Pure and Bayes-Nash Price of Anarchy for Generalized Second Price Auction

Renato Paes Leme
Department of Computer Science
Cornell University, Ithaca, NY
renatoppl@cs.cornell.edu

Éva Tardos
Department of Computer Science
Cornell University, Ithaca, NY
eva.tardos@cornell.edu

***Abstract*—The Generalized Second Price Auction has been the main mechanism used by search companies to auction positions for advertisements on search pages. In this paper we study the social welfare of the Nash equilibria of this game in various models. In the full information setting, socially optimal Nash equilibria are known to exist (i.e., the Price of Stability is 1). This paper is the first to prove bounds on the price of anarchy, and to give any bounds in the Bayesian setting.**

Our main result is to show that the price of anarchy is small assuming that all bidders play un-dominated strategies. In the full information setting we prove a bound of 1.618 for the price of anarchy for pure Nash equilibria, and a bound of 4 for mixed Nash equilibria. We also prove a bound of 8 for the price of anarchy in the Bayesian setting, when valuations are drawn independently, and the valuation is known only to the bidder and only the distributions used are common knowledge.

Our proof exhibits a combinatorial structure of Nash equilibria and uses this structure to bound the price of anarchy. While establishing the structure is simple in the case of pure and mixed Nash equilibria, the extension to the Bayesian setting requires the use of novel combinatorial techniques that can be of independent interest.

***Keywords*-game theory; price of anarchy; GSP; Sponsored Search Auction**

I. Introduction

Search engines and other online information sources use Sponsored Search Auctions, or AdWord auctions, to monetize their services via advertisements sold. These auctions allocate advertisement slots to companies, and companies are charged per click, that is, they are charged a fee for any user that clicks on the link associated with the advertisement. There has been much work in understanding various aspect of the auctions used in this context, see the survey of Lahaie et al. [8].

Here we consider Sponsored Search Auctions in a game theoretic context: consider the game played by advertisers in bidding for an advertisement slot. For each search word, advertisers can bid for showing their ad next to the search results. There are multiple slots for advertisements and slots higher on the page are more valuable (clicked on by more users). The bids are used to determine both the assignment of bidders to slots, and the fees charged. In the simplest model, the bidders are assigned to slots in order of bids, and the fee for each click is the bid occupying the next slot. This auction is called the *Generalized Second Price Auction* (GSP). More generally, positions and payments in the Generalized Second Price Auction depend also on the click-through rates associated with the bidders, the probability that the advertisement will get clicked on by the users if assigned to the best slot. This is the version of the Generalized Second Price Auction mechanism adopted by all search companies. Here we will focus on the basic model for simplicity of presentation, but our results extend to the standard model of separable click-though rates (see the full version of our paper [11]).

The Generalized Second Price Auction is a simple and natural generalization of the Vickrey auction [15] for a single slot (or single item). The Vickrey auction [15] for a single item, and its generalization, the Vickrey-Clarke-Groves Mechanism (VCG) [2], [5], make truthful behavior (when the advertisers reveal their true valuation) a dominant strategy, and make the resulting outcome maximize the social welfare. However, the Generalized Second Price Auction is neither truthful nor maximizes social welfare. In this paper we will consider the social welfare of the GSP auction outcomes. Our goal is to show that the intuition based on the similarity of GSP to the Vickrey auction is not so far from truth: we prove that the social welfare is within a small constant factor of the optimal in any Nash equilibrium under the mild assumption that the players use un-dominated strategies.

We consider both full information games when player valuations are fixed, and also consider the Bayesian setting when the values are independent random variables, the valuation is known only to the bidder, and only the distributions used are public knowledge.

In the case of the full information game Edelman et al. [3] and Varian [14] show that there exists Nash equilibria that are socially optimal (for both our

Supported in part by NSF grants CCF-0729006.
Supported in part by NSF grants CCF-0910940 and CCF-0729006, ONR grant N00014-98-1-0589, and a Yahoo! Research Alliance Grant.

simple model and the case of separable click-through rates). But there are Nash equilibria where the social welfare is arbitrarily smaller than the optimum even for the special case of the single item Vickrey auction. However these equilibria are unnatural, as some bid exceeds the players valuations, and hence the player takes unnecessary risk. We show that bidding above the valuation is dominated strategy, and define conservative bidders as bidders who won't bid above their valuations. Our results assume that players are conservative.

Our results: The main results of this paper are Price of anarchy bounds for pure, mixed and Bayesian Nash equilibria for the GSP game assuming conservative bidders. To motivate the conservative assumption, we observe that bidding above the player's valuation is dominated strategy in all settings.

For each setting, we exhibit a combinatorial structure of the Nash equilibria that can be of independent interest. To state this structure we need the following notation. For an advertiser k let v_k be the value of advertiser k for a click (a random variable in the Bayesian case). For a slot i, let $\pi(i)$ be the advertiser assigned to slot i in an equilibrium (a random variable, in the case of mixed Nash, or in the Bayesian setting).

- For the case of full information game, the social welfare in a pure Nash equilibrium with conservative bidders is at most a factor of 1.618 above the optimum. We achieve this bound via a structural characterization of such equilibria: for any two slots i and j, we show that in a Nash equilibrium with conservative bidders, we must have that
$$\frac{\alpha_j}{\alpha_i} + \frac{v_{\pi(i)}}{v_{\pi(j)}} \geq 1.$$
It is not hard to see that this structure implies that the assignment cannot be too far from the optimal: if two advertisers are assigned to positions not in their order of values, then either (i) the two advertisers have similar values for a click; or (ii) the click-through rates of the two slots are not very different, and hence in either case their relative order doesn't affect the social welfare very much.
- We also bound the quality of mixed Nash equilibria as a warm-up for the Bayesian setting. For a mixed Nash equilibrium $\pi(i)$ is a random variable, indicating the bidder assigned to slot i, and similarly let the random variable $\sigma(i)$ denote the slot assigned to bidder i. For notational convenience we number players and slots in order of decreasing valuation and click-through rates respectively. By this notation, bidder i should be assigned to slot i in the optimum. The inequality for pure Nash equilibria is derived by thinking about a pair of bidders that are assigned to slots in reverse order. Such pairs seem hard to define in the mixed case. Instead, we will consider bidder i and his optimal slot i, and get the following condition for mixed Nash equilibria
$$\frac{\mathbb{E}\alpha_{\sigma(i)}}{\alpha_i} + \frac{\mathbb{E}v_{\pi(i)}}{v_i} \geq \frac{1}{2},$$
We use this inequality to show that the social welfare of a mixed Nash equilibrium is at least one-fourth of the optimal social welfare.
- We prove a bound of 8 on the price of anarchy for the Bayesian setting, where the valuations v_k are drawn independently at random. We do this via a slightly more complicated structural property, showing that an expression similar to the one used in the case of mixed Nash must be at least 1/4th in expectation. However, establishing this inequality in the Bayesian setting in much harder. In the context of pure and mixed Nash, the inequality follows from the Nash property by considering a single deviation by a player, e.g., a player who would be assigned to slot i in the optimum, may want to try to bid high enough to take over slot i. In contrast, in the Bayesian case we obtain our structural result by considering many different bids, and combine the inequalities established by these bids to show the structure.
In the process we use a number of new techniques of independent interest. The bids we use for player i are twice the expected value of the minimal bid that takes slot k conditioned both on the value v_i and the fact that the optimal position for bidder i is k. We show via an interesting combinatorial argument using the max-flow min-cut theorem, that these bids decrease with k. Then we use a novel averaging technique (using linear programming) to combine the resulting inequalities.

Our results differ significantly from the existing work on the price of anarchy in a number of ways. Many of the known results can be summarized via a smoothness argument as observed by Roughgarden [12]. In contrast, it is easy to see that the GSP game is not smooth in the sense of [12] (see the full version of this paper [11] for an example). Second, most known price of anarchy results are for the case of full information games. The full information setting makes the strong assumption that all advertisers are aware of the valuations of all other players. In contrast, the Bayesian setting requires only the much weaker assumption that valuations are drawn from independent distributions, and these distributions are known to all players. The Bayesian game is a better model for real AdWord Auctions, since players submit a single bid that will be used in many auctions with different competitors, so players are, in fact, optimizing for a distribution of other players.

Related work: Sponsored search has been an active area of research in the last several years. Mehta et al. [10] considered AdWord auctions in the algorithmic context. Since the original models, there has been much work in the area, see the survey of Lahaie et al. [8]. Here we use the game theoretic model of the AdWord auctions of Edelman et al. [3] and Varian [14].

In the full information setting Edelman et al. [3] and Varian [14] show that the price of Stability for this game is l. More precisely, they consider a restricted class of Nash equilibria called Envy-free equilibria or Symmetric Nash Equilibria, and show that such equilibria exists, and all such equilibria are socially optimal. In this class of equilibria, an advertiser wouldn't be better off after switching his bids with the advertiser just above him. Note that this is a stronger requirement than Nash, as an advertiser cannot unilaterally switch to a position with higher click-through rate by simply increasing his bid. Edelman et al. [3] claim that if the bids eventually converge, they will converge to an envy-free equilibrium; otherwise some advertiser could increase his bid making the slot just above more expensive and therefore making the advertiser occupying it underbid him. They do not provide a formal game theoretical model that selects such equilibria.

Gomes and Sweeney [4] study the Generalized Second Price Auction in the Bayesian context. They show that, unlike the full information case, there may not exist symmetric or socially optimal equilibria in this model, and obtain sufficient conditions on click-through rates that guarantee the existence of a symmetric and efficient equilibrium.

Lahaie [7] considers the problem of quantifying the social efficiency of an equilibrium. He makes the strong assumption that the click-through rates α_i decay exponentially along the slots with a factor of $\frac{1}{\delta}$, and proves a price of anarchy bound of $\min\{\frac{1}{\delta}, 1 - \frac{1}{\delta}\}$. We make no assumptions on the click-through rates. Thompson and Leyton-Brown [13] study the efficiency loss of equilibria empirically in various models.

We assume that bidders are conservative, in the sense that no bidder is bidding above their own valuation. We can justify this assumption by noting that bidding above the valuation is a dominated strategy. Lucier and Borodin [9] and Christodoulou at al. [1] also use the conservative assumption to establish price-of-anarchy results in the context of combinatorial auctions.

The paper by Lucier and Borodin [9] on greedy auctions is also closely related to our work. They analyze the Price of Anarchy of the auction game induced by a Greedy Algorithm. They consider a general combinatorial auction setting, where a greedy algorithm is used for determining the allocation with payments computed using the critical price. They show via a type of smoothness argument (see [12]) that if the greedy algorithm is a c-approximation algorithm, then the Price of Anarchy of the resulting mechanism is $c + 1$ - for pure and mixed Nash and for Bayes-Nash equilibria. The Generalized Second Price mechanism is a type of greedy mechanism, but is not a combinatorial auction, and hence it does not fit the framework of Lucier and Borodin. The key to proving the $c + 1$ bound of Lucier and Borodin [9] is to consider possible deviating bids, such as a single minded bid for the slot in the optimal solution, or modifying a bid by changing it only on a single slot (the one allocated in the optimal solution). The combinatorial auction framework allows such complex bids; in contrast, the bids in GSP have limited expressivity, since a bid is a single number, and hence bidders cannot make single-minded declarations for a certain slot, or modify their bid only on one of the slots. Like the GSP game, many natural bidding languages have limited expressivity, since typically allowing arbitrary complex bids makes the optimization problem hard. The limited expressivity of the bidding language can increase the set of Nash equilibria (since there are fewer deviating bids to consider), so it is important to understand if such natural bidding languages result in greatly increased price of anarchy.

II. Preliminaries

We consider an auction with n advertisers and n slots (if there are fewer slots, add virtual slots with click-through rate zero). We model this auction as a game with n players, where each advertiser is one player. In the simple model the type of the advertisers is given by their valuation v_i, their value for one click. The strategy for each advertiser is a bid $b_i \in [0, \infty)$ which expresses the maximum he or she is willing to pay for a click.

The auction decides where to allocate each advertiser based on the bids. In the simple model, being assigned to the k-th slot results in α_k clicks and α_k is a monotone non-increasing sequence, i.e., $\alpha_1 \geq \alpha_2 \geq \ldots \geq \alpha_n$. The simple game proceeds as follows:

1) each advertiser submits a bid $b_i \geq 0$, which is the maximum he is willing to pay for a click
2) the advertisers are sorted by their bids (ties are broken arbitrarily). Call $\pi(k)$ the advertiser with the k-th highest bid
3) advertiser $\pi(k)$ is placed on slot k and therefore received α_k clicks
4) for each click, advertiser k pays $b_{\pi(k+1)}$, which is the next highest bid

The vector π is a permutation that indicates to which slot each player is assigned - it is determined by the set of bids (up to ties). We define the *utility* of a user i when occupying slot j as given by $u_i(b) = \alpha_j(v_i - b_{\pi(j+1)})$. We define the *social welfare* of this game as the total value that the bidders and the auctioneer get from playing it, which is: $\sum_j \alpha_j v_{\pi(j)}$. The goal of this paper is to bound the social welfare of the equilibria relative to the optimum. This measure is called the Price of Anarchy. We analyze the Price of Anarchy in three different settings of increasing complexity.

In the full version of our paper [11] we extend the results to the more general model of separable click-through rates, where the probability of clicking on an advertisement i displayed in slot j is $\alpha_j \gamma_i$. Now advertisers are assigned in order of the products $b_i \gamma_i$ (the expected total willingness of the bidder to pay: γ_i clicks at the rate of b_i), and the fee for a click is the critical value of the bid needed to keep the advertiser in his current slot. We get the simple model as a special case by assuming that $\gamma_i = 1$ for all bidders.

Pure Nash equilibrium: The valuation of each player is a fixed value v_i. We number the bidders (without loss of generality) so that $v_1 \geq v_2 \geq \ldots \geq v_n$. Each player chooses a pure strategy, i.e., a deterministic bid b_i. The bids $b = (b_1, \ldots, b_n)$ is a *Pure Nash Equilibrium* if no bidder can change his bid to increase his utility, i.e.:

$$u_i(b_i, b_{-i}) \geq u_i(b'_i, b_{-i}), \forall b'_i \in [0, \infty)$$

where b_{-i} denotes the vector of bids for bidders $j \neq i$.

To gain some intuition, suppose advertiser i is currently bidding b_i and occupying slot j. Changing his bid to something between $b_{\pi(j-1)}$ and $b_{\pi(j+1)}$ won't change the permutation π nor his payment. So, he could try to increase his utility by doing one of two things:

- increasing his bid to get a slot with a better click-through rate. If he wants to get a slot $k < j$ he needs to overbid advertiser $\pi(k)$, say by bidding $b_{\pi(k)} + \epsilon$. This way he gets slot k for the price $b_{\pi(k)}$ per click, getting utility $\alpha_k(v_i - b_{\pi(k)})$.
- decreasing his bid to get a worse but cheaper slot. If he wants to get slot $k > j$ he needs to bid below advertiser $\pi(k)$. This way he would get slot k for the price $b_{\pi(k+1)}$ per click, getting utility $\alpha_k(v_i - b_{\pi(k+1)})$.

Note the asymmetry between the two options. The symmetric (or envy free) equilibria studied by Edelman et al. [3] and Varian [14] satisfy the stronger symmetric condition that $\alpha_j(v_i - b_{\pi(j+1)}) \geq \alpha_k(v_i - b_{\pi(k+1)})$ for all k. Edelman et al. [3] and Varian [14] show that symmetric equilibria exist and have optimal welfare, hence the Price of Stability for this game is 1.

We are interested in bounding the *Pure Price of Anarchy*, which is the ratio $\sum_j \alpha_j v_j / \sum_j \alpha_j v_{\pi(j)}$, between the social welfare in the optimum and in the worst Nash equilibrium.

Mixed Nash equilibrium: The valuations v_i are still fixed and we can assume (without loss of generality) that $v_1 \geq \ldots \geq v_n$, but players pick a distribution over strategies. In a Mixed Nash equilibrium, each player chooses a random variable b_i for his bid such that the chosen random variable maximizes the expected utility for each player. In other words:

$$\mathbb{E}[u_i(b_i, b_{-i})] \geq \mathbb{E}[u_i(b'_i, b_{-i})], \forall b'_i, \forall i$$

where expectation is with respect to the distribution of bids. Now, the assignment π is a random variable determined by b and therefore the social welfare is also a random variable (even though the optimal welfare is fixed). The Price of Anarchy is the ratio: $\sum_j \alpha_j v_j / \mathbb{E}[\sum_j \alpha_j v_{\pi(j)}]$.

Bayes-Nash equilibrium: The partial information setting, using the framework of Harsanyi [6], provides a more realistic setting than the full information game. In this model the valuations v_i are drawn from independent distributions. The distributions used are common knowledge, but only player i is aware of his valuation v_i. (No assumptions are made about the distributions beyond independence). Now the strategy of a player i is to choose a bid (possibly at random) based on his own valuation v_i. Therefore, the strategy of player i is a bidding function $b_i(v_i)$ that associates for each valuation v_i a distribution of bids. A set of bidding functions is a *Bayes-Nash equilibrium* if for all $i, v_i, b'_i(v_i)$:

$$\mathbb{E}[u_i(b_i(v_i), b_{-i}(v_{-i}))|v_i] \geq \mathbb{E}[u_i(b'_i(v_i), b_{-i}(v_{-i})|v_i]$$

where expectations are taken over values and randomness used by players.

The Nash assignment π is a random variable, since it is dependent on the bids, which are random. The optimal allocation is also a random variable, let $\nu(k)$ be the slot occupied by player k in the optimal assignment. Therefore, ν is a random variable such that $v_i > v_j \Rightarrow \nu(i) < \nu(j)$. The optimal social welfare is therefore $\sum_j \alpha_{\nu(j)} v_j$. In this setting the quantity we want to bound is the *Bayes-Nash price of Anarchy*, which we define as the ratio: $\mathbb{E}[\sum_j \alpha_{\nu(j)} v_j] / \mathbb{E}[\sum_j \alpha_j v_{\pi(j)}]$.

A. *Equilibria with Low Social Welfare and Conservative Bidders*

Even for two slots the gap between the best and the worse Nash equilibrium can be arbitrarily large. For example, consider two slots with click-through-rates $\alpha_1 = 1$ and $\alpha_2 = 0$ and two advertisers with valuations $v_1 = 1$ and $v_2 = 0$. It is easy to check that the bids $b_1 = 0$ and $b_2 = 1$ are a Nash equilibrium where advertiser 1 gets the second slot and advertiser 2 gets the first slot. The social welfare in this equilibrium is 0 while the optimum is 1. The price of anarchy is therefore unbounded.

Notice, however, that this Nash equilibrium seems very artificial: the special case of GSP with $\alpha_1 = 1$ and $\alpha_2 = 0$ is the Vickrey auction, where truthful bidding of $b_i = v_i$ is a dominant strategy, yet in the equilibrium above the bids are not truthful. Advertiser 2 above is exposed to the risk of negative utility with no benefit: if advertiser 1 (or a new advertiser) adds a bid somewhere in the interval $(0, 1)$ this imposes a negative utility on advertiser 2.

More generally, for any bidder i, bidding above the valuation v_i (with any probability) is dominated by bidding v_i in any of the above models. We state the lemma here in the more general model of Bayesian games.

Lemma II.1 *Given a bidding function $b_i(v_i)$, a strategy in which $P(b_i(v_i) > v_i) > 0$ for some v_i is dominated by playing $b'_i(v_i) = \min\{v_i, b_i(v_i)\}$.*

We say that a player is *conservative* if he doesn't overbid, i.e., $P(b_i(v_i) \leq v_i) = 1$. We assume throughout the paper that players are conservative.

III. Pure Nash Equilibrium

Theorem III.1 *For 2 slots, if all advertisers are conservative, then the price of anarchy is exactly 1.25.*

Proof: To see that the price of anarchy is achievable consider two slots with $\alpha_1 = 1$ and $\alpha_2 = 1/2$, and two bidders with valuations $v_1 = 1$ and $v_2 = 1/2$, and note that the bids $b_1 = 0$ and $b_2 = 1/2$ form a Nash equilibrium. It is not hard to see that this is the worst case. See the full version [11] for more details. ∎

A. Weakly Feasible Assignments

Next we show that equilibria with conservative bidders satisfy the property mentioned in the introduction. We will call the assignments satisfying this property weakly feasible. In the next subsection we analyze the welfare properties of weakly feasible assignments.

Lemma III.2 *For any valuation v, click-through rates α and a Nash permutation π we have*

$$\frac{\alpha_j}{\alpha_i} + \frac{v_{\pi(i)}}{v_{\pi(j)}} \geq 1; \qquad (1)$$

in particular, $\frac{\alpha_j}{\alpha_i} \geq \frac{1}{2}$ or $\frac{v_{\pi(i)}}{v_{\pi(j)}} \geq \frac{1}{2}$.

Proof: If $j \leq i$ the inequality is obviously true. Otherwise consider the bidder $\pi(j)$ in slot j. Since it is a Nash equilibrium, the bidder in slot j is happy with his outcome and doesn't want to increase his bid to take slot i, so: $\alpha_j(v_{\pi(j)} - b_{\pi(j+1)}) \geq \alpha_i(v_{\pi(j)} - b_{\pi(i)})$ since $b_{\pi(j+1)} \geq 0$ and $b_{\pi(i)} \leq v_{\pi(i)}$ then: $\alpha_j v_{\pi(j)} \geq \alpha_i(v_{\pi(j)} - v_{\pi(i)})$ ∎

Inspired by the last lemma, given parameters α, v we say that permutation π is *weakly feasible* if inequality (1) holds for each i, j. The main result of this section follow from analyzing the price of anarchy ratio $\sum_j \alpha_j v_j / \sum_j \alpha_j v_{\pi(j)}$ over all weakly feasible permutations π.

B. Price of Anarchy Bound

Here we present the bound on the price of anarchy for weakly feasible permutations, and hence for GSP for conservative bidders.

Theorem III.3 *For conservative bidders, the price of anarchy for pure Nash equilibria of GSP is bounded by the golden ratio $\frac{1+\sqrt{5}}{2} \approx 1.618$.*

Proof: We will prove the bound by induction for all weakly feasible permutations. As a warm-up we will prove that the price of anarchy is bounded by 2, since the proof is easier and captures the main ideas.

We use induction on n. The case $n = 1$ is obvious. Consider parameters v, α and a weakly feasible permutation π. Let $i = \pi^{-1}(1)$ be the slot occupied by the advertiser with maximum valuation and $j = \pi(1)$ be the advertiser occupying the first slot. If $i = j = 1$ then we can apply the induction hypothesis right away. If not, inequality (1) tells us that $\frac{\alpha_i}{\alpha_1} \geq \frac{1}{2}$ or $\frac{v_j}{v_1} \geq \frac{1}{2}$. Suppose $\frac{\alpha_i}{\alpha_1} \geq \frac{1}{2}$ and consider an input with slot i and advertiser 1 deleted. The permutation π restricted to these $n - 1$ advertisers and $n - 1$ slots is still weakly feasible, so by the induction hypothesis:

$$\begin{aligned}\sum_{k \neq i} \alpha_k v_{\pi(k)} \\ &\geq \frac{1}{2}(\alpha_1 v_2 + ... + \alpha_{i-1} v_i + \alpha_{i+1} v_{i+1} + ... + \alpha_n v_n) \\ &\geq \frac{1}{2}(\alpha_2 v_2 + ... + \alpha_i v_i + \alpha_{i+1} v_{i+1} + ... + \alpha_n v_n)\end{aligned}$$

and therefore,

$$\sum_k \alpha_k v_{\pi(k)} = \alpha_i v_1 + \sum_{k \neq i} \alpha_k v_{\pi(k)} \geq \frac{1}{2}\alpha_1 v_1 + \frac{1}{2}\sum_{k>1} \alpha_k v_k$$

If $\frac{v_j}{v_1} \geq \frac{1}{2}$ we just do the same but deleting slot 1 and advertiser j from the input. This proves the bound of 2.

Next we sketch the proof of the improved bound. See the full version [11] for more details. As before, we prove the conclusion for all weakly feasible permutations. Let r_k be the worst price of anarchy for feasible permutations in a k slots auction. By the proof of Theorem III.1 we know that $r_2 = 1.25$. We will generate a recursion to bound r_k and then prove that the bound converges to the desired bound of $\frac{1+\sqrt{5}}{2}$.

Consider parameter α, v, a weakly feasible permutation π and let's assume $i = \pi^{-1}(1)$ and $j = \pi(1)$. If $i = j = 1$, the price of anarchy is bounded by r_{n-1}. If not, assume without loss of generality that $i \leq j$ (since inequality (1) is symmetric in α and v). Let $\beta = \frac{\alpha_1}{\alpha_i}$ and $\gamma = \frac{v_1}{v_j}$. We know that $\frac{1}{\beta} + \frac{1}{\gamma} \geq 1$. Following the outline of the previous proof we have:

$$\sum_k \alpha_k v_{\pi(k)} = \alpha_i v_1 + \sum_{k \neq i} \alpha_k v_{\pi(k)}$$

The first term is bounded by $\frac{1}{\beta}\alpha_1 v_1$. We bound the remaining terms as

$$\begin{aligned}\sum_{k \neq i} \alpha_k v_{\pi(k)} &\geq \frac{1}{r_{n-1}} \left(\sum_{k=2}^{i} \alpha_{k-1} v_k + \sum_{k=i+1}^{n} \alpha_k v_k \right) \\ &= \frac{1}{r_{n-1}} \left[\sum_{k=2}^{i} (\alpha_{k-1} - \alpha_k) v_k + \sum_{k>1} \alpha_k v_k \right] \\ &\geq \frac{1}{r_{n-1}} (\alpha_1 - \alpha_i) v_i + \frac{1}{r_{n-1}} \sum_{k>1} \alpha_k v_k\end{aligned}$$

By the assumption that $i \leq j$ we have $v_i \geq v_j = \frac{1}{\gamma} v_1 \geq \left(1 - \frac{1}{\beta}\right) v_1$, and we get

$$\sum_k \alpha_k v_{\pi(k)} \geq \left[\frac{1}{\beta} + \frac{1}{r_{n-1}} \left(1 - \frac{1}{\beta}\right)^2 \right] \alpha_1 v_1 + \\ + \frac{1}{r_{n-1}} \sum_{k>1} \alpha_k v_k$$

Symmetrically, we can remove slot 1 and advertiser j in the inductive step and get a similar equation. The bound for r_n is the maximum of the two. Finally, to get bound for r_n valid for all β we need to use the value of β that minimizes the resulting bound. We get the following recursion for r_n

$$r_n = \begin{cases} \left(1 - \frac{r_{n-1}}{4}\right)^{-1} & , r_{n-1} < \frac{4}{3} \\ \left(r_{n-1} - \sqrt{r_{n-1}^2 - r_{n-1}}\right)^{-1} & , r_{n-1} \geq \frac{4}{3} \end{cases}$$

To show that the sequence is bounded by $\varphi = \frac{1+\sqrt{5}}{2}$, note that if $r_{n-1} \leq \varphi$ then $r_n \leq \varphi$.

■

Remark: Proving matching upper and lower bounds for this problem remains an interesting open problem. The worse example of the Price of Anarchy the authors are aware of (in any of the models) is 1.259 (and it is for a pure Nash equilibrium for 3 players).

IV. MIXED NASH EQUILIBRIUM

As before, we assume that players are numbered such that $v_1 \geq \ldots \geq v_n$ and slots with click-through rates $\alpha_1 \geq \ldots \geq \alpha_n$. In a mixed Nash equilibrium the strategy of player i is a probability distribution on $[0, v_i]$ represented by a random variable b_i, and we assume that bidders are conservative: $P(b_i \leq v_i) = 1$.

Now the allocation, represented by the permutation π, is also a random variable. For notational convenience, let $\sigma = \pi^{-1}$. We begin by proving a bound similar to Lemma III.2 for mixed Nash and then using that to bound the price of anarchy. Note that by our notational assumption the position of bidder i in the optimal allocation is position i. The new inequality is different as it involves a bidder i and its location i in the optimal allocation, rather than two bidders that are allocated to "wrong relative positions".

Lemma IV.1 *If the random vector b is a mixed Nash equilibrium for GSP then for each player i:*

$$\frac{\mathbb{E}\alpha_{\sigma(i)}}{\alpha_i} + \frac{\mathbb{E}v_{\pi(i)}}{v_i} \geq \frac{1}{2} \tag{2}$$

Proof: We will consider whether player i benefits by deviating to the deterministic $b'_i = \min(v_i, 2\mathbb{E}b_{\pi(i)})$.

We claim that with probability at least $\frac{1}{2}$, this bid gets one of the slots of $\{1, \ldots, i\}$. If $b'_i = v_i$ then this happens for sure, as our conservative assumption guarantees that only the previous $i-1$ players can bid more. If $b'_i = 2\mathbb{E}b_{\pi(i)}$ then by Markov's inequality $P(b_{\pi(i)} \geq b'_i) \leq \frac{\mathbb{E}b_{\pi(i)}}{b'_i} = \frac{1}{2}$. Therefore we have that

$$\mathbb{E}\alpha_{\sigma(i)} v_i \geq \mathbb{E}u_i(b) \geq \mathbb{E}u_i(b'_i, b_{-i}) \geq \frac{1}{2}\alpha_i(v_i - b'_i) \geq \\ \geq \frac{1}{2}\alpha_i(v_i - 2\mathbb{E}b_{\pi(i)}) \geq \frac{1}{2}\alpha_i(v_i - 2\mathbb{E}v_{\pi(i)})$$

Now it is just a matter of rearranging the expression. ■

Theorem IV.2 *The Price of Anarchy for the mixed Nash equilibria of GSP with conservative bidders is at most 4.*

Proof: The proof is a simple application of Lemma IV.1 and some algebraic manipulation:

$$\mathbb{E}[\sum_i u_i(b)] = \frac{1}{2}\left[\mathbb{E}\sum_i \alpha_{\sigma(i)} v_i + \mathbb{E}\sum_i \alpha_i v_{\pi(i)}\right] = \\ = \frac{1}{2}\sum_i \alpha_i v_i \left(\frac{\mathbb{E}\alpha_{\sigma(i)}}{\alpha_i} + \frac{\mathbb{E}v_{\pi(i)}}{v_i}\right) \geq \frac{1}{4}\sum_i \alpha_i v_i$$

■

V. BAYES-NASH EQUILIBRIUM

Recall that in the Bayesian setting, the values v_i are independent random variables, their distributions are common knowledge, but the value v_i is only known to bidder i. A strategy for a player i is a bidding function $b_i(v_i)$ (or a probability distribution of such functions) where $b_i(v_i)$ is the player's bid when his value is v_i. As before, we will assume that $P(b_i(v_i) \leq v_i) = 1$, since overbidding is dominated strategy.

We will use π and $\sigma = \pi^{-1}$ to denote the permutation representing the allocation, and we will use ν to denote the random permutation (defined by v) such that player i occupies slot $\nu(i)$ in the optimal solution. The expected social welfare is $\mathbb{E}[\sum_i \alpha_i v_{\pi(i)}] = \mathbb{E}[\sum_i \alpha_{\sigma(i)} v_i]$ and the social optimum is given by $\mathbb{E}[\sum_i \alpha_{\nu(i)} v_i]$. The goal of this section is to bound the price of anarchy, the ratio of these two expectations.

Theorem V.1 *If a set of functions $b_1, \dots, b_n$ are a Bayes-Nash equilibrium in conservative strategies then:*

$$\mathbb{E}\left[\sum_i \alpha_i v_{\pi(i)}\right] \geq \frac{1}{8}\mathbb{E}\left[\sum_i \alpha_{\nu(i)} v_i\right]$$

that is, the Bayes-Nash Price of Anarchy in conservative strategies for GSP is bounded by 8*.*

The proof of the theorem is based on a structural characterization analogous to the one used for Pure and Mixed Nash equilibria in previous sections, but much harder to prove. The structural characterization for Mixed Nash (Lemma IV.1) can be written as $v_i \mathbb{E}\alpha_{\sigma(i)} + \alpha_i \mathbb{E} v_{\pi(i)} \geq \frac{1}{2}\alpha_i v_i$. The Bayesian structural characterization is obtained by taking expectation of this inequality (and losing a factor of 2). In the full information model, bidder i is assigned to slot i in the optimum by notation, and the inequality above uses this notational convenience. In the Bayesian setting, the optimal slot for a bidder is a random variable, so we cannot deterministically order bidders by valuation; instead we need to use a random variable $\nu(i)$ to denote the slot bidder i is assigned to in the optimum.

Lemma V.2 *If $\{b_i(\cdot)\}_i$ is a Bayes-Nash equilibrium of the GSP then for all i and for all v_i:*

$$v_i \mathbb{E}[\alpha_{\sigma(i)}|v_i] + \mathbb{E}[\alpha_{\nu(i)} v_{\pi(\nu(i))}|v_i] \geq \frac{1}{4} v_i \mathbb{E}[\alpha_{\nu(i)}|v_i]$$

The price of anarchy bound follows from the lemma.

Proof of Theorem V.1 :

$$\begin{aligned}
SW &= \frac{1}{2}\mathbb{E}\sum_i (\alpha_i v_{\pi(i)} + \alpha_{\sigma(i)} v_i) = \\
&= \frac{1}{2}\mathbb{E}\sum_i (\alpha_{\nu(i)} v_{\pi(\nu(i))} + \alpha_{\sigma(i)} v_i) = \\
&= \frac{1}{2}\mathbb{E}\left[\sum_i \mathbb{E}[\alpha_{\nu(i)} v_{\pi(\nu(i))}|v_i] + v_i \mathbb{E}[\alpha_{\sigma(i)}|v_i]\right] \geq \\
&\geq \frac{1}{8}\mathbb{E}\left[\sum_i v_i \alpha_{\nu(i)}\right]
\end{aligned}$$

■

The hard part of the proof is proving Lemma V.2. The main difficulty in the Bayesian setting is that the inequality is not established by a single deviating bid. The structural inequalities of Lemmas III.2 and IV.1 in the full information setting were obtained by considering a single deviation, e.g., for mixed Nash equilibria we considered a single bid just above $2\mathbb{E}b_{\pi(\nu(i))}$, as by Markov's inequality this value is above $b_{\pi(\nu(i))}$ with probability at least 1/2. In contrast, in the Bayesian setting, we obtain our structural result by considering deviations to different bids and then combining them using a novel averaging argument.

To define the deviating bids, consider the following notation: let $\pi^i(k)$ be the bidder occupying slot k in the case i didn't participate in the auction, i.e., $\pi^i(k) = \pi(k)$ if $\sigma(i) > \sigma(k)$ and $\pi^i(k) = \pi(k+1)$ otherwise. Note the following property of $\pi^i(k)$

Lemma V.3 *A deviating bid B by player i gets a slot k or above if and only if $B > b_{\pi^i(k)}$.*

For mixed equilibria in the full information setting, we considered the bid $2\mathbb{E}b_{\pi(\nu(i))}$. To extend this to the Bayesian setting, we will consider a sequence of bids, conditioned on the value of $\nu(i)$ defined as

$$B_k = \min\{v_i, 2\mathbb{E}[b_{\pi^i(k)}|v_i; \nu(i) = k]\}.$$

Notice that B_k is defined as a conditional expectation, so it is a function of v_i, and not a constant. We will drop the dependence on v_i from the notation as we are focusing on a single value v_i throughout the proof.

The proof of Lemma V.2, depends on two combinatorial results. The first is a structural property: we claim that the bids B_k are monotone in k for any fixed value of v_i. Showing this will allow us to argue that bid B_k not only has a good chance of taking slot k when $\nu(i) = k$, but also has a good chance of taking any other slot $k' > k$ when $\nu(i) = k'$, since $B_k \geq B_{k'}$.

Lemma V.4 *The expectation $\mathbb{E}[b_{\pi^i \nu(i)}|v_i, \nu(i) = k]$ is non-increasing in k for any fixed value v_i.*

We will prove the lemma above using flows and the max-flow min-cut theorem. The value B_k is defined as a conditional expectation assuming $\nu(i) = k$, while B_{k+1} is defined as a conditional expectation conditioning on a disjoint part of the probability space: assuming $\nu(i) = k+1$. To relate the two expectations we define a flow of probabilities from the probability space where $\nu(i) = k$ to the space where $\nu(i) = k+1$ that transfers the mass of probability with the property that the value $b_{\pi^i(\nu(i))}$ is non-increasing along the flow lines. This will prove that B_k, the expectation of $b_{\pi^i(\nu(i))}$ on the source side, is no bigger than B_{k+1}, the expectation of the same value on the sink side.

We combine the inequalities obtained by considering the different bids B_k using a novel "dual averaging argument", finding an average that will simultaneously guarantee that one average is not too low, and a different average is not to high. We combine the bids B_k via a probability distribution x (bidding B_k with probability x_k). The two inequalities of the lemma will guarantee that the resulting randomized bid, on one hand, gets a high enough number of clicks, and on the other hand, the resulting payment is not too large.

Lemma V.5 *Given any nonnegative values γ_k, B_k there is a probability distribution $x_k \geq 0$, $\sum_k x_k = 1$ such that*

$$\sum_k x_k \sum_{j=k}^{n} \gamma_j \geq \frac{1}{2} \sum_{j=1}^{n} \gamma_j$$

$$\sum_k x_k B_k \sum_{j=k}^{n} \gamma_j \leq \sum_{j=1}^{n} \gamma_j B_j$$

Before we prove these key lemmas, we show how to use them for proving the main Lemma V.2:

Proof of Lemma V.2 : As outlined above, we will consider n deviations for a player i at bids B_k for all possible slots k. Since the bidding functions are a Nash equilibrium, player i can't benefit from changing his strategy, and so each deviation will give us an inequality on the utility of player i. We will use Lemma V.5 to average the inequalities and get the claimed inequality.

Suppose bidder i deviates to $B_k = \min\{v_i, 2\mathbb{E}[b_{\pi^i(k)}|v_i; \nu(i) = k]\}$. Let α'_k be the random variable that means the click-through rate of the slot he occupies by bidding B_k. First we estimate the probability that by bidding B_k the player gets the slot k or better when $\nu(i) = k$. In the case $B_k = v_i$ this is trivially guaranteed, since only $\nu(i) - 1$ players have values above v_i and only these players can bid above v_i. If $B_k = 2\mathbb{E}[b_{\pi^i k}|v_i; \nu(i) = k]$, we use Lemma V.3, and Markov's inequality to get:

$$P(\alpha'_k \geq \alpha_k | v_i, \nu(i) = k) =$$
$$= P(B_k \geq b_{\pi^i(k)} | v_i, \nu(i) = k) \geq \frac{1}{2}.$$

Let $p_j = P(\nu(i) = j|v_i)$. Recall that by Lemma V.4 we have that $B_1 \geq B_2 \geq \ldots \geq B_n$, and hence the probability of bid B_k taking a slot j or better when $\nu(i) = j$ is also at least $1/2$ whenever $j \geq k$. The expected value of bidding B_k is at least $\mathbb{E}[\alpha'_k(v_i - B_k)|v_i]$, and the value for player i in the current solution is at most $v_i\mathbb{E}[\alpha_{\sigma(i)}|v_i]$. This leads to the following inequality.

$$v_i\mathbb{E}[\alpha_{\sigma(i)}|v_i] \geq \mathbb{E}[\alpha'_k(v_i - B_k)|v_i] =$$
$$= \sum_j p_j \mathbb{E}[\alpha'_k(v_i - B_k)|v_i, \nu(i) = j] \geq$$
$$\geq \sum_{j \geq k} \frac{1}{2} p_j \alpha_j (v_i - B_k).$$

Now we use the Lemma V.5 applied with B_k and $\gamma_k = p_k\alpha_k$. We can interpret x_k from the lemma as probabilities, and consider the deviating strategy of bidding B_k with probability x_k.

Combining the above inequalities with the coefficients x_k from the Lemma, we get that

$$v_i\mathbb{E}[\alpha_{\sigma(i)}|v_i] \geq \sum_k x_k \sum_{j \geq k} \frac{1}{2} p_j \alpha_j (v_i - B_k) \geq$$
$$\geq \frac{1}{4} v_i \sum_j \alpha_j p_j - \frac{1}{2} \sum_j p_j \alpha_j B_j \geq$$
$$\geq \frac{1}{4} v_i \mathbb{E}[\alpha_{\nu(i)}|v_i] - \mathbb{E}[\alpha_{\nu(i)} b_{\pi^i(\nu(i))}|v_i].$$

To get the claimed inequality, note that $b_{\pi^i(k)} \leq b_{\pi(k)} \leq v_{\pi(k)}$. ■

A. Proving that bids B_k are non-increasing

We will prove Lemma V.4 in several steps. First we prove bounds assuming all but a single player has a deterministic value, and we take expectations to get a conditional version. We define a probability flow from the probability space where $\nu(i) = k$ to the space where $\nu(i) = k + 1$ that transfers the mass of probability so that only a single value is changing along the flow edges, and hence by the first claim the value $b_{\pi^i(\nu(i))}$ is non-increasing along the flow lines. In transferring the probability mass we take advantage of the fact that the valuations are drawn from independent distributions.

Proof of Lemma V.4 : We want to prove that

$$\mathbb{E}[b_{\pi^i(k)}|v_i, \nu(i) = k] \geq \mathbb{E}[b_{\pi^i(k+1)}|v_i, \nu(i) = k + 1].$$

The value v_i is in position k in the optimum if exactly $n - k$ values are below v_i. Consider such a set S of agents, $i \notin S$, and the corresponding event:

$$A_S = \{v_j \leq v_i; \forall j \in S, v_j > v_i; \forall j \notin S\}.$$

The event $\nu(i) = k$ can now be stated as $\cup_{|S|=n-k} A_S$, and so what we are trying to prove is:

$$\mathbb{E}[b_{\pi^i(k)}|v_i, \cup_{|S|=n-k} A_S] \geq \mathbb{E}[b_{\pi^i(k+1)}|v_i, \cup_{|S'|=n-k-1} A_{S'}]$$

Consider a pair of sets $S' \subseteq S$, i.e., $S = S' \cup \{t\}$ for some agent $t \neq i$. The first claim is the following.

Claim V.6 *For sets S' and $S = S' \cup \{t\}$ for $t \neq i$,*

$$\mathbb{E}[b_{\pi^i(k)}|v_i, A_S] \geq \mathbb{E}[b_{\pi^i(k+1)}|v_i, A_{S'}]$$

To see this, notice that

$$\mathbb{E}[b_{\pi^i(k)}|v_i, A_S, \{v_j\}_{j \neq i,t}] \geq \mathbb{E}[b_{\pi^i(k+1)}|v_i, A_{S'}, \{v_j\}_{j \neq i,t}]$$

The conditioning on the two sides differs only by the value of bidder t. In identical conditioning and identical bids, the bid of position k is clearly higher than the bid of position $k+1$, and by letting one bidder (bidder t) change, we can't violate the above inequality. Taking the expectation over the valuations $\{v_j\}_{j\neq i,t}$ and the bids used (if the strategies are randomized) we get the inequality of Claim V.6.

To finish the proof of Lemma V.4, we would like to add the inequalities for different set pairs (S, S'). The next combinatorial lemma states that if the values v_i are drawn from independent distributions, then there is a "probability flow" $\lambda_{S,S'}$ that transfers the probability mass from $\cup_{|S|=n-k} A_S$ to $\cup_{|S'|=n-k-1} A_{S'}$ along the pairs $S' \subseteq S$. More formally, Lemma V.7 will show that there are coefficients $\lambda_{S,S'} \geq 0$ for $S' \subseteq S$ such that

$$\sum_S \lambda_{S,S'} = P(A_{S'} | v_i, \cup_{|T'|=n-k-1} A_{T'})$$

$$\sum_{S'} \lambda_{S,S'} = P(A_S | v_i, \cup_{|T|=n-k} A_T)$$

Taking linear combination of the inequalities (V.6) for set pair (S, S') with coefficients $\lambda_{S,S'}$ gives the claimed bound. ■

Lemma V.7 *If valuations are drawn from independent distributions, then there exists a probability flow* $\lambda_{S,S'} \geq 0$ *for set pairs* $S' \subseteq S$ *with* $|S'| = n-k-1$ *and* $|S| = n-k$ *such that the equations above hold.*

Proof: We will use the max-flow min-cut theorem to prove that the $\lambda_{S,S'}$ values exist. We characterize the probabilities $P(A_S | v_i, \cup_{|T|=n-k} A_T)$ using the independence assumption. Let $q_j = P(v_j \geq v_i)$, then we can write

$$P(A_S | v_i, \cup_{|T|=n-k} A_T) = \frac{\prod_{j\in S} q_j \prod_{j \notin S+i} (1-q_j)}{\sum_{|T|=n-k} \prod_{j\in T} q_j \prod_{j\notin T+i}(1-q_j)}$$

If we define $\phi_j = \frac{q_j}{1-q_j}$ and $\phi(S) = \prod_{j\in S} \phi_j$ then we can rewrite the above equation as

$$P(A_S | v_i, \cup_{|T|=n-k} A_T) = \frac{\phi(S)}{\sum_{|T|=n-k} \phi(T)}.$$

The existence of the $\lambda_{S,S'}$ values is equivalent to the existence a flow of value 1 in the following network: consider a bipartite graph where the left nodes are sources corresponding to sets S' with $|S'| = n-k-1$ with supply $\frac{\phi(S')}{\sum_{T'} \phi(T')}$ and the right nodes are sinks corresponding to sets S of $|S| = n-k$ with demand $\frac{\phi(S)}{\sum_T \phi(T)}$. We add an edge (S', S) with capacity ∞ if $S' \subseteq S$. We need to prove that the max-flow in this graph has flow value 1 (and then the flow values define $\lambda_{S',S}$). We use the max-flow/min-cut theorem (in this case, a weighted version of Hall's Theorem): there is a flow of value 1 if and only if for each collection of sets $A'_1, \ldots, A'_p$ of size $n-k-1$, the total supply, the flow that needs to leave the set, is at most the demand that is available at the neighbors of the set:

$$\sum_{i=1}^{p} \frac{\phi(A'_i)}{\sum_{S'} \phi(S')} \leq \sum_{A'_i \subseteq A, |A|=n-k} \frac{\phi(A)}{\sum_S \phi(S)}$$

which can be rewritten as

$$\sum_S \phi(S) \cdot \sum_i \phi(A'_i) \leq \sum_{A'_i \subseteq A, |A|=n-k} \phi(A) \cdot \sum_{S'} \phi(S')$$

Notice that both sides have sums of products of $2(n-k)-1$ terms of type ϕ_j. If we can prove that all terms in the LHS appear in the RHS with at least the same multiplicity we are done. We prove this based on a combinatorial construction.

The left-hand side consists of products of ϕ values for pairs of sets (S, A'_i). The right-hand side contains the products of ϕ values for pairs of the form $(S-j, A'_i+j)$ for some $j \in S \setminus A'_i$. We want to map each pair (S, A'_i) to a set $(S-j, A'_i+j)$ without collisions. If we can do this, it proves the claim. We say the pairs (S^1, A'_i) and (S^2, A'_j) are equivalent if $S^1 \cup A_i$ and $S^2 \cup A_j$ are the same (including multiplicities of the elements). Now we just need to map each equivalence class of elements in a collision-free manner. Lemma V.8 below shows a construction that satisfies the property for $t = \frac{1}{2}(|S \cup A'_i| - |S \cap A'_i| - 1)$: identify $(S \cup A'_i) \setminus (S \cap A'_i)$ with $[2t+1]$ and choose $j = f_t(A'_i \setminus S) \setminus A'_i$, where $[n] = \{1, \ldots, n\}$ and $\binom{S}{t} = \{T \subseteq S; |T| = t\}$. ■

Lemma V.8 *For all t there is a bijective function* $f_t : \binom{[2t+1]}{t} \to \binom{[2t+1]}{t+1}$ *such that* $S \subseteq f_t(S)$.

Proof: Consider a bipartite graph where the left nodes are $\binom{[2t+1]}{t}$ and the right nodes are $\binom{[2t+1]}{t+1}$ and there is an (A, B) edge if $A \subseteq B$. Notice this is a regular $k+1$-graph. Since all regular bipartite graphs have perfect matchings, the claim is proved. ■

B. Proving the dual averaging Lemma

Proof of Lemma V.5 : We want to prove that the following linear inequality system is feasible:

$$-\sum_k x_k \sum_{j=k}^{n} \gamma_j \leq -\frac{1}{2} \sum_{j=1}^{n} \gamma_j$$

$$\sum_k x_k B_k \sum_{j=k}^{n} \gamma_j \leq \sum_{j=1}^{n} \gamma_j B_j$$

$$\sum_k x_k = 1$$

$$x_k \geq 0$$

We will show that the dual system below is feasible and bounded which shows that the system above is feasible. The dual is

$$
\begin{aligned}
&\min -\phi\frac{1}{2}\sum_{j=1}^{n}\gamma_j + \psi\sum_{j=1}^{n}\gamma_j B_j + \xi \text{ s.t.}\\
&-\phi\left(\sum_{j=k}^{n}\gamma_j\right) + \psi B_k\left(\sum_{j=k}^{n}\gamma_j\right) + \xi \geq 0, \quad \forall k\\
&\phi,\psi \geq 0
\end{aligned}
$$

This linear program has a feasible solution for any $\phi, \psi \geq 0$ by setting ξ sufficiently high. So the linear program is the same as the following optimization problem:

$$
\begin{aligned}
\min_{\phi,\psi\geq 0} -\phi\frac{1}{2}\sum_{j=1}^{n}\gamma_j + \psi\sum_{j=1}^{n}\gamma_j B_j + \\
+ \max_k\left[\left(\sum_{j=k}^{n}\gamma_j\right)(\phi - \psi B_k)\right]
\end{aligned}
$$

Our goal is to prove that for any fixed $\gamma_k, B_k \geq 0$, for any values of $\phi, \psi \geq 0$ this is a non-negative expression, and establishing that it is bounded. We claim that for some value of k, the following must be non-negative:

$$
-\phi\frac{1}{2}\sum_{j=1}^{n}\gamma_j + \psi\sum_{j=1}^{n}\gamma_j B_j + \left(\sum_{j=k}^{n}\gamma_j\right)(\phi - \psi B_k)
$$

We will show this by summing the above expressions weighted by γ_k, and showing that the result is non-negative. Therefore, at least one of the summands must be non-negative. The sum is

$$
\begin{aligned}
\sum_k \gamma_k[-\phi\frac{1}{2}\sum_{j=1}^{n}\gamma_j + \psi\sum_{j=1}^{n}\gamma_j B_j + \\
+ \left(\sum_{j=k}^{n}\gamma_j\right)(\phi - \psi B_k)]
\end{aligned}
$$

And this expression is non-negative, since ϕ is multiplied by $\sum_k \sum_{j\geq k}\gamma_j\gamma_k - \frac{1}{2}\sum_k\sum_j \gamma_j\gamma_k$ which is ≥ 0 and ψ is multiplied by $\sum_k\sum_j \gamma_k\gamma_j B_j - \sum_k\sum_{j\geq k}\gamma_j\gamma_k B_k$, which is also ≥ 0.

■

ACKNOWLEDGMENTS

We would like to thank the anonymous reviewers and the participants of the AdAuction workshops of 2009 and 2010 for many helpful comments and suggestions.

REFERENCES

[1] G. Christodoulou, A. Kovács, and M. Schapira. Bayesian combinatorial auctions. In *ICALP '08: Proceedings of the 35th international colloquium on Automata, Languages and Programming, Part I*, pages 820–832, Berlin, Heidelberg, 2008. Springer-Verlag.

[2] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11(1), September 1971.

[3] Edelman, Benjamin, Ostrovsky, Michael, Schwarz, and Michael. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *The American Economic Review*, 97(1):242–259, March 2007.

[4] R. D. Gomes and K. S. Sweeney. Bayes-nash equilibria of the generalized second price auction. In *EC '09: Proceedings of the tenth ACM conference on Electronic commerce*, pages 107–108, New York, NY, USA, 2009. ACM.

[5] T. Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973.

[6] J. C. Harsanyi. Games with incomplete information played by "bayesian" players, i-iii. *Manage. Sci.*, 50(12 Supplement):1804–1817, 2004.

[7] S. Lahaie. An analysis of alternative slot auction designs for sponsored search. In *EC '06: Proceedings of the 7th ACM conference on Electronic commerce*, pages 218–227, New York, NY, USA, 2006. ACM.

[8] S. Lahaie, D. Pennock, A. Saberi, and R. Vohra. *Algorithmic Game Theory*, chapter Sponsored search auctions, pages 699–716. Cambridge University Press, 2007.

[9] B. Lucier and A. Borodin. Price of anarchy for greedy auctions. In *SODA '10*. ACM, 2010.

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

[11] R. Paes Leme and E. Tardos. Pure and bayes-nash price of anarchy for generalized second price auction, manuscript. http://www.cs.cornell.edu/~renatoppl/papers/paesleme_tardos_focs10_manuscript.pdf, 2010.

[12] T. Roughgarden. Intrinsic robustness of the price of anarchy. In *STOC '09: Proceedings of the 41st annual ACM symposium on Theory of computing*, pages 513–522, New York, NY, USA, 2009. ACM.

[13] D. R. M. Thompson and K. Leyton-Brown. Computational analysis of perfect-information position auctions. In *EC '09: Proceedings of the tenth ACM conference on Electronic commerce*, pages 51–60, New York, NY, USA, 2009. ACM.

[14] H. R. Varian. Position auctions. *International Journal of Industrial Organization*, 2006.

[15] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, 1961.