4.0 License.

WWW '22, April 25–29, 2022, Virtual Event, Lyon, France
 © 2022 Copyright held by the owner/author(s).
 ACM ISBN 978-1-4503-9096-5/22/04.
<https://doi.org/10.1145/3485447.3512062>

$\alpha^* = \frac{1+\sqrt{85}}{6} \approx 1.703$, this shows that the PoA of $\text{RAND}(\alpha^*, \frac{2}{5})$ is approximately 1.896. We also show (**Theorem 2**) that our analysis is tight for this range of α , i.e., there is an instance in which the welfare for $\text{RAND}(\alpha, \frac{2}{5})$ at equilibrium is precisely $\frac{2\alpha+6+\frac{2}{\alpha}}{2\alpha+1+\frac{2}{\alpha}}$. We note that when $\alpha = 1$ then the auction becomes identical to the second-price auction, and the PoA bound reduces to 2, which means that our result strictly generalizes the result in [1].

Thus we have a prior-free auction with equilibrium welfare strictly better than that of VCG ([1] provides a tight example with two bidders for the PoA of VCG). This is surprising, and highlights the difference between the ROI constrained auto-bidding setting and the standard quasi-linear setting, in which VCG achieves maximum welfare. Further, similar results in the quasi-linear setting for improving over the revenue of VCG depend on either prior knowledge of the distribution of values (*Bayesian auction design*) [15], or on evaluating over instances in which the values are picked from some unknown distribution from a well-formed class like regular distributions (*prior-independent auction design*, e.g., [7]). In comparison, our result is in the prior-free setting, the auction has no other input besides the bids and is evaluated on any set of values, yet improves on the welfare of VCG in equilibrium. We note that the welfare objective in our setting is Liquid welfare which captures both the desire and ability to pay. Even though the utility function is non-standard, the questions of incentive compatibility and mechanism design remain important: As shown in [1], it is only in a truthful auction that there is a simple (uniform) optimal bidding formula. We also restrict our attention to truthful auctions.

Our result above shows a strict improvement over the PoA of VCG in the two bidder setting. One can therefore ask for an auction which can get a PoA better than 2 more generally. While it is possible to extend the analysis in Theorem 1 to a fixed number of bidders and retain a PoA strictly less than 2 (we do not include this slightly generalized analysis here), we complement our positive result with a stark impossibility result in this setting. We show (**Theorem 5**) that as the number of bidders increases (in particular the number of bidders bidding per query increases), then the PoA of any (randomized) truthful anonymous auction is asymptotically no better than 2. This result closes the problem introduced in [1] for prior-free auctions (restricted to auctions without any additional information besides the bids).

Finally, we note that the model, auction design, and analytical results are in a stylistic and theoretical setting (see related work below), and is not meant to accurately represent all aspects and complexities of real-world practical implementations.

1.2 Related Work

As mentioned above, the closest related work is that in [1] and [6]. The work in [1] introduced the problem and provided the PoA analysis for VCG in a more general setting (with multiple ROI constraints); we restrict our analysis to the prototypical tCPA setting (but it applies directly to tROAS), and also to the single slot per query setting. In the context of the results in [6], it is interesting that we can also improve over the PoA of VCG without the additional value information. There has also been recent work on understanding the optimal mechanism design in Bayesian settings [4, 12] (with

the latter describing a few different models of the setting); in comparison our work does not consider advertiser incentives (e.g., to misreport the targets) but only the system response (as in [1, 6]).

The auction introduced here ($\text{RAND}(\alpha, p)$) is reminiscent of auctions studied in a few previous results. We will define $\text{RAND}(\alpha, p)$ in Sec. 2, but informally, it allocates to the higher bidder only if its bid is an $\alpha \geq 1$ factor greater than the other bidder; otherwise it allocates to both bidders randomly, with a higher probability to the higher bidder.

In [11], the authors introduced a prior-independent auction called (ϵ, δ) -inflated second-price auction, where the parameter $\delta \geq 1$ is similar to our α parameter, but used differently: with probability ϵ it runs a second-price auction, and with the rest of the probability the highest bidder wins only if its bid is greater than the next highest bid by a factor of $\delta \geq 1$ (otherwise the item is unallocated). For this auction, it was proved that it achieves a fraction strictly better than $\frac{n-1}{n}$ of the optimal revenue (in the standard quasi-linear utility setting, where n is the number of bidders), and an improved ratio of 0.512 was provided for two bidders. This result was generalized in [2] which introduced a family of prior-independent auctions called threshold-auctions, and proved stronger results for revenue in the prior-independent setting with a focus on the two bidder setting. In fact, $\text{RAND}(\alpha, p)$ lies in the family of threshold-auctions. One may expect that other similar members of threshold-auctions family may also improve over the PoA of VCG in the auto-bidding setting – although perhaps not the specific auctions studied in [2, 11]. We note that a similar auction called the ratio-auction was described in [14] in a different context.

The definition of liquid welfare used in the auto-bidding setting [1] is based on the liquid welfare definition introduced for the budgeted setting in [8] (see also [3]).

Finally, we note that prior to ROI based auto-bidding products such as tCPA and tROAS, a well-established and well-studied model was that of budgeted optimization – how to bid under simple budget constraints. There are several lines of work in this model, e.g., [5, 10] among others.

2 PRELIMINARIES

2.1 Auto-bidding and tCPA / tROAS constraints

In the general auto-bidding problem introduced in [1], each advertiser has a goal and a set of constraints. We will restrict our attention to tCPA, in which the advertiser’s goal is to maximize the number of conversions, (sales after the click on the ad) subject to a constraint which says that the average cost of the conversions is no more than an advertiser input target T . In other words, the goal in tCPA bidding is to maximize the conversions subject to the expected spend being at most T times the expected conversions. Thus the problem for a single bidder i is:

$$\begin{aligned} & \text{Maximize} \quad \sum_{j \in Q} x_{ij} \text{ctr}_{ij} v_{ij} & (1) \\ & \text{s.t.} \quad \sum_{j \in Q} x_{ij} \text{ctr}_{ij} \text{cpc}_{ij} \leq T(i) \cdot \sum_{j \in Q} x_{ij} \text{ctr}_{ij} v_{ij} \\ & \quad \forall j \in Q : 0 \leq x_{ij} \leq 1 \end{aligned}$$

Here, Q is the set of queries, the x_{ij} are the decision variables as to whether advertiser i should buy the slot on the j th query (we assume

a single slot per query for simplicity), ctr_{ij} is the click-through-rate for an ad of advertiser i on the j th query (the probability of a click given an impression), cpc_{ij} is the cost-per-click of the ad (determined by the auction from the other advertisers' bids and hence not known in advance), and $T(i)$ is advertiser i 's input target CPA. Finally v_{ij} is the value of a click on an ad for advertiser i on the j th query. For tCPA, v_{ij} is the conversion rate (the probability of a conversion given a click). This formulation easily extends to tROAS by taking v_{ij} to be the total value after the click, e.g., value of a sale.

Bidding formula. It was shown in [1] that the bidder can bid according to a simple bidding formula based on the optimal value of the variables of the dual of LP (1). If the auction is truthful, and we have the correct values of the optimal dual variables, then this bidding formula is optimal, i.e., the bidder i buys the optimal solution to LP (1). In the tCPA setting above, given the optimal dual variable $\gamma_i \geq 0$ for the tCPA constraint, the optimal bid for bidder i takes the simple form

$$bid(i, j) = \mu_i \cdot v_{ij}, \text{ where } \mu_i := 1 + \frac{1}{\gamma_i} \quad (2)$$

$\mu_i \geq 1$ is called the bid multiplier for the bidder A .

Equilibrium. An instance of the problem is a set of queries $j \in Q$ and a set of bidders $i \in A$, each with its own (private) tCPA constant $T(i)$. Each (i, j) pair has a ctr_{ij} and value v_{ij} . We are interested in the allocation outcome when all advertisers bid using the optimal bidding formula (2). Note that each bidder's bids depend on the dual optimal variable for its LP (1) in which the cpcs are determined in the auction based on the other bidders' bids. Thus we have to study the system at equilibrium.

DEFINITION 1. Consider a (possibly randomized) auction \mathcal{A} and an instance \mathcal{I} with set A of bidders and Q of queries. Let v_{ij} be the value of bidder i for query j in the instance, and consider a set of bid-multipliers $\{\mu_i\}_{i \in A}$, so that $bid(i, j) := \mu_i v_{ij}$ is the bid of bidder i for query j . Let $\{x_{ij}\}_{i \in A, j \in Q}$ be the (probabilistic) allocation achieved with bids $bid(i, j)$ under the auction \mathcal{A} , and let p_{ij} be the per-unit price (cpc) charged to bidder i for query j . For $\Delta, \gamma \geq 0$, the bids $\{bid(i, j)\}$ are said to be in a (Δ, γ) -equilibrium if the following hold.

- Each bidder satisfies its tCPA constraint up to a multiplicative factor of γ :

$$\forall i \in A: \sum_{j \in Q} p_{ij} ctr_{ij} x_{ij} \leq (1 + \gamma) T(i) \sum_{j \in Q} v_{ij} ctr_{ij} x_{ij}$$

- No bidder can deviate from its bid-multiplier unilaterally and gain more than a Δ amount of additive value while still satisfying its tCPA constraint (up to γ). More precisely, suppose bidder i changes its bid-multiplier to μ'_i and this changes its allocation to $\{x'_{ij}\}$ and the prices to $\{p'_{ij}\}$. Then $\forall i \in A$

$$\text{Either } \sum_j v_{ij} ctr_{ij} x'_{ij} < \sum_j v_{ij} ctr_{ij} x_{ij} + \Delta$$

(does not gain more than Δ),

$$\text{or } \sum_j p'_{ij} ctr_{ij} x'_{ij} > (1 + \gamma) \sum_j v_{ij} ctr_{ij} x'_{ij}$$

(violates the tCPA constraint by more than γ).

Our positive result on PoA (Sec. 3) will use an exact definition of equilibrium (with $\Delta = \gamma = 0$), while our impossibility result (Sec. 4) holds for any $\Delta, \gamma > 0$.

Liquid Welfare. Since the spend is constrained, we need to define welfare carefully. In this setting the appropriate definition of welfare is the Liquid Welfare (LW), first defined in [8] for the budgeted allocation case, and generalized to the auto-bidding setting in [1]. Liquid welfare captures both the willingness and the ability to pay for a given allocation. For a general auto-bidding problem (with potentially multiple constraints on spend), the liquid welfare for a given allocation is defined [1] as the sum over bidders, of the minimum value of the right hand sides of the constraints of the bidder's LP, in that allocation. The tCPA LP (1) has a single constraint, and the LW becomes equivalent to the sum of tCPA-weighted conversions:

DEFINITION 2. In the tCPA setting,

$$LW(\{x_{ij}\}) = \sum_{i \in A} T(i) \sum_{j \in Q} x_{ij} ctr_{ij} v_{ij} \quad (3)$$

where A is the set of advertisers, $T(i)$ is the target CPA for advertiser i , $\{x_{ij}\}$ denotes the allocation, and v_{ij} is the conversion rate of an ad of advertiser i on query j . Thus the liquid welfare is the total number of tCPA-weighted conversions.

Define $LW(i, \{x_{ij}\})$ as the contribution of advertiser i to the liquid welfare (Eq. 3), and let $spend(i, \{x_{ij}\})$ denote the total expected spend of advertiser i in the auction. We see from LP (1) that the latter is a lower bound on the former.

$$spend(i, \{x.\}) \leq LW(i, \{x.\}) \quad (4)$$

Price of Anarchy is defined as

$$PoA = \max_{\mathcal{I}} \max_{eq \in Eq(\mathcal{I})} \frac{LW(OPT(\mathcal{I}))}{LW(eq)}$$

where \mathcal{I} denotes an instance of the problem, $Eq(\mathcal{I})$ is the set of equilibria in instance \mathcal{I} , and $OPT(\mathcal{I})$ denotes the allocation with the highest liquid welfare for that instance.

2.2 Auction RAND(α, p)

We define auction RAND(α, p) for two bidders parameterized by $\alpha \geq 1$ and $p \in (0, 1/2]$. As mentioned in Sec. 1.2, this is a member of the *threshold-auctions* family from [2], and somewhat similar to the *bid-inflation* auction from [11] and the *ratio-auction* from [14].

DEFINITION 3. Let the bids be b_1, b_2 , and without loss of generality assume that $b_1 \geq b_2$. Then the allocation function is as follows:

- If $b_1 \geq \alpha \cdot b_2$, allocate to bidder 1.
- Else, allocate to bidder 1 w.p. $1 - p$, and to bidder 2 w.p. p .

The truthful prices fall out of the allocation function in the natural manner:

$$\text{If } b_1 \geq \alpha \cdot b_2, \text{ then } cost_1 = b_2 \left(\frac{p}{\alpha} + (1 - 2p) + p \cdot \alpha \right)$$

$$\text{Else, } E[cost_1] = b_2 \left(\frac{p}{\alpha} + (1 - 2p) \right), \text{ and } E[cost_2] = \frac{p \cdot b_1}{\alpha}$$

Note that RAND(1, p) is identical to a second-price auction for any p .

3 AN IMPROVED PRICE OF ANARCHY FOR THE TWO BIDDERS CASE

In this section we prove our main positive result, that the price of anarchy of $\text{RAND}(\alpha, p)$ in the two bidder setting, is strictly better than 2 for a choice of α and p , and hence is a strict improvement over VCG.

3.1 Techniques and Intuition

The proof for the PoA of 2 for VCG in [1] proceeds at a high level as follows: For any query, let the opt-bidder be the one that is allocated the query in the optimal allocation. Now, either (case 1) the opt-bidder got allocated in the equilibrium – in which case the optimal value is obtained for that query, or (case 2) another bidder was allocated – in which case its spend on the query is at least the opt-bidder’s bid, since that is a floor in the second price auction. Due to the ROI constraints, the total spend of an advertiser is also a lower bound on its contribution to liquid welfare, and therefore the spends in queries in case 2 help bound the welfare by the optimal welfare for those queries. Putting these two cases together, we get a ratio of 2 for VCG.

The trade-offs are different in $\text{RAND}(\alpha, p)$ (for $\alpha > 1$). On the positive side, if the auction allocates to the incorrect bidder, and if its bid is greater than that of the opt-bidder by more than a factor of α , then we see a gain in the attributions in case 2 above, since the spend is strictly greater than the opt-bidder’s bid (see Def. 3). However, this also comes with potential losses in attribution: When the opt-bidder is the highest but the second-highest bid is close to its bid then we only select the opt-bidder w.p. $1 - p < 1$, and the spend is also lower than the second-price auction. A similar dynamic plays out when the opt-bidder is not the highest, but is close to the highest, in which case you allocate to the opt-bidder w.p. p , but the spend is again lower than in the second-price auction.

With such opposing cases, one would expect to gain in the aggregate only if α and p are well-tuned to the underlying values – e.g., in Bayesian setting when the value distributions are known, or perhaps even in the prior-independent setting when there are underlying regular or MHR type distributions. Surprisingly, we show that for a range of α and p , the tradeoffs *always* work out in our favor, for any fixed set of values, i.e., in a prior-free manner.

3.2 Bounding the PoA

THEOREM 1. For $p = \frac{2}{5}$, and $\alpha \in [1, \frac{1+\sqrt{85}}{6}]$, the PoA of $\text{RAND}(\alpha, p)$ (defined for two bidders) is at most $\frac{2\alpha+6+\frac{2}{\alpha}}{2\alpha+1+\frac{2}{\alpha}}$. In particular, for $\alpha^* = \frac{1+\sqrt{85}}{6} \approx 1.703$, the PoA of $\text{RAND}(\alpha^*, \frac{2}{5})$ is at most a value ≈ 1.896 .

PROOF. Consider any instance \mathcal{I} consisting of two bidders, and a set of queries Q . Bidder $i \in \{0, 1\}$ has a tCPA of $T(i)$. The value of bidder i ’s ad on query $j \in Q$ is $v(i, j)$. For ease of exposition, we will assume that all click-through rates are 1.

Fix an optimal allocation OPT , and any equilibrium allocation EQ . For $j \in Q$, let $i^*(j)$ denote the bidder to which OPT allocates j , and let $i^o(j)$ denote the other bidder. Let $\text{ALGB}(j)$ be the bidder to which the auction allocates j in EQ .

For any $j \in Q$, define $\text{OPT}(j)$ as the contribution of j to the liquid welfare in the optimal allocation, i.e. (from Def. 2),

$$\text{OPT}(j) = T(i^*(j)) \cdot v(i^*(j), j)$$

Similarly, for any subset $Q' \subseteq Q$, define $\text{OPT}(Q')$ as the contribution of Q' to the liquid welfare in OPT :

$$\text{OPT}(Q') = \sum_{j \in Q'} \text{OPT}(j)$$

Firstly, note that, for all $j \in Q$, using the optimal bidding formula (Eq.2), and the fact that the optimal bid multiplier for each bidder is at least 1, we get

$$\begin{aligned} \text{bid}(i^*(j), j) &= \mu(i^*) \cdot T(i^*(j)) \cdot v(i^*(j), j) \\ &\geq T(i^*(j)) \cdot v(i^*(j), j) \\ &= \text{OPT}(j) \end{aligned} \quad (5)$$

Query partition: We partition Q into four parts based on the relative values of the bids of $i^*(j)$ and $i^o(j)$. For a query j in each part, we will compute, using the definition (3) of $\text{RAND}(\alpha, p)$:

- the probability that j is allocated to $i^*(j)$, and
- the expected spend on query j .

- Let Q_1 be the set of queries j such that

$$\text{bid}(i^*(j), j) \leq \frac{\text{bid}(i^o(j), j)}{\alpha}$$

For each $j \in Q_1$, $\text{RAND}(\alpha, p)$ allocates j to $i^o(j)$ w.p. 1:

$$\Pr[\text{ALGB}(j) = i^*(j)] = 0$$

Further, the expected spend on j is

$$\begin{aligned} E[\text{spend}(j)] &= \left(p \cdot \alpha + (1 - 2p) + \frac{p}{\alpha}\right) \cdot \text{bid}(i^*(j), j) \\ &\geq \left(p \cdot \alpha + (1 - 2p) + \frac{p}{\alpha}\right) \cdot \text{OPT}(j) \quad (\text{using Eq.5}) \end{aligned}$$

- Let Q_2 be the set of queries j such that

$$\frac{\text{bid}(i^o(j), j)}{\alpha} \leq \text{bid}(i^*(j), j) \leq \text{bid}(i^o(j), j)$$

For each $j \in Q_2$:

$$\Pr[\text{ALGB}(j) = i^*(j)] = p$$

The total expected spend on j is the sum of the expected spends for the two bidders which are:

$$E[\text{spend}(i^o(j), j)] = p \cdot \frac{\text{bid}(i^*(j), j)}{\alpha} + (1 - 2p) \cdot \text{bid}(i^*(j), j),$$

$$E[\text{spend}(i^*(j), j)] = p \cdot \frac{\text{bid}(i^o(j), j)}{\alpha} \geq p \cdot \frac{\text{bid}(i^*(j), j)}{\alpha},$$

where the last inequality is because $\text{bid}(i^o(j), j) \geq \text{bid}(i^*(j), j)$ in this case. Thus, using Eq.5, we get

$$E[\text{spend}(j)] \geq \left(\frac{2p}{\alpha} + (1 - 2p)\right) \cdot \text{OPT}(j)$$

- Let Q_3 be the set of queries j such that

$$\text{bid}(i^o(j), j) \leq \text{bid}(i^*(j), j) \leq \alpha \cdot \text{bid}(i^o(j), j)$$

Then for each $j \in Q_3$:

$$\Pr[\text{ALGB}(j) = i^*(j)] = 1 - p$$

The total expected spend is the sum of spends of the two bidders which are:

$$E[\text{spend}(i^o(j), j)] = p \cdot \frac{\text{bid}(i^*(j), j)}{\alpha}$$

and

$$\begin{aligned} E[\text{spend}(i^*(j), j)] &= \left(\frac{p \cdot \text{bid}(i^o(j), j)}{\alpha} + (1 - 2p) \cdot \text{bid}(i^o(j), j) \right) \\ &\geq \left(\frac{p \cdot \text{bid}(i^*(j), j)}{\alpha^2} + (1 - 2p) \cdot \frac{\text{bid}(i^*(j), j)}{\alpha} \right) \end{aligned}$$

where the last inequality is because $\text{bid}(i^o(j), j) \geq \frac{\text{bid}(i^*(j), j)}{\alpha}$ in this case. Thus, using Eq.5, we get

$$E[\text{spend}(j)] \geq \left(\frac{1-p}{\alpha} + \frac{p}{\alpha^2} \right) \cdot \text{OPT}(j)$$

(4) Let Q_4 be the set of queries j such that

$$\text{bid}(i^*(j), j) \geq \alpha \cdot \text{bid}(i^o(j), j)$$

Then for each $j \in Q_2$:

$$\text{Pr}[\text{ALGB}(j) = i^*(j)] = 1$$

while the expected spend can be 0 (if $\text{bid}(i^o(j), j) = 0$).

Bounding the welfare: Define the following constants:

$$\begin{aligned} m_1 &= 0 & s_1 &= p \cdot \alpha + (1 - 2p) + \frac{p}{\alpha} \\ m_2 &= p & s_2 &= \frac{2p}{\alpha} + (1 - 2p) \\ m_3 &= 1 - p & s_3 &= \frac{1-p}{\alpha} + \frac{p}{\alpha^2} \\ m_4 &= 1 & s_4 &= 0 \end{aligned}$$

From the calculations above, we have:

$$\forall k = 1..4 : \forall j \in Q_k, \text{Pr}[\text{ALGB}(j) = i^*(j)] = m_k \quad (6)$$

$$\forall k = 1..4 : \forall j \in Q_k, \text{spend}(j) \geq s_k \cdot \text{OPT}(j) \quad (7)$$

Now, for each case $k = 1 \dots 4$, we get the following two bounds on the welfare obtained in equilibrium.

Firstly, considering the probability that j is allocated to the optimal bidder i^* in the equilibrium, we get, using Eq. 6:

$$\begin{aligned} E[\text{LW}(Eq)] &\geq \sum_{k=1}^4 \sum_{j \in Q_k} \text{Pr}[\text{ALGB}(j) = i^*(j)] \cdot \text{OPT}(j) \\ &\geq \sum_{k=1}^4 m_k \cdot \text{OPT}(Q_k) \end{aligned} \quad (8)$$

Secondly,

$$\begin{aligned} E[\text{LW}(Eq)] &= \sum_{i \in A} E[\text{LW}(i, Eq)] \geq \sum_{i \in A} \text{spend}(i, Eq) \\ &= \sum_{k=1}^4 \sum_{j \in Q_k} E[\text{spend}(j)] \geq \sum_{k=1}^4 \sum_{j \in Q_k} s_k \cdot \text{OPT}(j) \\ &= \sum_{k=1}^4 s_k \text{OPT}(Q_k) \end{aligned} \quad (9)$$

where the first inequality follows from using Eq. 4 and the second inequality follows from Eq. 7.

Bounding the PoA via a factor revealing LP. Define the variables $x_k := \text{OPT}(Q_k)$, for $k = 1 \dots 4$. Normalizing $\text{OPT}(Q) = 1$, we get

$$\sum_k x_k = 1 \quad (10)$$

We can now bound the overall price of anarchy by minimizing $E[\text{LW}(Q)]$ over the constraints (8), (9), (10). Let z denote $E[\text{LW}(Q)]$, then we have the following linear program and its dual LP (with variables γ, β, δ):

$$\begin{array}{l|l} \text{Minimize } z & \text{Maximize } \delta \\ \text{s.t. } z - \sum_k m_k x_k \geq 0 & \text{s.t. } \forall k : -m_k \gamma - s_k \beta + \delta \leq 0 \\ z - \sum_k s_k x_k \geq 0 & \gamma + \beta \leq 1 \\ \sum_k x_k \geq 1 & \gamma, \beta, \delta \geq 0 \\ \forall k, x_k \geq 0 & \end{array}$$

Consider the following dual solution:

$$\delta = \frac{1}{\left(1 + \frac{1}{s_1}\right)}, \quad \gamma = \delta, \quad \beta = \frac{\delta}{s_1} \quad (11)$$

We show that this is a feasible solution for the choice of $^1 p = \frac{2}{5}$ and a range of α . For $k = 1$ and $k = 4$, the constraints in the dual LP are satisfied and in fact tight, since $m_1 = 0, m_4 = 1$, and $s_4 = 0$. Note also that

$$\gamma + \beta = \delta \left(1 + \frac{1}{s_1}\right) = 1$$

so the last constraint is satisfied and tight as well.

For $k = 2$, we need: $-m_2 \gamma - s_2 \beta + \delta \leq 0$. Plugging in the choice of $p = \frac{2}{5}$ and the values of $m_2, s_2, \gamma, \beta, \delta$, this reduces to:

$$3\alpha^2 - \alpha - 7 \leq 0,$$

which is satisfied for $\alpha \in [1, \frac{1+\sqrt{85}}{6}]$. Note that $\frac{1+\sqrt{85}}{6} \approx 1.703$.

For $k = 3$, we need: $-m_3 \gamma - s_3 \beta + \delta \leq 0$, which reduces to:

$$4\alpha^3 + 2\alpha^2 - 11\alpha - 10 \leq 0$$

This is satisfied for $\alpha \in [1, r]$, with $r \approx 1.8$.

Thus we see that the solution (11) is a feasible solution for the dual LP, for $\alpha \in [1, \frac{1+\sqrt{85}}{6}]$.

The value of the dual objective is $\delta = \frac{1}{\left(1 + \frac{1}{s_1}\right)} = \frac{2\alpha + 1 + \frac{2}{\alpha}}{2\alpha + 6 + \frac{2}{\alpha}}$. This

means that the minimum value of the primal LP, which is a lower bound on the value obtained in the auction equilibrium, is no less than this value. Since OPT is normalized to 1, this completes the proof the theorem. For $\alpha = \frac{1+\sqrt{85}}{6}$, the value of the dual objective is approx. 0.527, i.e., a price of anarchy of approx. 1.896.

For completeness, we note that there is a primal solution with the same value as this dual solution: $x_4 = z = \delta$, $x_1 = \frac{\delta}{s_1}$, and $x_2 = x_3 = 0$. \square

¹This choice of $p = \frac{2}{5}$ is found by a search and is almost optimal. We note that an arguably more natural choice of $p = \frac{1}{3}$ (which gives an even increase in allocation probability as the bid increases) is only slightly worse, giving a PoA of approx. 1.91.

3.3 A tight example

One may notice that the bound obtained above does not fully use the properties of the equilibrium – it does use the fact that tCPA constraints are satisfied (in Eq. 9), but does not use the property that no bidder has an incentive to change its bid-multiplier. So we may expect the bound to be very loose. Surprisingly, we show that there exists an instance in which the bound is tight, and the values of $x_k := OPT(Q_k)$ in an equilibrium in the instance are precisely those corresponding to an optimal LP solution in the proof above.




Figure 1: A tight example.

THEOREM 2. *The analysis in Theorem 1 is tight, i.e., there exists an instance with an equilibrium under $RAND(\alpha, p)$, in which the ratio of the optimal welfare to the welfare in the equilibrium is exactly $1 + \frac{1}{s_1}$, where $s_1 = p \cdot \alpha + (1 - 2p) + \frac{p}{\alpha}$. In particular, for $p = \frac{2}{5}$, the ratio is exactly $\frac{2\alpha+6+\frac{2}{\alpha}}{2\alpha+1+\frac{2}{\alpha}}$.*

PROOF. Consider the instance in Fig. 1. This instance is similar to the one used in [1] except for the value on edge $(B, 2)$ (reduced from 1 to $1/s_1$). We have two advertisers a, b and two queries 1, 2. The values for advertiser a are $v_1^a = 1$ and $v_2^a = \epsilon$ and for advertiser b are $v_1^b = 0$ and $v_2^b = \frac{1}{s_1}$. The click-through rates for all ad-query pairs are set to 1. For both the advertisers, we have a tCPA constraint with both $tCPA = 1$, so the constraint is that the spend should be at most the value.

Consider the bid multipliers:

$$\mu_A = \frac{\alpha}{\epsilon s_1} + 1, \mu_B = 1 \quad (12)$$

We will first show that this is an equilibrium:

- For item 2, A 's bid is $\frac{\alpha}{s_1} + \epsilon$, while B 's bid is $\frac{1}{s_1}$, so A wins item 2 with probability 1, and pays s_1 times B 's bid, equal to 1. A also gets item 1 for free since B 's bid is 0. Thus A 's spend is 1, which is less than its value, $1 + \epsilon$, for the two items obtained. Thus A 's tCPA constraint is satisfied, and since it gets both the items w.p. 1, it has no incentive to deviate from its bid of μ_A .
- B gets nothing, but it can not bid up (increase μ_B) to get item 2 with any probability and still satisfy its tCPA constraint:
 - If it bids very high to obtain item 2 with probability 1, then it has to pay s_1 times A 's bid which becomes $\alpha + \epsilon s_1$ which is more than the value of $\frac{1}{s_1}$ since $\alpha \geq 1$ and $s_1 \geq 1$.

- If it bids up to win the item with probability $1 - p$, then its spend is $bid(A) \left(\frac{p}{\alpha} + (1 - 2p) \right) > \frac{\alpha}{s_1} \left(\frac{p}{\alpha} + (1 - 2p) \right)$. But this is at least the expected value for B which is $\frac{1-p}{s_1}$ (since $\alpha \geq 1$).
 - If B bids up to win item 2 with probability p , then its spend is $p \cdot \frac{bid(A)}{\alpha} > \frac{p}{s_1}$, which is the expected value for B .
- Thus B does not want to defect from its bid of $\mu_B = 1$.

Thus the bid multipliers (12) are in equilibrium. In this equilibrium, the total value obtained is $1 + \epsilon$, while the optimal allocation (item 1 to A , item 2 to B) gets a value of $1 + \frac{1}{s_1}$, thus proving the theorem, as $\epsilon \rightarrow 0$. \square

4 AN IMPOSSIBILITY RESULT WITH MANY BIDDERS

In this section we will show an impossibility result when we have a large number of bidders bidding per query.

4.1 Intuition

Suppose we try to extend the proof in Sec. 3 to the case of many bidders per query (for some appropriate generalization of $RAND$ to more than two bidders). The problem that arises is that the contributions to the welfare from queries in classes Q_2 and Q_3 can become very small. To see this, note that a query in type Q_2 or Q_3 has the opt-bidder among the highest bidders within some margin (α). Now, the probability of the opt-bidder winning can be very low if there are a large number of bidders within the margin, since the anonymous auction has to fairly randomize between all of them. Thus, m_2 and m_3 can be close to 0. Now the bad scenario that could occur is that there is an instance and an equilibrium in which there are no queries in Q_1 ($x_1 = 0$), and a mix of queries, half in Q_4 (auction picked the opt-bidder) and half in Q_2 or Q_3 . If such a scenario arises in equilibrium then the welfare in any such generalization of $RAND$ would be no better than a half of the optimal welfare.

In fact, we show that such a scenario can be constructed in equilibrium for *any* randomized auction. We do so by creating an instance for any given randomized auction using the intuition for $RAND$ above. Queries will either be in Q_4 (i.e., the opt-bidder wins without any competition), or will be such that there are a large number of bidders with almost equal bids, and the opt-bidder is the lowest among them. In that case the opt-bidder has a very low chance of winning, but the overall spend among the other bidders is also not high enough to increase the welfare. The technical difficulty here is that the space of randomized auctions is very large, so we have to construct an instance that relies on very basic auction properties so that every claim holds for all auctions.

4.2 Anonymity and Max-Threshold

We start by defining two auction properties.

DEFINITION 4. *A randomized auction is said to be anonymous if its allocation does not depend on the identity of the bidders but only on the relative bid values.*

Our proof of the impossibility result holds only for anonymous auctions. Anonymity is often standard when deriving general results; note that anonymity precludes some known auctions such as those with personalized reserves (but these are not very applicable in the prior-free setting anyway). We will use the following (proof in Appendix A.2):

LEMMA 3. *In any anonymous auction, if there are k bidders, then the lowest bidder wins the item with probability at most $\frac{1}{k}$.*

Next, since we want to prove the result for all (anonymous) randomized auctions, we need to identify a technical property of any auction that we will use in our construction. This is a technical requirement used to consider cases of auctions that do not sell the item fully to a bidder even if it is the only bidder and bids as high as needed; intuitively such auctions can only have worse efficiency.

DEFINITION 5. *Consider the setting when there is exactly one bidder for the item and puts a bid of b . As b increases, the probability $P(b) = \Pr[\text{win with bid } b]$ of winning the item must increase in a truthful auction. Define the max-probability*

$$\pi^* = \lim_{b \rightarrow \infty} P(b)$$

be (the limit of) the highest probability with which the bidder can win. For simplicity of notation, we will assume that the limit is reached². Define the max-threshold as the lowest bid that achieves this highest probability.

$$M^* = \min\{b : P(b) = \pi^*\}$$

We will use the following (proof in Appendix A.2).

LEMMA 4. *For any (randomized) truthful auction with a max-probability of π^* and a max-threshold of M^* , the cost when there is a single bidder bidding $b \geq M^*$ is at most $\pi^* M^*$.*

4.3 Proving the bound

THEOREM 5. *For any randomized truthful anonymous auction, for any $\Delta, \gamma > 0$, there exists an instance with sufficiently many bidders n for which there exists a (Δ, γ) -equilibrium in which the total value is asymptotically $\frac{1}{2}$ of the optimal value for the instance, as $n \rightarrow \infty$.*

PROOF. Let the given the randomized auction \mathcal{A} have max-probability π^* and max-threshold M^* . For simplicity of notation going forward, if $M > 1$ then we re-scale all values in the following construction wlog so that $M = 1$. Note that we can assume $\pi^* \geq \frac{1}{2}$, otherwise we would be done, by taking an instance with one bidder and one item.

Constructing the instance. Consider the following instance. There are $2k$ bidders in total: k bidders A_0, \dots, A_{k-1} , and k bidders B_0, \dots, B_{k-1} . All bidders have a tCPA of 1, i.e., they are constrained to have their total cost to be at most the total value achieved (up to a factor of $1 + \gamma$). Correspondingly, there are $2k$ queries: k queries P_0, \dots, P_{k-1} and k queries Q_0, \dots, Q_{k-1} . We pick $a \geq 1, V \geq 1$ to be fixed later.

Values for P_i : For each $i \in [0, k-1]$, only A_i has a non-zero value for P_i , with $v(A_i, P_i) = \frac{a \cdot V}{\pi^*}$.

²For example, the standard assumption of consumer sovereignty implies that $\pi^* = 1$ and is achieved at a high enough bid.

Values for Q_i : The values of the bidders for the Q_i are in the following pattern. Firstly, for each $i \in [0, k-1]$, $v(B_i, Q_i) = V$, and $v(B_j, Q_i) = 0, \forall j \neq i$. Next, we define the k -long tuple

$$\tau^0 = (a \cdot V + \rho, a \cdot V + 2\rho, \dots, a \cdot V + k\rho)$$

Here $\rho > 0$ is some very small constant used only for tie-breaking. For $i \in [1, k-1]$, define τ^i to be the i^{th} rotation of τ^0 , i.e., $\tau^i = (a \cdot V + (i+1)\rho, a \cdot V + (i+2)\rho, \dots, a \cdot V + k\rho, a \cdot V + \rho, a \cdot V + 2\rho, \dots, a \cdot V + i\rho)$.

Now, for each $i \in [0, k-1]$, the values of the bidders $\{A_j\}$ for Q_i are determined as follows. Let $\epsilon > 0$ be a constant to be fixed later. The tuple of values

$$(v(A_0, Q_i), v(A_1, Q_i), \dots, v(A_{k-1}, Q_i))$$

is set to be equal to $\epsilon \cdot \tau^i$, where the latter multiplication is coordinate-wise (so, e.g., $v(A_0, Q_i) = \epsilon(a \cdot V + (i+1)\rho)$). Fig. 2 shows the instance for $k = 2$, i.e., with four bidders.




Figure 2: An instance with $k = 2$, i.e., 4 bidders and 4 queries. Here, $V > 1, a > 1, \epsilon > 0$. We take V to be large, $a \rightarrow 1$, and $\epsilon \rightarrow 0$. ρ can be considered a formal variable for tie-breaking.

Bid multipliers: For this instance, consider the following bid-multipliers: $\mu(A_i) = \frac{1}{\epsilon}, \forall i$, and $\mu(B_i) = 1, \forall i$.

Bids for P_i : Thus, for each i , $\text{bid}(A_i, P_i) = \frac{a \cdot V}{\epsilon \cdot \pi^*}$ and $\text{bid}(X, P_i) = 0, \forall X \neq A_i$.

Bids for Q_i : For each i , query Q_i gets the following bids: Firstly $\text{bid}(B_i, Q_i) = V$, and $\text{bid}(B_j, Q_i) = 0, \forall j \neq i$. Further, the tuple

$$(\text{bid}(A_0, Q_i), \text{bid}(A_1, Q_i), \dots, \text{bid}(A_{k-1}, Q_i)) = \tau^i.$$

The goal of the rest of the proof: To prove the theorem, we need to prove that for this instance and these bid-multipliers:

- (1) The welfare achieved in the auction with these bids is close to a half of the optimal possible welfare.
- (2) The bids form an equilibrium, i.e., for each bidder. (Def. 1):
 - (a) Their tCPA constraint is satisfied at these bids (up to a factor of $1 + \gamma$).
 - (b) They can not change their bid-multiplier and gain more than Δ additive value while staying within $(1 + \gamma)$ of the tCPA constraint.

Proving the result for any randomized auction. The challenge is to prove the above for any randomized auction. To do this, firstly, note that our instance is almost universal, i.e., it only depends on the given auction via the max-probability π^* and up to scaling of all numbers in the instance by the max-threshold M^* . Secondly, we rely only on basic properties that all anonymous randomized truthful auctions satisfy: (a) Lemma 4 which tackles the dependency

of π^* and M^* , (b) consequences of anonymity such as Lemma 3, and (c) simple properties following from truthfulness.

(1) Welfare: With these bids, the auction for P_i has only one bidder A_i and its bid is at least the max-threshold (which is at most 1), so from Def. 5 P_i is allocated to A_i w.p. π^* . Query Q_i gets allocated to the “correct” bidder B_i with probability at most $\frac{1}{k+1}$ since it is the lowest among $k+1$ bidders (Lemma 3). Thus the expected value obtained for the P_j is $a \cdot V$ each, and for each Q_j is at most $\frac{1}{k+1}V + \epsilon(a \cdot V + k\rho)$ (the latter term is an upper bound, being the highest possible value among the A_j for Q_i).

Thus the total value achieved in the auctions under these bids is (ignoring terms with ρ)³

$$\begin{aligned} \text{LW(EQ)} &\leq \pi^* \cdot \frac{a \cdot V}{\pi^*} \cdot k + \frac{k}{k+1}V + \epsilon \cdot a \cdot V \cdot k \\ &= \left(a + \frac{1}{k+1} + \epsilon \cdot a \right) \cdot V \cdot k. \end{aligned} \quad (13)$$

The optimal allocation allocates each P_i to A_i , and Q_i to B_i , giving a value of

$$\text{OPT} = \left(\frac{a}{\pi^*} + 1 \right) \cdot V \cdot k. \quad (14)$$

(2) Proving equilibrium: It remains to prove that this set of bids constitutes a (Δ, γ) -equilibrium.

(2a) Bidders A_i 's constraint: From Lemma 4: A_i is allocated P_i w.p. π^* and the cost for P_i is at most π^* since we re-scaled to set the max-threshold $M^* \leq 1$. Also, the per-unit price for A_i for any Q_j is at most its per-unit bid for Q_j . So we get (dropping the terms involving ρ):

$$\begin{aligned} \text{value}(A_i) &= v(A_i, P_i) \cdot P[P_i \text{ allocated to } A_i] \\ &+ \sum_{j=0}^{k-1} v(A_i, Q_j) \cdot P[Q_j \text{ allocated to } A_i] \\ &\geq a \cdot V + \sum_{j=0}^{k-1} \epsilon \cdot a \cdot V \cdot P[Q_j \text{ allocated to } A_i] \end{aligned} \quad (15)$$

$$\begin{aligned} \text{cost}(A_i) &\leq \pi^* + \sum_{j=0}^{k-1} \text{bid}(A_i, Q_j) \cdot P[Q_j \text{ allocated to } A_i] \\ &\leq 1 + \sum_{j=0}^{k-1} a \cdot V \cdot P[Q_j \text{ allocated to } A_i] \end{aligned} \quad (16)$$

To get a handle on the probabilities in Eqs. 15 and 16, we consider a **special setting** in which \mathcal{A} is given $k+1$ bidders with bids equal to the tuple

$$\beta = (V, a \cdot V + \rho, a \cdot V + 2\rho, \dots, a \cdot V + k\rho)$$

Let $(\chi_0, \chi_1, \dots, \chi_k)$ be the probabilities with which \mathcal{A} allocates the item to bidders $0, \dots, k$ in this setting. Clearly, $\sum_{s=0}^k \chi_s \leq 1$.

Given this property of \mathcal{A} , we can bound the total probability with which the Q_j are allocated in our instance. Note that each Q_j has $k+1$ bidders with a set of bids equal to the tuple β defined above. Hence each Q_j will be allocated according to the allocation

³The constant ρ can be taken to be very small, in fact it can be considered to be a formal variable, introduced purely for tie-breaking among the A_i . Thus, we will not include the ρ in the quantification of welfare and incentives.

$\{\chi\}$ defined above. Now due to the rotational symmetry across the Q_j , A_i is the $t_{ij} := ((i+j) \bmod k+1)^{\text{th}}$ lowest bidder for Q_j among the $\{A_s\}$ (while B_j is the lowest bidder for Q_j). Therefore, as a consequence of anonymity, A_i gets allocated Q_j with probability $\chi_{t_{ij}}$. As j ranges over $[0, k-1]$, t_{ij} ranges over $[1, k]$. Thus, we get:

$$\sum_{j=0}^{k-1} P[Q_j \text{ allocated to } A_i] = \sum_{s=1}^k \chi_s \leq 1 \quad (17)$$

That is, for all i , A_i is allocated an at most a unit expected total amount of any of the Q_j .

Thus, using Eq. 17 in Eqs. 15 and 16, we get:

$$\text{cost}(A_i) \leq \left(1 + \frac{1}{a \cdot V} \right) \text{value}(A_i) \quad (18)$$

(2b) Bounding the gain from defection for bidder A_i : Reducing the bid multiplier can not help since that can only reduce the value obtained. Now even if A_i increases its bid to infinity, the maximum value it can achieve is (ignoring the terms with ρ):

$$\pi^* \cdot v(A_i, P_i) + \sum_{j=0}^{k-1} v(A_i, Q_j) \leq a \cdot V + k \cdot \epsilon \cdot a \cdot V$$

Thus, using Eq. 15, the highest that A_i can gain in value by deviating from $\mu(A_i)$ (even ignoring its tCPA constraint) is

$$\text{Gain}(A_i) \leq k \cdot \epsilon \cdot a \cdot V \quad (19)$$

We provide the rest of the argument to prove the bidders B_i are also satisfied, and hence the bid-multipliers are at equilibrium, in Appendix A.1.

Setting the parameters. Now we can set our parameters to satisfy all the constraints for a (Δ, γ) -equilibrium with a welfare close to a half of the optimal welfare. From Eq. 18, we see that by taking $V > \frac{1}{\gamma}$, and noting that $a \geq 1$, A_i 's tCPA constraint is satisfied up to a $(1+\gamma)$ factor (the tCPA constraint for each B_i is strictly satisfied). From Eq. 20 (Appendix A.1), we see that by taking

$$k+1 > \frac{aV}{(a-1-\gamma)\Delta} > \frac{a}{(a-1-\gamma)\Delta\gamma}$$

we get that $\text{Gain}(B_i)$ is no more than Δ . From Eq. 19, we see that by taking $\epsilon \rightarrow 0$, small enough so that $k\epsilon aV < \Delta$, we get that $\text{Gain}(A_i)$ is no more than Δ . Finally, by taking $a \rightarrow 1$ (recall $a \geq 1$), we get from Eq. 13 and 14 that the value obtained via the bidding in any randomized auction is at most

$$\frac{a}{\frac{a}{\pi^*} + 1} + \frac{1}{(k+1)\left(\frac{a}{\pi^*} + 1\right)} + \frac{\epsilon a}{\frac{a}{\pi^*} + 1} \rightarrow \frac{\pi^*}{1 + \pi^*} \leq \frac{1}{2},$$

(as $a \rightarrow 1$, $\epsilon \rightarrow 0$ and $k \rightarrow \infty$), thus proving the theorem. \square

REFERENCES

- [1] Gagan Aggarwal, Ashwinkumar Badanidiyuru, and Aranyak Mehta. 2019. Autobidding with Constraints. In *Web and Internet Economics - 15th International Conference, WINE 2019, Proceedings (Lecture Notes in Computer Science, Vol. 11920)*. Springer, 17–30.
- [2] Amine Allouah and Omar Besbes. 2020. Prior-Independent Optimal Auctions. *Manag. Sci.* 66, 10 (2020), 4417–4432.
- [3] Yossi Azar, Michal Feldman, Nick Gravin, and Alan Roytman. 2017. Liquid Price of Anarchy. In *Algorithmic Game Theory - 10th International Symposium, SAGT 2017, L'Aquila, Italy, September 12–14, 2017, Proceedings*. 3–15.
- [4] Santiago R. Balseiro, Yuan Deng, Jieming Mao, Vahab S. Mirrokni, and Song Zuo. 2021. The Landscape of Auto-bidding Auctions: Value versus Utility Maximization. In *EC '21: The 22nd ACM Conference on Economics and Computation, 2021*. ACM, 132–133.

- [5] Santiago R. Balseiro and Yonatan Gur. 2019. Learning in Repeated Auctions with Budgets: Regret Minimization and Equilibrium. *Management Science* 65, 9 (September 2019), 3952–3968.
- [6] 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 (WWW '21)*. Association for Computing Machinery, 3965–3973.
- [7] Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. 2015. Revenue maximization with a single sample. *Games and Economic Behavior* 91 (2015), 318–333.
- [8] Shahar Dobzinski and Renato Paes Leme. 2014. Efficiency Guarantees in Auctions with Budgets. In *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part I*. 392–404.
- [9] Facebook 2022. Auto-bidding products support page. <https://www.facebook.com/business/help/1619591734742116>. Accessed: 2022-02-09.
- [10] 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.
- [11] Hu Fu, Nicole Immorlica, Brendan Lucier, and Philipp Strack. 2015. Randomization Beats Second Price as a Prior-Independent Auction. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, 2015*. ACM, 323.
- [12] Negin Golrezaei, Ilan Lobel, and Renato Paes Leme. 2021. Auction Design for ROI-Constrained Buyers. In *Proceedings of the Web Conference 2021 (WWW '21)*. 3941–3952.
- [13] Google 2022. Auto-bidding products support page. <https://support.google.com/google-ads/answer/2979071>. Accessed: 2022-02-09.
- [14] Jason D. Hartline and Tim Roughgarden. 2014. Optimal Platform Design. arXiv:1412.8518 [cs.GT]
- [15] Roger Myerson. 1981. Optimal Auction Design. *Mathematics of Operations Research* 6, 1 (1981), 58–73.

A APPENDIX

A.1 Missing proof of equilibrium in Theorem 5

(2a) Bidders B_i 's constraint: Firstly, for each i , B_i only bids on one query (Q_i), and its bid is equal to its value, since $\mu(B_i) = 1$. Since the per-unit cost is at most the bid, B_i 's tCPA constraint is satisfied no matter the probability with which Q_i is allocated to B_i .

(2b) Bounding the gain from defection for bidder B_i : To understand the incentive to deviate from $\mu(B_i) = 1$, note first that lowering its bid can only decrease its value obtained. Since B_i is the lowest bidder for Q_i , and the other bids are $a + j\rho$ for $j \in [1, k]$, by Lemma 3 we have that B_i is allocated Q_i with probability at most $\frac{1}{k+1}$. Now suppose B_i raises its bid-multiplier so that the bid for Q_i increases to some value $bid'(B_i, Q_i)$ and it wins Q_i with some increased probability p .

Since the auction is truthful, the total cost to B_i at a bid of $bid'(B_i, Q_i)$ can be computed using Myerson's lemma for truthful pricing (the area above the allocation curve). Consider the function \mathcal{F} which denotes the probability of allocating to B_i according to \mathcal{A} as B_i changes its bid, given that the other bidders $\{A_j\}$ bids are fixed at $\{aV + j\rho\}_{j \in [1, k]}$. Consider the curve at the bid $bid'(B_i, Q_i)$ giving a probability of p . The value to B_i is

$$value'(B_i) = p \cdot V$$

The cost to B_i is

$$\begin{aligned} cost'(B_i) &= \int_{b=0}^{bid'(B_i, Q_i)} (p - \mathcal{F}(b)) db \\ &\geq (p - \mathcal{F}(a \cdot V)) \cdot a \cdot V \geq \left(p - \frac{1}{k+1}\right) \cdot a \cdot V \end{aligned}$$

Here, first inequality follows because the allocation curve is non-decreasing, and $bid'(B_i, Q_i) \geq a \cdot V$, since the lowest bid among the other bidders is $a \cdot V + \rho$. The second inequality follows from




Figure 3: Bounding the cost for B_i as it raises its bid by the area of the rectangle in the figure.

Lemma 3, because at a bid of $a \cdot V$, B_i is still the lowest among $k + 1$ bidders. Fig. 3 gives a pictorial representation of this bound on the cost.

Since the tCPA constraint needs to be satisfied at the new bid up to a factor of $(1 + \gamma)$, we have (dropping the terms with ρ):

$$\begin{aligned} cost'(B_i) &\leq (1 + \gamma) \cdot value'(B_i) \\ &\Rightarrow \left(p - \frac{1}{k+1}\right) \cdot a \cdot V \leq (1 + \gamma) \cdot p \cdot V \Rightarrow p \leq \frac{a}{(a-1-\gamma)(k+1)} \end{aligned}$$

Thus we have a bound on the maximum possible gain that B_i can get by changing its bid while still respecting its tCPA constraint:

$$Gain(B_i) \leq \frac{a \cdot V}{(a-1-\gamma)(k+1)} \quad (20)$$

A.2 Proofs of lemmas from Sec. 4.2

PROOF. (of Lemma 3) Anonymity in a truthful auction (with a monotonic allocation function) implies fairness – a higher bidder can not get a lower probability of allocation. Fairness implies the property required. \square

PROOF. (of Lemma 4) This follows directly from Myerson's truthful pricing formula, where one can see that if the max-threshold bid is M^* and the max-probability is π^* , then the area above the allocation curve (at a bid at least M^*) is at most $\pi^* \cdot M^*$. The bound is tight when the auction is a second price auction with a reserve price equal to M^* (and $\pi^* = 1$). \square