

On the Efficiency of Equilibria in Generalized Second Price Auctions

Ioannis Caragiannis Christos Kaklamanis
Panagiotis Kanellopoulos Maria Kyropoulou
Department of Computer Engineering and Informatics
University of Patras and RACTI, Greece
{caragian,kakl,kanellop,kyropoul}@ceid.upatras.gr

ABSTRACT

In sponsored search auctions, advertisers compete for a number of available advertisement slots of different quality. The auctioneer decides the allocation of advertisers to slots using bids provided by them. Since the advertisers may act strategically and submit their bids in order to maximize their individual objectives, such an auction naturally defines a strategic game among the advertisers. In order to quantify the efficiency of outcomes in generalized second price auctions, we study the corresponding games and present new bounds on their price of anarchy, improving the recent results of Paes Leme and Tardos [16] and Lucier and Paes Leme [13]. For the full information setting, we prove a surprisingly low upper bound of 1.282 on the price of anarchy over pure Nash equilibria. Given the existing lower bounds, this bound denotes that the number of advertisers has almost no impact on the price of anarchy. The proof exploits the equilibrium conditions developed in [16] and follows by a detailed reasoning about the structure of equilibria and a novel relation of the price of anarchy to the objective value of a compact mathematical program. For more general equilibrium classes (i.e., mixed Nash, correlated, and coarse correlated equilibria), we present an upper bound of 2.310 on the price of anarchy. We also consider the setting where advertisers have incomplete information about their competitors and prove a price of anarchy upper bound of 3.037 over Bayes-Nash equilibria. In order to obtain the last two bounds, we adapt techniques of Lucier and Paes Leme [13] and significantly extend them with new arguments.

Categories and Subject Descriptors

F.2 [Theory of Computation]: Analysis of Algorithms and Problem Complexity—*Miscellaneous*; J.4 [Computer Applications]: Social and Behavioral Sciences—*Economics*

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. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC'11, June 5–9, 2011, San Jose, California, USA.

Copyright 2011 ACM 978-1-4503-0261-6/11/06 ...\$10.00.

General Terms

Algorithms, Economics, Theory

Keywords

Auctions, equilibria, generalized second price, price of anarchy

1. INTRODUCTION

Sponsored search auctions [11] have become extremely popular during the last decade as the main tool used by search engines and other information services in order to create income. A number of advertisement slots are available which can be thought of as ranked according to their significance (e.g., according to the number of clicks by viewers an advertisement assigned to a slot is expected to have). A sponsored search auction aims to allocate advertisement slots to advertisers. Advertisers cast their bids for the available slots and the auctioneer uses the bids in order to compute the allocation to slots and the fee each advertiser should pay for this service. The particular rule used in order to compute both the allocation and the payments defines a distinct auction.

We consider generalized second price auctions whose variations are widely used by leaders in the sponsored search industry such as Google and Yahoo! In generalized second price auctions, the advertisers are assigned the slots in decreasing order of their bids (i.e., the advertiser with the highest bid is assigned to the most significant slot, and so on) and each of them is required to pay the next highest bid (i.e., the advertiser assigned to the most significant slot will pay per click an amount equal to the bid of the advertiser assigned to the second significant slot, and so on).

Traditionally, Auction Theory (see [9] for an introduction) has given a central role to the requirement that an auction should impose truthful behavior by the potentially strategically acting participants. This requirement has several implications to the maximization of the social welfare, i.e., the quantity that includes both the auctioneer's revenue and the profit of the participants. Even though generalized second price auctions generalize the famous truthful Vickrey auction [21], they are known neither to be truthful [1] nor to guarantee social welfare maximization [11, 20] and this seems to come in sharp contrast to their recent success. In an attempt to provide a partial justification of this success, following previous work, we consider the following question:

how much can the strategic behavior of the advertisers affect the social welfare?

We address this question by considering the natural strategic games among the advertisers that generalized second price auctions define. Each advertiser (henceforth also called bidder) has her own valuation for each click and her utility from an allocation depends on the total valuation from the clicks on the slot she is assigned to minus her payment to the auctioneer. Hence, acting strategically means that each bidder aims to maximize her utility given the strategies (bids) of the other bidders. Such a behavior naturally leads to an equilibrium, i.e., a set of strategies from which no bidder has an incentive to deviate. We consider both the full information and the more realistic incomplete information (or Bayesian) setting. The corresponding equilibrium concepts of interest are the pure Nash, mixed Nash, correlated, and coarse correlated equilibria in the former and Bayes-Nash equilibria in the latter setting. In both cases, we make the reasonable assumption that bidders always cast conservative bids that do not exceed their valuations. We provide new social welfare guarantees for generalized second price auctions that are expressed as bounds on the price of anarchy of the corresponding games over equilibria of particular classes, improving previous results in the literature. In a nutshell, our results indicate that, despite the strategic behavior of the advertisers, the social welfare is always high.

Related work. The game-theoretic model we adopt in the current paper was proposed by [5] and [20] and was further used in the sequence of papers [6, 10, 13, 16, 19] (see also the surveys [11, 15]). Edelman et al. [5] and Varian [20] prove that Nash equilibria with optimal social welfare always exist for generalized second price auction games in the full information setting. In contrast, this is not the case for games in the Bayesian setting as proved in [6]. Lahaie [10] provides bounds on the social welfare of equilibria under strong assumptions for the click-through rates of the slots. Thompson and Leyton-Brown [19] study the efficiency of equilibria through experimentation.

Our work is closely related to, and improves the recent results of Paes Leme and Tardos [16] and Lucier and Paes Leme [13]. They consider conservative bidders and justify this assumption since bidders’ strategies are dominated otherwise. This is a natural assumption that is usually made in similar contexts as well, such as in combinatorial auctions (e.g., see [3, 4, 12]). The authors of [16] consider pure and mixed Nash equilibria in the full-information setting, for which they upper-bound the price of anarchy by 1.618 and 4, respectively. Their result for mixed Nash equilibria is valid for more general equilibrium classes as well. Furthermore, they present a tight (lower and upper) bound of $5/4$ for pure Nash equilibria and two bidders. For the incomplete information setting and Bayes-Nash equilibria, they show an upper bound of 8. The authors of [13] improve this last bound to 3.162, while (in the preliminary version of their paper) they present a tight bound of 1.259 on the price of anarchy over pure Nash equilibria for the case of three bidders (the lower bound has also been claimed in [16] without presenting the explicit construction).

Roughgarden [18] presents a sufficient condition for games, termed smoothness, so that the price of anarchy of a smooth game over pure Nash equilibria immediately extends also to mixed Nash, correlated, and coarse correlated equilibria as well. Smoothness arguments have been implicitly or explic-

itly used (see [18] and the references therein) in order to provide such “robust” bounds on the price of anarchy of several games. As observed in [16], generalized second price auctions do not correspond to smooth games.

Our results. We first consider pure Nash equilibria of generalized second price auction games under the full information setting (Section 3). We warm up by considering the case of three bidders (in Section 3.1), for which we present the upper bound of approximately 1.259 (by providing an alternative proof to the one in the preliminary version of [13]). This upper bound together with the bound of $5/4$ for two bidders from [16] serves as the base of the inductive proof of our more general result presented in Section 3.2: an extremely low upper bound of 1.282 on the price of anarchy in games with arbitrarily many bidders. This result implies that the number of bidders involved in the auction has almost no effect on the price of anarchy. The proof elaborates on techniques developed in [16] and follows by a detailed reasoning about the structure of equilibria and a novel relation of the price of anarchy to the objective value of a compact mathematical program.

Then, in Section 4, we consider the broad class of coarse correlated equilibria in the full information setting and prove an upper bound of 2.3102 on the price of anarchy, improving the upper bound of 4 from [16]. Clearly, this bound holds for more restricted equilibria classes, namely correlated and mixed Nash equilibria. In the incomplete information setting, we prove an upper bound of 3.037 on the price of anarchy over Bayes-Nash equilibria (Section 5), improving the bound of 3.162 from [13]. The proofs of these results explicitly take into account the bids of the bidders and bound their utility (in different ways) by considering several possible deviations. In order to obtain our bounds, we adapt techniques from [13] and significantly extend them with new arguments.

All our bounds hold (with appropriate adaptations in the proofs) in the more general model of separable click-through rates [11] in which the click-through rate of a slot depends on the bidder allocated to that slot as well. In order to keep the exposition simple, we do not consider this extension in the current text.

We begin with preliminary definitions in Section 2 and conclude with open problems in Section 6. Due to lack of space, some proofs have been omitted.

2. PRELIMINARIES

Before proceeding with the presentation of our results, we give some formal definitions. Throughout the paper, we consider generalized second price auctions with n bidders and n slots. In such an auction, each bidder i has a *valuation* v_i that denotes how much the bidder values a click on her ad. Each slot i has a non-negative *click-through rate* a_i that denotes the rate by which this slot is clicked by the viewers. Without loss of generality, we assume that bidders and slots are sorted according to their valuations and their click-through rates, so that $v_1 \geq v_2 \geq \dots \geq v_n \geq 0$ and $a_1 \geq a_2 \geq \dots \geq a_n \geq 0$. Given a vector of bids $\mathbf{b} = (b_1, b_2, \dots, b_n)$ with one bid per bidder, the generalized second price auction defines an assignment π according to which the bidder with the j -th highest bid is assigned to slot j (ties are broken arbitrarily). We denote by $\pi(j)$ the bidder assigned to slot j and by $\pi^{-1}(i)$ the slot to which bidder i is

allocated. Then, each bidder i is required to pay the bid of the bidder that is assigned to the slot below hers (i.e., slot $\pi^{-1}(i) + 1$) if any.

We study the strategic game among the bidders that is induced by such an auction and we refer to it as a *GSP auction game*. In this game, each bidder acts selfishly and aims to maximize her utility given the bids of the other bidders. Given a bid vector \mathbf{b} and the corresponding assignment π it induces, the *utility* of bidder i is

$$u_i(\mathbf{b}) = a_{\pi^{-1}(i)} (v_i - b_{\pi(\pi^{-1}(i)+1)}),$$

assuming that $b_{\pi(n+1)} = 0$. We assume that bidders are conservative, i.e., each bidder selects as her strategy a bid that does not exceed her valuation. A *pure strategy* for bidder i consists of a single bid $b_i \in [0, v_i]$, while a *mixed strategy* is a probability distribution over pure strategies. We say that a bid vector \mathbf{b} is a *mixed Nash equilibrium* for a GSP auction game if no bidder has an incentive to unilaterally deviate from her strategy in order to strictly increase her expected utility. I.e., for each bidder i and each alternative bid $b'_i \in [0, v_i]$, it holds that

$$\mathbb{E}[u_i(\mathbf{b})] \geq \mathbb{E}[u_i(b'_i, \mathbf{b}_{-i})],$$

where \mathbf{b}_{-i} denotes the bids of all bidders apart from i and the expectation is taken according to the randomness of the bids. If all bidders play pure strategies, then \mathbf{b} is a *pure Nash equilibrium*.

Correlated and coarse correlated equilibria can be viewed as generalizations of mixed Nash equilibria, where the bidders have a joint probability distribution instead of independent ones. Informally, in both settings, a mediator draws a bid vector from a publicly known distribution and secretly informs each bidder of her suggested strategy. If no bidder has an incentive to deviate from the suggested strategy, then this is a *correlated equilibrium* [2], while if no bidder has a pure strategy that she can always follow, irrespective of the outcome, and improve her expected utility, then this is a *coarse correlated equilibrium* [14] (see also [22]). More formally, there exists a joint probability distribution that draws the bid vector \mathbf{b} from the universe \mathcal{B} of all bid vectors (y_1, y_2, \dots, y_n) with $y_i \in [0, v_i]$. The expected utility of a bidder is then

$$\mathbb{E}[u_i(\mathbf{b})] = \sum_{y \in \mathcal{B}} \Pr[\mathbf{b} = y] \cdot u_i(y).$$

In a coarse correlated equilibrium, for any bidder i , there exists no alternative bid $b' \in [0, v_i]$ such that

$$\mathbb{E}[u_i(\mathbf{b})] < \mathbb{E}[u_i(b', \mathbf{b}_{-i})].$$

A correlated equilibrium is also a coarse correlated one.

The *social welfare* of an assignment π induced by a bid vector \mathbf{b} is then defined as

$$\mathcal{W}(\mathbf{b}) = \mathbb{E}\left[\sum_i a_i v_{\pi(i)}\right],$$

where the expectation is taken according to the randomness of the bids. Equivalently, we have

$$\mathcal{W}(\mathbf{b}) = \mathbb{E}\left[\sum_i a_{\pi^{-1}(i)} v_i\right] = \mathbb{E}\left[\sum_i u_i(\mathbf{b})\right] + \mathbb{E}\left[\sum_i a_i b_{\pi(i)+1}\right].$$

The *optimal social welfare* is

$$\text{OPT} = \sum_i a_i v_i.$$

We also consider the Bayesian setting [7] where bidders' valuations are random. In this setting, each bidder has a probability distribution on her valuation and her strategy depends on the actual valuation, i.e., a bid is now a function on the valuation. The optimal social welfare is defined as

$$\text{OPT} = \mathbb{E}\left[\sum_i a_{o(i)} v_i\right],$$

where $o(i)$ is the random variable denoting the slot that bidder i occupies in the optimal assignment and the expectation is taken according to the randomness in the bidders' valuations. In a *Bayes-Nash equilibrium*, for any bidder i , any possible value x for her valuation, and any alternative bid $b'_i(x)$, it holds that

$$\mathbb{E}[u_i(b_i(x), \mathbf{b}_{-i}(v_{-i})) | v_i = x] \geq \mathbb{E}[u_i(b'_i(x), \mathbf{b}_{-i}(v_{-i})) | v_i = x].$$

The social welfare is again

$$\mathcal{W}(\mathbf{b}) = \mathbb{E}\left[\sum_i a_i v_{\pi(i)}\right],$$

where the expectation is taken according to the randomness in the valuations and the bids.

The *price of anarchy* (introduced in [8]; see also [17]) of a game over a given class of equilibria is defined as the worst ratio of the optimal social welfare over the social welfare of an equilibrium (over all equilibria of the class), i.e., $\max_{\mathbf{b}} \text{OPT}/\mathcal{W}(\mathbf{b})$, where \mathbf{b} is restricted to equilibria. The price of anarchy for a class of games (over equilibria of a particular type) is the worst price of anarchy among the games in the class (over equilibria of the same type).

3. THE PRICE OF ANARCHY OVER PURE NASH EQUILIBRIA

In this section we present our results for pure Nash equilibria. We consider GSP auction games with n bidders with valuations $v_1 \geq \dots \geq v_n \geq 0$ and n slots with click-through rates $a_1 \geq \dots \geq a_n \geq 0$. We assume that neither all slots have the same click-through rate nor all bidders have the same valuation (in both cases, the price of anarchy is 1). We elaborate on the approach taken in [16] and use the notion of weakly feasible assignments defined therein.

DEFINITION 1. *An assignment π is weakly feasible if for each pair of bidders i, j , it holds that $a_{\pi^{-1}(i)} v_i \geq a_{\pi^{-1}(j)} (v_i - v_j)$.*

We refer to the inequalities in this definition as *weak feasibility conditions*. With some abuse of notation, we denote by $\mathcal{W}(\pi) = \sum_i a_i v_{\pi(i)}$ the social welfare of assignment π and use the term *efficiency* of assignment π to refer to the ratio $\text{OPT}/\mathcal{W}(\pi)$. As proved in [16], every pure Nash equilibrium corresponds to a weakly feasible assignment. Hence, the price of anarchy of a GSP auction game over pure Nash equilibria is upper-bounded by the worst-case efficiency among weakly feasible assignments.

DEFINITION 2. *An assignment π is called proper if for any two slots $i < j$ with equal click-through rates, it holds $\pi(i) < \pi(j)$.*

Clearly, for any non-proper weakly feasible assignment, we can construct a proper weakly feasible one with equal social welfare. Hence, in order to prove our upper bounds, we

essentially upper-bound the worst-case efficiency over proper weakly feasible assignments.

Given an assignment π , consider the directed graph $G(\pi)$ that has one node for each slot, and a directed edge for each bidder i that connects the node corresponding to slot i to the node corresponding to slot $\pi^{-1}(i)$. In general, $G(\pi)$ consists of a set of disjoint cycles and may contain self-loops.

DEFINITION 3. *An assignment π is called reducible if its directed graph $G(\pi)$ has more than one cycles. Otherwise, it is called irreducible.*

Given a reducible assignment π such that $G(\pi)$ has $c \geq 2$ cycles, we can construct c GSP auction subgames by considering the slots and the bidders that correspond to the nodes and edges of each cycle. Similarly, for $\ell = 1, \dots, c$, the restriction π^ℓ of π to the slots and bidders of the ℓ -th subgame is an assignment for this game. The next fact is implicit in [16].

FACT 4. *If assignment π is weakly feasible for the original GSP auction game, then π^ℓ is weakly feasible for the ℓ -th subgame as well, for $\ell = 1, \dots, c$. Then, the efficiency of π is at most the maximum efficiency among the assignments π^ℓ for $\ell = 1, \dots, c$.*

When considering irreducible weakly feasible assignments, we further assume that the index of the slot bidder 1 occupies is smaller than the index of the bidder that is assigned to slot 1. This is without loss of generality due to the following argument. Consider an irreducible weakly feasible assignment π for a GSP auction game with n bidders such that $\pi^{-1}(1) > \pi(1)$. We construct a new game with click-through rate $a'_i = v_i$ for slot i and valuation $v'_i = a_i$ for bidder i , for $i = 1, \dots, n$ and the assignment $\pi_* = \pi^{-1}$. Observe that $\pi_*^{-1}(1) = \pi(1) < \pi^{-1}(1) = \pi_*(1)$. Clearly, the optimal social welfare is the same in both games while the social welfare of π_* for the new game is $\mathcal{W}(\pi_*) = \sum_i a'_i v'_{\pi_*(i)} = \sum_i v_i a_{\pi^{-1}(i)} = \sum_i a_{\pi^{-1}(i)} v_i = \mathcal{W}(\pi)$. We can also prove the weak feasibility conditions for π_* in the new game for each i, j . In order to do so, consider the weak feasibility condition for π in the original game for bidders $\pi(j), \pi(i)$. It is $a_{\pi^{-1}(\pi(j))} v_{\pi(j)} \geq a_{\pi^{-1}(\pi(i))} (v_{\pi(j)} - v_{\pi(i)})$ and, equivalently, $v_{\pi(i)} a_{\pi^{-1}(\pi(i))} \geq v_{\pi(j)} (a_{\pi^{-1}(\pi(i))} - a_{\pi^{-1}(\pi(j))})$. By the definition of the click-through rates and the valuations in the new game and the definition of π_* , we obtain that $a'_{\pi_*^{-1}(i)} v'_i \geq a'_{\pi_*^{-1}(j)} (v'_i - v'_j)$ as desired.

We furthermore note that when $v_n = 0$, any proper weakly feasible assignment is reducible. This is obviously the case if all bidders with zero valuation use the last slots. Otherwise, consider a bidder i with non-zero valuation that is assigned a slot $\pi^{-1}(i) > \pi^{-1}(j)$ where j is a bidder with zero valuation. Since the assignment is proper, it holds that $a_{\pi^{-1}(i)} < a_{\pi^{-1}(j)}$. Then, we obtain a contradiction by the weak feasibility condition $a_{\pi^{-1}(i)} v_i \geq a_{\pi^{-1}(j)} (v_i - v_j)$ for bidders i, j .

3.1 GSP auction games with three bidders

We are ready to present our first result (Theorem 6). In the proof, we use the following technical lemma.

LEMMA 5. *Let $\zeta = 0.129567$. For any $\lambda \in [0, 1]$, it holds that $\sqrt{\lambda^3 + 1} \geq 1 - \zeta\lambda + \frac{\lambda^2}{2}$.*

PROOF. Since both parts of the inequality are non-negative for $\lambda \in [0, 1]$, it suffices to show that the function $f(\lambda) =$

$(\lambda^3 + 1) - \left(1 - \zeta\lambda + \frac{\lambda^2}{2}\right)^2$ is non-negative for $\lambda \in [0, 1]$.

Let $g(\lambda) = -\frac{\lambda^3}{4} + (1 + \zeta)\lambda^2 - (1 + \zeta^2)\lambda + 2\zeta$ and observe that $f(\lambda) = \lambda \cdot g(\lambda)$. The proof will follow by proving that $g(\lambda) \geq 0$ when $\lambda \in [0, 1]$. Observe that the derivative of g is strictly negative for $\lambda = 0$ and strictly positive for $\lambda = 1$. Hence, the minimum of g in $[0, 1]$ is achieved at the point $\lambda^* = \frac{4+4\zeta-2\sqrt{\zeta^2+8\zeta+1}}{3}$ where the derivative of g becomes zero. Straightforward calculations yield that $g(\lambda^*) > 0$ and the lemma follows. \square

THEOREM 6. *The price of anarchy over pure Nash equilibria of GSP auction games with three conservative bidders is at most 1.259134.*

PROOF. Consider a GSP auction game with three slots with click-through rates $a_1 \geq a_2 \geq a_3 \geq 0$ and three bidders with valuations $v_1 \geq v_2 \geq v_3 \geq 0$ and a proper weakly feasible assignment π of slots to bidders. We will prove the theorem by upper-bounding the efficiency of π by 1.259134. If π is reducible, then the efficiency is bounded by the efficiency for games with two bidders and the theorem follows by the upper bound of 5/4 proved in [16] for this case. So, in the following, we assume that π is irreducible; by the observation above, this implies that $v_3 > 0$. There are only two such assignments which are in fact symmetric: in the first, slots 1, 2, 3 are allocated to bidders 3, 1, 2, respectively, and in the second, slots 1, 2, 3 are allocated to bidders 2, 3, 1, respectively. Without loss of generality (see the discussion above), we assume that π is the former assignment.

Let β, γ, λ , and μ be such that $a_2 = \beta a_1$, $a_3 = \gamma a_1$, $v_2 = \lambda v_1$, and $v_3 = \mu v_1$. Clearly, it holds that $1 \geq \beta \geq \gamma \geq 0$ and $1 \geq \lambda \geq \mu > 0$. The social welfare of assignment π is $\mathcal{W}(\pi) = a_1 v_1 (\mu + \beta + \gamma\lambda)$ whereas the optimal social welfare is $\text{OPT} = a_1 v_1 (1 + \beta\lambda + \gamma\mu)$. Furthermore, since π is weakly feasible, the weak feasibility conditions for bidders 1 and 3 and bidders 2 and 3 are $a_2 v_1 \geq a_1 (v_1 - v_3)$ and $a_3 v_2 \geq a_1 (v_2 - v_3)$, respectively, i.e., $\beta \geq 1 - \mu$ and $\gamma \geq 1 - \frac{\mu}{\lambda}$. We are now ready to bound the efficiency of π . Let $\delta, \epsilon \geq 0$ be such that $\beta = 1 - \mu + \delta$ and $\gamma = 1 - \frac{\mu}{\lambda} + \epsilon$. We have

$$\begin{aligned} \frac{\text{OPT}}{\mathcal{W}(\pi)} &= \frac{1 + \beta\lambda + \gamma\mu}{\mu + \beta + \gamma\lambda} \\ &= \frac{1 + \lambda - \mu\lambda + \mu - \frac{\mu^2}{\lambda} + \delta\lambda + \epsilon\mu}{1 + \lambda - \mu + \delta + \epsilon\lambda} \\ &\leq \frac{1 + \lambda - \mu\lambda + \mu - \frac{\mu^2}{\lambda}}{1 + \lambda - \mu}. \end{aligned}$$

The inequality follows since $1 \geq \lambda \geq \mu > 0$ implies that $1 + \lambda - \mu\lambda + \mu - \frac{\mu^2}{\lambda} = 1 + \lambda - \mu + \mu(1 - \lambda) + \mu(1 - \mu/\lambda) \geq 1 + \lambda - \mu \geq 1$ and $\delta + \epsilon\lambda \geq \delta\lambda + \epsilon\mu \geq 0$.

For $\mu \in [0, 1]$, this last expression is maximized for the value of μ that makes its derivative with respect to μ equal to zero, i.e., $\mu = -\sqrt{\lambda^3 + 1} + \lambda + 1$. By substituting μ , we obtain that

$$\begin{aligned} \frac{\text{OPT}}{\mathcal{W}(\pi)} &\leq \frac{\lambda^2 + \lambda + 2 - 2\sqrt{\lambda^3 + 1}}{\lambda} \\ &\leq 1 + 2\zeta \\ &= 1.259134, \end{aligned}$$

where $\zeta = 0.129567$ and the second inequality follows by Lemma 5. \square

3.2 GSP auction games with many bidders

In this section, we prove our main result for pure Nash equilibria.

THEOREM 7. *The price of anarchy over pure Nash equilibria of GSP auction games with conservative bidders is at most $\frac{61+7\sqrt{217}}{128} \approx 1.28216$.*

In our proof, we will need the following technical lemma.

LEMMA 8. *Let $r = \frac{61+7\sqrt{217}}{128} \approx 1.28216$ and $f(\beta, \gamma, \lambda, \mu) = \mu + \beta(1 - \frac{\lambda}{r}) + \gamma(\lambda - \mu) - \frac{1}{r}$. Then, the objective value of the mathematical program*

$$\begin{aligned} & \text{minimize} && f(\beta, \gamma, \lambda, \mu) \\ & \text{subject to} && \beta \geq 1 - \mu \\ & && \gamma \geq 1 - \mu/\lambda \\ & && 1 \geq \lambda \geq \mu > 0 \\ & && 1 \geq \beta, \gamma \geq 0 \end{aligned}$$

is non-negative.

PROOF. Since $\mu \leq \lambda \leq 1$, we have that $f(\beta, \gamma, \lambda, \mu)$ is non-decreasing in β and γ . Using the first two constraints, we have that the objective value of the mathematical program is at least

$$f\left(1 - \mu, 1 - \frac{\mu}{\lambda}, \lambda, \mu\right) = 1 - \frac{1}{r} + \lambda - \frac{\lambda}{r} - \mu\left(2 - \frac{\lambda}{r}\right) + \frac{\mu^2}{\lambda},$$

which is minimized for $\mu = \lambda - \frac{\lambda^2}{2r}$ to

$$f\left(1 - \lambda + \frac{\lambda^2}{2r}, \frac{\lambda}{2r}, \lambda, \lambda - \frac{\lambda^2}{2r}\right) = 1 - \frac{1}{r} - \frac{\lambda}{r} + \frac{\lambda^2}{r} - \frac{\lambda^3}{4r^2}.$$

In order to complete the proof it suffices to show that the function $g(\lambda) = 1 - \frac{1}{r} - \frac{\lambda}{r} + \frac{\lambda^2}{r} - \frac{\lambda^3}{4r^2}$ is non-negative for $\lambda \in [0, 1]$. Observe that $g(\lambda)$ is a polynomial of degree 3 and, hence, it has at most one local minimum. Also observe that the derivative of $g(\lambda)$ is $-\frac{1}{r} + \frac{2\lambda}{r} - \frac{3\lambda^2}{4r^2}$ which is strictly negative for $\lambda = 0$ and strictly positive for $\lambda = 1$. Hence, its minimum in $[0, 1]$ is achieved at the point $\lambda^* = \frac{4r-2\sqrt{4r^2-3r}}{3}$ where the derivative becomes zero. Straightforward calculations yield that $g(\lambda^*) = 0$ and the lemma follows. \square

PROOF OF THEOREM 7. In order to prove the theorem, we will prove that the worst-case efficiency among weakly feasible assignments of any GSP auction game is at most $r = \frac{61+7\sqrt{217}}{128} \approx 1.28216$. We use induction. As the base of our induction, we use the fact that GSP auction games with one, two, or three bidders have worst-case efficiency among weakly feasible assignments at most 1.28216. For a single bidder, the claim is trivial. For two bidders, it follows by [16], and for three bidders, it follows by the proof of Theorem 6. Let $n \geq 4$ be an integer. Using the inductive hypothesis that the worst-case efficiency among weakly feasible assignments of any GSP auction game with at most $n-1$ bidders is at most r , we will show that this is also the case for any GSP auction game with n bidders.

Consider a GSP auction game with n bidders with valuations $v_1 \geq v_2 \geq \dots \geq v_n \geq 0$ and n slots with click-through rates $a_1 \geq a_2 \geq \dots \geq a_n \geq 0$ and let π be a proper weakly feasible assignment. If π is reducible, the claim follows by

Fact 4 and the inductive hypothesis. So, in the following, we assume that π is irreducible; this implies that $v_n > 0$. Let j be the bidder that is assigned slot 1 and i_1 be the slot assigned to bidder 1. Without loss of generality, we assume that $i_1 < j$ since the other case is symmetric; see the discussion at the beginning of Section 3. Also, let i_2 be the slot assigned to bidder i_1 . By our assumptions, the integers $j, 1, i_1,$ and i_2 are different.

We will show that

$$\begin{aligned} \mathcal{W}(\pi) \geq & a_1 v_j + a_{i_1} (v_1 - \frac{v_{i_1}}{r}) + a_{i_2} (v_{i_1} - v_j) - \frac{a_1 v_1}{r} \\ & + \frac{\text{OPT}}{r}. \end{aligned} \quad (1)$$

Assuming that (1) holds, we apply Lemma 8 with $\beta = a_{i_1}/a_1$, $\gamma = a_{i_2}/a_1$, $\lambda = v_{i_1}/v_1$, and $\mu = v_j/v_1$. Clearly, the last two constraints of the mathematical program in Lemma 8 are satisfied. Also, observe that the weak feasibility conditions for bidders 1 and j and bidders i_1 and j in assignment π are $a_{i_1} v_1 \geq a_1 (v_1 - v_j)$ and $a_{i_2} v_{i_1} \geq a_1 (v_{i_1} - v_j)$, respectively, i.e., $\beta \geq 1 - \mu$ and $\gamma \geq 1 - \mu/\lambda$ and the first two constraints of the mathematical program in Lemma 8 are satisfied as well. Now, using inequality (1) and Lemma 8, we have that

$$\mathcal{W}(\pi) \geq f\left(\frac{a_{i_1}}{a_1}, \frac{a_{i_2}}{a_1}, \frac{v_{i_1}}{v_1}, \frac{v_j}{v_1}\right) \cdot a_1 v_1 + \frac{\text{OPT}}{r} \geq \frac{\text{OPT}}{r}$$

and the proof follows.

It remains to prove inequality (1). We distinguish between three cases depending on the relative order of $j, i_1,$ and i_2 ; in each of these cases, we further distinguish between two subcases.

Case I.1: $1 < i_1 < j < i_2$ and $a_j \leq a_{i_2} r$. Consider the restriction of the original game that consists of the bidders different than $j, 1,$ and i_1 and the slots different than $1, i_1,$ and i_2 . Let π' be the restriction of π to the bidders and slots of the new game. Clearly, this assignment is weakly feasible for the new game since the weak feasibility conditions for π' are just a subset of the corresponding conditions for π (for the original game). Also, note that the optimal assignment for the restricted game assigns bidder k to slot k for $k = 2, \dots, i_1 - 1, i_1 + 1, \dots, j - 1, i_2 + 1, \dots, n$ and bidder $k + 1$ to slot k for $k = j, \dots, i_2 - 1$. By the inductive hypothesis, we know that the efficiency of π' is at most r . Hence, we can bound the social welfare of π as

$$\begin{aligned} \mathcal{W}(\pi) &= a_1 v_j + a_{i_1} v_1 + a_{i_2} v_{i_1} + \sum_{k \notin \{1, i_1, i_2\}} a_k v_{\pi(k)} \\ &= a_1 v_j + a_{i_1} v_1 + a_{i_2} v_{i_1} + \mathcal{W}(\pi') \\ &\geq a_1 v_j + a_{i_1} v_1 + a_{i_2} v_{i_1} + \frac{1}{r} \left(\sum_{k=2}^{i_1-1} a_k v_k \right. \\ &\quad \left. + \sum_{k=i_1+1}^{j-1} a_k v_k + \sum_{k=j}^{i_2-1} a_k v_{k+1} + \sum_{k=i_2+1}^n a_k v_k \right) \end{aligned}$$

$$\begin{aligned}
&\geq a_1 v_j + a_{i_1} v_1 + a_{i_2} v_{i_1} + \frac{1}{r} \left(\sum_{k=2}^{i_1-1} a_k v_k + \sum_{k=i_1+1}^{j-1} a_k v_k \right. \\
&\quad \left. + \sum_{k=j+1}^{i_2} a_k v_k + \sum_{k=i_2+1}^n a_k v_k \right) \\
&= a_1 v_j + a_{i_1} v_1 + a_{i_2} v_{i_1} + \frac{1}{r} \left(\sum_{k=1}^n a_k v_k - a_1 v_1 \right. \\
&\quad \left. - a_{i_1} v_{i_1} - a_j v_j \right) \\
&= a_1 v_j + a_{i_1} v_1 + a_{i_2} v_{i_1} - \frac{1}{r} (a_1 v_1 + a_{i_1} v_{i_1} + a_j v_j) \\
&\quad + \frac{\text{OPT}}{r} \\
&\geq a_1 v_j + a_{i_1} (v_1 - \frac{v_{i_1}}{r}) + a_{i_2} (v_{i_1} - v_j) - \frac{a_1 v_1}{r} + \frac{\text{OPT}}{r}
\end{aligned}$$

and inequality (1) follows. The first inequality follows by the inductive hypothesis and the definition of the optimal assignment for the restricted game. The second inequality follows since $a_k \geq a_{k+1}$ for $k = j, \dots, i_2 - 1$. The last inequality follows since $a_j \leq a_{i_2} r$.

Case I.2: $1 < i_1 < j < i_2$ and $a_j > a_{i_2} r$. Consider the restriction of the original game that consists of the bidders different than j and 1 and the slots different than 1 and i_1 . Let π' be the restriction of π to the bidders and slots of the new game. Again, this assignment is weakly feasible for the new game since the weak feasibility conditions for π' are just a subset of the ones for π (for the original game). Also, note that the optimal assignment for the restricted game assigns bidder k to slot k for $k = 2, \dots, i_1 - 1, j + 1, \dots, n$ and bidder $k - 1$ to slot k for $k = i_1 + 1, \dots, j$. By the inductive hypothesis, we know that the efficiency of π' is at most r . Hence, we can bound the social welfare of π as

$$\begin{aligned}
&\mathcal{W}(\pi) \\
&= a_1 v_j + a_{i_1} v_1 + \sum_{k \notin \{1, i_1\}} a_k v_{\pi(k)} \\
&= a_1 v_j + a_{i_1} v_1 + \mathcal{W}(\pi') \\
&\geq a_1 v_j + a_{i_1} v_1 + \frac{1}{r} \left(\sum_{k=2}^{i_1-1} a_k v_k + \sum_{k=i_1+1}^j a_k v_{k-1} \right. \\
&\quad \left. + \sum_{k=j+1}^n a_k v_k \right) \\
&= a_1 v_j + a_{i_1} v_1 + \frac{1}{r} \left(\sum_{k=1}^n a_k v_k - a_1 v_1 - a_{i_1} v_{i_1} \right. \\
&\quad \left. + \sum_{k=i_1+1}^j a_k (v_{k-1} - v_k) \right) \\
&\geq a_1 v_j + a_{i_1} v_1 + \frac{1}{r} \left(\sum_{k=1}^n a_k v_k - a_1 v_1 - a_{i_1} v_{i_1} \right. \\
&\quad \left. + a_j \sum_{k=i_1+1}^j (v_{k-1} - v_k) \right)
\end{aligned}$$

$$\begin{aligned}
&= a_1 v_j + a_{i_1} v_1 - \frac{1}{r} (a_1 v_1 + a_{i_1} v_{i_1} + a_j v_j - a_j v_{i_1}) \\
&\quad + \frac{\text{OPT}}{r} \\
&> a_1 v_j + a_{i_1} (v_1 - \frac{v_{i_1}}{r}) + a_{i_2} (v_{i_1} - v_j) - \frac{a_1 v_1}{r} + \frac{\text{OPT}}{r}
\end{aligned}$$

and inequality (1) follows. The first inequality follows by the inductive hypothesis and the definition of the optimal assignment for the restricted game. The second inequality follows since $a_k \geq a_j$ and $v_{k-1} - v_k \geq 0$ for $k = i_1 + 1, \dots, j$. The last inequality follows since $a_j > a_{i_2} r$.

Case II.1: $1 < i_1 < i_2 < j$ and $v_{i_2} \leq v_j r$. Consider the restriction of the original game that consists of the bidders different than j , 1, and i_1 and the slots different than 1, i_1 , and i_2 . Let π' be the restriction of π to the bidders and slots of the new game. Clearly, this assignment is weakly feasible for the new game since the weak feasibility conditions for π' are just a subset of the ones for π (for the original game). Also, note that the optimal assignment for the restricted game assigns bidder k to slot k for $k = 2, \dots, i_1 - 1, i_1 + 1, \dots, i_2 - 1, j + 1, \dots, n$ and bidder $k - 1$ to slot k for $k = i_2 + 1, \dots, j$. By the inductive hypothesis, we know that the efficiency of π' is at most r . Hence, we can bound the social welfare of π as

$$\begin{aligned}
&\mathcal{W}(\pi) \\
&= a_1 v_j + a_{i_1} v_1 + a_{i_2} v_{i_1} + \sum_{k \notin \{1, i_1, i_2\}} a_k v_{\pi(k)} \\
&= a_1 v_j + a_{i_1} v_1 + a_{i_2} v_{i_1} + \mathcal{W}(\pi') \\
&\geq a_1 v_j + a_{i_1} v_1 + a_{i_2} v_{i_1} + \frac{1}{r} \left(\sum_{k=2}^{i_1-1} a_k v_k \right. \\
&\quad \left. + \sum_{k=i_1+1}^{i_2-1} a_k v_k + \sum_{k=i_2+1}^j a_k v_{k-1} + \sum_{k=j+1}^n a_k v_k \right) \\
&\geq a_1 v_j + a_{i_1} v_1 + a_{i_2} v_{i_1} + \frac{1}{r} \left(\sum_{k=2}^{i_1-1} a_k v_k \right. \\
&\quad \left. + \sum_{k=i_1+1}^{i_2-1} a_k v_k + \sum_{k=i_2+1}^j a_k v_k + \sum_{k=j+1}^n a_k v_k \right) \\
&= a_1 v_j + a_{i_1} v_1 + a_{i_2} v_{i_1} + \frac{1}{r} \left(\sum_{k=1}^n a_k v_k - a_1 v_1 \right. \\
&\quad \left. - a_{i_1} v_{i_1} - a_{i_2} v_{i_2} \right) \\
&= a_1 v_j + a_{i_1} v_1 + a_{i_2} v_{i_1} - \frac{1}{r} (a_1 v_1 + a_{i_1} v_{i_1} \\
&\quad + a_{i_2} v_{i_2}) + \frac{\text{OPT}}{r} \\
&\geq a_1 v_j + a_{i_1} (v_1 - \frac{v_{i_1}}{r}) + a_{i_2} (v_{i_1} - v_j) - \frac{a_1 v_1}{r} + \frac{\text{OPT}}{r}
\end{aligned}$$

and inequality (1) follows. The first inequality follows by the inductive hypothesis and the definition of the optimal assignment for the restricted game. The second inequality follows since $v_{k-1} \geq v_k$ for $k = i_2 + 1, \dots, j$. The last inequality follows since $v_{i_2} \leq v_j r$.

Case II.2: $1 < i_1 < i_2 < j$ and $v_{i_2} > v_j r$. Consider the restriction of the original game that consists of the bidders

different than 1 and i_1 and the slots different than i_1 and i_2 . Let π' be the restriction of π to the bidders and slots of the new game. Again, this assignment is weakly feasible for the new game since the weak feasibility conditions for π' are just a subset of the ones for π (for the original game). Also, note that the optimal assignment for the restricted game assigns bidder k to slot k for $k = i_2 + 1, \dots, n$, bidder $i_1 + 1$ to slot $i_1 - 1$, and bidder $k + 1$ to slot k for $k = 1, \dots, i_1 - 2, i_1 + 1, \dots, i_2 - 1$. By the inductive hypothesis, we know that the efficiency of π' is at most r . Hence, we can bound the social welfare of π as

$$\begin{aligned}
& \mathcal{W}(\pi) \\
&= a_{i_1} v_1 + a_{i_2} v_{i_1} + \sum_{k \notin \{i_1, i_2\}} a_k v_{\pi(k)} \\
&= a_{i_1} v_1 + a_{i_2} v_{i_1} + \mathcal{W}(\pi') \\
&\geq a_{i_1} v_1 + a_{i_2} v_{i_1} + \frac{1}{r} \left(\sum_{k=1}^{i_1-2} a_k v_{k+1} + a_{i_1-1} v_{i_1+1} \right. \\
&\quad \left. + \sum_{k=i_1+1}^{i_2-1} a_k v_{k+1} + \sum_{k=i_2+1}^n a_k v_k \right) \\
&= a_{i_1} v_1 + a_{i_2} v_{i_1} + \frac{1}{r} \left(\sum_{k=1}^n a_k v_k + \sum_{k=1}^{i_1-2} (a_k - a_{k+1}) v_{k+1} \right. \\
&\quad \left. + (a_{i_1-1} - a_{i_1+1}) v_{i_1+1} \right. \\
&\quad \left. + \sum_{k=i_1+1}^{i_2-1} (a_k - a_{k+1}) v_{k+1} - a_1 v_1 - a_{i_1} v_{i_1} \right) \\
&\geq a_{i_1} v_1 + a_{i_2} v_{i_1} + \frac{1}{r} \left(\sum_{k=1}^n a_k v_k + \sum_{k=1}^{i_1-2} (a_k - a_{k+1}) v_{i_2} \right. \\
&\quad \left. + (a_{i_1-1} - a_{i_1+1}) v_{i_2} \right. \\
&\quad \left. + \sum_{k=i_1+1}^{i_2-1} (a_k - a_{k+1}) v_{i_2} - a_1 v_1 - a_{i_1} v_{i_1} \right) \\
&= a_{i_1} v_1 + a_{i_2} v_{i_1} - \frac{1}{r} (a_1 v_1 + a_{i_1} v_{i_1} + a_{i_2} v_{i_2} - a_1 v_{i_2}) \\
&\quad + \frac{\text{OPT}}{r} \\
&> a_1 v_j + a_{i_1} (v_1 - \frac{v_{i_1}}{r}) + a_{i_2} (v_{i_1} - v_j) - \frac{a_1 v_1}{r} + \frac{\text{OPT}}{r}
\end{aligned}$$

and inequality (1) follows. The first inequality follows by the inductive hypothesis and the definition of the optimal assignment for the restricted game. The second inequality follows since $a_k - a_{k+1} \geq 0$ and $v_{k+1} \geq v_{i_2}$ for $k = 1, \dots, i_1 - 2, i_1 + 1, \dots, i_2 - 1$ and $a_{i_1-1} - a_{i_1+1} \geq 0$ and $v_{i_1+1} \geq v_{i_2}$. The last inequality follows since $v_{i_2} > v_j r$, and $a_1 > a_{i_2}$.

Due to lack of space, the third case where $1 < i_2 < i_1 < j$ is omitted. \square

4. COARSE CORRELATED EQUILIBRIA

In this section, we prove our upper bound on the price of anarchy over coarse correlated equilibria. Unlike the proof of previous upper bounds (e.g., in [16]), our proof is based on taking into account the bids in the analysis. We consider a GSP auction game with n slots with click-through rates $a_1 \geq a_2 \geq \dots \geq a_n \geq 0$ and n bidders with valuations

$v_1 \geq v_2 \geq \dots \geq v_n \geq 0$. Let \mathbf{b} denote the random bid vector drawn according to a probability distribution that corresponds to a coarse correlated equilibrium. We denote by π the random assignment induced by \mathbf{b} . Also, let $b_{\pi(i)}$ be the random variable denoting the i -th highest bid among all bidders.

We can extend the main argument in the proof of Lemma 6 in [13] (adapted to coarse correlated equilibria) in order to obtain the following lower bound on the expected utility of each bidder at a coarse correlated equilibrium.

LEMMA 9. *Consider a coarse correlated equilibrium and bidder i . For any $\beta > 0$, it holds that*

$$\mathbb{E}[u_i(\mathbf{b})] \geq \beta(1 - e^{-1/\beta}) a_i v_i - \beta \mathbb{E}[a_i b_{\pi(i)}].$$

We are now ready to prove the main result of this section. In our proof, we combine the above lemma together with a stronger property concerning the bidder with the highest valuation.

THEOREM 10. *The price of anarchy over coarse correlated equilibria of GSP auction games with conservative bidders is at most 2.3102.*

PROOF. We prove the theorem by lower-bounding the expected utility of each bidder at a coarse correlated equilibrium. Let $\beta \geq 1$ be a parameter to be fixed later. Consider bidder 1 and her deviation to the bid v_1 . Then, bidder 1 would always be allocated slot 1 and pay a_1 times the highest bid among the remaining bidders (which is upper bounded by $b_{\pi(1)}$). Hence,

$$\begin{aligned}
\mathbb{E}[u_1(\mathbf{b})] &\geq \mathbb{E}[u_1(v_1, \mathbf{b}_{-1})] \\
&\geq a_1 v_1 - \mathbb{E}[a_1 b_{\pi(1)}] \\
&\geq \beta(1 - e^{-1/\beta}) (a_1 v_1 - \mathbb{E}[a_1 b_{\pi(1)}]), \quad (2)
\end{aligned}$$

where the last inequality follows since $v_1 \geq b_{\pi(1)}$ and since $\beta \geq 1$ implies that $\beta(1 - e^{-1/\beta}) \leq 1$.

Now, we will use the fact that the social welfare is the sum of the expected utilities of the bidders plus the total bids paid. We have

$$\begin{aligned}
\beta \mathcal{W}(\mathbf{b}) &= \mathcal{W}(\mathbf{b}) + (\beta - 1) \mathcal{W}(\mathbf{b}) \\
&= \mathbb{E}[u_1(\mathbf{b})] + \mathbb{E}[\sum_{i \geq 2} u_i(\mathbf{b})] + \mathbb{E}[\sum_i a_i b_{\pi(i+1)}] \\
&\quad + (\beta - 1) \mathbb{E}[\sum_i a_i v_{\pi(i)}] \\
&\geq \beta(1 - e^{-1/\beta}) a_1 v_1 - \beta(1 - e^{-1/\beta}) \mathbb{E}[a_1 b_{\pi(1)}] \\
&\quad + \beta(1 - e^{-1/\beta}) \sum_{i \geq 2} a_i v_i - \beta \sum_{i \geq 2} \mathbb{E}[a_i b_{\pi(i)}] \\
&\quad + \sum_{i \geq 2} \mathbb{E}[a_i b_{\pi(i)}] + (\beta - 1) \mathbb{E}[a_1 v_{\pi(1)}] \\
&\quad + (\beta - 1) \sum_{i \geq 2} \mathbb{E}[a_i v_{\pi(i)}] \\
&\geq \beta(1 - e^{-1/\beta}) \text{OPT} + (\beta e^{-1/\beta} - 1) \mathbb{E}[a_1 b_{\pi(1)}],
\end{aligned}$$

where the first inequality follows by the lower bounds on the bidders' utilities in Lemma 9 and inequality (2) and the fact that $a_i \geq a_{i+1}$ for $i = 1, \dots, n - 1$, and the second inequality follows by the definition of OPT and since $\beta \geq 1$ and $v_{\pi(i)} \geq b_{\pi(i)}$ for any i .

By setting β such that $\beta e^{-1/\beta} = 1$, we obtain that the price of anarchy $\text{OPT}/\mathcal{W}(\mathbf{b})$ is at most $\beta/(\beta - 1)$. This occurs for $\beta \approx 1.76322$, where the ratio becomes 2.3102. \square

5. THE BAYESIAN SETTING

In this section, we prove our upper bound on the price of anarchy in the Bayesian setting. In our proof, we consider a GSP auction game with n slots with click-through rates $a_1 \geq a_2 \geq \dots \geq a_n \geq 0$ and n bidders with random valuations $v_1, v_2, \dots, v_n \geq 0$. Let \mathbf{b} denote the bid functions of the bidders at a Bayes-Nash equilibrium. We denote by π the random assignment induced by \mathbf{b} and by $o(i)$ the random variable that indicates the slot bidder i occupies in the optimal assignment; therefore, $o^{-1}(j)$ stands for the bidder that occupies slot j in the optimal assignment. Also, let $b_{\pi(i)}$ be the random variable denoting the i -th highest bid among all bidders.

We will prove our main result by appropriately combining two lower bounds for the expected utilities of the bidders. In the proof of the first lower bound (Lemma 12), we will need the following technical lemma.

LEMMA 11. *For any ξ and any non-negative X, Y , it holds that*

$$(X - Y)^2 \geq \left(\xi - \frac{\xi^2}{4} \right) X^2 - \xi XY.$$

PROOF. Let $f(X, Y) = (X - Y)^2 - (\xi - \xi^2/4)X^2 + \xi XY$. It suffices to prove that $f(X, Y) \geq 0$ for any ξ . Then, $f(X, Y) = (1 - \xi + \xi^2/4)X^2 + Y^2 - (2 - \xi)XY = ((1 - \xi/2)X - Y)^2 \geq 0$. \square

LEMMA 12. *Consider a Bayes-Nash equilibrium. Then, for any γ and δ , it holds that*

$$\begin{aligned} \sum_i \mathbb{E}[u_i(\mathbf{b})] &\geq (\gamma/2 - \gamma^2/8) \text{OPT} - \gamma/2 \sum_i \mathbb{E}[a_i b_{\pi(i)}] \\ &\quad + (\delta - \delta^2/4 - \gamma/2 + \gamma^2/8) \mathbb{E}[a_1 v_{o^{-1}(1)}] \\ &\quad - (\delta - \gamma/2) \mathbb{E}[a_1 b_{\pi(1)}]. \end{aligned}$$

PROOF. By the definition of the Bayes-Nash equilibrium, it holds that any bidder i can not increase her expected utility when her valuation is $v_i = x$, by deviating to any bid $b'_i < x$. We will argue about all possible slots that bidder i can be assigned to, when having valuation x and bidding b'_i . Observe that for any bid $b'_i < x$ such that $b'_i > b_{\pi(1)}$, bidder i is allocated to the first slot when bidding b'_i and pays at most $a_1 b_{\pi(1)}$. Similarly, for any bid $b'_i < x$ such that $b'_i > b_{\pi(j)}$, bidder i is allocated to slot j (or a higher one) and pays at most b'_i per click. Let A_x^{ij} denote the event that $v_i = x$ and $o(i) = j$ and B_x^{ij} denote the event that $o(i) = j$ given that $v_i = x$. Since the bid b_i maximizes the expected utility of bidder i , we have that

$$\begin{aligned} &\mathbb{E}[u_i(\mathbf{b})|v_i = x] \\ &\geq \mathbb{E}[u_i(b'_i, \mathbf{b}_{-i})|v_i = x] \\ &= \sum_{j=1}^n \mathbb{E}[u_i(b'_i, \mathbf{b}_{-i})|A_x^{ij}] \cdot \Pr[B_x^{ij}] \end{aligned}$$

$$\begin{aligned} &\geq a_1 \left(x - \mathbb{E}[b_{\pi(1)}|A_x^{i1}] \right) \cdot \Pr[b_{\pi(1)} < b'_i|A_x^{i1}] \cdot \Pr[B_x^{i1}] \\ &\quad + \sum_{j=2}^n a_j \left(x - b'_i \right) \cdot \Pr[b_{\pi(j)} < b'_i|A_x^{ij}] \cdot \Pr[B_x^{ij}] \\ &= a_1 \left(x - \mathbb{E}[b_{\pi(1)}|A_x^{i1}] \right) \cdot \left(1 - \Pr[b_{\pi(1)} \geq b'_i|A_x^{i1}] \right) \\ &\quad \cdot \Pr[B_x^{i1}] \\ &\quad + \sum_{j=2}^n a_j \left(x - b'_i \right) \cdot \left(1 - \Pr[b_{\pi(j)} \geq b'_i|A_x^{ij}] \right) \cdot \Pr[B_x^{ij}]. \end{aligned}$$

Using this inequality we have

$$\begin{aligned} &x \mathbb{E}[u_i(\mathbf{b})|v_i = x] \\ &= \int_0^x \mathbb{E}[u_i(\mathbf{b})|v_i = x] db'_i \\ &\geq a_1 \int_0^x \left(x - \mathbb{E}[b_{\pi(1)}|A_x^{i1}] \right) \cdot \left(1 - \Pr[b_{\pi(1)} \geq b'_i|A_x^{i1}] \right) db'_i \\ &\quad \cdot \Pr[B_x^{i1}] \\ &\quad + \sum_{j=2}^n a_j \int_0^x \left(x - b'_i \right) \cdot \left(1 - \Pr[b_{\pi(j)} \geq b'_i|A_x^{ij}] \right) db'_i \\ &\quad \cdot \Pr[B_x^{ij}] \\ &= a_1 \left(x^2 - 2x \mathbb{E}[b_{\pi(1)}|A_x^{i1}] + \mathbb{E}^2[b_{\pi(1)}|A_x^{i1}] \right) \cdot \Pr[B_x^{i1}] \\ &\quad + \sum_{j=2}^n a_j \left(\frac{x^2}{2} - x \mathbb{E}[b_{\pi(j)}|A_x^{ij}] + \frac{1}{2} \mathbb{E}[b_{\pi(j)}^2|A_x^{ij}] \right) \cdot \Pr[B_x^{ij}] \\ &\geq a_1 \left(x - \mathbb{E}[b_{\pi(1)}|A_x^{i1}] \right)^2 \cdot \Pr[B_x^{i1}] \\ &\quad + \frac{1}{2} \sum_{j=2}^n a_j \left(x - \mathbb{E}[b_{\pi(j)}|A_x^{ij}] \right)^2 \cdot \Pr[B_x^{ij}], \end{aligned}$$

where the second equality holds since A_x^{i1} implies $b_{\pi(1)} \leq x$ and, hence, $\mathbb{E}[b_{\pi(1)}|A_x^{i1}] = \int_0^x \Pr[b_{\pi(1)} \geq b'_i|A_x^{i1}] db'_i$, and, furthermore, since $\int_0^\infty z \Pr[Z \geq z] dz = \frac{1}{2} \mathbb{E}[Z^2]$, for any random variable Z that takes non-negative values. The second inequality holds since $\mathbb{E}[X^2] \geq \mathbb{E}^2[X]$ for any non-negative random variable X .

We will now use the above inequality and will apply Lemma 11 to its right-hand side (with $X = x$, $Y = \mathbb{E}[b_{\pi(j)}|A_x^{ij}]$ and ξ equal to γ and δ) in order to lower-bound $\mathbb{E}[u_i(\mathbf{b})|v_i = x]$. We will also exploit the fact that $x = \mathbb{E}[v_i|A_x^{ij}]$, for any $j = 1, \dots, n$, when $\Pr[B_x^{ij}] > 0$. We have

$$\begin{aligned} &\mathbb{E}[u_i(\mathbf{b})|v_i = x] \\ &\geq \left((\delta - \delta^2/4) \mathbb{E}[a_1 v_i|A_x^{i1}] - \delta \mathbb{E}[a_1 b_{\pi(1)}|A_x^{i1}] \right) \cdot \Pr[B_x^{i1}] \\ &\quad + \frac{1}{2} \sum_{j=2}^n \left((\gamma - \gamma^2/4) \mathbb{E}[a_j v_i|A_x^{ij}] - \gamma \mathbb{E}[a_j b_{\pi(j)}|A_x^{ij}] \right) \\ &\quad \cdot \Pr[B_x^{ij}]. \end{aligned}$$

We can now bound the unconditional expected utility of bidder i using this last inequality. We have

$$\begin{aligned} &\mathbb{E}[u_i(\mathbf{b})] \\ &= \int_0^\infty \mathbb{E}[u_i(\mathbf{b})|v_i = x] \cdot \Pr[v_i = x] dx \end{aligned}$$

$$\begin{aligned}
&\geq (\delta - \delta^2/4) \int_0^\infty \mathbb{E}[a_1 v_i | A_x^{i1}] \cdot \Pr[B_x^{i1}] \cdot \Pr[v_i = x] dx \\
&\quad - \delta \int_0^\infty \mathbb{E}[a_1 b_{\pi(1)} | A_x^{i1}] \cdot \Pr[B_x^{i1}] \cdot \Pr[v_i = x] dx \\
&\quad + (\gamma/2 - \gamma^2/8) \\
&\quad \cdot \sum_{j=2}^n \int_0^\infty \mathbb{E}[a_j v_i | A_x^{ij}] \cdot \Pr[B_x^{ij}] \cdot \Pr[v_i = x] dx \\
&\quad - \gamma/2 \sum_{j=2}^n \int_0^\infty \mathbb{E}[a_j b_{\pi(j)} | A_x^{ij}] \cdot \Pr[B_x^{ij}] \cdot \Pr[v_i = x] dx \\
&= (\delta - \delta^2/4) \mathbb{E}[a_1 v_i | o(i) = 1] \cdot \Pr[o(i) = 1] \\
&\quad - \delta \mathbb{E}[a_1 b_{\pi(1)} | o(i) = 1] \cdot \Pr[o(i) = 1] \\
&\quad + (\gamma/2 - \gamma^2/8) \sum_{j=2}^n \mathbb{E}[a_j v_i | o(i) = j] \cdot \Pr[o(i) = j] \\
&\quad - \gamma/2 \sum_{j=2}^n \mathbb{E}[a_j b_{\pi(j)} | o(i) = j] \cdot \Pr[o(i) = j] \\
&= (\gamma/2 - \gamma^2/8) \mathbb{E}[a_{o(i)} v_i] - \gamma/2 \cdot \mathbb{E}[a_{o(i)} b_{\pi(o(i))}] \\
&\quad + ((\delta - \delta^2/4) - (\gamma/2 - \gamma^2/8)) \\
&\quad \cdot \mathbb{E}[a_1 v_i | o(i) = 1] \cdot \Pr[o(i) = 1] \\
&\quad - (\delta - \gamma/2) \mathbb{E}[a_1 b_{\pi(1)} | o(i) = 1] \cdot \Pr[o(i) = 1],
\end{aligned}$$

where the second equality holds since

$$\begin{aligned}
&\int_0^\infty \mathbb{E}[Z | A_x^{ij}] \cdot \Pr[B_x^{ij}] \cdot \Pr[v_i = x] dx \\
&= \mathbb{E}[Z | o(i) = j] \cdot \Pr[o(i) = j],
\end{aligned}$$

for any non-negative random variable Z .

By summing over all bidders, we have

$$\begin{aligned}
&\sum_i \mathbb{E}[u_i(\mathbf{b})] \\
&\geq (\gamma/2 - \gamma^2/8) \sum_i \mathbb{E}[a_{o(i)} v_i] - \gamma/2 \sum_i \mathbb{E}[a_{o(i)} b_{\pi(o(i))}] \\
&\quad + ((\delta - \delta^2/4) - (\gamma/2 - \gamma^2/8)) \\
&\quad \cdot \sum_i \mathbb{E}[a_1 v_i | o(i) = 1] \cdot \Pr[o(i) = 1] \\
&\quad - (\delta - \gamma/2) \sum_i \mathbb{E}[a_1 b_{\pi(1)} | o(i) = 1] \cdot \Pr[o(i) = 1] \\
&= (\gamma/2 - \gamma^2/8) \text{OPT} - \gamma/2 \sum_i \mathbb{E}[a_i b_{\pi(i)}] \\
&\quad + (\delta - \delta^2/4 - \gamma/2 + \gamma^2/8) \mathbb{E}[a_1 v_{o^{-1}(1)}] \\
&\quad - (\delta - \gamma/2) \mathbb{E}[a_1 b_{\pi(1)}],
\end{aligned}$$

and the lemma follows. \square

The next lemma provides a second lower bound on the sum of expected utilities. Its proof (omitted due to space constraints) extends the arguments in [13]. Besides the use of a technical lemma in order to relate the quantities to parameter β , a subtle (and rather surprising) technical point is that the bound is obtained by ignoring possible gains the bidders may have when allocated to the slot with the highest click-through rate. This point is crucial in order to obtain our improved bound.

LEMMA 13. Consider a Bayes-Nash equilibrium. Then, for any $\beta > 0$, it holds that

$$\begin{aligned}
\sum_i \mathbb{E}[u_i(\mathbf{b})] &\geq \beta(1 - e^{-1/\beta}) \text{OPT} - \beta \sum_i \mathbb{E}[a_i b_{\pi(i)}] \\
&\quad - \beta(1 - e^{-1/\beta}) \mathbb{E}[a_1 v_{o^{-1}(1)}] + \beta \mathbb{E}[a_1 b_{\pi(1)}].
\end{aligned}$$

We are now ready to prove our main result concerning the price of anarchy for Bayes-Nash equilibria. The proof follows by appropriately taking into account the lower bounds on the sum of expected utilities that are proven in Lemmas 12 and 13.

THEOREM 14. The price of anarchy over Bayes-Nash equilibria of GSP auction games with conservative bidders is at most 3.037.

PROOF. Consider a Bayes-Nash equilibrium. We will use the fact that the social welfare is the sum of the expected utilities of the bidders plus the total bids paid. Let $\beta > 0$, $\mu \in [0, 1]$, and non-negative γ, δ and λ be parameters to be fixed later. We have

$$\begin{aligned}
&(1 + \lambda) \mathcal{W}(\mathbf{b}) \\
&= \mu \mathbb{E}[\sum_i u_i(\mathbf{b})] + (1 - \mu) \mathbb{E}[\sum_i u_i(\mathbf{b})] + \mathbb{E}[\sum_i a_i b_{\pi(i+1)}] \\
&\quad + \lambda \mathbb{E}[\sum_i a_i v_{\pi(i)}] \\
&\geq \mu \mathbb{E}[\sum_i u_i(\mathbf{b})] + (1 - \mu) \mathbb{E}[\sum_i u_i(\mathbf{b})] \\
&\quad + (1 + \lambda) \mathbb{E}[\sum_{i \geq 2} a_i b_{\pi(i)}] + \lambda \mathbb{E}[a_1 b_{\pi(1)}] \\
&\geq \mu \left(\beta(1 - e^{-1/\beta}) \text{OPT} - \beta \sum_i \mathbb{E}[a_i b_{\pi(i)}] \right. \\
&\quad \left. - \beta(1 - e^{-1/\beta}) \mathbb{E}[a_1 v_{o^{-1}(1)}] + \beta \mathbb{E}[a_1 b_{\pi(1)}] \right) \\
&\quad + (1 - \mu) \left(\left(\frac{\gamma}{2} - \frac{\gamma^2}{8} \right) \text{OPT} - \frac{\gamma}{2} \sum_i \mathbb{E}[a_i b_{\pi(i)}] \right. \\
&\quad \left. + \left(\delta - \frac{\delta^2}{4} - \frac{\gamma}{2} + \frac{\gamma^2}{8} \right) \mathbb{E}[a_1 v_{o^{-1}(1)}] \right. \\
&\quad \left. - \left(\delta - \frac{\gamma}{2} \right) \mathbb{E}[a_1 b_{\pi(1)}] \right) + (1 + \lambda) \mathbb{E}[\sum_{i \geq 2} a_i b_{\pi(i)}] \\
&\quad + \lambda \mathbb{E}[a_1 b_{\pi(1)}] \\
&= \left(\mu \beta(1 - e^{-1/\beta}) + (1 - \mu) \left(\frac{\gamma}{2} - \frac{\gamma^2}{8} \right) \right) \text{OPT} \\
&\quad + \left(\lambda + 1 - \mu \beta - (1 - \mu) \frac{\gamma}{2} \right) \sum_{i \geq 2} \mathbb{E}[a_i b_{\pi(i)}] \\
&\quad + \left((1 - \mu) \left(\delta - \frac{\delta^2}{4} - \frac{\gamma}{2} + \frac{\gamma^2}{8} \right) - \mu \beta(1 - e^{-1/\beta}) \right) \\
&\quad \cdot \mathbb{E}[a_1 v_{o^{-1}(1)}] + (\lambda - (1 - \mu) \delta) \mathbb{E}[a_1 b_{\pi(1)}].
\end{aligned}$$

The first inequality follows since $a_i \geq a_{i+1}$ for $i = 1, \dots, n-1$ and $v_{\pi(i)} \geq b_{\pi(i)}$, and the second one follows by applying Lemmas 12 and 13. By setting $\lambda = 0.76$, $\mu = 0.2$, $\beta = 5.05$, $\gamma = 1.86$, and $\delta = 0.95$, we have that the second, third and fourth term in the right-hand side of the above inequality are non-negative and the ratio $\text{OPT}/\mathcal{W}(\mathbf{b})$ is then bounded by 3.037 as desired. \square

6. OPEN PROBLEMS

Our work leaves several open problems. Still, there is a small gap between the upper and lower bounds on the price of anarchy over pure Nash equilibria. Even though we have found several weakly feasible assignments with efficiency higher than 1.259 in games with more than three bidders, none of them corresponds to a pure Nash equilibrium. So, it is interesting to prove or disprove whether the worst price of anarchy over pure Nash equilibria is obtained in GSP auction games with just three bidders. A proof of such a statement would require an explicit accounting of the bids and should not be based on weakly feasible assignments.

The case of more general equilibria is even more challenging. Here, there is no known lower bound besides the one for pure Nash equilibria. Computing tight bounds for mixed Nash, correlated, or coarse correlated equilibria are challenging open problems. Is the price of anarchy over equilibria in some of these classes worse than the price of anarchy over pure Nash equilibria? The same question applies to the Bayesian setting as well. The fact (observed in [16]) that GSP auction games are not smooth games according to the definition in [18] does not preclude a negative answer. On the positive side, our new upper bound for pure Nash equilibria might make the search for a game with a strictly worse non-pure Nash equilibrium easier.

7. REFERENCES

- [1] G. Aggarwal, A. Goel, and R. Motwani. Truthful auctions for pricing search keywords. In *Proceedings of the 7th ACM Conference on Electronic Commerce (EC)*, pages 1–7, 2006.
- [2] R. J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1):67–96, 1974.
- [3] K. Bhawalkar and T. Roughgarden. Welfare guarantees for combinatorial auctions with item bidding. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2011.
- [4] G. Christodoulou, A. Kovács, and M. Shapira. Bayesian combinatorial auctions. In *Proceedings of the 35th International Colloquium on Automata, Languages and Programming (ICALP)*, pages 820–832, 2008.
- [5] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertizing and the generalized second-price auction: selling billions of dollars worth of keywords. *The American Economic Review*, 97(1):242–259, 2007.
- [6] R. Gomes and K. Sweeney. Bayes-Nash equilibria of the generalized second price auction. In *Proceedings of the 10th ACM Conference on Electronic Commerce (EC)*, page 107, 2009.
- [7] J. C. Harsanyi. Games with incomplete information played by Bayesian players, parts i-iii. *Management Science*, 14:159–182, 330–334, 486–502, 1967-1968.
- [8] E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 404–413, 1999.
- [9] V. Krishna. *Auction Theory*. Elsevier Science, 2002.
- [10] S. Lahaie. An analysis of alternative slot auction designs for sponsored search. In *Proceedings of the 7th ACM Conference on Electronic Commerce (EC)*, pages 218–227, 2006.
- [11] S. Lahaie, D. Pennock, A. Saberi, and R. Vohra. Sponsored search auctions. In N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, editors, *Algorithmic Game Theory*, pages 699–716. 2007.
- [12] B. Lucier and A. Borodin. Price of anarchy for greedy auctions. In *Proceedings of the 21th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 537–553, 2010.
- [13] B. Lucier and R. Paes Leme. GSP auctions with correlated types. In *Proceedings of the 12th Annual ACM Conference on Electronic Commerce (EC)*, 2011. Preliminary version: Improved social welfare bounds for GSP at equilibrium. arXiv:1011.3268, 2010.
- [14] H. Moulin and J. P. Vial. Strategically zero-sum games: the class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory*, 7:201–221, 1978.
- [15] S. Muthukrishnan. Internet ad auctions: insights and directions. In *Proceedings of the 35th International Colloquium on Automata, Languages and Programming (ICALP)*, pages 14–23, 2008.
- [16] R. Paes Leme and E. Tardos. Pure and Bayes-Nash price of anarchy for generalized second price auctions. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 735–744, 2010.
- [17] C. H. Papadimitriou. Algorithms, games and the Internet. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 749–753, 2001.
- [18] T. Roughgarden. Intrinsic robustness of the price of anarchy. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC)*, pages 513–522, 2009.
- [19] D. R. M. Thompson and K. Leyton-Brown. Computational analysis of perfect-information position auctions. In *Proceedings of the 10th ACM Conference on Electronic Commerce (EC)*, pages 51–60, 2009.
- [20] H. Varian. Position auctions. *International Journal of Industrial Organization*, 25:1163–1178, 2007.
- [21] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, 1961.
- [22] H. P. Young. *Strategic Learning and its Limits*. Oxford University Press.