

Simultaneous Auctions Are (Almost) Efficient

Michal Feldman
The Hebrew University of
Jerusalem & Harvard
University
michal.feldman@huji.ac.il

Nick Gravin
Nanyang Technological
University
ngravin@pmail.ntu.edu.sg

Hu Fu
Cornell University
Dept. of Computer Science
hufu@cs.cornell.edu

Brendan Lucier
Microsoft Research New
England
brlucier@microsoft.com

ABSTRACT

Simultaneous item auctions are simple and practical procedures for allocating items to bidders with potentially complex preferences. In a simultaneous auction, every bidder submits independent bids on all items simultaneously. The allocation and prices are then resolved for each item separately, based solely on the bids submitted on that item. We study the efficiency of Bayes-Nash equilibrium (BNE) outcomes of simultaneous first- and second-price auctions when bidders have complement-free (a.k.a. subadditive) valuations. While it is known that the social welfare of every pure Nash equilibrium (NE) constitutes a constant fraction of the optimal social welfare, a pure NE rarely exists, and moreover, the full information assumption is often unrealistic. Therefore, quantifying the welfare loss in Bayes-Nash equilibria is of particular interest. Previous work established a logarithmic bound on the ratio between the social welfare of a BNE and the expected optimal social welfare in both first-price auctions (Hassidim et al. [11]) and second-price auctions (Bhawalkar and Roughgarden [2]), leaving a large gap between a constant and a logarithmic ratio. We introduce a new proof technique and use it to resolve both of these gaps in a unified way. Specifically, we show that the expected social welfare of any BNE is at least $1/2$ of the optimal social welfare in the case of first-price auctions, and at least $1/4$ in the case of second-price auctions.

Categories and Subject Descriptors

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

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

STOC'13, June 1–4, 2013, Palo Alto, California, USA.

Copyright 2013 ACM 978-1-4503-2029-0/13/06 ...\$15.00.

Keywords

Simultaneous auction; First-price auction; Second-price auction; Bayes Nash Equilibrium; Price of anarchy

1. INTRODUCTION

The central problem in algorithmic mechanism design is to determine how best to allocate resources among individuals, while respecting both computational constraints and the individual incentives of the participants. Much of the theoretical work in this field to date has focused on solving such problems *truthfully*. In a truthful mechanism, the participants reveal their preferences in full to a central orchestrator, who then distributes the resources in a way that incentivizes truthful revelation. Such an approach has theoretical appeal, but truthful mechanisms tend to be complex and are rarely used in practice. Instead, it is common to forego truthfulness and use simpler mechanisms. Canonical examples of such auctions are the generalized second price (GSP) auctions for online advertising [7, 24], and the ascending price auction for electromagnetic spectrum allocation [18]. Given that such simple auctions are used in practice, it is of crucial importance to determine how they actually perform when used by rational agents.

Consider the problem of resolving a *combinatorial auction*. In such a problem there is a large set M of m objects for sale, and n potential buyers. Each buyer has a private value function $v_i: 2^M \rightarrow \mathbb{R}_{\geq 0}$ mapping sets of objects to their associated values. The goal of the market designer is to decide how to allocate the objects among the buyers to maximize the overall social efficiency. One approach would be to elicit the valuation function from each bidder, then attempt to solve the resulting optimization problem. However, the valuation function is an object of exponential size, and this approach leads inevitably to large communication and computation complexity overheads. It is not surprising, therefore, that in existing online marketplaces such as eBay, buyers do not express their (potentially complex) preferences directly; rather, each item is auctioned independently, and a buyer is forced to bid separately on individual items. This approach is simple and natural, and relieves the burden of expressing a potentially complex valuation function. On the other hand, this limited expressiveness could potentially lead to inefficient outcomes. This begs the question: how well does the outcome of simultaneous item auctions approximate the socially optimal allocation?

In order to evaluate the performance of non-truthful mechanisms, we take the economic viewpoint that self-interested agents will apply bidding strategies at equilibrium, so that no agent can unilaterally improve his outcome by changing his strategy. We apply a quantitative approach, and ask how well the performance at equilibrium approximates the socially optimal outcome. Since there may potentially be multiple equilibria, we will bound the performance in the worst case over equilibria. Put another way, our approach is to use the *price of anarchy* as a performance measure for the analysis of mechanisms.

The fact that equilibria of simultaneous auctions might not be socially optimal was first observed by Bikhchandani [4], who studied the complete information¹ setting. As he states:

“Simultaneous sealed bid auctions are likely to be inefficient under complete information and hence, also under the more realistic assumption of incomplete information about buyer reservation values.”

Our goal is to bound the extent of this inefficiency in the incomplete information setting. To this end, we model incomplete information using the standard Bayesian framework. In this model, the buyers’ valuations are assumed to be drawn independently from (not necessarily identical) distributions. This product distribution is commonly known to all of the participants; we think of this as representing the public’s aggregate beliefs about the buyers in the market. While the distributions are common knowledge, each agent’s true valuation is private. This Bayesian model generalizes the full-information model of Nash equilibrium, which implicitly supposes that the type profile is known by all participants. Note that while the agents are aware of the type distribution, the mechanism (which applies simultaneous item auctions) is prior-free and hence agnostic to this information.

Pricing and Efficiency in Simultaneous Auctions.

We consider separately the case in which items are sold via first-price auctions (in which the player who bids highest wins and pays his bid), and the case of second-price auctions² (in which the winning bidder pays the second-highest bid). The differences between first and second-price simultaneous auctions have received significant attention in the recent literature. For example, a pure Nash equilibrium of our mechanism with simultaneous first-price auctions is equivalent to a Walrasian equilibrium [4, 11], and therefore must obtain the optimal social welfare [16]. On the other hand, every pure Nash equilibrium for simultaneous second-price auctions is equivalent to a Conditional equilibrium, and hence obtains at least half of the optimal social welfare [9]. While these constant factor bounds are appealing, their power is marred by the fact that pure equilibria do not exist in general. In fact, based on the equivalence results above, their existence is quite restrictive (e.g., for simultaneous first-price auctions, existence is guaranteed for an

extremely restrictive family of valuations, called *gross substitutes* [10]). Moreover, pure Nash equilibria rely on the very strong assumption of full information, which is rare in practice.

Can we hope for such constant-factor bounds to hold for general Bayes-Nash equilibria? For general valuations the answer is no. Consider, for example, the case of a buyer who has a very large value for the set of all objects for sale, but no value for any strict subset. In this case, any positive bid carries great risk: the buyer might win some items but not others, leaving him with negative utility. It therefore seems that complements do not synergize well with item bidding, and indeed it has been shown by Hassidim et al. [11] that the price of anarchy (with respect to mixed equilibria) in a first-price auction can be as high as $\Omega(\sqrt{m})$ when bidders’ valuations exhibit complementarities. The same lower bound can be easily extended to the case of second-price auctions.³

Our main result is that *the presence of complements is the only barrier to a constant price of anarchy*. We show that when buyer valuations are complement-free (a.k.a. subadditive), the (Bayesian) price of anarchy of the simultaneous item auction mechanism is at most a constant, in both the first- and second-price auctions.

For first-price auctions, we show that any Bayes-Nash equilibrium yields at least half of the optimal social welfare. This improves upon the previously best-known bound of $O(\log n)$ due to Hassidim et al. [11], where n is the number of bidders.

RESULT 1: [BPoA ≤ 2 in simultaneous first-price auctions.] *When buyers have subadditive valuations, the Bayesian price of anarchy of the simultaneous first-price item auction mechanism is at most 2.*

For simultaneous Vickrey auctions, it is not possible to bound the worst-case performance at equilibrium, even when there is only a single object for sale. This impossibility is due to arguably unnatural equilibria in which certain players grossly overreport their values, prompting others to bid nothing. To circumvent this issue one must impose an assumption that agents avoid such “overbidding” strategies. In the *strong no-overbidding assumption*, used by Christodoulou et al. [6] and Bhawalkar and Roughgarden [2], it is assumed that each agent i chooses bids so that, for every set of objects S , the sum of the bids on S is at most $v_i(S)$. We show that under this assumption, the Bayesian price of anarchy for simultaneous Vickrey auctions is at most 4. This improves upon the previously best-known bound of $O(\log n)$ due to Bhawalkar and Roughgarden [2].

RESULT 2: [BPoA ≤ 4 in simultaneous second-price auctions.] *When buyers have subadditive valuations, the Bayesian price of anarchy of the simultaneous Vickrey auction mechanism is at most 4, under the strong no-overbidding assumption.*

The strong no-overbidding assumption is quite strong, as it must hold for *every* set of items. A somewhat weaker as-

¹In a complete (or full) information setting, it is assumed that the bidders’ valuations are commonly known to all participants

²Second-price item auctions are also known as Vickrey auctions; we will use these terms interchangeably.

³As explained in the sequel, to obtain meaningful results in second-price auctions one needs to impose *no-overbidding* assumptions on the bidding strategies, defined formally in Section 2.3. The $\Omega(\sqrt{m})$ lower bound extends to the case of second-price auctions under the *weak no-overbidding* assumption. The alternative *strong no-overbidding* assumption is meaningless in the case of complements, as it precludes item bidding altogether.

sumption, referred to as *weak no-overbidding*, requires that the no overbidding condition holds only in expectation over the distribution of sets won by a player at equilibrium. That is, agents are said to be *weakly no-overbidding* if they apply strategies such that expected value of each agent’s winnings is at least the expected sum of his winning bids [9]. Roughly speaking, weak no-overbidding supposes that agents are generally averse to winning sets with bids that are higher than their true values. However, unlike strong no-overbidding, it does not preclude strategies in which an agent overbids on sets that he does not expect to win, i.e. in order to more accurately express his willingness to pay for other sets. For an expanded discussion of the no-overbidding assumptions, see Section 6.

Notably, the BNE outcomes under the two no-overbidding assumptions are incomparable; while the weak assumption is more permissive, and thus enables a richer set of behaviors in equilibrium, it also introduces new ways to deviate from the prescribed equilibrium. Therefore, a constant bound on the Bayesian PoA under the weakly no-overbidding assumption does not follow directly. Nevertheless, we show that the bound of 4 on the Bayesian PoA extends also to the case of weakly no-overbidding agents.

Bhawalkar and Roughgarden [2] showed that, under the strong no-overbidding assumption, the Bayesian PoA of the simultaneous Vickrey auction is strictly greater than 2, and furthermore the price of anarchy is $\Omega(n^{1/4})$ when agent values are allowed to be correlated. In the full version of the paper we show that similar results hold also under the weak no-overbidding assumption, proving bounds strictly greater than 2 and $\Omega(n^{1/6})$, respectively.

Our constant bounds hold for subadditive bidders, whereas constant bounds on Bayesian price of anarchy were previously known only for the subclass of fractionally subadditive (i.e. XOS) valuations [6]. Previous work that attempted to bound the BPoA for subadditive valuations [2, 11] provided constant bounds for XOS valuations, then used the logarithmic factor separation between XOS and subadditive valuations to establish a logarithmic upper bound on the BPoA for subadditive valuations. While it seems plausible to use the close relation between XOS and subadditive valuations, any analysis that follows this trajectory would encounter this inevitable logarithmic gap. The challenge, therefore, is in developing a new proof technique for subadditive valuations, which does not go through XOS valuations. This is the approach taken in this work.

It should be noted that subadditive valuations are more expressive than their XOS counterparts, and obtaining price of anarchy bounds for subadditive valuations is significantly more challenging. In particular, for XOS valuations, a player who aims to win a certain set S has a natural choice of bid: the additive valuation that determines his value for set S . For subadditive valuations, there is no such notion of a natural bid aimed at representing one’s value for a particular set, and hence even determining how best to bid on a certain set of interest is a non-trivial task.

Related Work.

Combinatorial auctions is a canonical subject of study in algorithmic mechanism design (see [19] and references therein for the large body of literature on this subject). While most previous work focuses on the design of truthful mechanisms,

we follow the more recent literature on the analysis of simple and practical (albeit not truthful) auctions.

Following the rich literature on the *price of anarchy* (PoA) [see, e.g., [12, 22], for references], Christodoulou et al. [6] pioneered the study of the *Bayesian price of anarchy* (BPoA) and applied it to item-bidding auctions. They bounded the BPoA by 2 in simultaneous second-price auctions with XOS valuations, which are equivalent to fractionally subadditive valuations [8]. The same bound was extended to the more general class of subadditive valuations by Bhawalkar and Roughgarden [2], and later to general valuations by Fu et al. [9], albeit only with respect to *pure* equilibria (when they exist). The price of anarchy was studied also in simultaneous first-price auctions by Hassidim et al. [11], who showed a pure PoA of 1 for general valuations⁴, and a constant BPoA for XOS valuations. The effect of the underlying single-item auction on the PoA was further studied by Bhawalkar and Roughgarden [3].

For both first- and second-price simultaneous auctions, the BPoA for subadditive valuations was not previously known to be better than $O(\log n)$. Previous techniques applied the constant bounds for XOS valuations, using the $O(\log n)$ separation between XOS and subadditive valuations [see e.g. [2]].

Studies on PoA and BPoA have provided insights into other settings, e.g. auctions employing greedy algorithms [14], Generalized Second Price Auctions [20, 15, 5], uniform-price multi-unit auctions [17], and network formation settings [1].

The *smoothness* technique for Bayesian games, developed by Roughgarden [21] and Syrgkanis [23], provides a method for extending bounds on pure PoA to Bayesian PoA. However, to the best of our knowledge, our approach does not fall within this framework. Roughly speaking, the smoothness framework requires that each player can find a good “default” strategy given his type, which is independent of the opponents’ strategy selections. However, subadditive valuations do not seem to admit such bids,⁵ and indeed the strategies we consider in our analysis depend heavily on the distribution of strategies applied by all players at equilibrium.

Organization of the paper.

We introduce the necessary background and notation in Section 2. Our analysis then proceeds in two parts. In the first part, Section 3, we consider a single-player game in which the player, a subadditive buyer, must determine how best to bid on a set of objects against a distribution over price vectors. We show that, for every distribution for which the expected sum of prices is not too large, the buyer has a bidding strategy that guarantees a high expected utility (compared to the player’s value for the set of all objects).

In the second part of our analysis for the first-price (Section 4) and Vickrey (Section 5) auctions, we show that every Bayesian-Nash equilibrium must have high expected social welfare. We do this by considering deviations in which an agent uses the bidding strategy from the single-player game described

⁴Pure Nash equilibria rarely exist in this case though, as they are shown to be equivalent to Walrasian equilibria of the corresponding two-sided market.

⁵We note that one can apply the technique on XOS valuations, but because of the $O(\log n)$ separation between XOS and subadditive valuations [see e.g. [2]] this gives only a logarithmic bound.

in Section 3, applied to some subset of the objects. This subset of objects is chosen randomly: agent i draws a new profile of types for his opponents from the type distribution, then considers bidding for the set he would be allocated under this “virtual” type profile. At a BNE, agent i cannot benefit from such a randomized deviation; we show that this implies that the social welfare at equilibrium is at least a constant times the optimal welfare.

2. PRELIMINARIES

2.1 Auctions and Equilibria

Combinatorial Auctions.

In a combinatorial auction, m items are sold to n bidders. Each bidder has a private combinatorial valuation captured by a set function $v : 2^{[m]} \rightarrow \mathbb{R}_+$ over different bundles $S \subseteq [m]$. Throughout the paper we assume the valuations are *monotone*, i.e. for every subset $S \subseteq T \subseteq [m]$ it holds that $v(S) \leq v(T)$. In a *Bayesian* (partial-information) setting, the bidders’ valuation profile \mathbf{v} is drawn from a commonly known product distribution⁶ $\mathcal{F} = \mathcal{F}_1 \times \dots \times \mathcal{F}_n$. The outcome of an auction consists of an allocation $\mathbf{X} = (X_1, \dots, X_n) \in 2^{[m] \times n}$, where X_i is the bundle of items allocated to bidder i , and payments made by each bidder. The *social welfare* of an allocation is $\sum_{i \in [n]} v_i(X_i)$. For any given valuation profile \mathbf{v} , we let $(\text{OPT}_1^{\mathbf{v}}, \dots, \text{OPT}_n^{\mathbf{v}})$ denote the welfare-maximizing assignment for profile \mathbf{v} .

Simultaneous Item-Bidding Auctions.

In a simultaneous item-bidding auction, each bidder simultaneously submits a vector of bids, one for each item. The outcome of the auction is then determined item by item according to the bids placed on each item. In this paper we study two forms of such auctions: *simultaneous first price auctions* and *simultaneous second price auctions*.⁷ In both auctions, each item is allocated to the bidder who has placed the highest bid on it (breaking ties arbitrarily but consistently). In a (simultaneous) first price auction, the winner of each item pays his bid on that item, while in a (simultaneous) second price auction, the winner of each item pays the second highest bid on that item. We now give a more formal description of this process.

We generally write $b_i(j)$ to denote the bid of player i on item j , and \vec{b}_i for the vector of bids placed by bidder i . Alternatively, we may think of agent i ’s bid b_i as an additive function $b_i(S) = \sum_{j \in S} b_i(j)$ that corresponds⁸ to the bid-vector \vec{b}_i . Given a sequence of bid profiles $\mathbf{b} = (b_1, \dots, b_n)$, we write $W_i(\mathbf{b})$ for the set of items won by bidder i , and $\vec{p}_i \in \mathbb{R}_+^m$ the vector of payments made by bidder i on the items. In this notation, the first- and second-price auctions

⁶Whenever an expectation is taken with respect to valuations, it will be assumed that they are drawn from these corresponding distributions.

⁷The word “simultaneous” is often omitted, as we study only simultaneous (in contrast to sequential) auctions.

⁸There is an easy equivalence between an additive function $a(S) := \sum_{j \in S} a(\{j\})$ and its concise vector description $\vec{a} = (a(\{1\}), \dots, a(\{m\}))$. We will use functional and vector representations interchangeably as the situation demands.

can be summarized as follows:

First-price	Vickrey
won set: $W_i(\mathbf{b}) = \{j \in [m] \mid b_i(j) > b_k(j), \forall k \neq i\}$	
payment: $p_i(j) =$	
$\begin{cases} b_i(j), & j \in W_i(\mathbf{b}) \\ 0, & j \notin W_i(\mathbf{b}) \end{cases}$	$\begin{cases} \max_{k \neq i} b_k(j), & j \in W_i(\mathbf{b}) \\ 0, & j \notin W_i(\mathbf{b}) \end{cases}$

We assume bidders have quasi-linear utilities, i.e. the *utility* of bidder i for a given bid profile \mathbf{b} is given by $u_i(\mathbf{b}) = v_i(W_i(\mathbf{b})) - p_i(W_i(\mathbf{b}))$.

A Single Bidder’s Perspective on Bidding.

In both first- and second-price auctions, the set of items won by a bidder i bidding b_i is determined solely by a coordinate-wise comparison between b_i and the largest bid placed by the other bidders. Let $\varphi_i(\mathbf{b}_{-i})$ be the vector whose j -th component is $\max_{k \neq i} b_k(j)$. It is often convenient to write $W(b_i, \mathbf{b}_{-i})$ as $W(b_i, \vec{p})$ where $\vec{p} = \varphi_i(\mathbf{b}_{-i})$. We think of \vec{p} as the vector of prices perceived by bidder i : in the second price auction, the bidder pays the price on an item if his bid exceeds it; and in the first price auction the bidder pays his own bid on such an item, and \vec{p} is the minimum such winning bid. It is in this light that we often write $\varphi_i(\mathbf{b}_{-i})$ as prices \vec{p} when this causes no confusion. We will also shorten the notation $v(W(b, \vec{p}))$ to $v(b, \vec{p})$, meaning the value obtained when bidding b against perceived prices \vec{p} .

Strategies and Equilibria.

Buyers select their bids strategically in order to maximize utility. The bidding behavior of a buyer given its valuation is described by a *strategy*. A strategy s_i maps each valuation v_i to a distribution over bid vectors; we interpret $s_i(v_i)$ as the (possibly randomized) set of bids placed by bidder i when his type is v_i .

Definition 1. (Bayes-Nash Equilibrium) A profile of strategies $\mathbf{s} = (s_1(v_1), \dots, s_n(v_n))$ is in *Bayes-Nash equilibrium* (BNE) for distribution \mathcal{F} if, for every buyer i , type v_i , and bidding strategy \tilde{s}_i ,

$$\mathbf{E}_{\mathbf{v}_{-i}} \left[\begin{array}{c} \mathbf{E} \\ \mathbf{b}_{-i} \sim \mathbf{s}(\mathbf{v}_{-i}), \\ \mathbf{b}_i \sim s_i(v_i) \end{array} [u_i(b_i, \mathbf{b}_{-i})] \right] \geq \mathbf{E}_{\mathbf{v}_{-i}} \left[\begin{array}{c} \mathbf{E} \\ \mathbf{b}_{-i} \sim \mathbf{s}(\mathbf{v}_{-i}), \\ \vec{b}_i \sim \tilde{s}_i \end{array} [u_i(\vec{b}_i, \mathbf{b}_{-i})] \right].$$

Given Fubini’s Theorem, we can shorten the condition as follows (such shorthand forms are used throughout the paper):

$$\mathbf{E}_{\mathbf{v}_{-i}, \mathbf{b} \sim \mathbf{s}(\mathbf{v})} [u_i(\mathbf{b})] \geq \mathbf{E}_{\mathbf{v}_{-i}, \mathbf{b} \sim \mathbf{s}(\mathbf{v}), \vec{b}_i \sim \tilde{s}_i} [u_i(\vec{b}_i, \mathbf{b}_{-i})]. \quad (1)$$

Definition 2. (Bayesian Price of Anarchy) Given an auction type (either first- or second-price), the *Bayesian price of anarchy* (BPoA) is the worst-case ratio between the expected optimal welfare and the expected welfare at a BNE and is given by

$$\max_{\substack{(\mathcal{F}, \mathbf{s}): \\ \mathbf{s} \text{ a BNE for } \mathcal{F}}} \frac{\mathbf{E}_{\mathbf{v}} [\sum_i v_i(\text{OPT}_i^{\mathbf{v}})]}{\mathbf{E}_{\mathbf{v}, \mathbf{b} \sim \mathbf{s}(\mathbf{v})} [\sum_i v_i(W_i(\mathbf{b}))]}.$$

For second price auctions we will consider BPoA under natural restrictions on the strategies used by the bidders. In

such cases, the maximum in Definition 2 is taken with respect to BNE under that restricted class of strategies. We note that a BNE is guaranteed to exist as long as the space of valuations and potential bids is discretized, say with all values expressed as increments of some $\epsilon > 0$. A more detailed discussion of BNE existence appears in the full version of the paper.

2.2 Subadditive Valuations

We focus on valuations that are complement-free in the following general sense:

Definition 3. A set function $v : 2^{[m]} \rightarrow \mathbb{R}_+$ is *subadditive* if, for any subsets $S_1, S_2 \subset [m]$,

$$v(S_1) + v(S_2) \geq v(S_1 \cup S_2).$$

The class of subadditive functions strictly includes a hierarchy of more restrictive complement-free functions such as submodular and gross substitute functions (see 13 for definitions and discussions). Among these, the XOS functions, as defined below, have a particular kinship with subadditive functions. XOS literally means XOR (taking the maximum) of OR's (taking sums), and this class of valuations is known to be equivalent to the class of *fractionally subadditive* functions [8].

Definition 4. A function $v : 2^{[m]} \rightarrow \mathbb{R}_+$ is said to be *XOS* if there exists a collection of additive functions $a_1(\cdot), \dots, a_k(\cdot)$ (that is, $a_i(S) := \sum_{j \in S} a_i(\{j\})$ for every set $S \subset [m]$), such that for each $S \subseteq [m]$, $v(S) := \max_{1 \leq i \leq k} a_i(S)$.

One of the characterizations of XOS functions uses the following definition.

Definition 5. A function $f(\cdot)$ is said to be *dominated* by a set function $g(\cdot)$ if for any subset $S \subseteq [m]$, $f(S) \leq g(S)$. We say that a vector $\vec{a} = (a_1, \dots, a_m)$ is dominated by a set function $v(\cdot)$, if as an additive function $a(\cdot)$ is dominated by $v(\cdot)$.

It is not too difficult to observe that $v(\cdot)$ is XOS if and only if for every set $T \subset [m]$ there is an additive function $a(\cdot)$ dominated by $v(\cdot)$ such that $a(T) = v(T)$.

For a general subadditive function $v(\cdot)$, it can be the case that any additive function $a(\cdot)$ dominated by $v(\cdot)$ has $\Omega(\log(m))$ gap from $v([m])$, i.e. $\Omega(\log(m))a([m]) \leq v([m])$, (See 2 for such an example) and a logarithmic factor is also an upper bound. Previous work that attempted to bound the BPoA for subadditive valuations [2, 11] provided constant bounds for XOS valuations, then used the logarithmic factor separation between XOS and subadditive valuations to establish a logarithmic upper bound on the BPoA for subadditive valuations. In order to establish a constant bound for subadditive valuations, we turn to a different technique.

2.3 Overbidding

It is well known that in second price auctions, even with only a single item, the price of anarchy can be infinite when bidders are not restricted in their bids.⁹ To exclude such

⁹A canonical example is two bidders who value the item at 0 and a large number h , respectively, but the first bidder bids $h + 1$ and the second bidder bids 0.

pathological cases, previous literature [e.g. 6, 2] has made the following *no-overbidding* assumption standard:¹⁰

Definition 6. A bidder is *strongly no-overbidding* if his bid $b(\cdot)$ is dominated by his valuation $v(\cdot)$.

In other words, a bidder is guaranteed to derive non-negative utility, no matter what are the prices in the market. Thus strong no overbidding is a strong risk-aversion assumption on the buyers. One may also consider less risk concerned bidders—in the following we generalize a weaker assumption of no-overbidding introduced by Fu et al. [9].

Definition 7. Given a price distribution \mathcal{D} , a bidder is said to be *weakly no-overbidding* if his bid vector b satisfies $\mathbf{E}_{p \sim \mathcal{D}}[v(W(b, p))] \geq \mathbf{E}_{p \sim \mathcal{D}}[b(W(b, p))]$, where $W(b, p)$ denotes the subset of items he wins when he bids b at price p , i.e., $W(b, p) = \{j \in [m] \mid b(j) \geq p(j)\}$.

We will bound BPoA under both weakly and strongly no-overbidding assumptions for simultaneous second price auctions (note that these sets are incomparable).

3. BIDDING UNDER UNCERTAIN PRICES

As discussed in Section 2, a bidder in a simultaneous auction faces the problem of maximizing his utility in presence of uncertain prices (which are the largest bids placed by other bidders). While this maximization problem is intricate, we show in this section particular bidding strategies that result in utilities comparable with the bidder's value of the whole bundle minus the expected total prices. In other words, given a price distribution \mathcal{D} , it is desired to have a bidding strategy b such that

$$\mathbf{E}_{p \sim \mathcal{D}}[v(b, p)] - b([m]) \geq \alpha v([m]) - \mathbf{E}_{p \sim \mathcal{D}}[p([m])], \quad (2)$$

for some constant α . Such bidding strategies are key ingredients of the BPoA proofs in later sections, and may also be of independent interest.

For fixed prices, achieving (2) is trivial, even for $\alpha = 1$; indeed, given a price vector \vec{p} , by bidding according to $b = p$, a bidder obtains $v(b, p) - b([m]) = v([m]) - p([m])$. The case in which prices are drawn at random is more intricate, and is the subject of the remainder of this section.

LEMMA 1 (Bidding against a price distribution) For any distribution \mathcal{D} of prices p and any subadditive valuation $v(\cdot)$ there exists a bid b_0 such that

$$\mathbf{E}_{p \sim \mathcal{D}}[v(b_0, p)] - b_0([m]) \geq \frac{1}{2}v([m]) - \mathbf{E}_{p \sim \mathcal{D}}[p([m])]. \quad (3)$$

PROOF. We show a random bidding strategy that guarantees the desired inequality in expectation, and infer the existence of a bid, drawn from the suggested distribution, that achieves the same inequality. Consider a bid that is drawn according to the exact same distribution as the prices. It holds that

$$\begin{aligned} \mathbf{E}_{b \sim \mathcal{D}}[\mathbf{E}_{p \sim \mathcal{D}}[v(b, p)]] &= \mathbf{E}_{p \sim \mathcal{D}}[\mathbf{E}_{b \sim \mathcal{D}}[v(b, p)]] \\ &= \frac{1}{2} \mathbf{E}_{b \sim \mathcal{D}}[\mathbf{E}_{p \sim \mathcal{D}}[v(b, p) + v(p, b)]] \\ &\geq \frac{1}{2} \mathbf{E}_{b \sim \mathcal{D}}[\mathbf{E}_{p \sim \mathcal{D}}[v([m])]] \\ &= \frac{1}{2}v([m]), \end{aligned}$$

¹⁰We note that such no-overbidding assumptions were also made in other contexts [e.g. 14, 20].

where the inequality follows from subadditivity (which guarantees that $v(b, p) + v(p, b) \geq v([m])$ for every p and b). Using the last inequality, it follows that

$$\begin{aligned} \mathbf{E}_{b \sim \mathcal{D}} [\mathbf{E}_{p \sim \mathcal{D}} [v(b, p)] - b([m])] &\geq \frac{1}{2}v([m]) - \mathbf{E}_{b \sim \mathcal{D}} [b([m])] \\ &= \frac{1}{2}v([m]) - \mathbf{E}_{p \sim \mathcal{D}} [p([m])]. \end{aligned}$$

Since a bid drawn from \mathcal{D} satisfies (3) in expectation, there must exist a bid b_0 satisfying (3). \square

Safe Bidding Under Uncertainty

As noted in Section 2.3, in order to obtain any meaningful bound on BPoA for second price auctions, one needs to assume that bidders are not overbidding. Unfortunately, Lemma 1 is not concerned with such requirements. This problem is addressed in Lemma 3, where it is shown that a strongly no-overbidding strategy analogous to that in Lemma 1 always exists.

Notably, when the no-overbidding requirement is imposed, the existence of a bid satisfying (2) is nontrivial even for the case in which the prices are fixed. The following lemma, rephrased from 2, establishes its existence:

LEMMA 2 (follows from Lemma 3.3 in 2) For a given price vector p and any subadditive valuation $v(\cdot)$ there exists a bid b dominated by $v(\cdot)$ such that

$$v(b, p) - b([m]) \geq v([m]) - p([m]).$$

We now turn to analyze the case of random prices.

LEMMA 3 (No Overbidding Against Price Distributions) For any distribution \mathcal{D} of prices p and any subadditive valuation $v(\cdot)$ there exists a bid b_0 dominated by $v(\cdot)$ such that

$$\mathbf{E}_{p \sim \mathcal{D}} [v(b_0, p)] - b_0([m]) \geq \frac{1}{2}v([m]) - \mathbf{E}_{p \sim \mathcal{D}} [p([m])]. \quad (4)$$

PROOF. Let q be any price vector in the support of the distribution \mathcal{D} . Let $T \subseteq [m]$ be a maximal set such that $v(T) \leq q(T)$. We consider a *truncated* price vector \tilde{q} , which is set to 0 on the coordinates corresponding to T , and coincides with q on the coordinates corresponding to $[m] \setminus T$.

We first observe that \tilde{q} is dominated by $v(\cdot)$. Indeed, for any set $R \subset [m] \setminus T$ it holds that $v(R) > q(R)$, since otherwise

$$v(R \cup T) \leq v(R) + v(T) \leq q(R) + q(T) = q(R \cup T),$$

in contradiction to the fact that T is a maximal set satisfying $v(T) \leq q(T)$.

We next establish that for any bid b , it holds that

$$v(b, q) + q([m]) \geq v(b, \tilde{q}) + \tilde{q}([m]). \quad (5)$$

Indeed, we have $W(b, \tilde{q}) \subseteq W(b, q) \cup T$. Therefore, $v(b, \tilde{q}) \leq v(b, q) + v(T)$ due to subadditivity of $v(\cdot)$. Now (5) follows by observing that $q([m]) - \tilde{q}([m]) = q(T) \geq v(T)$.

We next define the distribution $\tilde{\mathcal{D}} := \{\tilde{q} \mid q \sim \mathcal{D}\}$, which consists of truncated prices drawn from \mathcal{D} . Equation (5) now extends for any bid b to

$$\mathbf{E}_{p \sim \mathcal{D}} [v(b, p) + p([m])] \geq \mathbf{E}_{\tilde{p} \sim \tilde{\mathcal{D}}} [v(b, \tilde{p}) + \tilde{p}([m])]. \quad (6)$$

Recall that each $\tilde{q} \sim \tilde{\mathcal{D}}$ is dominated by $v(\cdot)$, therefore, bidding any b drawn from $\tilde{\mathcal{D}}$ satisfies the strongly no-overbidding requirement. Furthermore, by applying (6) to each $b \sim \tilde{\mathcal{D}}$ we get

$$\begin{aligned} \mathbf{E}_{b \sim \tilde{\mathcal{D}}} [\mathbf{E}_{p \sim \mathcal{D}} [v(b, p) + p([m])]] &\geq \mathbf{E}_{b \sim \tilde{\mathcal{D}}} [\mathbf{E}_{\tilde{p} \sim \tilde{\mathcal{D}}} [v(b, \tilde{p}) + \tilde{p}([m])]] \\ &= \mathbf{E}_{b \sim \tilde{\mathcal{D}}} [\mathbf{E}_{\tilde{p} \sim \tilde{\mathcal{D}}} [v(b, \tilde{p})]] + \mathbf{E}_{b \sim \tilde{\mathcal{D}}} [b([m])] \\ &\geq \frac{1}{2}v([m]) + \mathbf{E}_{b \sim \tilde{\mathcal{D}}} [b([m])], \end{aligned}$$

where the last inequality follows in a manner similar to the proof of Lemma 1. The assertion of the lemma follows. \square

4. BPOA OF FIRST PRICE AUCTIONS

In this section we apply the bidding strategy from Lemma 1 to bound the Bayesian price of anarchy of simultaneous first-price auctions.

THEOREM 4. *In a simultaneous first-price auction with subadditive bidders, the Bayesian price of anarchy is at most 2.*

PROOF. We begin with a brief outline of the proof. Our plan is to fix a Bayes-Nash equilibrium and then consider, for each agent, a potential deviating strategy. This deviation will use the bidding strategy from Lemma 1, applied to some subset of the objects. To determine which subset to bid upon, each agent i will do the following: given her own value v_i , she will draw a “virtual” type profile \mathbf{v}_i^* for the other agents from distribution \mathcal{F} , and then bid upon the set that she would be assigned in the optimal allocation for (v_i, \mathbf{v}_i^*) . To determine how to bid upon this set, she draws a second type profile for the other agents, \mathbf{v}_{-i} , as dictated by Lemma 1. At BNE, agent i cannot benefit from such a randomized deviation; this implies a bound on the expected utility of each agent at equilibrium (inequality (8)). By taking a sum over all agents and using linearity of expectation to disentangle the random variables \mathbf{v} and \mathbf{v}^* , we show that this implies the social welfare at equilibrium is at least a constant times the optimal welfare.

We now proceed with the details, beginning with notation. Fix type distributions $\mathcal{F} = \prod_{i=1}^n \mathcal{F}_i$ and let \mathbf{s} be a BNE for \mathcal{F} . Fix an agent i and an arbitrary subadditive valuation v_i . Fix an arbitrary \mathbf{v}_{-i} , and let $\mathbf{v} = (v_i, \mathbf{v}_{-i})$. Fix an arbitrary \mathbf{v}_{-i}^* , and let $\mathbf{v}^* = (v_i, \mathbf{v}_{-i}^*)$. Recall that $(\text{OPT}_1^*, \dots, \text{OPT}_n^*)$ is the welfare-optimal allocation for \mathbf{v}^* .

Recall that each bid profile \mathbf{b}_{-i} induces a price vector $\varphi_i(\mathbf{b}_{-i})$ on bidder i . Let \tilde{p} be equal to $\varphi_i(\mathbf{b}_{-i})$ on OPT_i^* and 0 elsewhere. Let \mathcal{D} be the distribution over these price vectors $\tilde{p} = \tilde{p}(\mathbf{b}_{-i})$, where $\mathbf{b}_{-i} \sim \mathbf{s}(\mathbf{v})$. That is, \mathcal{D} is precisely the distribution over the maximum bids on the items in OPT_i^* , excluding the bid of player i . Note that \mathbf{v}^* , which is different from \mathbf{v} , was used only to determine the set OPT_i^* , whereas \mathbf{v} determines the distribution of prices over the items in OPT_i^* . Much of the following proof involves handling and, to some extent, disentangling the two. By replacing $[m]$ by OPT_i^* in Lemma 1, there exists a bid vector b_i' over the objects in OPT_i^* such that, thinking now of p as an additive function,

$$\begin{aligned} \mathbf{E}_{p \sim \mathcal{D}} [v_i(b_i', p)] - b_i'(\text{OPT}_i^*) &\geq \frac{1}{2}v_i(\text{OPT}_i^*) - \mathbf{E}_{p \sim \mathcal{D}} [p(\text{OPT}_i^*)]. \quad (7) \end{aligned}$$

Since \mathbf{s} forms a BNE, we have that

$$\begin{aligned} \mathbf{E}_{\substack{\mathbf{v}_{-i}, \\ \mathbf{b} \sim \mathbf{s}(\mathbf{v})}} [u_i(\mathbf{b})] &\geq \mathbf{E}_{\substack{\mathbf{v}_{-i}, \\ \mathbf{b} \sim \mathbf{s}(\mathbf{v})}} [u_i(b_i', \mathbf{b}_{-i})] \\ &= \mathbf{E}_{\substack{\mathbf{v}_{-i}, \\ \mathbf{b} \sim \mathbf{s}(\mathbf{v})}} [v_i(b_i', \varphi_i(\mathbf{b}_{-i}))] - \mathbf{E}_{\substack{\mathbf{v}_{-i}, \\ \mathbf{b} \sim \mathbf{s}(\mathbf{v})}} [b_i'(W_i(b_i', \mathbf{b}_{-i}))] \\ &\geq \mathbf{E}_{p \sim \mathcal{D}} [v_i(b_i', p)] - b_i'(\text{OPT}_i^*), \end{aligned}$$

where the last inequality follows from the definition of \mathcal{D} and the fact that $W_i(b_i', \mathbf{b}_{-i}) \subseteq \text{OPT}_i^*$ for all \mathbf{b}_{-i} . Applying (7) and the definition of $p \sim \mathcal{D}$, we conclude that

$$\begin{aligned} \mathbf{E}_{\substack{\mathbf{v}_{-i}, \\ \mathbf{b} \sim \mathbf{s}(\mathbf{v})}} [u_i(\mathbf{b})] & \tag{8} \\ &\geq \frac{1}{2} v_i(\text{OPT}_i^*) - \mathbf{E}_{\substack{\mathbf{v}_{-i}, \\ \mathbf{b}_{-i} \sim \mathbf{s}_{-i}(\mathbf{v}_{-i})}} \left[\sum_{j \in \text{OPT}_i^*} \max_{k \neq i} b_k(j) \right]. \end{aligned}$$

Taking the sum over all i and expectations over all $v_i \sim \mathcal{F}_i$ and $\mathbf{v}_{-i}^* \sim \mathcal{F}_{-i}$, we conclude that

$$\begin{aligned} \sum_i \mathbf{E}_{\substack{\mathbf{v}, \mathbf{v}_{-i}^*, \\ \mathbf{b} \sim \mathbf{s}(\mathbf{v})}} [u_i(\mathbf{b})] &\geq \frac{1}{2} \sum_i \mathbf{E}_{v_i, \mathbf{v}_{-i}^*} [v_i(\text{OPT}_i^*)] & \tag{9} \\ &\quad - \sum_i \mathbf{E}_{\substack{\mathbf{v}, \mathbf{v}_{-i}^*, \\ \mathbf{b}_{-i} \sim \mathbf{s}_{-i}(\mathbf{v}_{-i})}} \left[\sum_{j \in \text{OPT}_i^*} \max_{k \neq i} b_k(j) \right]. \end{aligned}$$

Let us consider each of the three terms of (9) in turn. The LHS is equal to $\mathbf{E}_{\mathbf{v}, \mathbf{b} \sim \mathbf{s}(\mathbf{v})} [\sum_i u_i(\mathbf{b})]$, as \mathbf{v}_{-i}^* does not appear inside the expectation. The first term on the RHS is equal to $\frac{1}{2} \mathbf{E}_{\mathbf{v}} [\sum_i v_i(\text{OPT}_i^*)]$, by relabeling \mathbf{v}_{-i}^* by \mathbf{v}_{-i} . For the final term on the RHS of (9), we note that

$$\begin{aligned} &\sum_i \mathbf{E}_{\substack{\mathbf{v}, \mathbf{v}_{-i}^*, \\ \mathbf{b}_{-i} \sim \mathbf{s}_{-i}(\mathbf{v}_{-i})}} \left[\sum_{j \in \text{OPT}_i^*} \max_{k \neq i} b_k(j) \right] \\ &\leq \sum_i \mathbf{E}_{\substack{\mathbf{v}, \mathbf{v}_{-i}^*, \hat{v}_i, \\ \mathbf{b} \sim \mathbf{s}(\hat{v}_i, \mathbf{v}_{-i})}} \left[\sum_{j \in \text{OPT}_i^*} \max_k b_k(j) \right] \\ &= \mathbf{E}_{\mathbf{v}, \mathbf{b} \sim \mathbf{s}(\mathbf{v})} \left[\sum_j \max_k b_k(j) \right], \end{aligned}$$

where the first inequality follows due to the fact we take a maximum over a larger set, and the last equality follows from the fact that OPT_i^* imposes a partition over $[m]$, and by relabeling. We note a subtlety: in the first line we select a bid vector \mathbf{b} with respect to $(\hat{v}_i, \mathbf{v}_{-i})$, rather than (v_i, \mathbf{v}_{-i}) , so that \mathbf{b} is independent of the partition $(\text{OPT}_1^*, \dots, \text{OPT}_n^*)$. Applying these simplifications to the terms of (9), we conclude that

$$\begin{aligned} \mathbf{E}_{\mathbf{v}, \mathbf{b} \sim \mathbf{s}(\mathbf{v})} \left[\sum_i u_i(\mathbf{b}) \right] &\geq \frac{1}{2} \mathbf{E}_{\mathbf{v}} \left[\sum_i v_i(\text{OPT}_i^*) \right] & \tag{10} \\ &\quad - \mathbf{E}_{\mathbf{v}, \mathbf{b} \sim \mathbf{s}(\mathbf{v})} \left[\sum_j \max_k b_k(j) \right]. \end{aligned}$$

Since we are in a first-price auction setting, it holds that

$$\begin{aligned} \mathbf{E}_{\mathbf{v}, \mathbf{b} \sim \mathbf{s}(\mathbf{v})} \left[\sum_i u_i(\mathbf{b}) \right] &= \mathbf{E}_{\mathbf{v}, \mathbf{b} \sim \mathbf{s}(\mathbf{v})} \left[\sum_i v_i(W_i(\mathbf{b})) \right] \\ &\quad - \mathbf{E}_{\mathbf{v}, \mathbf{b} \sim \mathbf{s}(\mathbf{v})} \left[\sum_j \max_k b_k(j) \right]. \end{aligned}$$

Equation (10) therefore implies that

$$\mathbf{E}_{\mathbf{v}, \mathbf{b} \sim \mathbf{s}(\mathbf{v})} \left[\sum_i v_i(W_i(\mathbf{b})) \right] \geq \frac{1}{2} \mathbf{E}_{\mathbf{v}} \left[\sum_i v_i(\text{OPT}_i^*) \right],$$

which yields the desired result. \square

Remark: In the full version of the paper we show that the upper bound does not carry over to the case where the bidders' valuations are correlated. Specifically, a polynomial lower bound of $\Omega(n^{1/6})$ is given on the Bayesian price of anarchy for this case. The construction is based on a lower bound due to Bhawalkar and Roughgarden [2] for second-price auctions.

5. BPOA OF SECOND PRICE AUCTIONS

We now turn to the case of simultaneous second-price auctions. We show that the Bayesian price of anarchy of such an auction is always at most 4 for subadditive bidders, assuming that bidders' valuations are independent and bidders select strategies that satisfy either the strong or weak no-overbidding assumption.

THEOREM 5. *In simultaneous second-price auctions where bidders have subadditive valuations, and every bidder is either strongly or weakly no-overbidding, the Bayesian price of anarchy is at most 4.*

PROOF. Fix type distributions \mathcal{F} and let \mathbf{s} be a BNE for \mathcal{F} . We can then derive inequality (10) in precisely the same way as in the proof of Theorem 4 (using now Lemma 3 instead of Lemma 1); we then have that

$$\begin{aligned} \mathbf{E}_{\mathbf{v}, \mathbf{b} \sim \mathbf{s}(\mathbf{v})} \left[\sum_i u_i(\mathbf{b}) \right] &\geq \frac{1}{2} \mathbf{E}_{\mathbf{v}} \left[\sum_i v_i(\text{OPT}_i^*) \right] & \tag{11} \\ &\quad - \mathbf{E}_{\mathbf{v}, \mathbf{b} \sim \mathbf{s}(\mathbf{v})} \left[\sum_j \max_k b_k(j) \right]. \end{aligned}$$

Note that $\mathbf{E}_{\mathbf{v}, \mathbf{b} \sim \mathbf{s}(\mathbf{v})} [\sum_i v_i(W_i(\mathbf{b}))] \geq \mathbf{E}_{\mathbf{v}, \mathbf{b} \sim \mathbf{s}(\mathbf{v})} [\sum_i u_i(\mathbf{b})]$. Also, since each agent i is assumed to be strongly or weakly no overbidding, it holds that

$$\begin{aligned} \mathbf{E}_{\mathbf{v}, \mathbf{b} \sim \mathbf{s}(\mathbf{v})} \left[\sum_j \max_k b_k(j) \right] &= \mathbf{E}_{\mathbf{v}, \mathbf{b} \sim \mathbf{s}(\mathbf{v})} \left[\sum_i \sum_{j \in W_i(\mathbf{b})} b_i(j) \right] \\ &\leq \mathbf{E}_{\mathbf{v}, \mathbf{b} \sim \mathbf{s}(\mathbf{v})} \left[\sum_i v_i(W_i(\mathbf{b})) \right]. \end{aligned}$$

Equation (11) therefore implies that

$$\begin{aligned} \mathbf{E}_{\mathbf{v}, \mathbf{b} \sim \mathbf{s}(\mathbf{v})} \left[\sum_i v_i(W_i(\mathbf{b})) \right] &\geq \frac{1}{2} \mathbf{E}_{\mathbf{v}} \left[\sum_i v_i(\text{OPT}_i^*) \right] \\ &\quad - \mathbf{E}_{\mathbf{v}, \mathbf{b} \sim \mathbf{s}(\mathbf{v})} \left[\sum_i v_i(W_i(\mathbf{b})) \right], \end{aligned}$$

as required. \square

Bhawalkar and Roughgarden [2] showed that the Bayesian price of anarchy of second price auctions can be strictly worse than the pure price of anarchy when bidders are strongly no overbidding. In what follows we give an example showing that such a gap exists also when bidders are weakly no overbidding. We note that this gap is not implied by the example given by Bhawalkar and Roughgarden since the strategy profile in their example is not a BNE under the weaker no overbidding notion (as can be easily verified). The full analysis of the example appears in the full version of the paper; the following is a sketch.

EXAMPLE 1 (Bayesian price of anarchy can be strictly larger than 2 when bidders are weakly no overbidding and have subadditive valuations) Consider an instance with 2 bidders and 6 items, where the set of items is divided into two sets, of 3 items each, denoted S_1 and S_2 . Throughout, we shall present the example with parameters a and b for ease of presentation. The lower bound is obtained by substituting $a = 0.06$ and $b = 0.85$. In what follows, we describe the valuation function of bidder 1; bidder 2's valuation is symmetric w.r.t. the sets S_1 and S_2 . Bidder 1's valuation over the items in S_1 is additive with respective values (over the 3 items) of (a, a, b) , (b, a, a) or (a, b, a) , each with probability $1/3$. Bidder 1's valuation over the items in S_2 is 2 if she gets all three items, and 1 for any non-empty strict subset of S_2 . Bidder 1's valuation for an arbitrary subset T the maximum of her value for $T \cap S_1$ and her value for $T \cap S_2$. One can verify that this is indeed a subadditive valuation function.

We claim that the profile in which each bidder i bids her true (additive) valuation on S_i and 0 on all other items is a Bayes-Nash equilibrium with weakly no overbidding bidders for the specified parameter values. The full proof is deferred to the full version of the paper, where it is shown that the only beneficial deviations break the weakly no-overbidding assumption. Under this bidding profile, each bidder derives a utility of $2a + b$, amounting to a social welfare of $2(2a + b) = 1.94$. In contrast, if bidder 1 is allocated S_2 and bidder 2 is allocated S_1 , then each bidder derives a utility of 2, amounting to a social welfare of 4. Consequently, the Bayesian price of anarchy is $4/1.94 > 2.061$.

6. NO OVERBIDDING: A DISCUSSION

In our analysis of the BPoA of second-price auctions we have adopted either the strong version or the weak version of the no-overbidding assumption. A few conceptual remarks are in order.

We can think of no-overbidding assumptions as representing a form of risk aversion. The strong no-overbidding assumption guarantees to the bidder a non-negative utility, independent of the behavior of the other players; i.e., even if the other players behave in an arbitrary way. The weak no-overbidding assumption, in contrast, guarantees to the bidder a non-negative utility only if the other bidders behave "as expected". However, when the other bidders behave as expected, the bidder is guaranteed a non-negative utility even if the auction changes, ex-post, from a second-price auction to a first-price auction.

Let us give an example to illustrate the difference between the two assumptions. Consider an instance of a simultaneous second-price auction with two bidders and two items, say $\{a, b\}$. The first bidder is unit-demand; with probability

1 his valuation is such that he has value 1 for any non-empty subset of the items. The second bidder's valuation is additive, and distributed as follows: with probability $1/2$ she values a for 0.9 and b for 1.1, and with the remaining probability $1/2$ she values a for 1.1 and b for 0.9. In this instance, since the second bidder's valuation is additive it is a dominant strategy for her to bid her true value on each item. The best response for the first bidder is then to bid between 0.9 and 1 on each item: this guarantees that he wins one of the items and pays 0.9. This profile of strategies then forms a BNE for this instance. This bidding strategy of player 1 does not satisfy the strong no-overbidding assumption: it requires that he indicate a value of at least 1.8 for the set $\{a, b\}$, which is larger than his true value 1. However, it does satisfy the weak no-overbidding assumption given the behavior of bidder 2, since bidder 1 expects to win only one item (of value 1) with a bid of 0.9.

The above example illustrates a situation in which the best response of a player is permitted by weak no-overbidding but excluded by strong no-overbidding. There also exist cases in which a best response is excluded by the weak no-overbidding assumption as well. Ex. 1 is one such case: the players can improve their utilities, but only by applying strategies that violate weak no-overbidding. A direction for future research would be to determine whether there is a weaker restriction on strategies that never excludes best-responses, yet still guarantees a constant price of anarchy bound.

An additional subtle point is in order. The use of no-overbidding assumptions in Vickrey auctions and GSP auctions [20, 15] was justified by the fact that overbidding is weakly dominated: any overbidding strategy can be converted to a no-overbidding strategy that performs at least as well, regardless of the behavior of the other agents. For the case of simultaneous item auctions, our no-overbidding assumption cannot be relaxed to the assumption that bidders avoid such *dominated strategies*. In particular, there exists an instance of a second-price auction with a Bayes-Nash equilibrium in which all bidders play undominated strategies, and the Bayesian price of anarchy is $\Omega(n)$. For example, consider an instance with n unit-demand bidders and n items, where every bidder $i = 1, \dots, n-1$ values each of item i and item n at $1 - \epsilon$ (for some $\epsilon > 0$), and bidder n values all items $1, \dots, n-1$ at 1 (and has no value for item n). One can easily verify that, for bidder n , to bid 1 on all the first $n-1$ items is an undominated strategy (while it obviously breaks the strong no overbidding requirement). Consider the strategy profile in which bidder n bids according to this strategy, and each of bidders $i = 1, \dots, n-1$ bids 0 on item i and $1 - \epsilon$ on item n . This is a Bayes-Nash equilibrium in undominated strategies in a second-price auction, which gives social welfare $2 - \epsilon$, compared to the optimal social welfare, which is roughly n .

References

- [1] Noga Alon, Yuval Emek, Michal Feldman, and Moshe Tennenholtz. Bayesian ignorance. In *PODC*, pages 384–391, 2010.
- [2] Kshipra Bhawalkar and Tim Roughgarden. Welfare guarantees for combinatorial auctions with item bidding. In *SODA*, pages 700–709, 2011.

- [3] Kshipra Bhawalkar and Tim Roughgarden. Simultaneous single-item auctions. In *WINE*, 2012.
- [4] Sushil Bikhchandani. Auctions of heterogeneous objects. *Games and Economic Behavior*, 26(2): 193–220, 1999.
- [5] Ioannis Caragiannis, Panagiotis Kanellopoulos, Christos Kaklamanis, and Maria Kyropoulou. On the efficiency of equilibria in generalized second price auctions. In *EC*, pages 81–90, 2011.
- [6] George Christodoulou, Annamária Kovács, and Michael Schapira. Bayesian combinatorial auctions. In *ICALP*, pages 820–832, 2008.
- [7] Benjamin Edelman, Michael Ostrovsky, Michael Schwarz, Thank Drew Fudenberg, Louis Kaplow, Robin Lee, Paul Milgrom, Muriel Niederle, and Ariel Pakes. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97, 2005.
- [8] Uriel Feige. On maximizing welfare when utility functions are subadditive. *SIAM J. Comput.*, 39(1): 122–142, 2009.
- [9] Hu Fu, Robert Kleinberg, and Ron Lavi. Conditional equilibrium outcomes via ascending price processes with applications to combinatorial auctions with item bidding. In *EC*, page 586, 2012.
- [10] Faruk Gul and Ennio Stacchetti. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory*, 87(1):95–124, 1999.
- [11] Avinatan Hassidim, Haim Kaplan, Yishay Mansour, and Noam Nisan. Non-price equilibria in markets of discrete goods. In *EC*, pages 295–296, 2011.
- [12] Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *STACS*, pages 404–413, 1999.
- [13] Benny Lehmann, Daniel J. Lehmann, and Noam Nisan. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior*, 55 (2):270–296, 2006.
- [14] Brendan Lucier and Allan Borodin. Price of anarchy for greedy auctions. In *SODA*, pages 537–553, 2010.
- [15] Brendan Lucier and Renato Paes Leme. Gsp auctions with correlated types. In *EC*, pages 71–80, 2011.
- [16] J. Mamer and S. Bikhchandani. Competitive equilibrium in an exchange economy with indivisibilities. *Journal of Economic Theory*, 74: 385–413, 1997.
- [17] Evangelos Markakis and Orestis Telelis. Uniform price auctions: Equilibria and efficiency. In *SAGT*, pages 227–238, 2012.
- [18] Paul Milgrom. Putting auction theory to work: The simultaneous ascending auction. *Journal of Political Economy*, 108:245–272, 1998.
- [19] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.
- [20] Renato Paes Leme and Éva Tardos. Pure and bayes-nash price of anarchy for generalized second price auction. In *FOCS*, pages 735–744, 2010.
- [21] Tim Roughgarden. The price of anarchy in games of incomplete information. In *EC*, pages 862–879, 2012.
- [22] Tim Roughgarden and Éva Tardos. Introduction to the inefficiency of equilibria. In Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani, editors, *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.
- [23] Vasilis Syrgkanis. Bayesian games and the smoothness framework. *CoRR*, abs/1203.5155, 2012.
- [24] Hal R. Varian. Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178, 2007.