

Simultaneous Auctions are (almost) Efficient

Michal Feldman*

Hu Fu†

Nick Gravin‡

Brendan Lucier§

September 24, 2012

Abstract

Simultaneous item auctions are simple procedures for allocating items to bidders with potentially complex preferences over different item sets. In a simultaneous auction, every bidder submits 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. Such procedures occur in practice (e.g. eBay) but are not truthful. We study the efficiency of Bayesian Nash equilibrium (BNE) outcomes of simultaneous first- and second-price auctions when bidders have complement-free (a.k.a. subadditive) valuations. We show that the expected social welfare of any BNE is at least $\frac{1}{2}$ of the optimal social welfare in the case of first-price auctions, and at least $\frac{1}{4}$ in the case of second-price auctions. These results improve upon the previously-known logarithmic bounds, which were established by Hassidim et al. (2011) for first-price auctions and by Bhawalkar and Roughgarden (2011) for second-price auctions.

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 (Edelman et al., 2005; Varian, 2007), and the ascending price auction for electromagnetic spectrum allocation (Milgrom, 1998). Given that such simple auctions are used in practice, it is of crucial importance to determine how they actually perform when used by rational (and strategic) 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

*Hebrew University of Jerusalem, and Harvard University, michal.feldman@huji.ac.il

†Dept. of Computer Science, Cornell University, hufu@cs.cornell.edu

‡Nanyang Technological University, ngravin@pmail.ntu.edu.sg

§Microsoft Research New England, brlucier@microsoft.com

Work was done while first three authors were visiting Microsoft Research New England.

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, 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 (1999), 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 item 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 (Bikhchandani, 1999; Hassidim et al., 2011), and therefore must obtain the optimal social welfare (Mamer and Bikhchandani, 1997). On the other hand, every pure Nash equilibrium for second-price auctions is equivalent to a Conditional equilibrium, and hence obtains at least half of the optimal social welfare (Fu et al., 2012). While

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.

these constant factor bounds are appealing, their power is marred by the fact that pure equilibria do not exist in general.

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. (2011) 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 auctions³.

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. (2011), 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. (2008) and Bhawalkar and Roughgarden (2011), 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.

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 assumption, referred to as *weak no-overbidding*, requires the 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 (Fu et al., 2012). 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

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.

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 Appendix C.

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. We show that the bound of 4 on the Bayesian PoA extends also to the case of weakly no-overbidding agents.

Bhawalkar and Roughgarden (2011) showed that, under the strong no-overbidding assumption, the Bayesian price of anarchy 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. 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 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 (Christodoulou et al., 2008). 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 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 Works Combinatorial auctions is a canonical subject of study in algorithmic mechanism design (see Nisan et al., 2007 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., Roughgarden and Tardos, 2007, for references), Christodoulou et al. (2008) 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 functions (Feige, 2009). The same bound was extended to the more general class of subadditive valuations by Bhawalkar and Roughgarden (2011), and later to general valuations by Fu et al. (2012), albeit only with respect to *pure* equilibria (when they exist). The pure PoA was studied also in simultaneous first-price auctions by Hassidim et al. (2011), who showed a pure PoA of 1 for general valuations⁴.

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 known bounds for XOS valuations, using the $O(\log n)$ separation between XOS and subadditive valuations (see e.g. Bhawalkar and Roughgarden, 2011).

Studies on PoA and BPoA have provided insights into other settings, e.g. auctions employing greedy algorithms (Lucier and Borodin, 2010), Generalized Second Price Auctions (Paes Leme and Tardos, 2010; Lucier and Paes Leme, 2011; Caragiannis et al., 2011), and also game-theoretic settings that are not related to auctions, such as network formation games (Alon et al., 2010).

The *smoothness* technique for Bayesian games, developed by Roughgarden (2012) and Syrgkanis (2012), provides a method for extending bounds on pure PoA to Bayesian PoA. However, to the

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.

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 bids⁵, 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 Bayes-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 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 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 decided 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

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. Bhawalkar and Roughgarden, 2011), this gives only a logarithmic bound.

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.

pays his bid on that item, and 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 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:	$p_i(j) = \begin{cases} b_i(j), & j \in W_i(\mathbf{b}) \\ 0, & j \notin W_i(\mathbf{b}) \end{cases}$	$p_i(j) = \begin{cases} \max_{k \neq i} b_k(j), & j \in W_i(\mathbf{b}) \\ 0, & j \notin W_i(\mathbf{b}) \end{cases}$

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 (Bayesian Nash Equilibria). 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 ,

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

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

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

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.

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 optimal expected 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 is given in Appendix D.

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 Lehmann et al., 2006 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 (Feige, 2009).

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 Bhawalkar and Roughgarden, 2011 for such an example) and a logarithmic factor is also an upper bound. Previous work that attempted to bound the BPoA for subadditive valuations (Bhawalkar and Roughgarden, 2011; Hassidim et al., 2011) 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.

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 bids⁹. To exclude such pathological cases, previous literature (e.g. Christodoulou et al., 2008; Bhawalkar and Roughgarden, 2011) has made the following *no-overbidding* assumption standard¹⁰:

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. (2012).

Definition 7. Given a price distribution \mathcal{D} , a bidder is *weakly no-overbidding* if his bid vector b satisfies that $\mathbf{E}_{p \sim \mathcal{D}}[v(W(b, p))] \geq \mathbf{E}_{p \sim \mathcal{D}}[b(W(b, p))]$, 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.

3 Bidding Strategies 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 be of interest on their own.

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

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

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.

¹⁰We note that such no-overbidding assumptions were also made in other contexts (e.g. Lucier and Borodin, 2010; Paes Leme and Tardos, 2010).

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

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

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

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

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

3.1 No-Overbidding Strategies Under Uncertain Prices

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 3 is not concerned with such requirements. This problem is addressed in Lemma 5, where it is shown that a strongly no-overbidding strategy analogous to that in Lemma 3 always exists.

Notably, when the no-overbidding requirement is imposed, the existence of a bid satisfying (2) is already nontrivial when the prices are fixed. The following lemma, rephrased from Bhawalkar and Roughgarden (2011), establishes its existence:

Lemma 4 (Lemma 3.3 in Bhawalkar and Roughgarden, 2011). *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 must now analyze the case when prices are drawn randomly.

Lemma 5 (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 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 consist 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 3. The assertion of the lemma follows \square

4 Price of Anarchy for First Price Auctions

In this section we apply the bidding strategy from Lemma 3 to bound the Bayesian price of anarchy of simultaneous first-price auctions.

Theorem 6. *In a simultaneous first-price auction with subadditive bidders, the Bayesian price of anarchy is at most 2.*

Proof. Fix type distributions \mathcal{F} and let \mathbf{s} be a BNE for \mathcal{F} . Choose some agent i and an arbitrary sub-additive valuation v_i . Fix an arbitrary \mathbf{v}_{-i}^* , and let $\mathbf{v}^* = (v_i, \mathbf{v}_{-i}^*)$. Recall that $(\text{OPT}_1^{\mathbf{v}^*}, \dots, \text{OPT}_n^{\mathbf{v}^*})$ is the welfare-optimal allocation for \mathbf{v}^* .

Recall that each bid profile \mathbf{b}_{-i} defines for bidder i a price vector $\varphi_i(\mathbf{b}_{-i})$. Let \vec{p} be equal to $\varphi_i(\mathbf{b}_{-i})$ on $\text{OPT}_i^{\mathbf{v}^*}$ and 0 elsewhere. Let \mathcal{D} be the distribution over these price vectors $\vec{p} = \vec{p}(\mathbf{b}_{-i})$, where $\mathbf{b} \sim \mathbf{s}(\mathbf{v}^*)$. That is, \mathcal{D} is precisely the distribution over the maximum bids on the items in $\text{OPT}_i^{\mathbf{v}^*}$, excluding the bid of player i . By Lemma 3 (and replacing $[m]$ there by $\text{OPT}_i^{\mathbf{v}^*}$), there exists a bid vector b_i' over the objects in $\text{OPT}_i^{\mathbf{v}^*}$ such that, thinking now of p as an additive function,

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

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

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

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

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

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

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

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 $\frac{1}{2} \mathbf{E}_{\mathbf{v}} [\sum_i v_i(\text{OPT}_i^{\mathbf{y}})]$, by relabeling \mathbf{v}_{-i}^* by \mathbf{v}_{-i} . For the final term on the RHS of (9), we note that

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

where the first inequality is due to taking a maximum over a larger set, and the last equality follows since $\text{OPT}_i^{\mathbf{y}^*}$ form a partition of $[m]$ (and by relabeling). We note a subtlety: in the first line we select 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 $(\text{OPT}_1^{\mathbf{y}^*}, \dots, \text{OPT}_n^{\mathbf{y}^*})$. Applying these simplifications to the terms of (9), we conclude that

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

Since we are in a first-price auction, we have that $\mathbf{E}_{\mathbf{v}, \mathbf{b} \sim \mathbf{s}(\mathbf{v})} [\sum_i u_i(\mathbf{b})] = \mathbf{E}_{\mathbf{v}, \mathbf{b} \sim \mathbf{s}(\mathbf{v})} [\sum_i v_i(W_i(\mathbf{b}))] - \mathbf{E}_{\mathbf{v}, \mathbf{b} \sim \mathbf{s}(\mathbf{v})} [\sum_j \max_k b_k(j)]$. Equation (10) therefore implies

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

which yields the desired result. \square

Remark: In Appendix B, we show that the upper bound does not carry over to the case where the bidders' valuations are correlated. In particular, a polynomial lower bound of $\Omega(n^{1/6})$ is given on the Bayesian price of anarchy for this case. The construction is based heavily upon a lower bound due to Bhawalkar and Roughgarden (2011) for second-price auctions.

5 Price of Anarchy for 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 select strategies that satisfy either the strong or weak no-overbidding assumption.

Theorem 7. *In simultaneous second price auctions where bidders have subadditive valuations independently drawn and each of them is 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 6 (using now Lemma 5 instead of Lemma 3); we then have that

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

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,

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

Equation (11) therefore implies

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

as required. \square

Bhawalkar and Roughgarden 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 the following 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.

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 Bayesian equilibrium with weakly no overbidding bidders for the specified parameter values. The full proof is deferred to the appendix, 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 $4/1.94 > 2.061$.

References

- Alon, N., Emek, Y., Feldman, M., and Tennenholtz, M. (2010). Bayesian ignorance. In *Proceedings of the 29th ACM SIGACT-SIGOPS symposium on Principles of distributed computing*, PODC '10, pages 384–391, New York, NY, USA. ACM.
- Bhawalkar, K. and Roughgarden, T. (2011). Welfare guarantees for combinatorial auctions with item bidding. In *SODA*, pages 700–709.
- Bikhchandani, S. (1999). Auctions of heterogeneous objects. *Games and Economic Behavior*, 26(2):193–220.
- Caragiannis, I., Kanellopoulos, P., Kaklamani, C., and Kyropoulou, M. (2011). On the efficiency of equilibria in generalized second price auctions. In *EC'11*.
- Christodoulou, G., Kovács, A., and Schapira, M. (2008). Bayesian combinatorial auctions. In *35th International Colloquium on Automata, Languages and Programming, 35th International Colloquium*, pages 820–832.
- Edelman, B., Ostrovsky, M., Schwarz, M., Fudenberg, T. D., Kaplow, L., Lee, R., Milgrom, P., Niederle, M., and Pakes, A. (2005). Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97.
- Feige, U. (2009). On maximizing welfare when utility functions are subadditive. *SIAM J. Comput.*, 39(1):122–142.
- Fu, H., Kleinberg, R., and Lavi, R. (2012). Conditional equilibrium outcomes via ascending price processes with applications to combinatorial auctions with item bidding. In *ACM Conference on Electronic Commerce*, page 586.
- Hassidim, A., Kaplan, H., Mansour, Y., and Nisan, N. (2011). Non-price equilibria in markets of discrete goods. In *ACM Conference on Electronic Commerce*, pages 295–296.
- Jackson, M. O., Simon, L. K., Swinkels, J. M., and Zame, W. R. (2002). Communication and equilibrium in discontinuous games of incomplete information. *Econometrica*, 70(5):1711–1740.
- Lehmann, B., Lehmann, D. J., and Nisan, N. (2006). Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior*, 55(2):270–296.

- Lucier, B. and Borodin, A. (2010). Price of anarchy for greedy auctions. In *21st ACM-SIAM Symposium on Discrete Algorithms*, pages 537–553.
- Lucier, B. and Paes Leme, R. (2011). Gsp auctions with correlated types. In *EC'11*.
- Mamer, J. and Bikhchandani, S. (1997). *Journal of Economic Theory*, 74:385–413.
- Milgrom, P. (1998). Putting auction theory to work: The simultaneous ascending auction. *Journal of Political Economy*, 108:245–272.
- Nisan, N., Roughgarden, T., Tardos, E., and Vazirani, V. V. (2007). *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA.
- Paes Leme, R. and Tardos, É. (2010). Pure and bayes-nash price of anarchy for generalized second price auction. In *FOCS*, pages 735–744.
- Roughgarden, T. (2012). The price of anarchy in games of incomplete information. In *ACM Conference on Electronic Commerce*, pages 862–879.
- Roughgarden, T. and Tardos, E. (2007). Introduction to the inefficiency of equilibria. In Nisan, N., Roughgarden, T., Tardos, E., and Vazirani, V. V., editors, *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA.
- Simon, L. K. and Zame, W. R. (1990). Discontinuous games and endogenous sharing rules. *Econometrica*, 58(4):861–72.
- Syrkkanis, V. (2012). Bayesian games and the smoothness framework. *CoRR*, abs/1203.5155.
- Varian, H. R. (2007). Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178.

A A Proof of Example 1

In this section we prove that the strategy profile in Ex. 1 is a Bayesian equilibrium with weakly no overbidding bidders. To establish this, we need to show that every beneficial deviation breaks the weakly no overbidding assumption. Notably, since weak no-overbidding is required for every bid in the support, it is sufficient to consider only pure deviations. By symmetry, it suffices to consider only deviations by bidder 1. Finally, it suffices to consider only deviations in which bidder 1 bids either 0, a , or b on each item in S_2 and 0 on all items in S_1 ; this is because bidder 1 obtains value from either S_1 or S_2 but not both, and without deviating bidder 1 obtains all of S_1 at no cost.

The following table includes all the possible bids (in the rows), and their respective expected values, expected payments and expected bids (in the columns). For clarity of presentation, we present the expressions in parametric forms, and write the corresponding values for $a = 0.06$ and $b = 0.85$ in brackets.

Deviation	$\mathbf{E}_{p \sim \mathcal{D}}[v(W(b, p))]$	$\mathbf{E}_{p \sim \mathcal{D}}[p(W(b, p))]$	$\mathbf{E}_{p \sim \mathcal{D}}[b(W(b, p))]$
(a, a, a)	1	$2a$ [0.12]	$2a$ [0.12]
(a, a, b)	$\frac{1}{3} \cdot 2 + \frac{2}{3} \cdot 1 = \frac{4}{3}$	$\frac{1}{3}(2a + b) + \frac{2}{3}(2a) = 2a + \frac{1}{3}b$ [0.4033..]	$\frac{1}{3}(2a + b) + \frac{2}{3}(a + b) = \frac{4}{3}a + b$ [0.93]
(a, b, b)	$\frac{2}{3} \cdot 2 + \frac{1}{3} \cdot 1 = \frac{5}{3}$	$\frac{2}{3}(2a + b) + \frac{1}{3}(2a) = 2a + \frac{2}{3}b$ [0.6866..]	$\frac{2}{3}(a + 2b) + \frac{1}{3}(2b) = \frac{2}{3}a + 2b$ [1.74]
(b, b, b)	2	$2a + b$ [0.97]	$3b$ [2.55]
$(a, 0, 0)$	$\frac{2}{3} \cdot 1 + \frac{1}{3} \cdot 0 = \frac{2}{3}$	$\frac{2}{3}a$ [0.04]	$\frac{2}{3}a$ [0.04]
$(a, a, 0)$	1	$\frac{1}{3}(2a) + \frac{2}{3}a = \frac{4}{3}a$ [0.08]	$\frac{1}{3}(2a) + \frac{2}{3}a = \frac{4}{3}a$ [0.08]
$(b, 0, 0)$	1	$\frac{2}{3}a + \frac{1}{3}b$ [0.2833..]	b [0.85]

It is evident from the table that for deviations (a, b, b) and (b, b, b) , $\mathbf{E}[v(W(b, p))] < \mathbf{E}[b(W(b, p))]$, and therefore they do not satisfy weakly no overbidding. For each of the remaining deviations, the obtained expected utility (which equals $\mathbf{E}[v(W(b, p))] - \mathbf{E}[p(W(b, p))]$) is smaller than the current expected utility (which equals $2a + b = 0.97$). We conclude that the strategy profile in the example is a Bayesian equilibrium with weakly no overbidding bidders, as required.

B A Lower Bound for Correlated Valuations

In this section we give a polynomial lower bound, $\Omega(n^{1/6})$, on the Bayesian price of anarchy for first-price auctions with subadditive valuations, when the valuation distributions are correlated among the bidders. In fact, our example will hold even when all valuations are unit demand. The construction is based heavily upon a lower bound due to Bhawalkar and Roughgarden (2011) for second-price auctions.

Example 2 (High price of anarchy for correlated valuations and weakly no-overbidding players). There are $n + (n + 1)\sqrt{n}$ items and $3n$ players. Players occur in triples. Each triple contains one player of type I and two players of type II . A valuation from the correlated distribution \mathcal{D} is drawn as follows. First, a set T of \sqrt{n} items are selected at random; we will refer to these items as the common pool. Next, n of the remaining items are selected at random and labelled a_1, \dots, a_n ; we refer to these as the reserve items. Finally, the remaining $n\sqrt{n}$ items are partitioned into sets S_1, \dots, S_n , each of size \sqrt{n} ; we refer to these as the mock pools. Reserve item a_i and mock pool S_i are matched with the i th triple of players.

Given the labelling of the items, the player valuations are as follows. There are two possibilities for the valuation profile; an atypical case that occurs with probability $p = \frac{1}{n^{1/6}}$, and a typical case that occurs with the remaining probability $1 - p$. In the typical case, each player of type II has value $n^{-1/6}$ for the corresponding reserve item a_i , and each player of type I has value 1 for any non-empty subset of the common pool plus the corresponding reserve item, $T \cup \{a_i\}$. In the atypical case, each player of type II has the zero valuation, and each player of type I , say from triple i , has value 1 for any non-empty subset of the corresponding mock pool plus reserve item, $S_i \cup \{a_i\}$.

First note that we can assume in a Bayes-Nash equilibrium that each player of type II always bids $n^{-1/6}$ on his reserve item, in the typical case. Bidding more than $n^{-1/6}$ leads to negative utility if he wins, and bidding less than $n^{-1/6}$ allows the other type II bidder in the triple to obtain positive utility by winning the item with a bid less than $n^{-1/6}$. Thus both agents of type II in a triple will bid $n^{-1/6}$, causing both to have utility 0. In the atypical case, each type II bidder trivially bids 0 on all items.

A player of type I , when bidding in equilibrium, cannot distinguish between the typical and atypical cases; he always sees a set of $\sqrt{n} + 1$ items for which he has value, and each item is equally likely to be the reserve item. Note that it has to bid at least $n^{-1/6}$ on the reserve item in order to win it in the typical case. Suppose that the player bids at least $n^{-1/6}$ on some number k of the $\sqrt{n} + 1$ items. Then if the valuation profile is atypical the player will win all k items, and pay at least $k \cdot n^{-1/6}$. The expected payment of the player is therefore at least $pkn^{-1/6} = kn^{-1/3}$. If $k > n^{1/3}$ then the expected payment of the player is greater than 1, and hence his expected utility is negative, contradicting the assumption of Bayes-Nash equilibrium. We therefore conclude that $k \leq n^{1/3}$. Each player of type I will therefore win its reserve item in the typical case with probability at most $k/(\sqrt{n} + 1) < n^{-1/6}$.

We conclude that the social welfare of any Bayes-Nash equilibrium \mathbf{s} satisfies

$$\mathbf{E}_{\mathbf{v}, \mathbf{b} \sim \mathbf{s}(\mathbf{v})} \left[\sum_i v_i(W_i(\mathbf{b})) \right] \leq pn + (1-p)(n \cdot n^{-1/6} + n \cdot n^{-1/6} \cdot 1 + \sqrt{n} \cdot 1) = O(n^{5/6})$$

where the expression for the typical case includes bounds on the value obtained by the type II bidders, the value of the type I bidders who win reserve items, and the value of the type I bidders who win items from the common pool, respectively. Since the optimal social welfare is at least n in each case, the price of anarchy is at least $\Omega(n^{1/6})$.

C 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 also excluded by the weak no-overbidding assumption. 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, but yet still guarantees a constant price of anarchy bound.

The use of no-overbidding assumptions in Vickrey auctions and GSP auctions (Paes Leme and Tardos, 2010; Lucier and Paes Leme, 2011) 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 Bayesian 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 Bayesian 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 .

D Existence of Equilibria

The simultaneous auction games we consider have continuous type spaces (i.e. valuations) and continuous (pure) strategy spaces (i.e. potential bids). In general, equilibria may not exist in such infinite games, even when the strategy space is compact. As a toy example, consider a game in which each bidder declares a value from $[0, 1]$, and whoever declares the largest value strictly less than 1 wins; such a game does not admit any mixed equilibria. The existence of equilibria in infinite games is an involved topic, a full discussion of which falls outside the scope of this paper; we hope only to give a brief discussion of relevant issues and results.

Consider first a variant of our auction game in which agent types and bids are discretized and bounded. That is, suppose that all values lie in $[0, 1]$, and moreover that there is some $\epsilon > 0$ such that for each agent i and every set of items S , $v_i(S)$ can be expressed as $\epsilon \times k_i(S)$ for some integer $k_i(S) \geq 0$. Furthermore, each agent is restricted to placing bids from $[0, 1]$, each of which must be a multiple of ϵ . In this restricted game, a Bayes-Nash equilibrium always exists. To see this, note that the set of (pure) strategies is finite: it is the set of all functions mapping agent types (a finite set) to bid vectors (also finite). We can interpret the (Bayesian) game of incomplete information as the following normal-form game: each agent selects a bidding function ex ante, and the payoffs correspond to the expected payoffs in the Bayesian game under the commonly known distribution of types. Since the strategy space is finite, Nash's result implies the existence of a mixed Nash equilibrium of this normal-form game, which corresponds precisely to a Bayes-Nash equilibrium of the original game.

Let us turn now to the more general setting of continuous agent valuations and bids, say

constrained to lie in $[0, 1]$. We can, of course, approximate the continuous setting via discretizations to an ϵ -grid with arbitrarily small choice of ϵ . It is tempting to take the limit $\epsilon \rightarrow 0$, but the existence of an equilibrium in each approximate game does not necessarily imply the existence of an equilibrium for the limit case, even in settings of complete information. Consider, for example, the sale of a single object via first-price auction between two bidders, where the first bidder has value $1/2$ and the second bidder has value 1 , in which ties are broken in favor of bidder 1. The natural equilibrium in this case is for bidder 1 to bid $1/2$ and bidder 2 to bid slightly higher, but no fixed bid of bidder 2 is optimal: any bid of the form $1/2 + \delta$ is strictly worse than $1/2 + \delta/2$ for any $\delta > 0$. Note that Nash's theorem does not imply the existence of equilibria in this case because utilities are discontinuous in the strategy space: the utility of a player bidding on a single item is discontinuous at the bid value of the highest-bidding competitor.

The issue in the above example lies in the choice of tie-breaking rule. If ties were broken in favor of bidder 2, there would indeed exist an equilibrium in which each bidder bids $1/2$. It turns out that this is not coincidental: for the specific case of mixed Nash equilibria in games of complete information, a result due to Simon and Zame (1990) implies that non-existence of equilibrium is always due to the choice of tie-breaking rule. Applied in the context of simultaneous item auctions, their result implies that for any profile of agent types, there exists a tie-breaking rule (i.e. manner of distributing items for which multiple players declare the same bid) such that a mixed Nash equilibrium exists. Importantly, the tie-breaking rule used may depend on the types of the agents.

For settings of incomplete information the situation is less clear. The work of Simon and Zame (1990) has been extended to Bayes-Nash equilibria under some restrictions on agent utilities, such as in Jackson et al. (2002). They demonstrate that in certain auction settings, one can incentivize agents to reveal their types for the purpose of implementing a tie-breaking rule that guarantees the existence of an equilibrium. However, their approach relies crucially on bidder utilities being affine functions of the auction outcome, which is not necessarily the case for combinatorial auctions with non-additive agent valuations.

To the best of our understanding, for the particular case of simultaneous item auctions for agents with subadditive valuations, prior work does not imply a manner of selecting tie-breaking rules so that BNE always exist. We therefore view our results as pertaining most directly to discrete approximations of the auction with continuous agent types, and leave for future work the task of establishing (or disproving) BNE existence in general.