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

Renato Paes Leme * Éva Tardos †

Abstract

Generalized Second Price Auction, also knows as Ad Word auctions, and its variants has been the main mechanism used by search companies to auction positions for sponsored search links. In this paper we study the social welfare of the Nash equilibria of this game. It is known that socially optimal Nash equilibria exists (i.e., that the Price of Stability for this game is 1). This paper is the first to prove bounds on the price of anarchy.

Our main result is to show that under some mild assumptions the price of anarchy is small. For pure Nash equilibria we bound the price of anarchy by 1.618, assuming all bidders are playing un-dominated strategies. For mixed Nash equilibria we prove a bound of 4 under the same assumption. We also extend the result to the Bayesian setting when bidders valuations are also random, and prove a bound of 8 for this case.

Our proof exhibits a combinatorial structure of Nash equilibria and use 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.

* renatoppl@cs.cornell.edu. Dept. of Computer Science, Cornell University, Ithaca, NY 14853. Supported in part by NSF grants CCF-0729006.

† eva@cs.cornell.edu. Dept. of Computer Science, Cornell University, Ithaca, NY 14853 Supported in part by NSF grants CCF-0910940 and CCF-0729006, ONR grant N00014-98-1-0589, and a Yahoo! Research Alliance Grant.

1 Introduction

Search engines and other online information sources use sponsored search auction, or AdWord auctions, to monetize their services. These actions allocate advertisement slots to companies, and companies are charged pay per click, that is, they are charged a fee for any user that clicks on the link associated with the advertisement. Mehta, Saberi, Vazirani, and Vazirani [9, 10] considered AdWord auctions in the algorithmic context, studying the problem of assigning AdWords to advertisers online as the word shows up in a search query. Since the introduction of the model, there has been much work in the area, see the survey of Lahaie et al [7].

Here we consider AdWords in a game theoretic context: consider the game played by advertisers in bidding for an AdWord. The bids are used to determine both the assignment of bidders to slots, and also the fees charged. The bidders are assigned to slots in order of bids, and the fee for each click is decided by variant of the so-called **Generalized Second Price** Auction (GSP), a simple generalization of the well-known Vickrey auction [14] for a single item (or a single advertising slot). The Vickrey auction [14] for a single item, and its generalization, the Vickrey-Clarke-Groves Mechanism (VCG) [3, 5], make truthful behavior (when the advertisers reveal their true valuation) dominant strategy, and make the resulting outcome maximize the social welfare.

Generalized Second Price Auction, the mechanism adopted by all search companies, is a simple and natural generalization of the Vickrey auction for a single slot, but it is neither truthful nor maximizes social welfare. In this paper we will consider the social welfare of the GSP auction outcomes. Our goal in this paper is to show that the intuition based on the similarity of GSP to the truthful 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 mild assumption that the players use un-dominated strategies.

We consider both full information games when player valuations are known, and also consider the Bayesian setting when the values are independent random variables. 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 [11]. In contrast, we show in Appendix B that the GSP game is not smooth in the sense of [11]. Second, most known price of anarchy results are for the case of full information games. The full information setting assumes that all advertisers are aware of the valuations of all other players. This is a very strong assumption and is not realistic. In contrast, the Bayesian setting requires only much weaker assumption that valuations are drawn from independent distributions, and these distributions are known to the other players. Proving the price of anarchy bound for the Bayesian setting requires the use of novel combinatorial techniques.

We use a standard model of separable click-though rates: where the probability of clicking on an advertisement j displayed in slot i is $\alpha_i\gamma_j$, i.e., the probability is a product of two separable components: depending on the slot, and on the advertiser respectively. To simplify the presentation, for the main part of the paper, we will focus on the simple case when $\gamma_j = 1$ for all j, that is, the probability of a click depends only on the slot. In Appendix A we show that it is easy to extend our results to the model with separable click-through rates.

For both our simple model, and the case of separable click-through rates, it is known that there exists Nash equilibria that are socially optimal [4, 13], i.e., that the price of stability is 1. It is not hard to give simple examples of Nash equilibria where the social welfare is arbitrarily smaller than the optimum. 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 a 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 players valuation is dominated strategy in all settings.

For each setting we exhibit a combinatorial structure of 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 a Nash equilibrium (a random variable, in the case of mixed Nash, or in the Bayesian setting).

- For the case of pure Nash equilibria the social welfare in a 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 bids, 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. 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 in order of decreasing valuation, and number slots in order of decreasing click rates. By this notation, bidder i should be assigned to slot i in the optimal solution. We derive the structure of pure Nash equilibria by thinking about a pair of bidders that are assigned to slots in reverse order. Such pairs are harder to define in the mixed case. Instead, we will consider bider i and his optimal slot i, and get the following condition for mixed Nash equilibria.
$$\frac{\mathbb{E}\alpha_{\sigma(i)}}{\alpha_i} + \frac{\mathbb{E}v_{\pi(i)}}{v_i} \geq \frac{1}{2},$$
and use this inequality to show that the social welfare of a mixed Nash equilibrium is at least an fourth fraction of the social welfare of the socially optimal assignment. Note that an analogous inequality with bound 1 holds also for pure Nash, which implies a bound of 2 on the pure price of anarchy. We used a different characterization above to be able to prove the stronger bound/

- We prove a bound of 8 on the price of anarchy for the Bayesian setting, and the valuations v_k are 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 simple 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, we are not aware of a single deviating bid that can help establish a useful structure. Instead, we obtain our structural result by considering many different bids, and show that the inequalities established by the different bids can be combined 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 closely related to $2\mathbb{E}[b_{\pi k}|v_i; \nu(i) = k]$, twice the expected value of the bid in slot k where the expectation is conditioned both on the value v_i of player i, and the fact that its optimal position is k. These expectations as defined here depend on player i both through the conditioning and though the position bidder i gets in the permutation π. This dual dependence makes it hard to prove any properties of them. For example, it would be natural to assume that for any value v_i the values monotone increase with k, but that is not always the case. To get around this problem, we will use instead slightly smaller values for the bids. We show via an interesting combinatorial argument using the max-flow min-cut theorem, that the modified values do increase with k. Then we use a novel averaging technique (using linear programming) to combine the resulting inequalities to establish the simple structural property.

Related work Sponsored search has been a very active area of research in the last several years. Mehta et al. [10, 9] 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 [7] for a general introduction. Here we use the game theoretic model of the AdWord auctions of Edelman et al [4] and Varian [13], for a truthful auction see Aggarwal et all [1].

We use the model of separable click-though rates, where the click through rate for bidder j in slot i can be expressed in a simple product form $\gamma_j \alpha_i$. For these models Edelman et al [4] and Varian [13] show that the price of Stability for this game is 1, that is, there exists Nash equilibria that are socially optimal.

Lahaie [6] also considers the problem of quantifying the social efficiency of an equilibrium. He makes the strong assumption that click-through-rate α_i decays exponentially along the slots with a factor of $\frac{1}{\delta}$, and proves a price of anarchy of $\min\{\frac{1}{\delta}, 1-\frac{1}{\delta}\}$. In this paper, we make no assumptions on the click-through-rates. Thompson and Leyton-Brown [12] 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 his valuation is a dominated strategy. Lucier and Borodin [8] and Christodoulou at al [2] also use the conservative assumption to establish price-of-anarchy results in the context of combinatorial auctions. Without any additional requirement Nash equilibria, even in the case of the single item Vickrey auction, can have social welfare that is arbitrarily bad compared to the optimal social welfare. However, we show that Nash equilibria of conservative bidders is within small constant factor of the optimum.

The paper by Lucier and Borodin [8] on greedy auctions is also closely related to our work. They analyze the Price of Anarchy of the auction game induced by Greedy Mechanisms. The consider in a general combinatorial auction setting, greedy algorithms with payments are computed using the critical price. They show via a type of smoothness argument (see [11]) that of 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, and further the GSP game does not satisfy the smoothness condition. The key to proving the $c + 1$ bound of Lucier and Borodin [8] is to consider possible bids, such as a single minded bid for the slot in the optimal solution, or modifying a bit 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, as 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, as typically allowing arbitrary complex bids makes the optimization problem hard. However, the limited expressivity of the bidding language can increase the set of Nash equilibria (as there are fewer deviating bids to consider). It is important to understand if such natural bidding languages result in greatly increased price of anarchy.

2 Preliminaries

We consider an auction with n advertisers and n slots (if there are less slots than advertisers, consider additional virtual slots with click-through-rate zero). We model this auction as a game with n players, where each advertiser is one player. The types of the advertisers are given by their valuation v_i, which expresses their value for one click. The strategy for each advertiser is a bid $b_i \in [0, \infty)$.

There are n slots and based on the bids, we decide where to allocate each advertiser. In the most simple model, the k-th slot contains α_k clicks and α_k is a monotone non-increasing sequece, i.e., $\alpha_1 \geq \alpha_2 \geq \ldots \geq \alpha_n$. We prove our results for this simple model, but they extend naturally to the more realistic model of separable click-through-rates, as we show in Appendix A. The game proceeds as follows:

1. each advertiser submits a bid $b_i \geq 0$, which is his declared value for a click
2. the advertiser 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 completely determined by the set of bids. 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 the bidders get from playing it, which is: $\sum_j \alpha_j v_{\pi(j)}$. Here in this paper we are concerned about bounding the social welfare in an equilibrium of this game relative to the optimal. This measure is called Price of Anarchy. We analyze the Price of Anarchy in three different settings of increasing complexity:

- **Pure Nash equilibrium:** The valuation of each player v_i is a fixed value. We consider without loss of generality that $v_1 \geq v_2 \geq \ldots \geq v_n$. Each player chooses a pure strategy, i.e., a deterministic bid b_i. We say that a set of bids $b = (b_1, \ldots, b_n)$ is a **Pure Nash Equilibrium** if any bidder can change his bid an increase his utility, i.e.:

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

  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 and therefore won't change the allocation 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 would get 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 $\alpha_k(v_i - b_{\pi(k+1)})$.

  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 optimal and in a Nash equilibrium.

- **Mixed Nash equilibrium:** The valuation v_i are still fixed and we can assume (without loss of generality) that $v_1 \geq \ldots \geq v_n$, but each player is allowed to pick a distribution over strategies. We can think that each player chooses a random variable b_i and the Nash equilibrium means that the chosen random variable maximizes the expected utility. In other words:

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

  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 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:** In a more realistic model, the player's don't know the valuations of other players, but they have beliefs about it. This is modeled as follows: the valuation v_i are drawn from independent distributions. A player chooses a bid (possibly in a randomized fashion) based on his own valuation. 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 said to be a Bayes-Nash equilibrium if:
$$\mathbb{E}[u_i(b_i(v_i), b_{-i}(v_{-i}))|v_i] \geq \mathbb{E}[u_i(b_i'(v_i), b_{-i}(v_{-i})|v_i], \forall b_i'(v_i) \quad \forall v_i$$
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, and we define it by ν: let $\nu(k)$ be the slot occupied by player i 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, given by the ratio: $\mathbb{E}[\sum_j \alpha_{\nu(j)} v_j]/\mathbb{E}[\sum_j \alpha_j v_{\pi(j)}]$

2.1 Equilibria with Low Social Welfare

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 $\alpha_2 = r$ and two advertisers with valuations $v_1 = 1$ and $v_2 = 0$. It is easy to check that the bids $b_1 = 0$ and $b_2 = 1 - r$ 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 r while the optimal is 1. The price of anarchy is therefore $1/r$. Since r can be any number from 0 to 1, the gap between the optimal and the worse Nash can be arbitrarily large.

Notice however that this Nash equilibrium seems very artificial: advertiser 2 is exposed to the risk of negative utility: if advertiser 1 (or another advertiser) adds a bid somewhere between 0 and $1-r$ this imposes a negative utility on advertiser 2. For advertiser 2, bidding anything greater than zero is clearly a dominated strategy. It is not hard to see that for any v_i, bidding $b_i > v_i$ is dominated by the bid v_i.

3 Pure Nash Equilibrium

We say pure bid b_i for advertiser i is *conservative* if $b_i \leq v_i$. Next we show that non-conservative bids are dominated. We say that a strategy b_i is *dominated* if there is some b_i' such that $u_i(b_i, b_{-i}) \leq u_i(b_i', b_{-i})$ for all b_{-i} and for at least one value of b_{-i} it holds strictly.

Lemma 3.1 *A bid $b_i > v_i$ is dominated by $b_i' = v_i$.*

Proof. Decreasing the bid from b_i to v_i changes allocation and or payment only if there is a bidder $j \neq i$ with $v_i < b_j < v_j$. However, in this case, bidder i getting negative utility for each click, and with bid v_i, he cannot get negative utility. ∎

Given the parameters α, v, we say that b is a **conservative bidder equilibrium** if it is a Nash equilibrium and $b_i \leq v_i$ for all bidders i.

Theorem 3.2 *For 2 slots, if all advertisers are conservative, then the price of anarchy is exactly 1.25.*

Proof. We can suppose without loss of generality (by scaling α and v) that $\alpha_1 = 1$, $\alpha_2 = r$ and $\alpha_1 v_1 + \alpha_2 v_2 = 1$. In any non-optimal Nash equilibrium $b_1 \leq b_2$ and by the Nash condition $r(v_1 - 0) \geq 1(v_1 - b_2)$ and by the conservative condition $b_2 \leq v_2$. Substituting $v_1 = 1 - r v_2$ in those two expressions and combining them to eliminate the b_2 term we get: $v_2 \geq \frac{1-r}{1-r(r-1)}$. Therefore the social welfare in any non-optimal Nash is $\alpha_1 v_2 + \alpha_2 v_1 = 1 v_2 + r(1 - r v_2) \geq 1 + r(1-r) \leq 1.25$. ∎

3.1 Weakly Feasible Assignments

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

The equations (**??**) are not very easy to work with, since they are not very symmetric and they depend on b. We propose a cleaner form of representing an equilibrium that just uses α, v and the permutation π. Although it is a weaker property it still captures most of the trade-offs:

1. if values v_i are very close then the order of the bidders doesn't influence the social welfare that much
2. if values v_i are very well separated, then permutations that would produce a bad social welfare are not feasible because they violate Nash constraints

Lemma 3.3 *Given v, α and a Nash permutation π, if $i < j$ and $\pi(i) > \pi(j)$ then:*

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

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

Proof. Since it is a Nash equilibrium bidder in slot j is happy with his condition and don'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 equation 1 holds for each $i < j$, $\pi(i) > \pi(j)$. From Lemma 3.3 we know that:

Corollary 3.4 *Given α, v, any permutation corresponding to a Nash equilibrium with conservative bids is weakly feasible.*

Our main results 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.

3.2 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. We prove it is bounded by 1.618. We will prove this bound for weakly feasible permutations and it will automatically be deduced to a bound for feasible permutations. Notice that the weakly feasible permutation nicely capture the fact that if advertisers i and j are in the "wrong relative position" (i.e. different to the one in the optimal) then either their values are close (within a factor of 2) or their click-through-rates are close (within a factor of 2). The proof of the 1.618 factor can be found in Appendix C.

Theorem 3.5 *For conservative bidders, the price of anarchy for pure Nash equilibria of GSP is bounded by $\frac{1+\sqrt{5}}{2} \approx 1.618$.*

Proof. 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 show this by induction on n. For 2 advertisers and 2 slots we know that the worst possible social welfare for a weakly feasible permutation is at most a 1.25 times the optimum. So, now we need to prove the inductive step. Consider parameters v, α and a weakly feasible permutation π. Let $i = \pi^{-1}(1)$ be the slot occupied by the advertiser of higher value and $j = \pi(1)$ be the advertiser occupying the first slot. If $i = j = 1$ then we can apply the inductive hypothesis right away. If not, equation 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. This input has $n-1$ advertisers and $n-1$ slots and the permutation π restricted to those is still weakly feasible, so by the inductive hypothesis:

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

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 finishes the bound of 2.

The above analysis is not tight. For example, when the ratio α_i/α_j is equal to 1/2, our inequality states that we also have $v_i/v_j \geq 1/2$. We prove the stronger bound in Appendix C by a more careful analysis using the full strength of the inequality. As before, we prove the conclusion for all weakly feasible permutations. We define a sequence of values r_k so that for k slots social welfare is at least an r_k times the optimum. We know that $r_2 = 1.25$, and use a similar but more careful induction proof to set up a recursion of r_k, and then show (in Appendix C) that r_k converges to the desired bound of $\frac{1+\sqrt{5}}{2}$. ■

4 Mixed Nash equilibrium

All our results so far dealt with Pure Nash equilibria of the GSP. Here we prove a bound on the Price of Anarchy of 4 for the mixed Nash equilibria. Consider the same setting: players with valuations $v_1 \geq \ldots \geq v_n$ and slots with click-through-rates $\alpha_1 \geq \ldots \geq \alpha_n$. Now, the strategy of player i is a probability distribution on $[0, v_i]$ represented by a random variable b_i.

Lemma 4.1 *A randomized bid b_i where $P(b_i > v_i) > 0$ is dominated by $b'_i = \min(v_i, b_i)$.*

Now, the allocation, represented by the permutation π is also a random variable. For notational convenience, let $\sigma = \pi^{-1}$. A random vector $b = (b_1, \ldots, b_n)$ is a mixed Nash equilibrium if for each deterministic bid b'_i: $\mathbb{E}u_i(b_i, b_{-i}) \geq \mathbb{E}u_i(b'_i, b_{-i})$. We begin by proving a bound similar to Lemma 3.3 for mixed Nash and then using that to prove a weaker bound. Note that the bound 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 4.2 *If the random vector b is a mixed Nash equilibrium for GSP then for each player i:*

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

Proof. Bidder i by our notation has the ith highest valuation, and hence would be in the ith slot in the optimal assignment. We will consider whether player i benefits by deviating to the deterministic $b'_i = \min(v_i, 2\mathbb{E}b_{\pi(i)})$, where $\mathbb{E}b_{\pi(i)}$ is the expected value of the bid that gets slot i.

We claim that with probability at least $\frac{1}{2}$, the bidder gets one of the slots of $\{1, \ldots, i\}$. If $b'_i = v_i$ then it for sure gets at least the i-th slot as our conservative assumption guarantees that only the previous $i-1$ players can bid more. If $b'_i = 2\mathbb{E}b_{\pi(i)}$ 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:

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

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

Theorem 4.3 *The Price of Anarchy for the mixed Nash equilibria of GSP with conservative bidders is* ≤ 4*.*

Proof. The proof is a simple application of Lemma 4.2 and some algebraic manipulation:

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

■

5 Bayes-Nash equilibrium

Recall that in the Bayesian setting, the values v_i are independent random variables and their distributions are public knowledge. 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 players bid when his value is v_i. For simplicity of presentation we will consider only pure bidding functions, but our results extend to randomized bids also. As before, we will assume that the bid $b_i(v_i)$ is in the range $[0, v_i]$. Here we are assuming that players don't overbid (as overbidding is dominated strategy).

As before, 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.

5.1 Price of Anarchy bound

Theorem 5.1 *If a set of bids* $b_1, \ldots, b_n$ *are a Bayes Nash equilibrium in conservative strategies then:*

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

in other words, GSP has a Bayes-Nash Price of Anarchy in conservative strategies bounded by 8*.*

The proof of the theorem is based on a structural characterization analogous to the one used for Mixed Nash equilibria, which is similar in form to the previous ones, but much harder to prove. We can state the inequality used to bound the price of anarchy for mixed Nash as $v_i\mathbb{E}\alpha_{\sigma(i)} + \alpha_i \mathbb{E}v_{\pi(i)} \geq \frac{1}{2}v_i\alpha_i$. To extend this expression we need to take expectations over valuations.

Player i knows his own valuation v_i, so in considering alternate bids for the player, we need to consider expectations conditioned on the value v_i. To be able to deal with the conditioning of the player i on his own valuation v_i we make the terms on the expression independent of i. To do this, 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 $\sigma(i) > \sigma(k)$ $\pi^i(k) = \pi(k+1)$ otherwise. Now we can state the claimed inequality. Note that the bound is stronger than suggested by the case of mixed Nash, as $b_{\pi^i(\nu(i))} \geq b_{\pi(\nu(i))}$.

Lemma 5.2 *If $\{b_i(\cdot)\}_i$ is a Bayes-Nash equilibrium of the GSP then:*

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

As before the price of anarchy bound follows easily from the lemma.

Proof of Theorem 5.1 : The proof follows the lines of the proof of theorem 4.3:

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

recalling that $\pi^i(k)$ is either $\pi(k)$ or $\pi(k+1)$ implying the inequality in the second line. ■

The hard part of the proof is proving Lemma 5.2. The main difficulty in the Bayesian setting is that the inequality is not established by a single deviating bid. In the base of pure Nash we considered the bidder in slot j bidding $b' = b_{\pi(i)}$ for a pair of slots $i < j$. In the case of mixed Nash, we considered bidder i bidding $b' = \min(v_i, 2\mathbb{E}b_{\pi(i)})$. In both cases the structural inequality followed by considering a single deviation. In contrast, we obtain our structural result by considering many different bids, and using a novel averaging argument, and a structural property to show that the inequalities established by the different bids can be combined to show a simple structure.

It would be natural to consider a set of bids $2\mathbb{E}[b_{\pi(k)}|v_i; \nu(i) = k]$ twice the expected value of the bid in slot k where the expectation is conditioned on the value v_i of player i, and the fact that its optimal position is k. The bids as defined depend on player i both through the conditioning and though the position bidder i gets in the permutation i, which makes it hard to prove any properties of them. For example, these bids may not be monotone functions of k. To get around this problem, we will use instead the slightly smaller bids

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

which gets rid of this second dependence. Notice that B_k is defined as a conditional expectation, so it is a function of v_i. Notice therefore that it is a bidding function and not a constant function.

The proof of Lemma 5.2, depends on two combinatorial results. The first, is a structural property: we claim that the bids B_k as now defined are monotone in k. This will allows us to argue that bid B_k has a good chance of taking a slot $k' > k$ when $\nu(i) = k'$, as $B_k \geq B_{k'}$.

Lemma 5.3 *Given bidding functions b_i, $\mathbb{E}[b_{\pi^i \nu(i)}|v_i, \nu(i) = k]$ in non-increasing in k*

We will prove the lemma above using a combinatorial argument and max-flow min-cut theorem.

To be able to combine the inequalities we get by a considering the different bids B_k we use a novel way to combine inequalities via a 'dual averaging argument'. The combination will simultaneously guarantee that one average is not too low, and a different average is not to hight. We expect that this Lemma 5.4, which prove using linear programming duality can have other applications.

Lemma 5.4 *Given any positive values γ_k and B_k There are $x_k \geq 0$, $\sum_k x_k = 1$ 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 5.2:

Proof of Lemma 5.2 : As outlined above we will consider n deviation for player i at bids B_k for all possible slots k. Since it is a Nash equilibrium player i can't benefit from changing his strategy each will give us an inequality on the utility. We the will use Lemma 5.4 to average them to get the claimed inequality. Consider bidder i deviates to $B_k = \max\{v_i, 2\mathbb{E}[b_{\pi^i(k)}|v_i; \nu(i) = k]\}$. Notice that by Lemma 5.3 we have that $B_1 \geq B_2 \geq \ldots \geq B_n$. Let $p_k = P(\nu(i) = k|v_i)$ and let α' be the random variable that means the click-through-rate of the slot he occupies by bidding B_k, then:

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

where the last inequality will follow by estimating the probability that by bidding B_k the player gets the slot j or better when $\nu(i) = j$ for some $j > k$. In the case $B_k = v_i$ it is trivial. If $B_k = 2\mathbb{E}[b_{\pi^i k}|v_i; \nu(i) = k]$, then we use that $B_j \geq B_k$ by Lemma 5.3, and we get:

$$P(\alpha' \geq \alpha_j | v_i, \nu(i) = j) = P(B_k \geq b_{\pi^i(j)}|v_i, \nu(i) = j) \geq P(B_j \geq b_{\pi^i(j)}|v_i, \nu(i) = j) \geq \frac{1}{2}$$

by Markov's Inequality. Now, we use the Lemma 5.4 applied with B_k and $\gamma_k = p_k \alpha_k$. Using the x_k from the Lemma, we get:

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

■

5.2 Proving that bids B_k are non-increasing in k

We will prove Lemma 5.3 in several steps. First we prove bounds assuming all but a single player has a deterministic value, and take expectations to get a conditional version. Then we use a max-flow min-cut argument to average these bounds get the claimed overall bound.

Proof of Lemma 5.3 : 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:

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

Take a set $S' \subseteq S$, i.e., $S = S' \cup \{t\}$ for some agent $t \neq i$. The first claim is that:

Claim 5.5 *For a set 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 the bid of position k is clearly higher than the bid of position $k+1$, and by letting one bidder change, we can't violate the above inequality. Taking the expectation over $\{v_j\}_{j\neq i,t}$ we get the inequality in of Claim 5.5.

To finish the proof of Lemma 5.3, we would like to add the inequalities for different set pairs (S, S'). The next combinatorial lemma states that there are values $\lambda_{S,S'} \geq 0$ for $S' \subseteq S$ such that:

$$\sum_S \lambda_{S,S'} = P(A_{S'}|v_i, \cup_{|S'|=n-k-1} A_{S'}) \text{ and } \sum_{S'} \lambda_{S,S'} = P(A_S|v_i, \cup_{|S|=n-k} A_S)$$

Taking the a linear combination of the inequalities (5.5) for set pair (S, S') with coefficients $\lambda_{S,S'}$ lets the claimed bound. ■

Lemma 5.6 *There exists values $\lambda_{S,S'} \geq 0$ 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 network flows to prove that the $\lambda_{S,S'}$ values exist. Before we do that, consider the following characterization of $P(A_S|v_i, \cup_{|S|=n-k} A_S)$: let $p_j = P(v_j \geq v_i)$, then we can write:

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

If we define $\phi_j = \frac{p_j}{1-p_j}$ and $\phi(S) = \prod_{j\in S} \phi_j$ then we can rewrite:

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

The existence of the $\lambda_{S,S'}$ is equivalent to the existence of a max-flow in the following network: consider a bipartite graph with left nodes corresponding to sets S' of $|S'| = n-k-1$ and inflow $\frac{\phi(S')}{\sum_{S'} \phi(S')}$ and the right nodes corresponding to sets S of $|S| = n-k$ and outflow $\frac{\phi(S)}{\sum_S \phi(S)}$. We add an edge (S', S) if $S' \subseteq S$ with capacity ∞. We need to prove that there max-flow in this graph has size 1 (and then the flow values define $\lambda_{S',S}$). We use the min-cut/max-flow theorem (in this case, this is a weighted version of Hall's Theorem): there is a max-flow of size 1 if and only if for each collection of sets $A'_1, \ldots, A'_p$ of size $n-k-1$, it is the case that he total flow that needs to enter the set is at least as big as the outflow 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 the at least same multiplicity we are done. We prove it based on a combinatorial construction.

The products on the two sides consist products of ϕ values for pairs of sets (S, A_i) and $(S-j, A_j+j)$ respectively. We want to map each pair (S, A_i) to $(S-j, A_j+j)$ 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, just need to map each equivalence class of elements in a collision-free manner. The Lemma 5.7 below shows that the following construction satisfies the property: take $t = \frac{1}{2}(|S \cup A_i| - |S \cap A_i| - 1)$, identify $(S \cup A_i) \setminus (S \cap A_i)$ with $[2t+1]$ and choose take $j = f_t(A_i \setminus S) \setminus A_i$. ■

Lemma 5.7 *For all t, then there is a bijective function $f_t : \binom{[2t+1]}{t} \to \binom{[2t+1]}{t+1}$ such that $S \subseteq f_t(S)$, where $[n] = \{1, \ldots, n\}$ and $\binom{S}{t} = \{T \subseteq S; |T| = t\}$.*

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, which prove the claim. ■

5.3 Proving the dual averaging Lemma

Proof of Lemma 5.4 : We want to prove that the following linear programming problem is feasible:

$$
\begin{aligned}
&\max 0 \text{ s.t.} \\
&\quad -\sum_k x_k \sum_{j=k}^n \gamma_j \leq -\frac{1}{2} \sum_{j=1}^n \gamma_j \\
&\quad \sum_k x_k B_k \sum_{j=k}^n \gamma_j \leq \sum_{j=1}^n \gamma_j B_j \\
&\quad \sum_k x_k = 1 \\
&\quad x_k \geq 0
\end{aligned}
$$

Verifying that this program is feasible is the same as verifying that the dual is feasible and bounded. 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 problem has a solution for any $\phi, \psi \geq 0$ by setting ξ sufficiently low. So the linear program is the same as the following optimization problem:

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

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 its bounded. We claim that one of the following must be non-negative for sum value of k:

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

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

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

And this expression is nonnegative, as 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.

■

References

[1] G. Aggarwal, A. Goel, and R. Motwani. Truthful auctions for pricing search keywords. In *EC '06: Proceedings of the 7th ACM conference on Electronic commerce*, pages 1–7, New York, NY, USA, 2006. ACM.

[2] 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.

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

[4] 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.

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

[6] 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.

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

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

[9] A. Mehta, A. Saberi, U. V. Vazirani, and V. V. Vazirani. Adwords and generalized on-line matching. In *FOCS*, pages 264–273, 2005.

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

[11] 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.

[12] 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.

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

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

Appendix A: Extension to separable click-through-rates

So far, we have considered that the click-through-rates of advertiser i placed on slot j depends only on the slot in which he is placed. A more general model called **separable click-through-rates** assumes it depends on a product of two factors: one depending on the bidder and other depending on the slot. Let's say that if advertiser i is placed on slot j, it will get click-through-rate $\gamma_i \alpha_j$ where γ_i is some "quality factor" attributed to each advertiser. The generalization of Second Price Auction for this setting ranks the advertisers in order of $\gamma_i b_i$ and charges an advertiser the minimum value it needed to bid to conserve his position. For example, if π is the permutation defined by sorting $\gamma_i b_i$ (i.e, $\pi(k)$ is the advertiser with the k^{th} highest value of $\gamma_i b_i$) then we charge advertiser $\pi(j)$ the amount of: $b_{\pi(j+1)} \gamma_{\pi(j+1)} / \gamma_{\pi(j)}$.

In this setting the utility of bidder i assigned to slot j is $u_i = \gamma_i \alpha_j \left(v_i - \frac{b_{\pi(j+1)} \gamma_{\pi(j+1)}}{\gamma_i} \right)$ and the social welfare is given by $\sum_k \alpha_k \gamma_{\pi(k)} v_{\pi(k)}$. Consider that $\alpha_1 \geq \ldots \geq \alpha_n$ and that $\gamma_1 v_1 \geq \ldots \geq \gamma_n v_n$. The definition of Nash equilibrium is analogous. Notice we can obtain a result very similar with Lemma 3.3 just by repeating the same calculations for this model:

Lemma 5.8 *Given v, α, γ and a feasible permutation π (a permutation from a Nash equilibrium) in the separable click-through-rate model, if $i < j$ and $\pi(i) > \pi(j)$ then:*

$$\frac{\alpha_j}{\alpha_i} + \frac{\gamma_{\pi(i)} v_{\pi(i)}}{\gamma_{\pi(j)} v_{\pi(j)}} \geq 1 \tag{2}$$

Proof. Since advertiser $\pi(j)$ can't increase his utility by taking slot i, we have that:

$$\gamma_{\pi(j)} \alpha_j \left(v_{\pi(j)} - \frac{b_{\pi(j+1)} \gamma_{\pi(j+1)}}{\gamma_{\pi(j)}} \right) \geq \gamma_{\pi(j)} \alpha_i \left(v_{\pi(j)} - \frac{b_{\pi(i)} \gamma_{\pi(i)}}{\gamma_{\pi(j)}} \right)$$

using that $b_{\pi(j+1)} \geq 0$ and $b_{\pi(i)} \leq v_{\pi(i)}$ we get the desired result. ■

Similarly, the structure characterization for Mixed and Bayes-Nash equilibria can be rephrased in the context of separable click-through-rates, effectively replacing the values v_i for bider i with the product $v_i \gamma_i$ in all expressions.

Appendix B: GSP is not a smooth game

This paper gave the first $O(1)$ bound for the Price of Anarchy of the GSP. As Roughgarden points out in [11], games studied so far (as congestion games, facility location, valid utility games, ...) have their Price of Anarchy proof based on a smoothness argument. In this section we note that this proof is essentially different from all previous Price of Anarchy analysis as the GSP game is not smooth.

A game is said to be (λ, μ)-smooth if the following property holds:

$$\sum_i u_i(s_i^*, s_{-i}) \geq \lambda SW(s^*) - \mu SW(s)$$

for all possible strategies s, s^*, where u_i are utilities of each player and SW is the social welfare function which is given by $SW = \sum_i u_i$. To model GSP as one of this games, we consider a game of $n+1$ players - the n advertisers and the search engine. Each advertiser has one value v_i and its strategies are bids in $[0, v_i]$, still supposing them conservative. The search engine has only one strategy, which is "run GSP", and its utility are the payments it receives. The search engine is clearly not really playing the game, it is just there to make the social welfare the sum of the utilities. The following theorem shows that GSP is not a smooth game:

Theorem 5.9 *Conservative GSP is not (λ, μ)-smooth for any parameters λ, μ.*

Proof. Consider the game with 2 slots with click-through-rates 1 and α and two advertisers with values 1 and v. Let $s = (b_1, b_2)$ and $s^* = (b_3, b_4)$ where $1 > v > b_2 > b_3 > b_4 > b_1$. For this case, the expression $\sum_i u_i(s_i^*, s_{-i}) \geq \lambda SW(s^*) - \mu SW(s)$ becomes:

$$b_1 + \alpha(1-0) + 1(v - b_1) \geq \lambda(1 + \alpha v) - \mu(v + \alpha)$$

Simplifying we get:

$$(1+\mu)(\alpha + v) \geq \lambda(1 + \alpha v)$$

Since α and v are parameters, for any λ, μ, we can make them arbitrarily small violating the inequality for any $\lambda > 0$, ■

Appendix C: 1.618 bound for the Pure Price of Anarchy

Theorem 5.10 *For conservative bidders, the price of anarchy for pure Nash equilibria is bounded by $\frac{1+\sqrt{5}}{2} \approx 1.618$.*

Proof. As before, we prove the conclusion for all weakly feasible permutations. We define a sequence of values r_k so that we can prove that for k slots social welfare is at least an r_k fraction of the optimum, and prove that r_k converges to the desired bound. Let $r_2 = 1.25$ and suppose we have $r_2, r_3, ..., r_{n-1}$ and that this property holds for them. Let's calculate some "small" value of r_n so that the property still holds.

Again, consider parameter α, v, a weakly feasible permutation π and let's assume $i = \pi^{-1}(1)$ and $j = \pi(1)$. If $i = j = 1$, this is an easy case and it is straightforward to see that in this case the price of anarchy can be bounded by r_{n-1}. If not, assume without loss of generality that $i \leq j$ (since equation 1 is symmetric in α and v we can just interchange the roles of them in the proof if $i > j$). 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 lines of the proof of the last theorem we have:

$$\begin{aligned}
\sum_k \alpha_k v_{\pi(k)} &= \alpha_i v_1 + \sum_{k \neq i} \alpha_k v_{\pi(k)} \geq \frac{1}{\beta}\alpha_1 v_1 + \frac{1}{r_{n-1}} \left(\sum_{k=2}^{i} \alpha_{k-1} v_k + \sum_{k=i+1}^{n} \alpha_k v_k \right) \geq \\
&= \frac{1}{\beta}\alpha_1 v_1 + \frac{1}{r_{n-1}} \left[\sum_{k=2}^{i} (\alpha_{k-1} - \alpha_k) v_k + \sum_{k>1} \alpha_k v_k \right] \geq \\
&\geq \frac{1}{\beta}\alpha_1 v_1 + \frac{1}{r_{n-1}} (\alpha_1 - \alpha_i) v_i + \frac{1}{r_{n-1}} \sum_{k>1} \alpha_k v_k
\end{aligned}$$

Now, we can use $i \leq j$ to say: $v_i \geq v_j = \frac{1}{\gamma} v_1 \geq \left(1 - \frac{1}{\beta}\right) v_1$.

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

So, we would like to find some r_n such that we can say that $\sum_k \alpha_k v_{\pi(k)} \geq \frac{1}{r_n} \sum_k \alpha_k v_k$ for all $\beta \geq 1$, so we would like to have: $\frac{1}{r_n} \leq \min\left\{\frac{1}{r_{n-1}}, \frac{1}{\beta} + \frac{1}{r_{n-1}}\left(1-\frac{1}{\beta}\right)^2\right\}$ for any $\beta \geq 1$. But notice some other bound we can get is:

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

by following the lines of the proof of last theorem, but removing slot 1 and advertiser j in the inductive step. So another alternative is to get: $\frac{1}{r_n} \leq \min\left\{\frac{1}{r_{n-1}}, 1-\frac{1}{\beta}\right\}$ for every $\beta \geq 1$. So if we can get $1/r_n$ bounded by the maximum of those two quantities, we are done. Summarizing that, we need:

$$r_n \geq \max\left\{r_{n-1}, \left[\max\left\{1-\frac{1}{\beta}, \frac{1}{\beta} + \frac{1}{r_{n-1}}\left(1-\frac{1}{\beta}\right)^2\right\}\right]^{-1}\right\}$$

for all $\beta \geq 1$.

Now we need to evaluate for which value of $\frac{1}{\beta} \in (0,1]$ we have the minimum for $\max\left\{1-\frac{1}{\beta}, \frac{1}{\beta} + \frac{1}{r_{n-1}}\left(1-\frac{1}{\beta}\right)^2\right\}$. The minimum can be in two points: the minimum of the quadratic function or the intersection between those two functions. They intersect at $\frac{1}{\beta} = -r + 1 + \sqrt{r^2 - r}$ (where r stands for r_{n-1}) and the quadratic minimum is at $1 - \frac{1}{2}r$. So, for $r \geq \frac{4}{3}$, the minimum occurs in the intersection and for $r < \frac{4}{3}$, it occurs in the quadratic minimum. So:

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

since we want the smallest possible ratio. This allows to define r_k recursively from $r_2 = 1.25$ and it is easy to see that the sequence monotonically converges to the fixed point of that function which is the golden ration $\varphi = \frac{1+\sqrt{5}}{2} \approx 1.618$. This happens because the function that maps r_{n-1} to r_n is non-decreasing and has a fixed point in φ, so if $r_{n-1} \leq \varphi$ then $r_n \leq \varphi$.

■