Bayes-Nash Price of Anarchy for GSP

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@cs.cornell.edu

ABSTRACT

Generalized Second Price Auction, also known as Ad Word auction, and its variants have 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 under the Bayes-Nash solution concept (i.e., in a partial information setting). In this model, the value of each player for one click is drawn independently from a distribution. Each player knows his own value but he knows only the probability distribution of the other players values. We compare the expected social welfare in a Bayes-Nash equilibrium with the expected value of the optimal social welfare. We obtain a bound of 8 for the Bayes-Nash Price of Anarchy of GSP.

Our proof exhibits a combinatorial structure of Nash equilibria and use this structure to bound the price of anarchy. Our proof of this structural property uses novel combinatorial techniques that can be of independent interest.

1. INTRODUCTION

Search engines and other online information sources use sponsored search auction, or AdWord auctions, to monetize their services. 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. Since the introduction of the model, there has been much work in the area, see the survey of Lahaie et al [9].

Here we consider AdWords as a 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 [15] for a single item (or a single advertising slot). The Vickrey auction [15] for a single item makes truthful behavior (when the advertisers reveal their true valuation) dominant strategy, and make the resulting outcome maximize the social welfare. This truthful auction is extended to complex scenarios by the Vickrey-Clarke-Groves Mechanism (VCG) [3, 6] which also maximizes 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, and prove that it is within a small constant factor of the optimal.

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, we will focus on the case when $\gamma_j = 1$ for all j, that is, the probability of a click depends only on the slot. It is easy to extend our results to the model with separable click-through rates by considering the product $\gamma_j v_j$ in place of the values v_j for each player j.

In the full information model, it is known that there exists Nash equilibria that are socially optimal [4, 14] in both our simple model, and in the case of separable click-through rates, 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. This implies that the player takes unnecessary risk, as bidding above the valuation is a dominated strategy. We define conservative bidders as bidders who won't bid above their valuations. Paes Leme and Tardos [10] proved a bound of 1.618 to the Price of Anarchy for Pure Nash equilibria assuming that players are conservative.

In this paper, we consider the partial information setting of a Bayesian game (see Harsanyi [7]). In that setting, players have beliefs about other players valuations, i.e., that valuations of different players are drawn from independent distri-

*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.

bution, and the players know only the probability distribution of their opponents, not the valuations. The Bayesian setting models real AdWord Auctions better as players submit a single bid that will be used in many auctions with different competitors (e.g., depending on the exact search query), so players are, in fact, optimizing over a distribution of other players. Gomes and Sweeney [5] study the Generalized Second Price Auction in the Bayesian context. They show that, unlike the full information case, there may not exists socially optimal Nash equilibria in this model if the click-through rates of different slots are sufficiently close.

Our results. The main result of this paper is a constant Price of Anarchy bound for the Bayes-Nash equilibria for the GSP game assuming player's valuations are drawn from independent (though not necessarily identical) distributions, and bidders are conservative bidders. To motivate the conservative assumption, we observe that bidding above the players valuation is dominated strategy also in the Bayesian setting. We prove that the gap between the expectations of the optimal and any Bayes-Nash equilibrium is at most a factor of 8.

We exhibit a combinatorial structure of Bayes-Nash equilibria that can be of independent interest. To derive this structural characterization we use novel combinatorial techniques. We focus on a player i, and consider a set of multiple deviations obtained by conditioning on the position of this bidder in the optimal allocation. We show that these bids are monotone in the position, use a novel dual averaging technique to combine the Nash inequalities obtained for the separate deviations.

Our results differ significantly from the existing work on the price of anarchy in that many of the known results can be summarized via a smoothness argument, as observed by Roughgarden [12]. In contrast, we show that the GSP game is not smooth in the sense of [12].

Related work. Sponsored search has been a very active area of research in the last several years. See the survey of Lahaie et al [9] for a general introduction. Here we use the game theoretic model of the AdWord auctions of Edelman et al [4] and Varian [14], for a truthful auction see Aggarwal et all [1].

In the full information model, Edelman et al [4] and Varian [14] show that the Price of Stability for this game is 1 in separable click-through rates, that is, there exists Nash equilibria that are socially optimal. 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-rates by simply increasing their bid. Edelman et al [4] 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 [5] study Generalized Second Price Auction in the Bayesian context. They show that, unlike the full information case, there may not exists socially optimal Nash equilibria in this model, and obtain sufficient conditions on click-through rates that guarantee existence a symmetric and socially optimal equilibrium. In [5], all the valuations are drawn iid from the same distribution. In contrast, we consider all equilibria (not only the symmetric ones), prove bounds on the price of anarchy, and do not assume that the players distributing are identical.

Lahaie [8] was the first to try to quantify the social efficiency of an equilibrium in the worst case setting. 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 $\min\{\frac{1}{\delta}, 1 - \frac{1}{\delta}\}$. Later, Paes Leme and Tardos [10] give an 1.618 bound for the Price of Anarchy, without any assumptions on the click-through-rate structure. Thompson and Leyton-Brown [13] study the efficiency loss of equilibria empirically in various models.

In this paper, 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 [11] 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 [11] on greedy auctions is also closely related to our work. They analyze the Price of Anarchy of the auction game induced by Greedy Mechanisms. They consider a general combinatorial auction setting: greedy algorithms with payments are computed using the critical price. They show via a type of smoothness argument (see [12]) 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. The key to proving the $c + 1$ bound of Lucier and Borodin [11] 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 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. Many natural bidding languages have limited expressivity (like the GSP game), as allowing arbitrary complex bids typically 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 a 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 sequence, 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 using the same argument used in [10]. 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). Let $\pi(k)$ denote 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 $\pi(k)$ pays $b_{\pi(k+1)}$, which is the next highest bid, where we use the notation of $b_{\pi(n+1)} = 0$.

The vector π is a permutation that indicates which advertiser is assigned to each slot - 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 and the auctioneer get from playing it, which is: $\sum_j \alpha_j v_{\pi(j)}$. In this paper we are concerned with bounding the expected social welfare in an equilibrium of this game relative to the optimal welfare.

We consider the partial information setting of Harsanyi [7] where the players don't know the valuations of other players, only the distributions. We assume that the valuation v_i are drawn from independent (but not necessarily identical) 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:

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

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, and we define it by ν: let $\nu(i)$ 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 $\sum_j \alpha_{\nu(j)} v_j$. The quantity we want to bound is the Bayes-Nash price of Anarchy given by the ratio:

$$\text{Bayes-Nash PoA} = \mathbb{E}\left[\sum_j \alpha_{\nu(j)} v_j\right] / \mathbb{E}\left[\sum_j \alpha_j v_{\pi(j)}\right]$$

Overbidding is a dominated strategy. Even in the full information case (when the distributions over the valuations are fixed), and even for just two bidders, the gap between the best and the worse Nash equilibrium can be arbitrarily large as shown in [10]. This difficult exists even in truthful mechanisms, say the simple single-slot Vickrey auction: a player with zero value can make a very large bid and a bidder with high value bids zero and this is an equilibrium, even though the high-bidding player is not gaining by doing that. A way of overcoming it is by supposing that players do not play dominated strategies. It is not hard to see that bidding above the valuation is a dominated strategy.

LEMMA 2.1. *A bidding function $b_i(v_i)$ 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 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. As stated by the lemma above, this is a weaker assumption than that players don't play dominates strategies.

3. BAYES-NASH PRICE OF ANARCHY

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

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

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 Pure Nash equilibria in [10], but much harder to prove. A special case of the characterization of [10] (for $j = \sigma(\nu(i))$) can be written as:

$$v_i \alpha_{\sigma(i)} + \alpha_{\nu(i)} v_{\pi(\nu(i))} \geq v_i \alpha_{\nu(i)}. \tag{1}$$

We prove the following Bayesian version of this inequality, where all expectations are conditioned on the valuation v_i of player i, and are over the valuation of other players, and the random choices made in bidding:

LEMMA 3.2. *If $\{b_i(\cdot)\}_i$ is a Bayes-Nash equilibrium of the GSP game using conservative bids then:*

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

The price of anarchy bound follows easily from the lemma.

Proof of Theorem 3.1 :

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

■

The hard part of the proof is proving Lemma 3.2. The main difficulty in the Bayesian setting is that the inequality is not established by a single deviating bid. The structural inequality (1) for pure Nash in the full information setting was obtained by considering a single deviation: player i bidding just above $b_{\pi(\nu(i))}$, the bid that is allocated the position bidder i occupies in the optimum. By the conservative assumption $b_{\pi(\nu(i))} \leq v_{\pi(\nu(i))}$, and then we get the claimed inequality from the Nash property that the value of the deviating bid $\alpha_{\nu(i)}(v_i - b_{\pi(\nu(i))})$ is no bigger than the Nash value for player i which is at most $v_i \alpha_{\sigma(i)}$. We can get a similar characterization for mixed Nash equilibria (loosing a factor of 2) by considering the 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 3.3. *A deviating bid B by player i can get a slot k or better if and only if $B \geq b_{\pi^i(k)}$.*

To extend the bid $2\mathbb{E}b_{\pi(\nu(i))}$ 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 function. We will drop the dependence on v_i for notational convenience.

The proof of Lemma 3.2, 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 for any fixed value of v_i. This will allows 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 $\nu(i) = k'$, as $B_k \geq B_{k'}$.

LEMMA 3.4. *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 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 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.

To be able to combine the inequalities we get by a considering the different bids B_k we use 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. We expect that this Lemma 3.5, which we prove using linear programming duality, can have other applications.

LEMMA 3.5. *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 3.2:

Proof of Lemma 3.2 : As outlined above we will consider n deviation for player i at bids B_k for all possible slots k. Since b is a Nash equilibrium, a player i cannot benefit from changing his/her strategy, each alternative bid will give us an inequality on the utility. We will use Lemma 3.5 to average them to get the claimed inequality.

Consider 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, as 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 3.3, 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 3.4 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 $v_i\mathbb{E}[\alpha_{\sigma(i)}|v_i]$. Using the above bound, this leads to the following inequality.

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

Now, we use the Lemma 3.5 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:

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

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

3.1 Proving that bids B_k are non-increasing

We will prove Lemma 3.4 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 use these to define a probability flow from the probability space where $\nu(i) = k$ to the space where $\nu(i) = k + 1$ so that the value $b_{\pi^i \nu(i)}$ is non-increasing along the flow lines. In transferring the probability pass we take advantage of the fact that the valuations are drawn from independent distributions.

Proof of Lemma 3.4 : 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:

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

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

CLAIM 3.6. *For a set S', and $S = S' \cup \{t\}$ for $t \neq i$*

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

PROOF. 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 (bidder t) change, we can't violate the above inequality. Taking the expectation over $\{v_j\}_{j \neq i,t}$ we get the inequality of Claim 3.6. ∎

To finish the proof of Lemma 3.4, 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, we need to show that there are coefficients $\lambda_{S,S'} \geq 0$ for $S' \subseteq S$ such that:

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

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

Taking the a linear combination of the inequalities (3.6) for set pair (S, S') with coefficients $\lambda_{S,S'}$ lets the bound claimed by Lemma 3.4. ∎

LEMMA 3.7. *If valuations are drawn from independent distributions, 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:

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

The existence of the $\lambda_{S,S'}$ 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' of size $|S'| = n-k-1$ with supply $\frac{\phi(S')}{\sum_{T'}\phi(T')}$ and the right nodes are sinks corresponding to sets S of size $|S| = n-k-1$ with demand $\frac{\phi(S)}{\sum_{T}\phi(T)}$, where the sums are over sets T' of size $n-k$ and sets T of size $n-k$ respectively. We add an edge (S', S) if $S' \subseteq S$ with capacity ∞. 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 min-cut/max-flow theorem (in this case, this is a weighted version of Hall's Theorem): there is a flow of size 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 as big as 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 the at least same multiplicity we are done. We prove it 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 $(S-j, A'_i + j)$ for $j \in S \setminus A'_i$. We want to map each pair (S, A'_i) to $(S-j, A'_i+j)$ 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, just need to map each equivalence class of elements in a collision-free manner. The Lemma 3.8 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 $j = f_t(A'_i \setminus S) \setminus A'_i$. ■

LEMMA 3.8. *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, the claim is proved. ■

3.2 Proving the dual averaging Lemma

Proof of Lemma 3.5 : We want to prove that the following linear programming problem is feasible:

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

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 high. 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 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

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

And this expression is non-negative, 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.

■

4. GSP IS NOT A SMOOTH GAME

As Roughgarden points out in [12], most of the games studied so far (as congestion games, facility location, valid utility games, etc) 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 4.1. *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 $\lambda > 0$, ■

5. 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] 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.

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

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

[8] 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.

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

[10] R. P. Leme and E. Tardos. Sponsored search equilibria for conservative bidders. In *AdAuctions Workshop*, 2009.

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