

GSP Auctions with Correlated Types

Brendan Lucier^{*}
Dept of Computer Science
University of Toronto
blucier@cs.toronto.edu

Renato Paes Leme[†]
Dept of Computer Science
Cornell University
renatopl@cs.cornell.edu

ABSTRACT

The Generalized Second Price (GSP) auction is the primary method by which sponsored search advertisements are sold. We study the performance of this auction in the Bayesian setting for players with correlated types. Correlation arises very naturally in the context of sponsored search auctions, especially as a result of uncertainty inherent in the behaviour of the underlying ad allocation algorithm. We demonstrate that the Bayesian Price of Anarchy of the GSP auction is bounded by 4, even when agents have arbitrarily correlated types. Our proof highlights a connection between the GSP mechanism and the concept of smoothness in games, which may be of independent interest.

For the special case of uncorrelated (i.e. independent) agent types, we improve our bound to $2(1 - 1/e)^{-1} \approx 3.16$, significantly improving upon previously known bounds. Using our techniques, we obtain the same bound on the performance of GSP at coarse correlated equilibria, which captures (for example) a repeated-auction setting in which agents apply regret-minimizing bidding strategies. Moreover, our analysis is robust against the presence of irrational bidders and settings of asymmetric information, and our bounds degrade gracefully when agents apply strategies that form only an approximate equilibrium.

Categories and Subject Descriptors

J.4 [Social and Behavioral Sciences]: Economics; F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems

General Terms

Algorithms, Economics, Theory

^{*}Part of this work was performed while the author was visiting Microsoft Research, New England.

[†]Supported in part by NSF grant CCF-0729006. Part of this work was performed while the author was visiting Microsoft Research, New England.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

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

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

Keywords

Sponsored search auctions, Price of anarchy

1. INTRODUCTION

The sale of advertising space is the primary source of revenue for many providers of online services. This is due, in part, to the fact that providers can tailor advertisements to the preferences of individual users. A search engine, for example, can choose to display ads that synergize well with a query being searched. However, such dynamic provision of content complicates the process of selling ad space to potential advertisers. The now-standard method has advertisers place bids – representing the amount they would be willing to pay per click – which are resolved in an automated auction whenever ads are to be displayed.

By far the most popular bid-resolution method currently in use is the Generalized Second Price auction (GSP), a generalization of the well-known Vickrey auction. In the GSP, there are multiple ad “slots” of varying appeal (i.e. slots at the top of the page are more effective). Advertisers are assigned slots in order of their bids, with the highest bidders receiving the best slots; each advertiser then pays an amount equal to the bid of the next-highest bidder. While simple to understand and use, the GSP has some notable drawbacks: unlike the Vickrey auction it is not truthful, and it does not generally guarantee the most efficient outcome (i.e. the outcome that maximizes social welfare). Nevertheless, the use of GSP has been extremely successful in practice. This begs the question: *are there theoretical properties of the Generalized Second Price auction that would explain its prevalence?*

Since the GSP auction is invoked every time a user queries a keyword of interest, it is best characterized as a repeated auction in which players repeatedly bid for ad slots. However, modeling equilibrium strategies in a repeated game of this nature is notoriously difficult; optimal bidding strategies must account for threatened behaviour in future rounds, optimal exploration of the bidding space, and so on. A common simplification used in the literature is to focus on auctions for a single keyword, and suppose that agents will quickly learn each others’ valuations and reach a stationary equilibrium. Under this assumption, the stationary equilibrium would correspond naturally to a Nash equilibrium in the full information, one-shot version of the GSP auction [7]. It has therefore become common practice to study pure, full-information equilibria of the one-shot game, as an approximation to expected behaviour in the more general repeated game [6, 24, 20].

In reality, however, the set and types of players can vary significantly between rounds of a GSP auction. The reason for this is that each query is unique, in the sense that it is defined not only by the set of keywords invoked but also the time the query was performed, the location and history of the user and many other factors. This *context* is taken into account by an underlying *ad allocation algorithm*, which is controlled by the search engine. The ad allocation algorithm not only selects which players will participate in an auction instance, but also assigns a *quality factor* to each player: a score that measures how likely that participant’s ad will be clicked for that query. These quality factors are then used to scale the bids of the advertisers. The effective bid and effective type of a player are therefore random variables, which can be thought as the original values times some exogenous quality score. Thus, even if agents converge to a stationary bidding pattern, it is not necessarily the case that this implies full certainty about agent types in each round.

The uncertainty about the types of other agents can be modelled in a Bayesian, partial information setting. Bayes-Nash equilibria for GSP have been studied previously [8, 19], but these works require that valuations are drawn from independent distributions. However, the independence assumption fails to capture the source of randomness due to query context and the ad allocation algorithm, which inherently introduces correlations among agent values. Our main contribution is to extend the analysis of quality of outcomes of Bayesian GSP to cases where player values are arbitrarily correlated, thereby capturing sources of uncertainty that occur in practical sponsored search auctions.

Our techniques are general enough to apply even when our model is extended to capture various additional features of agent behaviour. For instance, there are different types of players in the AdMarkets, which may have differing degrees of information about their competitors. Some smaller players (such as individual advertisers) might be clueless about the valuations of the other players and expected behaviour of quality scores, while others (say bidding agencies or large companies with web advertising departments) may have a much better understanding of how individual rounds of the auction will proceed. Even among those, there are various levels of market research one can perform. Our bound on social efficiency in the Bayesian model holds in settings with this natural asymmetry in information.

One of the fundamental assumptions in auction analysis is that all players are perfectly rational, i.e., utility optimizers. However, in reality (and especially in large online settings), it is natural to assume that some fraction of the players participating in an AdAuction might have particularly unsophisticated bidding strategies. In fact, some agents may not even play at equilibrium in the single-shot approximation of the GSP auction, or may only be able to converge to an approximate equilibrium. We discuss the robustness of our bounds to the presence of players bidding with limited (or no) rationality.

We also consider an equilibrium model suited to long-run bidder behaviour in GSP, where it is not required that agents converge to a stable equilibrium. Namely, returning to the repeated-game form of GSP, we consider situations in which agents play in order to minimize *regret*. Such strategies are not necessarily in equilibrium in the repeated game, but capture the intuition that agents attempt to learn beneficial bidding strategies over time. Roughly speaking, such

a model assumes that agents observe the bidding patterns of others and modify their own bids in such a way that their long-term performance approaches that of a single optimal strategy chosen in hindsight. It is well-known from learning theory that such regret minimization is easy to achieve via simple bidding techniques [11]. We demonstrate that our techniques apply to this setting and bound the Price of Total Anarchy [2], which is the ratio between the social welfare of the optimal allocation and the average social welfare obtained by GSP when agents minimize regret over a sufficiently long number of rounds.

Results.

Our main result is a bound on the social welfare obtained at Bayes-Nash equilibrium for the GSP auction. Specifically, we show that the Bayesian Price of Anarchy for GSP is at most 4 for correlated valuations. At the heart of this proof lies a structural characterization of the GSP auction that has a flavor similar to the characterization of smooth games introduced by Roughgarden [23], even though it is known that GSP is not a smooth game. Furthermore, if valuations can be assumed to be independent, we can improve this bound to $2(1 - 1/e)^{-1} \approx 3.164$. This improves upon the previous best-known bounds of 8 for BNE and 4 for (mixed) NE [19].

Perhaps just as important as the improved bounds, however, is the straightforward and robust nature of our analysis. In particular, our results extend to give the same bound for coarse correlated equilibria (see Section 4.2), which implies that the social welfare when agents play repeatedly in order to minimize total regret is within a $2(1 - 1/e)^{-1}$ factor of the optimal social welfare. Also, we prove that the factor of 4 for agents with correlated distributions continues to hold even if agents have asymmetric access to distributional information, in the form of exogenously provided signals.

Moreover, our results are resilient against the presence of irrational agents, in the following sense. Suppose that, in addition to the rational participants in the auction, there is also some set of agents who apply arbitrary strategies. We can view these as irrational participants who do not understand how to bid strategically. Note that, in such a setting, it is not possible for an auction to guarantee a fraction of the social welfare obtainable from the irrational bidders; after all, a bidder with very large value may decide (irrationally) to bid 0 and effectively not participate in the auction. What we *can* show, however, is that the presence of the irrational bidders does not interfere with the auction’s ability to approximate the welfare obtainable from the rational bidders. That is, the ratio of the optimal social welfare *of the rational bidders* to the total social welfare obtained at any BNE is at most $2(1 - 1/e)^{-1} \approx 3.164$. This result requires an assumption on the play of the irrational bidders, which is that no player bids more than his true value. We feel that this is a reasonable assumption, as overbidding is a dominated strategy that is easily avoided; we therefore view the irrational bidders as novice or uninformed participants who would avoid dominated strategies, rather than truly adversarial agents. Finally, if rational agents are limited in their rationality in that they can be assumed only to reach a δ -approximate equilibrium, our welfare bounds continue to hold with the addition of an extra factor of $(1 + \delta)$.

Our results hold for a standard model of separable click-through rates, where the probability that a user clicks on

an advertisement j in slot i is of the form $\alpha_i \gamma_j$. That is, it is a product of two separable components: one for the advertiser, and one for the slot. For ease of exposition, we will focus on the special case that $\gamma_j = 1$ for all j . However, we note that our results extend to the more general case of separable click-through rates.

Related work.

In recent years there has been a surge of work on algorithmic mechanism design for sponsored search, beginning with Mehta et al. [17, 16]. See the survey of Lahaie et al [13] for an overview of subsequent developments. The GSP model applied in this manuscript is due to Edelman et al [6] and Varian [24].

The work most closely related to ours is that of Paes Leme and Tardos, who also study equilibria of GSP [19]. They give upper bounds on the Price of Anarchy in pure, mixed, and Bayesian strategies; achieving bounds of 1.618, 4, and 8, respectively. Our main result is a simplification and strengthening of their results for the mixed and Bayesian cases, as well as an extension to different but related solution concepts.

The study of Price of Anarchy for non-truthful auction mechanisms (specially in the Bayesian setting) is a quite successful research line initiated by Christodoulou, Kovács and Schapira [4] and developed in Borodin and Lucier [15], Lucier [14], Paes Leme and Tardos [19] and most recently in the work of Bhawalkar and Roughgarden [1]. To the best of our knowledge, this is the first paper in that research line whose Price of Anarchy bounds hold when player valuation are drawn from a correlated distribution. In the study of optimal truthful mechanism design, correlated valuations have recently been studied by Papadimitriou and Pierrakos [22] and Dobzinski, Fu and Kleinberg [5].

Edelman et al [6] and Varian [24] study Envy-free equilibria of GSP (a special case of Nash equilibrium) in the full information setting. They demonstrate that such equilibria exist, and that all such equilibria are socially optimal. Gomes and Sweeney [8] study the Generalized Second Price Auction in the Bayesian context. They show that, unlike the full information case, there may not exist symmetric or socially optimal equilibria in this model, and obtain sufficient conditions on click-through-rates that guarantee the existence of a symmetric and efficient equilibrium. Lahaie [12] also considers the problem of bounding the social welfare obtained at equilibrium, but restricts attention to the special case that click-through-rate α_i decays exponentially along the slots with a factor of $\frac{1}{\delta}$. Under this assumption, Lahaie proves a price of anarchy of $\min\{\frac{1}{\delta}, 1 - \frac{1}{\delta}\}$.

The study of regret-minimization goes back to the work of Hannan on repeated two-player games [9]. Kalai and Vempala [11] extend the work of Hannan to online optimization problems, and Kakade et al [10] further extend to settings of approximate regret minimization. Blum et al [2] apply regret-minimization to the study of inefficiency in repeated games, coining the phrase “price of total anarchy” for the worst-case ratio between the optimal objective value and the average objective value when agents minimize regret.

Many of the above solution concepts have been tied together by Roughgarden [23], who showed that for a class of *smooth* games the price of anarchy, Bayesian price of anarchy, and price of total anarchy are identical. The smoothness criterion is quite strong, and does not apply in general to the

GSP mechanism. However, we note that the methods used to analyze mechanisms at Nash Equilibrium share a common flavour with the results for smooth games. As part of the proof of our main result, we isolate a property related to smoothness that encapsulates many of the insights that drives our bounds.

In an independent work, Caragiannis et al [3] also improve the bounds in [19]. The authors prove a bound of 1.282 for the Full Information setting, almost matching the 1.259 lower bounds, and improve and extend the techniques of Paes Leme and Lucier [21] (a previous arXiv version of the present paper), showing bounds of 2.310 for correlated equilibria in the Full Information setting and 3.037 for Bayes-Nash equilibria with independent types.

2. PRELIMINARIES

We consider an auction with n advertisers and n slots¹. Each advertiser i has a private type v_i , representing his or her value per click received. The sequence $\mathbf{v} = (v_1, \dots, v_n)$ is referred to as the *type profile*. We will write \mathbf{v}_{-i} for \mathbf{v} excluding the i th entry, so that $\mathbf{v} = (v_i, \mathbf{v}_{-i})$.

An *outcome* is an assignment of advertisers to slots. An outcome can be viewed as a permutation π with $\pi(k)$ being the player assigned to slot k . When advertiser i is assigned to the k -th slot, he gets $\alpha_k \gamma_i$ clicks, where $\alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_n$ are called *click-through-rates* and γ_i are called *quality factors*. Without loss of generality and for clarity of exposition, we consider for the rest of the paper $\gamma_i = 1$ for all players.

A mechanism for this auction elicits a bid $b_i \in [0, \infty)$ from each agent i , which is interpreted as a type declaration, and returns an assignment as well as a price p_i per click for each agent. We assume $b_i \leq v_i$, since overbidding is clearly a dominated strategy. If advertiser i is assigned to slot j , his *utility* is $\alpha_j(v_i - p_i)$, which is the number of clicks received times profit per click. The *social welfare* of outcome π is $SW(\pi, \mathbf{v}) = \sum_j \alpha_j v_{\pi(j)}$, the total value of the solution for the participants. The optimal social welfare is $OPT(\mathbf{v}) = \max_{\pi} SW(\pi, \mathbf{v})$.

We focus on a particular mechanism, the Generalized Second Price auction, which works as follows. Given bid profile \mathbf{b} , the auction sets $\pi(k)$ to be the advertiser with the k th highest bid (breaking ties arbitrarily). That is, GSP assigns slots with higher click-through-rate to agents with higher bids. Payments are then set according to $p_i = b_{\pi(\pi^{-1}(i)+1)}$. That is, the payment of the k th highest bidder is precisely the bid of the next-highest bidder (where we take $b_{n+1} = 0$). We will write $u_i(\mathbf{b})$ for the utility derived by agent i from the GSP when agents bid according to \mathbf{b} .

For the remainder of the paper, we will write $\pi(\mathbf{b}, j)$ to be the player assigned to slot j by GSP when the agents bid according to \mathbf{b} . We will also write $\sigma(\mathbf{b}, i)$ for the slot assigned to bidder i by GSP, again when agents bid according to \mathbf{b} . In other words, $\sigma(\mathbf{b}, \cdot) = \pi^{-1}(\mathbf{b}, \cdot)$. We write $\pi^i(\mathbf{b}_{-i}, j)$ to be the player that would be assigned to slot j by GSP if agent i did not participate in the auction. We will write $\nu(\mathbf{v})$ for the optimal assignment of slots to bidders for value profile \mathbf{v} , so that $\nu(\mathbf{v}, i)$ is the slot that would be allocated to agent i in the optimal assignment².

¹we handle unequal numbers of slots and advertisers by adding virtual slots with click-through-rate zero or virtual advertisers with zero value per click.

²We note that, since GSP makes the optimal assignment for

2.1 Pure and Mixed Nash Equilibrium

A (*pure*) strategy for agent i is a function $b_i : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ that maps each private value to a declared bid. A *mixed strategy* maps a private value to a distribution over bids, corresponding to a randomized declaration. Given a value profile \mathbf{v} , we say that strategy profile \mathbf{b} is a mixed Nash equilibrium if, for all i and all alternative strategies $b'_i(\cdot)$,

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

That is, each agent i maximizes his utility by bidding according to strategy $b_i(\cdot)$. We say this is a pure Nash equilibrium if, in addition, all strategies are pure. We define the (*mixed*) *Price of Anarchy* to be the worst-case ratio between social welfare in the optimum and expected social welfare in GSP across all valuation profiles and all mixed Nash equilibria:

$$\sup_{\mathbf{v}, \mathbf{b}(\cdot)_{NE}} \frac{OPT(\mathbf{v})}{\mathbb{E}_{\mathbf{b}}[SW(\pi(\mathbf{b}(\mathbf{v})), \mathbf{v})]}.$$

2.2 Bayesian setting

In a Bayesian setting, we suppose that each agent's type is drawn from a publicly known (possibly correlated) distribution, i.e., $\mathbf{v} \sim \mathbf{F}$. We then say that strategy profile \mathbf{b} is a Bayes-Nash equilibrium for distributions \mathbf{F} if, for all i , all v_i , and all alternative strategies b'_i ,

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

That is, each agent maximizes his expected utility by bidding in accordance with strategy $b_i(\cdot)$, where expectation is taken over the distribution of the other agents' types and any randomness in their strategies. We define the *Bayes-Nash Price of Anarchy* to be the worst-case ratio between social welfare in the optimum and social welfare in GSP across all distributions and all Nash equilibria:

$$\max_{\mathbf{F}, \mathbf{b}(\cdot)_{BNE}} \frac{\mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[OPT(\mathbf{v})]}{\mathbb{E}_{\mathbf{v} \sim \mathbf{F}, \mathbf{b}(\mathbf{v})}[SW(\pi(\mathbf{b}(\mathbf{v})), \mathbf{v})]}.$$

2.3 Signals and Information Asymmetry

We define an extension of the setting above, incorporating a Bayesian version of information asymmetry. In this model, each player's type consists of a value v_i and a signal s_i drawn from an arbitrary signal space. The signals can be thought of as privately-gained insight that refines agent i 's conditional distribution over the space of other agents' types. It is publicly known that (\mathbf{v}, \mathbf{s}) comes from a certain distribution \mathbf{F}' , which can be arbitrarily correlated. As a special case, one could choose for the value of a player to be a function of his signal (or even the vector of signals).

The presence of signals captures the notion that some agents might have a better potential to infer the other players' valuations than others, or may be endowed with privileged information. We do note, however, that agents are aware of the bidding strategies $\mathbf{b}(\cdot, \cdot)$ and the distribution \mathbf{F}' , so that agents can rationalize about the effects of signals upon the bidding behaviour of their opponents.

In this model, a strategy is a bidding function mapping (v_i, s_i) to a distribution of possible bids. The bid profile is

given a bid declaration, we actually have that $\nu(\mathbf{v}, i)$ and $\sigma(\mathbf{v}, i)$ are identically equal. We define ν mainly for use when emphasizing the distinction between an optimal assignment for a value profile and the assignment that results from a given bid profile.

a Bayes-Nash in the asymmetric information model if:

$$\begin{aligned} \mathbb{E}_{\mathbf{v}_{-i}|v_i, s_i}[u_i(b_i(v_i, s_i), \mathbf{b}_{-i}(\mathbf{v}_{-i}, s_{-i}))] &\geq \\ \mathbb{E}_{\mathbf{v}_{-i}|v_i, s_i}[u_i(b'_i(v_i, s_i), \mathbf{b}_{-i}(\mathbf{v}_{-i}, s_{-i}))] &\end{aligned}$$

2.4 Repeated Auctions and Regret Minimization

We now consider the GSP auction in a repeated setting. In this model, the GSP auction is run $T \geq 1$ times with the same slots and agents. The private value profile \mathbf{v} of the agents does not change between rounds, but the agents are free to change their bids. We write b_i^t for the bid of agent i on round t . We refer to $D = (b^1, \dots, b^T)$ as a *declaration sequence*. We will write $\pi(D)$ for the sequence of permutations generated by GSP on input sequence D . The average social welfare generated by GSP is then $SW(\pi(D), \mathbf{v}) = \frac{1}{T} \sum_t SW(\pi(\mathbf{b}^t), \mathbf{v})$.

The full range of equilibria in such a repeated game is very rich, so we restrict ourselves to a particular non-equilibrium form of play that nevertheless captures the intuition that agents learn appropriate bidding strategies over time, without necessitating convergence to a stationary equilibrium.

We say that declaration sequence $D = (b^1, \dots, b^T)$ *minimizes external regret* for agent i if, for any fixed declaration b_i ,

$$\sum_t u_i(b_i^t, b_{-i}^t) \geq \sum_t u_i(b_i, b_{-i}^t) + o(T).$$

That is, as T grows large, the utility of agent i approaches the utility of the optimal fixed strategy in hindsight. The *Price of Total Anarchy* [2] is the worst-case ratio between social welfare in the optimum and the average social welfare obtained by GSP across all declaration sequences that minimize external regret for all agents. That is, the price of total anarchy is

$$\lim_{T \rightarrow \infty} \max_{\mathbf{v}, D} \frac{OPT(\mathbf{v})}{SW(\pi(D), \mathbf{v})}$$

where the maximum is taken over declaration sequences that minimize external regret for all agents.

We note that the above model can be generalized to account for player uncertainty between rounds. For instance, one could assume that there is a large pool Ω of players, and on each round t the ad allocation algorithm (modeled as a randomized exogenous process) selects a subset of players $S_t \subseteq \Omega$ to participate in the auction that round. In this case, each player $i \in \Omega$ has a fixed valuation v_i , but agents are necessarily uncertain of the valuations (i.e. the identities) of their opponents on each round. The notion of regret minimization extends readily to such partial-information settings: if we write $\tau(T, i)$ for the set of rounds $t \leq T$ in which player i is selected, we say that the bid sequence b_i^t minimizes external regret if, for all b'_i ,

$$\sum_{t \in \tau(T, i)} u_i(b_i^t, b_{-i}^t) \geq \sum_{t \in \tau(T, i)} u_i(b'_i, b_{-i}^t) + o(|\tau(T, i)|).$$

Such a model more directly captures the fact that an auction participant cannot perfectly predict the set of competitors she will face on any given round. A more thorough discussion of such generalized models appears in the full version of this paper.

3. CORRELATED PRICE OF ANARCHY

Our main result is a bound of 4 on the Bayes-Nash PoA for GSP, which holds even if we allow agent types to be arbitrarily correlated. The same result holds for the Bayesian model of information asymmetry as well - extending the proof to this case is straightforward - but for clarity of exposition, we will give the proof without considering signals.

Theorem 1 *The Correlated Bayesian Price of Anarchy of GSP is at most 4.*

Our proof of Theorem 1 is related to the concept of smoothness due to Roughgarden [23]. Write $SW(s)$ for the social welfare generated by strategy profile s for some (implicitly defined) game. Roughgarden defines a game as (λ, μ) -smooth if, for all pairs of strategy profiles \mathbf{s}, \mathbf{s}' , we have

$$\sum_i u_i(s'_i, s_{-i}) \geq \lambda SW(s') - \mu SW(s).$$

Roughly speaking, smoothness captures the property that if strategy profile s' results in a significantly larger social welfare than another strategy profile s , then this gap in welfare can be captured by the marginal increases in the individual agents' utilities when unilaterally switching their strategies from s to s' .

Paes Leme and Tardos showed in [18] that GSP auction is not smooth for any choice of parameters. The intuition is that in a mechanism game, welfare is determined not by strategy profiles directly, but through the intermediate space of outcomes. Thus, even if a bidding strategy s' yields an outcome that is much more efficient than s , it may be that the outcome does not change when any single agent deviates from s to s' . Nevertheless, we now note that GSP satisfies a conceptually weaker but related property: there is a *particular* strategy profile s' that maximizes social welfare, and for which a smoothness-like inequality holds for any other strategy profile s .

Definition 2 (Semi-Smooth Games) *We say that a game is (λ, μ) -semi-smooth if there exists some s' maximizing the social welfare such that, for any strategy profile s ,*

$$\sum_i u_i(s'_i, s_{-i}) \geq \lambda SW(s') - \mu SW(s).$$

As we now demonstrate, the GSP auction corresponds to an semi-smooth game.

Lemma 3 *GSP in the full information setting is $(\frac{1}{2}, 1)$ -semi-smooth.*

PROOF. Fix valuation profile \mathbf{v} and agent i . Consider any bid profile \mathbf{b} and define $b'_i = \frac{v_i}{2}$. We claim that

$$u_i\left(\frac{v_i}{2}, \mathbf{b}_{-i}\right) + \alpha_i v_{\pi(\mathbf{b}, i)} \geq \frac{1}{2} \alpha_i v_i.$$

Summing the above expression for all i gives us the desired result. To see that the above expression holds, consider two cases:

1. By switching his bid from $b_i(v_i)$ to $v_i/2$, player i wins some slot $j \leq i$. In this case $u_i(\frac{v_i}{2}, \mathbf{b}_{-i}) \geq \frac{1}{2} \alpha_i v_i$.
2. Otherwise, the player who wins slot i under bidding profile \mathbf{b} is bidding more than $v_i/2$. In this case $\alpha_i v_{\pi(i)} \geq \frac{1}{2} \alpha_i v_i$.

We now prove Theorem 1 using the semi-smoothness of GSP. Notice that the statement of semi-smoothness doesn't place any restrictions on strategy profile \mathbf{b} , and in particular does not require that it be a Nash equilibrium. So, in the Bayesian setting, for any type profile \mathbf{v} and for any $\mathbf{b}(\mathbf{v})$, Lemma 3 implies that

$$\sum_i u_i(b'_i(v_i), \mathbf{b}_{-i}(\mathbf{v}_{-i})) \geq \frac{1}{2} OPT(\mathbf{v}) - SW(\pi(\mathbf{b}(\mathbf{v})), \mathbf{v}).$$

for $b'_i(v_i) = v_i/2$.

A proof of Theorem 1 is now straightforward by taking expectations and using the Bayes-Nash inequality for a given Bayes-Nash equilibrium strategy profile \mathbf{b} :

$$\begin{aligned} \mathbb{E}[SW(\pi(\mathbf{b}(\mathbf{v})), \mathbf{v})] &= \sum_i \mathbb{E}_{v_i} \mathbb{E}_{\mathbf{v}_{-i}|v_i} [u_i(b_i(v_i), \mathbf{b}_{-i}(\mathbf{v}_{-i}))] \\ &\geq \sum_i \mathbb{E}_{v_i} \mathbb{E}_{\mathbf{v}_{-i}|v_i} [u_i(b'_i(v_i), \mathbf{b}_{-i}(\mathbf{v}_{-i}))] \\ &\geq \frac{1}{2} \mathbb{E}[OPT(\mathbf{v})] - \mathbb{E}[SW(\pi(\mathbf{b}(\mathbf{v})), \mathbf{v})]. \end{aligned}$$

which proves the Price of Anarchy of 4.

4. INDEPENDENT VALUATIONS

We now wish to tighten the bound obtained in Theorem 1 beyond 4, for the special case that player valuations are not correlated. Let us begin by describing the high-level idea, and the way in which it relies on type independence. We can rephrase the main idea of Theorem 1 as follows: at equilibrium, for each bidder i , it must be that the expected outcome for bidder i is no worse than the expected outcome that results from applying any alternative strategy. In particular, this must be true for the strategy that prescribes a bid equal to half of bidder i 's true value. This, in turn, implies that the expected bid for slots much better than the expected allocation to agent i must be at least $v_i/2$, as otherwise agent i would have incentive to deviate. However, this suggests that, on average, a bidder with value at least $v_i/2$ will obtain such high-valued slots, which makes up for any loss in social welfare due to those slots not being assigned to agent i .

The key to the above argument was in identifying an alternative strategy for agent i to consider - namely, bidding half of his true value - and using the fact that this alternative strategy either gives player i a high utility or the player currently in slot i is paying a high price. This strategy has the strong property that it does not depend on the expected bids of the other agents; it is in this way that the argument is related to smoothness. However, more generally, we may be able to obtain tighter bounds by reasoning about strategies that are motivated by other agents' behaviour at equilibrium. In particular, it may be that a certain gap in social welfare implies the existence of some agent with some beneficial deviating strategy, but it might not be the case that a single strategy leads to an improvement in utility in every instance.

Indeed, what we will show is that if, for a given fixed valuation profile, the expected winning bid for some slot i is sufficiently low, then there must exist a deviating strategy for agent i that improves his welfare. The precise nature

of this strategy will depend on the particular strategies employed by the other agents. Note that this argument does not immediately imply a bound on the expected welfare of GSP, since the existence of a deviating strategy for agent i conditions on agent i 's type, whereas the expected social welfare is in expectation over all types. It is here that we use independence to convert conditional bounds into bounds on the overall expected social welfare.

In the remainder of the section we will first formalize the above argument, then demonstrate that the analysis extends to coarse correlated equilibria, outcomes of learning strategies, and the presence of irrational agents.

4.1 Performance at Bayes-Nash Equilibrium

We are now ready to prove the following theorem.

Theorem 4 *The Bayesian Price of Anarchy of GSP is at most $2(1 - 1/e)^{-1} \approx 3.164$, assuming that agent types are independently distributed.*

The proof of Theorem 4 proceeds in two steps. We first show that a structural property of bidding profiles implies a bound on the social welfare obtained by GSP (Lemma 5). We then show that this structural property holds at all BNE of the GSP (Lemma 6).

Lemma 5 *Suppose that $\mathbf{v} \sim \mathbf{F}$ and agents apply strategy profile $\mathbf{b}(\cdot)$. Suppose further that the following is true:*

$$\mathbb{E}_{\mathbf{v}_{-i}}[\alpha_{\sigma(\mathbf{b}(\mathbf{v}),i)}v_i] + \mathbb{E}_{\mathbf{v}_{-i}}[\alpha_k b_{\pi^i(\mathbf{b}_{-i}(\mathbf{v}_{-i}),k)}] \geq \gamma \alpha_k v_i \quad (1)$$

for all slots k , players i , and values v_i . Then

$$\mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[SW(\pi(\mathbf{b}(\mathbf{v})), \mathbf{v})] \geq \frac{1}{2} \gamma \mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[OPT(\mathbf{v})].$$

Lemma 6 *At any BNE of GSP, (1) holds with $\gamma = 1 - \frac{1}{e}$.*

Lemma 5 and Lemma 6 immediately imply Theorem 4. Note that it is in Lemma 5 that we make use of the assumption that agent types are independent.

Proof of Lemma 5 : Fix some value profile \mathbf{v} . For notational convenience, let Γ be the induced distribution on bid profiles $\mathbf{b} = \mathbf{b}(\mathbf{v})$ when $\mathbf{v} \sim \mathbf{F}$. Then for any player i , value v_i , and slot k , if we write $b_i = b_i(v_i)$, then we can express (1) as:

$$\mathbb{E}_{\mathbf{b}_{-i}}[\alpha_{\sigma(\mathbf{b},i)}v_i] + \mathbb{E}_{\mathbf{b}_{-i}}[\alpha_k b_{\pi^i(\mathbf{b}_{-i},k)}] \geq \gamma \alpha_k v_i.$$

This inequality essentially states that, under the bidding strategies employed by the other agents, the expected value of the slot obtained by agent i plus the expected bid required to obtain slot k is at least γ times the value agent i would obtain from slot k . Note that \mathbf{v}_{-i} does not appear in this expression; bids \mathbf{b}_{-i} are taken to be drawn from induced distribution Γ_{-i} . Now, recalling that $\nu(\mathbf{v}, i)$ is the slot assigned to player i in the optimal assignment for values \mathbf{v} , we can take $k = \nu(\mathbf{v}, i)$ in the above inequality. We then have

$$\mathbb{E}_{\mathbf{b}_{-i}}[\alpha_{\sigma(\mathbf{b},i)}v_i] + \mathbb{E}_{\mathbf{b}}[\alpha_{\nu(\mathbf{v},i)}b_{\pi^i(\mathbf{b}_{-i},\nu(\mathbf{v},i))}] \geq \gamma \alpha_{\nu(\mathbf{v},i)}v_i$$

for all \mathbf{v} and all i , where we used the fact that strategy $b_i(\cdot)$ does not appear in the second term. Summing over all i

and taking expectation over \mathbf{v} , we have

$$\begin{aligned} & \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} \left[\sum_i \mathbb{E}_{\mathbf{b}_{-i}}[\alpha_{\sigma(\mathbf{b},i)}v_i] \right] + \\ & \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} \left[\sum_i \mathbb{E}_{\mathbf{b}}[\alpha_{\nu(\mathbf{v},i)}b_{\pi^i(\mathbf{b}_{-i},\nu(\mathbf{v},i))}] \right] \geq \quad (2) \\ & \geq \gamma \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} \left[\sum_i \alpha_{\nu(\mathbf{v},i)}v_i \right] \end{aligned}$$

Consider each of the three expectations in (2). For the third term, we note

$$\mathbb{E}_{\mathbf{v} \sim \mathbf{F}} \left[\sum_i \alpha_{\nu(\mathbf{v},i)}v_i \right] = \mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[OPT(\mathbf{v})].$$

For the first term, linearity of expectation and independence of agent types implies

$$\begin{aligned} \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} \left[\sum_i \mathbb{E}_{\mathbf{b}_{-i}}[\alpha_{\sigma(\mathbf{b},i)}v_i] \right] &= \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} \left[\sum_i \alpha_{\sigma(\mathbf{b}(\mathbf{v}),i)}v_i \right] \\ &= \mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[SW(\pi(\mathbf{b}(\mathbf{v})), \mathbf{v})]. \end{aligned}$$

For the second expectation, notice that (again using type independence):

$$\begin{aligned} & \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} \mathbb{E}_{\mathbf{b} \sim \Gamma} \left[\sum_i \alpha_{\nu(\mathbf{v},i)}b_{\pi^i(\mathbf{b}_{-i},\nu(\mathbf{v},i))} \right] \leq \\ & \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} \mathbb{E}_{\mathbf{b} \sim \Gamma} \left[\sum_i \alpha_{\nu(\mathbf{v},i)}b_{\pi(\mathbf{b},\nu(\mathbf{v},i))} \right] \\ & = \mathbb{E}_{\mathbf{b} \sim \Gamma} \left[\sum_k \alpha_k b_{\pi(b,k)} \right] \end{aligned}$$

which is at most $\mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[SW(\pi(\mathbf{v}), \mathbf{v})]$.

We therefore conclude from (2) that $2\mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[SW(\pi(\mathbf{v}), \mathbf{v})] \geq \gamma \mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[OPT(\mathbf{v})]$, completing the proof. \blacksquare

Proof of Lemma 6 : Fix slot k , player i , and value v_i . We wish to show that

$$\mathbb{E}_{\mathbf{v}_{-i}}[\alpha_{\sigma(\mathbf{b}(\mathbf{v}),i)}v_i] + \mathbb{E}_{\mathbf{v}_{-i}}[\alpha_k b_{\pi^i(\mathbf{b}_{-i}(\mathbf{v}_{-i}),k)}] \geq \gamma \alpha_k v_i. \quad (3)$$

The intuition is that if agent i 's expected welfare at equilibrium and the expected bid required to obtain slot k are both very low compared to agent i 's value for slot k , then agent i should attempt to deviate to a strategy aimed at obtaining slot k . The complicating factor is that the bid needed to win slot k is only known to be low in expectation, so it is not immediately clear how agent i should optimally attempt to win slot k . We must therefore bound the size of a gap in welfare that implies the existence of a single strategy that provides a benefit to player i in expectation.

For notational convenience we write $W = \mathbb{E}_{\mathbf{v}_{-i}}[\alpha_{\sigma(\mathbf{b}(\mathbf{v}),i)}v_i]$ note that W is precisely the expected welfare of agent i given that his type is v_i . First note that if $\alpha_k v_i < W$ then (3) is trivial. So we will consider $\alpha_k \geq W/v_i$. We will prove that

$$\begin{aligned} W + \mathbb{E}_{\mathbf{v}_{-i}}[\alpha_k b_{\pi^i(\mathbf{b}_{-i}(\mathbf{v}_{-i}),k)}] &\geq \\ &\geq \alpha_k v_i - W \cdot \log \frac{\alpha_k v_i}{W} \end{aligned}$$

and then dividing everything by $\alpha_k v_i$ and using that $\frac{\log(x)}{x} \leq \frac{1}{e}$ we arrive at inequality (3).

Consider any bid b'_i for agent i such that $b_i < b'_i < v_i$. Then for each slot k , since bid b_i is utility-maximizing for agent i , the utility of bidding b'_i satisfies

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

Also, if agent i bids b'_i and moreover it is true that $b_{\pi^i(\mathbf{b}_{-i}, k)} < b'_i$, then agent i will win a slot with at least α_k clicks. Thus,

$$\mathbb{E}_{\mathbf{v}_{-i}}[u_i(b'_i, \mathbf{b}_{-i})] \geq (v_i - b'_i)\alpha_k \mathbb{P}_{\mathbf{v}_{-i}}[b_{\pi^i(\mathbf{b}_{-i}, k)} < b'_i].$$

Combining these two inequalities and substituting $z = v_i - b'_i$, we get

$$\mathbb{P}_{\mathbf{v}_{-i}}[v_i - b_{\pi^i(\mathbf{b}_{-i}, k)} > z] \leq \frac{W}{z \cdot \alpha_k}.$$

We are now able to estimate the expected value of $v_i - b_{\pi^i(\mathbf{b}_{-i}, k)}$. Since $v_i - b_{\pi^i(\mathbf{b}_{-i}, k)} \leq v_i$ with probability 1, we have

$$\begin{aligned} \mathbb{E}[v_i - b_{\pi^i(\mathbf{b}_{-i}, k)}] &= \int_0^{v_i} \mathbb{P}_{\mathbf{v}_{-i}}[v_i - b_{\pi^i(\mathbf{b}_{-i}, k)} > z] dz \\ &\leq \int_0^{W/\alpha_k} 1 dz + \int_{W/\alpha_k}^{v_i} \frac{W}{\alpha_k z} dz \\ &\leq \frac{W}{\alpha_k} + \frac{W}{\alpha_k} \left(\log v_i - \log \frac{W}{\alpha_k} \right). \end{aligned}$$

Multiplying both sides by α_k and rearranging gives the required inequality. \blacksquare

4.2 Coarsely Correlated bids and Price of Total Anarchy

Notice that the proof of the previous section applies even in cases where agent bids are coarsely correlated. In such a case, we can consider a common source of randomness \mathcal{R} and each bidding function to be a function $b_i(v_i, r)$, where $r \sim \mathcal{R}$. We call a profile of bidding functions a coarse correlated equilibrium if:

$$\begin{aligned} \mathbb{E}_{\mathbf{v}_{-i} \sim \mathbf{F}_{-i}}[u_i(b_i(v_i, r), \mathbf{b}_{-i}(\mathbf{v}_{-i}, r))] &\geq \\ \mathbb{E}_{\mathbf{v}_{-i} \sim \mathbf{F}_{-i}}[u_i(b'_i(v_i, r), \mathbf{b}_{-i}(\mathbf{v}_{-i}, r))], \forall i, v_i, r \end{aligned}$$

We still suppose $v_i \sim F_i$ where F_i are independent distributions. In this case, F and \mathcal{R} induce a distribution Γ on the bids.

Adapting Lemma 6 to this context is straightforward. Now, to adapt Lemma 5, we observe that the proof follows as before, except that we must argue that

$$\mathbb{E}_{\mathbf{b}_{-i} \sim \Gamma_{-i}}[\alpha_{\sigma(v, i)} b_{\pi^i(\mathbf{b}_{-i}, \sigma(v, i))}] = \mathbb{E}_{\mathbf{b} \sim \Gamma}[\alpha_{\sigma(v, i)} b_{\pi^i(\mathbf{b}_{-i}, \sigma(v, i))}].$$

However, this follows because the marginal distribution of $b \sim \Gamma$ restricted to $-i$ is precisely $b_{-i} \sim \Gamma_{-i}$, as the shared randomness from \mathcal{R} is not affected by restriction to a marginal distribution.

We now note that the result from the previous section implies a bound on the price of total anarchy for GSP. This follows because, whenever bidding sequence $D = (\mathbf{b}^1, \dots, \mathbf{b}^T)$ minimizes regret for all agents, the bidding strategy with shared randomness $b_i(v_i, t) = b_i^t(v_i)$ for $t \in [T]$ is a coarse correlated equilibrium. Lemmas 5 and 6 therefore imply that, for all \mathbf{v} ,

$$\mathbb{E}_{t \in [T]}[SW(\pi(\mathbf{b}^t(\mathbf{v})), \mathbf{v})] \geq \frac{1}{2}(1 - 1/e)OPT(\mathbf{v})$$

which implies that the price of total anarchy is bounded by $2(1 - 1/e)^{-1}$. Thus, even when agents apply strategies that minimize regret but do not necessarily converge to a stationary equilibrium, we still obtain bounds on the average social welfare obtained by GSP over many rounds.

4.3 Irrational Agents

We now consider a setting in which, in addition to the n advertisers who bid rationally, there are m ‘‘irrational’’ advertisers who cannot be assumed to bid at equilibrium. Write N for the set of rational advertisers, and M for the set of irrational advertisers. Note that we still think of the irrational advertisers as being true players, who still receive value per click. The irrational bidders simply cannot be assumed to apply optimal strategies; for example, they may not have experience with the GSP auction, or not know about historical bidding patterns.

Given an outcome π (which is an assignment of these $n + m$ bidders to $n + m$ slots), the definition of social welfare is unchanged: it is $SW(\pi, \mathbf{v}) = \sum_{i \in N \cup M} v_i \alpha_{\pi(i)}$. We define the social welfare of bidders in N to be precisely that: $SW_N(\pi, \mathbf{v}) = \sum_{i \in N} v_i \alpha_{\pi(i)}$. The optimal social welfare for bidders in N is $OPT_N(\mathbf{v}) = \max_{\pi} SW_N(\pi, \mathbf{v})$.

We wish to show that the total social welfare obtained by GSP is a good approximation to $OPT_N(\mathbf{v})$ when the players in N play at equilibrium and the players in M play arbitrarily. That is, the addition of irrational players does not degrade the social welfare guarantees of GSP had they not participated. In order to make this claim, we must impose a restriction on the behaviour of the irrational players: that they do not overbid. In other words, we require that $b_i(v_i) \leq v_i$ for all $i \in M$ and all v_i . We feel this is a natural restriction: overbidding is easily seen to be a dominated strategy (i.e. any strategy that bids higher than v_i is dominated by a strategy that lowers such bids to be at most v_i). Moreover, it is arguable that inexperienced bidders would bid conservatively, and not risk a large payment with no gain.³

Given that irrational agents do not overbid, we note that our BPoA bounds go through in this setting almost without change. In particular, our structural property (1) continues to hold for all agents in N .

Lemma 7 Equation (1) holds with $\gamma = 1 - \frac{1}{e}$ for all $i \in N$.

PROOF. This proof follows the proof of Lemma 6 without change. Note that in that proof we used only the fact that the bidding strategy of agent i is a best response, so the fact that other agents may not bid at equilibrium does not affect the argument. \blacksquare

The corresponding version of Lemma 5 then follows from (1) just as in the setting without irrational agents, with only minor modifications.

Lemma 8 If (1) holds for all $i \in N$, then

$$\mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[SW(\pi(\mathbf{b}(\mathbf{v})), \mathbf{v})] \geq \frac{1}{2} \gamma \mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[OPT_N(\mathbf{v})].$$

³This relies on the simplifying assumption that all bidders have knowledge of their own private valuations. Admittedly, this requires a certain level of sophistication and may be difficult to attain in practice. Our argument is thus limited to imperfect strategy choice given perfect knowledge of types. It remains open to extend this analysis to agents who may misunderstand their own valuations.

PROOF. Precisely as in the proof of Lemma 5, we obtain

$$\begin{aligned} & \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} \left[\sum_{i \in N} \mathbb{E}_{\mathbf{b}_{-i}} [\alpha_{\sigma(\mathbf{b}, i)} v_i] \right] + \\ & \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} \left[\sum_{i \in N} \mathbb{E}_{\mathbf{b}} [\alpha_{\nu(\mathbf{v}, i)} b_{\pi^i(\mathbf{b}, \nu(\mathbf{v}, i))}] \right] \geq \\ & \geq \gamma \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} \left[\sum_{i \in N} \alpha_{\nu(\mathbf{v}, i)} v_i \right] \end{aligned} \quad (4)$$

where we note that the summations are over agents in N . Then, as in Lemma 5,

$$\mathbb{E}_{\mathbf{v} \sim \mathbf{F}} \left[\sum_{i \in N} \alpha_{\nu(\mathbf{v}, i)} v_i \right] = \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} [OPT_N(\mathbf{v})].$$

and

$$\begin{aligned} \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} \left[\sum_{i \in N} \mathbb{E}_{\mathbf{b}_{-i}} [\alpha_{\sigma(\mathbf{b}, i)} v_i] \right] &= \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} [SW_N(\pi(\mathbf{b}(\mathbf{v})), \mathbf{v})] \\ &\leq \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} [SW(\pi(\mathbf{b}(\mathbf{v})), \mathbf{v})]. \end{aligned}$$

For the second expectation in (4), we require a slight deviation from the previous argument. We notice that

$$\begin{aligned} & \mathbb{E}_{\mathbf{v} \sim \mathbf{F}, \mathbf{b} \sim \Gamma} \left[\sum_{i \in N} \alpha_{\nu(\mathbf{v}, i)} b_{\pi^i(\mathbf{b}_{-i}, \nu(\mathbf{v}, i))} \right] \\ & \leq \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} \mathbb{E}_{\mathbf{b} \sim \Gamma} \left[\sum_{i \in N} \alpha_{\nu(\mathbf{v}, i)} b_{\pi(\mathbf{b}, \nu(\mathbf{v}, i))} \right] \\ & = \mathbb{E}_{\mathbf{b} \sim \Gamma} \left[\sum_{k: \exists i \in N, \nu(\mathbf{v}, i) = k} \alpha_k b_{\pi(\mathbf{b}, k)} \right] \\ & \leq \mathbb{E}_{\mathbf{b} \sim \Gamma} \left[\sum_k \alpha_k v_{\pi(\mathbf{b}, k)} \right] = \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} [SW(\pi(\mathbf{v}), \mathbf{v})]. \end{aligned}$$

We therefore conclude from (4) that $2\mathbb{E}_{\mathbf{v} \sim \mathbf{F}} [SW(\pi(\mathbf{v}), \mathbf{v})] \geq \gamma \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} [OPT_N(\mathbf{v})]$, completing the proof. \blacksquare

Together, Lemma 7 and Lemma 8 imply that the Bayesian Price of Anarchy of GSP is at most $2(1 - 1/e)^{-1}$ even in the presence of irrational bidders. Following the comments in Section 4.2, we can apply the same argument to obtain a matching bound on the Price of Total Anarchy with irrational bidders.

4.4 Approximate Equilibria

It is commonly assumed that rational bidders apply strategies at equilibrium. However, it may be that due to limits on rationality or indifference between small differences in utility, agents converge only on an approximate Bayes-Nash equilibrium. Given type distribution \mathbf{F} , we say that strategy profile \mathbf{b} is an ϵ -BNE if, for all agents i and all types v_i ,

$$\mathbb{E}_{\mathbf{v}_{-i}} [u_i(b_i(v_i), \mathbf{b}_{-i}(\mathbf{v}_{-i}))] \geq (1 - \epsilon) \mathbb{E}_{\mathbf{v}_{-i}} [u_i(b_i'(v_i), \mathbf{b}_{-i}(\mathbf{v}_{-i}))].$$

Note that our choice of the multiplicative definition of approximate equilibrium, motivated by the fact that we elected not to scale values to lie in $[0, 1]$.

We now claim that our bound for social welfare at equilibrium depends on the Bayes-Nash equilibrium condition in a continuous way, so that the bound degrades gracefully as we increase the degree to which a bidding strategy only

approximates an equilibrium. Specifically, if agents apply strategy profile \mathbf{b} which forms an ϵ -BNE, for $0 \leq \epsilon < 1 - 1/e$, then the expected social welfare of GSP is within a factor of $2(1 - \epsilon - 1/e)^{-1}$.

The proof requires only a minor modification to the proof of Theorem 4. Lemma 5 is not affected by this change. The statement of Lemma 6 is then modified to $\gamma = 1 - \epsilon - 1/e$. To obtain this bound, we need only note that the inequality due to \mathbf{b} being a BNE becomes

$$W \geq (1 - \epsilon) \mathbb{E}_{\mathbf{v}_{-i}} [u_i(b_i', \mathbf{b}_{-i})],$$

from which we conclude (in the same way as in the proof of Lemma 6) that

$$(1 - \epsilon) \mathbb{E}[v_i - b_{\pi^i(\mathbf{b}_{-i}, k)}] \leq \frac{W}{\alpha_k} + \frac{W}{\alpha_k} \left(\log v_i - \log \frac{W}{\alpha_k} \right).$$

Multiplying through by α_k and rearranging yields

$$\mathbb{E}[\alpha_{\sigma(\mathbf{b}(\mathbf{v}), i)} v_i] + \mathbb{E}[\alpha_k b_{\pi^i(\mathbf{b}_{-i}, k)}] \leq (1 - \epsilon) \alpha_k v_i - W \cdot \log \frac{\alpha_k v_i}{W}$$

which implies the desired result, using that $\frac{\log(x)}{x} \leq \frac{1}{e}$.

Acknowledgments: We thank Éva Tardos for the many fruitful discussions. We also thank Silvio Micali and Costis Daskalakis for pointing us to the asymmetric information model.

5. REFERENCES

- [1] K. Bhawalkar and T. Roughgarden. Welfare guarantees for combinatorial auctions with item bidding. In *SODA 2011*, 2011.
- [2] A. Blum, M. Hajiaghayi, K. Ligett, and A. Roth. Regret minimization and the price of total anarchy. In *STOC08*, 2008.
- [3] I. Caragiannis, P. Kanellopoulos, C. Kaklamanis, and M. Kyropoulou. On the efficiency of equilibria in generalized second price auctions. In *EC'11*, 2011.
- [4] 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.
- [5] S. Dobzinski, H. Fu, and R. Kleinberg. Optimal auctions with correlated bidders are easy. In *STOC 2011*, 2011.
- [6] B. Edelman, M. Ostrovsky, and M. Schwarz. 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.
- [7] B. Edelman and M. Schwarz. Optimal auction design and equilibrium selection in sponsored search auctions. In *Working Paper*, 2010.
- [8] 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.
- [9] J. Hannan. Approximation to bayes risk in repeated plays. In M. Dresher, A. Tucker, and P. Wolfe, editors, *Contributions to the Theory of Games*, volume 4. Princeton University Press, 1957.

- [10] S. Kakade, A. Kalai, and K. Ligett. Playing games with approximation algorithms. In *STOC07*, 2007.
- [11] A. Kalai and S. Vempala. Efficient algorithms for online decision problems. *J. Comput. Syst. Sci.*, 2005.
- [12] 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.
- [13] S. Lahaie, D. Pennock, A. Saberi, and R. Vohra. *Algorithmic Game Theory*, chapter Sponsored search auctions, pages 699–716. Cambridge University Press, 2007.
- [14] B. Lucier. Beyond equilibria: Mechanisms for repeated combinatorial auctions. In *ICS10: Innovations in Computer Science 2010*, 2010.
- [15] B. Lucier and A. Borodin. Price of anarchy for greedy auctions. In *SODA '10*. ACM, 2010.
- [16] A. Mehta, A. Saberi, U. V. Vazirani, and V. V. Vazirani. Adwords and generalized on-line matching. In *FOCS*, pages 264–273, 2005.
- [17] A. Mehta, A. Saberi, U. V. Vazirani, and V. V. Vazirani. Adwords and generalized online matching. *J. ACM*, 54(5), 2007.
- [18] R. Paes Leme and Éva Tardos. Bayes-nash price of anarchy for gsp. In *AdAuctions Workshop 2010*, 2010.
- [19] R. Paes Leme and Éva Tardos. Pure and bayes-nash price of anarchy for generalized second price auctions. In *FOCS10: 51st Annual IEEE Symposium on Foundations of Computer Science*, 2010.
- [20] R. Paes Leme and Éva Tardos. Sponsored search equilibria for conservative bidders. In *AdAuctions Workshop 2009*, 2010.
- [21] R. Paes Leme and B. Lucier. Improved social welfare bounds for gsp at equilibrium. In *arXiv*, 2011.
- [22] C. Papadimitriou and G. Pierrakos. On optimal single-item auctions. In *STOC 2011*, 2011.
- [23] 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.
- [24] H. R. Varian. Position auctions. *International Journal of Industrial Organization*, 2006.