Welfare Guarantees for Combinatorial Auctions with Item Bidding

Kshipra Bhawalkar* Tim Roughgarden†

September 2, 2011

Abstract

We analyze the price of anarchy (POA) in a simple and practical non-truthful combinatorial auction when players have subadditive valuations for goods. We study the mechanism that sells every good in parallel with separate second-price auctions. We first prove that under a standard "no overbidding" assumption, for every subadditive valuation profile, every pure Nash equilibrium has welfare at least 50% of optimal — i.e., the POA is at most 2. For the incomplete information setting, we prove that the POA with respect to Bayes-Nash equilibria is strictly larger than 2 — an unusual separation from the full-information model — and is at most $2 \ln m$, where m is the number of goods.

1 Introduction

Theoretical work in algorithmic mechanism design has championed mechanisms with *dominant strategies*, where each player has a "foolproof" strategy that is optimal no matter what strategies the other players choose. The appeal of such mechanisms is clear: they impose little decision-making burden on the participants and have a predictable outcome. For example, in the Vickrey (second-price) single-item auction, truthful bidding is a dominant strategy for each player. Under the relatively weak assumption that bidders play these dominant strategies, the Vickrey auction is guaranteed to allocate the item to the bidder who values it the most.

Dominant strategies are expendable in practice, and many important real-world mechanisms do not have them. One recurring theme is that designing for dominant strategies leads to mechanisms that are overly "complex" for practical implementations. For instance, in sponsored search auctions, the Vickrey-Clarke-Groves (VCG) mechanism has an intuitively more complicated payment formula than that of the Generalized Second Price (GSP) auction that is used in practice (see e.g. [7]). The combinatorial auctions that are used in practice — for example, government auctions to sell wireless spectrum — are carefully crafted to achieve a number of economic and computational design goals, but the existence of dominant strategies is not one of them [16].

If mechanisms without dominant strategies are important to analyze, how should we do it? An obvious idea is to use an equilibrium concept to predict a mechanism's plausible outcomes, and to identify the performance of a mechanism with that of its equilibria. For example, a typical goal in classical mechanism design is to prove that one or all equilibria of a mechanism implement a given social choice rule. Here, we take a quantitative approach and ask how well the equilibrium

*Department of Computer Science, Stanford University, 460 Gates Building, 353 Serra Mall, Stanford, CA 94305. Supported by a Stanford Graduate Fellowship. Email: `kshipra@cs.stanford.edu`.

†Department of Computer Science, Stanford University, 462 Gates Building, 353 Serra Mall, Stanford, CA 94305. Supported in part by NSF grant CCF-1016885, an ONR Young Investigator Award, an ONR PECASE Award, an AFOSR MURI grant, and an Alfred P. Sloan Fellowship. Email: `tim@cs.stanford.edu`.

performance of a mechanism *approximates* that of a socially optimal outcome. Thus our work applies the "price of anarchy" framework to the analysis of mechanisms without dominant strategies.

1.1 Combinatorial Auctions with Item Bidding

Our most basic results are for the following model. There are n players and m goods. Each player i has a private valuation v_i that describes its value for each of the $2^m - 1$ non-empty bundles of goods. Every valuation v_i is *subadditive*, meaning that $v_i(S \cup T) \leq v_i(S) + v_i(T)$ for every pair S, T of subsets of goods. The objective is to maximize the social welfare $\sum_{i=1}^{n} v_i(S_i)$ over all allocations $S_1, \ldots, S_n$ of the goods to the players. This goal can be achieved using the VCG mechanism, in which each player has a dominant strategy that is to truthfully reveal its entire valuation to the mechanism. There are numerous challenges to implementing this mechanism (see e.g. [1]), such as the exponential amount of information that each player needs to communicate up front. By contrast, in a *combinatorial auction with item bidding*, each player submits one bid for each good — m bids instead of $2^m - 1$ — and each good is sold separately using a Vickrey auction.

There are at least three different motivations for studying combinatorial auctions with item bidding. The most obvious one is as a simple and practical alternative to combinatorial auctions with dominant strategies. The second, noted in [5], is that combinatorial auctions with item bidding are, in effect, already being used: a bidder that is trying to buy a bundle of goods in parallel on eBay (say) is implicitly participating in such an auction. The third motivation is theoretical, and suggests a general approach to welfare maximization when players have high-dimensional type spaces. Suppose players' valuations lie in the set V. The VCG mechanism has dominant strategies and is welfare-maximizing, but if V is very rich then asking a player to describe fully its valuation is unrealistic. A natural alternative is to accept bids only from a smaller set $W \subseteq V$ of easily describable valuations and to run the VCG mechanism *as if* bidders' true valuations lie in W. Truthfulness is obviously lost — the true valuations may lie in $V \setminus W$ — but there is hope for good performance at equilibrium if W is a "good enough approximation" of V. Our results substantiate this hope for the case where V and W are the sets of subadditive and additive valuations, respectively — with additive valuations, the VCG mechanism is equivalent to a separate Vickrey auction for each good.

1.2 Our Results

We prove worst-case bounds on the welfare of equilibria in combinatorial auctions with item bidding when players have subadditive valuations. While there is no hope of proving a non-trivial bound on the welfare of *every* equilibrium, even for the single-item Vickrey auction[1], our guarantees apply to every equilibrium with "no overbidding", meaning that the sum of a player's bids on a subset of goods is at most the player's value for that subset.[2]

Our first result is for the full information model, in which players' valuations are commonly known.

Result 1 The welfare of every pure Nash equilibrium with no overbidding in a combinatorial auction with item bidding and subadditive valuations is at least 1/2 times that of the welfare-maximizing allocation.

This bound is tight in the worst case.

[1] Consider two bidders with valuations 1 and 0 and equilibrium bids 0 and 1, respectively.

[2] Similar assumptions are made (by necessity) and economically justified in the related works discussed in Section 1.4.

Our second result is for the incomplete information model, where each player's valuation is drawn independently from a (player-specific) distribution over subadditive valuations. The distributions are commonly known but the actual valuations are private. Note that the full information model is the special case in which every distribution is a point mass.

We show that there can be Bayes-Nash equilibria with no overbidding with expected welfare less than half of the expected welfare of an optimal allocation. Nevertheless, we prove the following guarantee.

Result 2 The expected welfare of every Bayes-Nash equilibrium with no overbidding in a combinatorial auction with item bidding and subadditive valuations is at least $1/2\mathcal{H}_m$ times that of the welfare-maximizing allocation, where $\mathcal{H}_m = 1 + \frac{1}{2} + \frac{1}{3} + \cdots + \frac{1}{m} \approx \ln m$ is the mth Harmonic number.

We also extend this second guarantee to the expected welfare of every coarse correlated equilibrium[3] — and hence to the special cases of mixed Nash and correlated equilibria — in the full information model.

Finally, we prove that both of our bounds degrade gracefully as we relax the subadditivity and "no overbidding" constraints. For example, if we assume only that every player overbids by at most some constant factor, then our upper bounds grow only by constant factors.

1.3 Our Techniques

We prove our first result in two parts. The interface between the two parts is the following single-player "game", which essentially models the best-response problem faced by a player in a combinatorial auction with item bidding. An adversary assigns each good j a non-negative price $p(j)$. The player chooses a bid vector a, subject to no overbidding with respect to its valuation v. The player wins good j if and only if $a(j) \geq p(j)$. The player's payoff in this game under strategy a is, by definition, $v(S) + \sum_{j \notin S} p(j)$, where S is the set of goods that it wins.

The first part of our proof is a reduction, from proving a guarantee for equilibria in combinatorial auctions with item bidding to proving a guarantee for this single-player game. Let V be a class of valuations. Suppose that for every valuation $v \in V$ and price vector p, the player has a countering bid vector a with payoff at least $v(M)/\alpha$, where M is the set of all goods. (Note that if prices are high, the player cannot win all the goods without overbidding.) We prove that, in this case, every pure Nash equilibrium with player valuations in V and no overbidding has welfare within a 2α factor of optimal. The second part of our proof provides, for every subadditive valuation v and price vector p, a bid vector a with payoff at least $v(M)$.

Our second result requires a different approach. One concrete reason is that, as we show, the guarantee of 2 for pure Nash equilibria does not hold for Bayes-Nash equilibria. Intuitively, the proof approach above fails because, in the incomplete information model, players best respond to *distributions* over price vectors rather than to *fixed* price vectors. Our solution is to reverse the order of the player and the adversary in the single-player game above: a player with a valuation $v \in V$ first chooses a bid vector a without overbidding, and then the adversary picks a price vector p to minimize the player's payoff. We prove that if for every $v \in V$ the player has a bid vector that guarantees a payoff of at least $v(M)/\beta$ for every price vector, then every Bayes-Nash equilibrium with player valuations in V and no overbidding has expected welfare within a 2β factor of optimal. This argument applies also to coarse correlated equilibria in the full information model. Finally,

[3]Equivalently, to sequences of repeated play in which every player has vanishing time-averaged regret — see the "price of total anarchy" defined by Blum et al. [3].

we prove that $\beta \leq \mathcal{H}_m$ for the class of subadditive functions. Our proof uses a greedy Set Cover-style argument to show that all subadditive functions are "$\mathcal{H}_m$-approximations" of fractionally subadditive functions, a class of valuations for which $\beta = 1$.[4]

1.4 Related Work

Our work is closely related to that of Christodoulou, Kovács, and Schapira [5], who were the first to study what we are calling combinatorial auctions with item bidding, and who proved special cases of our two main results for submodular valuations.[5] Their approximation guarantee for Bayes-Nash equilibria is a factor of 2 (like for pure Nash equilibria), in contrast to our bound of $2 \ln m$.

Subadditive valuations are much more expressive, and hence technically more challenging, than submodular valuations. For example, for the problem of computing a welfare-maximizing allocation, a simple greedy algorithm produces a 2-approximation with submodular valuations [12], while the only known polynomial-time constant-factor approximation algorithm with subadditive valuations is the ingenious LP rounding scheme by Feige [8]. Thus our first result, that every pure Nash equilibrium with no overbidding is already a 2-approximation with subadditive valuations, is arguably surprising.[6] Also, our lower bound that rules out an approximation guarantee of 2 for Bayes-Nash equilibria with subadditive valuations offers a formal separation between the submodular and subadditive cases. This lower bound also implies that no proof of our first result can follow the "smoothness paradigm" of [17] and hence must be novel in some sense. In contrast, the proofs in [5] can be recast as smoothness proofs in the sense of [17].

Since the work of Christodoulou, Kovács, and Schapira [5], several other papers applied the price of anarchy framework to mechanisms without dominant strategies [2, 4, 9, 11, 13, 14, 15]. We describe a few closest to our work. First, Lucier and Borodin [14] study conditions under which an α-approximation algorithm for welfare maximization can be combined with the VCG payment rule to obtain a mechanism with equilibria that are (close to) α-approximate. The results in [14] do not seem to have any interesting implications for combinatorial auctions with subadditive valuations (with item bidding or otherwise). A series of papers [4, 13, 15] prove that, under a no overbidding assumption, the welfare of all Nash and Bayes-Nash equilibria in the GSP sponsored search auction are within a constant factor of optimal. Like the present work, most of the proofs in these papers do not follow the smoothness paradigm in [17], and the bounds for pure Nash, mixed Nash and Bayes-Nash equilibria are different. However, these works do not give a lower bound that proves an unavoidable separation between these guarantees (as we do here).

After the conference version of this work [**?**], Hassidim, Kaplan, Mansour and Nisan [11] studied combinatorial auctions with item bidding where each good is sold using a *first-price* auction. Among many other results, they prove two that are related to ours here: pure Nash equilibria in their model have optimal welfare (when they exist); and for bidders with subadditive valuations, the expected welfare of every Bayes-Nash equilibrium is $\Omega(1/\log m)$ times that of an optimal allocation, where m is the number of goods.

[4] A similar result to this last step was proved earlier by Dobzinski [6].

[5] Their results hold for the somewhat more general class of XOS (i.e., "XOR-of-OR-of-singletons" [12]) or, equivalently, "fractionally subadditive" [8] valuations. These are a strict subclass of the subadditive valuations [**?**].

[6] To be clear, our work does not seem to imply any new algorithmic results for the problem.

2 Preliminaries

2.1 Combinatorial Auctions

In a combinatorial auction (CA), there is a set of n players and a set M of m goods (or items). Each player i has a *valuation* $v_i : 2^M \to \mathbb{R}^+$ that describes its value for each subset of the goods. We always assume that $v_i(\emptyset) = 0$ and $v_i(S) \leq v_i(T)$ for all i and $S \subseteq T$. The *social welfare* $SW(\mathbf{X})$ of an allocation $\mathbf{X} := \{X_1, X_2, \ldots, X_n\}$ of the goods to the players is $\sum_{i=1}^n v_i(X_i)$. For a valuation profile $\mathbf{v} = \{v_1, v_2, \ldots, v_n\}$, we denote the welfare-maximizing outcome by $OPT(\mathbf{v})$.

We consider CAs with *item bidding*. Each player i submits m bids, one for each good. Letting $\mathbf{b} = (b_1, b_2, \ldots, b_n)$ denote the bid profile, each good j is allocated to the highest bidder i at a price equal to the second-highest bid $\max_{k \neq i} b_k(j)$ for that item. Ties are broken arbitrarily. We use $X_i(\mathbf{b})$ to denote the goods allocated to player i in bid profile $\mathbf{b}$ and $SW(\mathbf{b}) = \sum_{i=1}^n v_i(X_i(\mathbf{b}))$ the social welfare of the resulting allocation. Player i's *utility* from bid profile $\mathbf{b}$ is defined as

$$u_i(\mathbf{b}) = v_i(X_i(\mathbf{b})) - \sum_{j \in X_i(\mathbf{b})} \max_{k \neq i} b_k(j).$$

As noted in the Introduction, no welfare guarantees are possible, even for the single-item Vickrey auction, without restricting players from using bids that are much larger than their valuations. For a valuation v and parameter $\gamma \geq 1$, call a bid vector b *γ-conservative* if

$$\sum_{j \in S} b(j) \leq \gamma \cdot v(S) \tag{1}$$

for every subset $S \subseteq M$ of goods. The strategy space of a *γ-conservative player* is, by definition, the set of bid vectors that are γ-conservative for its valuation. We sometimes call such bid vectors *feasible* for the player.

2.2 The Full Information Model

Players generally have no dominant strategies in a CA with item bidding, so we consider several types of equilibria. In the simplest *full information model*, players' valuations are commonly known. For a fixed valuation profile $\mathbf{v} = \{v_1, v_2, \ldots, v_n\}$, a bid profile $\mathbf{b} = \{b_1, b_2, \ldots, b_n\}$ is a *pure Nash equilibrium* if $u_i(\mathbf{b}) \geq u_i(b'_i, \mathbf{b}_{-i})$ for every player i and (feasible) deviation b'_i, where $(b'_i, \mathbf{b}_{-i})$ denotes the bid profile when player i bids b'_i and all other players bid according to $\mathbf{b}$.

The *pure price of anarchy (POA)* is the ratio of the social welfare of the optimal outcome and that of the worst pure Nash equilibrium:

$$\text{pure POA} := \max_{\mathbf{b}\text{: a pure Nash eq.}} \frac{SW(OPT(\mathbf{v}))}{SW(\mathbf{b})}. \tag{2}$$

If there is no pure Nash equilibrium (Example :pure-examples), then the POA is undefined.

We next review three standard and more permissive equilibrium concepts for the full information model: mixed Nash, correlated, and coarse correlated equilibria. A *mixed strategy* for player i is a distribution B_i over bid vectors. Let $\mathbf{B}$ denote the product of the players' distributions and $\mathbf{B}_{-i}$ the analogous product for the players other than i. The mixed profile $\{B_1, B_2, \ldots, B_n\}$ is a *mixed Nash equilibrium* if for each player i and bid vector b'_i,

$$\mathop{\mathbb{E}}_{\mathbf{b} \sim \mathbf{B}}[u_i(\mathbf{b})] \geq \mathop{\mathbb{E}}_{\mathbf{b_{-i}} \sim \mathbf{B}_{-i}}[u_i(b'_i, \mathbf{b}_{-i})].$$

We note that, by linearity, the equilibrium conditions above only need to be stated for deviations to pure strategies. Also, every pure Nash equilibrium is also a mixed Nash equilibrium.

A distribution $\mathbf{B}$ over bid profiles — which now need not be a product distribution — is a *correlated equilibrium* if

$$\mathop{\mathbb{E}}_{\mathbf{b}\sim\mathbf{B}}[u_i(\mathbf{b})|b_i] \geq \mathop{\mathbb{E}}_{\mathbf{b}\sim\mathbf{B}}[u_i(b_i', \mathbf{b}_{-i})|b_i]$$

for every player i, bid vector b_i belonging to a profile $\mathbf{b}$ in the support of $\mathbf{B}$, and feasible bid vector b_i'. Such a distribution $\mathbf{B}$ is a *coarse correlated equilibrium* if

$$\mathop{\mathbb{E}}_{\mathbf{b}\sim\mathbf{B}}[u_i(\mathbf{b})] \geq \mathop{\mathbb{E}}_{\mathbf{b}\sim\mathbf{B}}[u_i(b_i', \mathbf{b}_{-i})]$$

for every player i and bid b_i'. Every mixed Nash equilibrium is a correlated equilibrium and every correlated equilibrium is a coarse correlated equilibrium. For each of these three equilibrium concepts, the corresponding price of anarchy is defined analogous to (2), where the denominator is the worst-case expected social welfare of an equilibrium. See Young [19], for example, for more details on and interpretations of these three equilibrium concepts.

2.3 The Incomplete Information Model

When players' valuations are not commonly known, the standard model is a *Bayesian game* [10]. Let V_i denote the possible valuations of player i. For each player i, there is a commonly known prior distribution D_i over the valuations in V_i. Let $\mathcal{D} = D_1 \times D_2 \cdots \times D_n$ denote the induced product distribution over the set $\mathcal{V} = V_1 \times V_2 \cdots \times V_n$ of all valuation profiles. A mixed strategy for player i is then a mapping B_i of each valuation $v_i \in V_i$ to a distribution $B_i(v_i)$ over bid vectors feasible for v_i. A *mixed Bayes-Nash equilibrium* is a strategy B_i for each player i such that no player i can increase its expected utility by a unilateral deviation:

$$\mathop{\mathbb{E}}_{\mathbf{v}_{-i}\sim\mathcal{D}_{-i},\mathbf{b}\sim\mathbf{B}(\mathbf{v})}[u_i(\mathbf{b})|v_i] \geq \mathop{\mathbb{E}}_{\mathbf{v}_{-i}\sim\mathcal{D}_{-i},\mathbf{b}_{-i}\sim\mathbf{B}_{-i}(\mathbf{v}_{-i})}[u_i(b_i', \mathbf{b}_{-i})|v_i]$$

for every valuation $v_i \in V_i$ and every (pure) deviation b_i' that is feasible for v_i. The expected optimal social welfare is $\mathbb{E}_{\mathbf{v}\sim\mathcal{D}}[SW(OPT(\mathbf{v}))]$ and the POA is the ratio of this quantity and the worst expected social welfare of a Bayes-Nash equilibrium:

$$\text{Bayes-Nash POA} := \max_{\mathbf{B}:\text{a Bayes-Nash eq.}} \frac{\mathbb{E}_{\mathbf{v}\sim\mathcal{D}}[SW(OPT(\mathbf{v}))]}{\mathbb{E}_{\mathbf{v}\sim\mathcal{D},\mathbf{b}\sim\mathbf{B}(\mathbf{v})}[SW(\mathbf{b})]}.$$

We have defined the incomplete information model only for bidders with independent, not necessarily identical, valuation distributions. While the model extends to correlated valuation distributions in the natural way, the Bayes-Nash POA can be quite large for such distributions (Section 4.4).

2.4 Existence of Equilibria

There are CAs with item bidding and subadditive valuation profiles with no pure Nash equilibrium exists.[7]

[7]When every player's valuation is submodular, or more generally XOS, there is always a pure Nash equilibrium [5].

Example 2.1 (Pure Nash Equilibria Need Not Exist) Suppose there are two 1-conservative players and three items, with the players' subadditive valuation functions defined as follows:

$$\text{Player 1} : v_1(M) = 3.2, v_1(S) = 1.6 \text{ for } S \neq M, \emptyset;$$
$$\text{Player 2} : v_2(M) = 2, v_2(S) = 1 \text{ for } S \neq M, \emptyset.$$

Suppose there is a pure Nash equilibrium. First suppose that the second player wins some items in this equilibrium. Then, the first player does not win the whole set M and its utility is at most 1.6. Let x, y, z denote the second player's respective bids for the three items. Since the player is 1-conservative, each pair of of these bids sums to at most 1, and $x + y + z \leq 1.5$. Then, the first player can submit a 1-conservative bid of the form $x + \epsilon, y + \epsilon, z + \epsilon$ for small enough ϵ, obtaining utility at least $3.2 - 1.5 > 1.6$.

Second, suppose that the first player wins all three goods in the Nash equilibrium, with bids x, y, z. Since the player uses a 1-conservative bid, either x or y (x, say) is less than 1. The second player then has a 1-conservative bid on x that increases its utility, from zero to a positive amount.

Existence of mixed equilibria, in both the complete and incomplete information models, is a subtle issue. XXX NEED TO FINISH

3 Upper Bound on the Pure POA

This section presents our first result, an upper bound on the pure POA in CAs with item bidding. Recall from (1) that a γ-conservative player with valuation v only uses bid vectors b such that $\sum_{j \in S} b(j) \leq \gamma \cdot v(S)$ for every subset S of goods. Also, a valuation v is *α-subadditive* if for all subsets $S_1, S_2, \ldots, S_k$ of goods,

$$v(S_1 \cup S_2 \cup \cdots \cup S_k) \leq \alpha \left[v(S_1) + v(S_2) + \cdots + v(S_k) \right]. \tag{3}$$

Note that a valuation is subadditive in the usual sense if and only if it is 1-subadditive.

Theorem 3.1 (Upper Bound on the Pure POA) *For every CA with item bidding that has γ-conservative players with α-subadditive valuations and at least one pure Nash equilibrium, the pure POA is at most $\alpha(\gamma + 1)$.*

In particular, the pure POA is at most 2 when players have subadditive valuations and do not overbid on any subset. Before proving this theorem, we show by example that this upper bound is tight in the worst case for subadditive valuations for all values of γ.

Example 3.2 (Theorem 3.1 Is Tight) Consider two goods (A and B) and two players with the following valuations:

$$\begin{aligned} v_1(\{A\}) &= 1, & v_2(\{A\}) &= \gamma + 1, \\ v_1(\{B\}) &= \gamma + 1, & v_2(\{B\}) &= 1, \\ v_1(\{A, B\}) &= \gamma + 1 & v_2(\{A, B\}) &= \gamma + 1. \end{aligned}$$

These valuations are submodular and hence subadditive. The optimal allocation — with goods A and B assigned to players 2 and 1, respectively — has social welfare $2(\gamma + 1)$.

On the other hand, suppose players 1 and 2 bid $(\gamma, 0)$ and $(0, \gamma)$, respectively. These bids are γ-conservative. The social welfare of the corresponding allocation is 2. A simple case analysis shows that this bid profile is a Nash equilibrium, and hence the pure POA of this example is at least $\gamma + 1$.

We now turn to the proof of Theorem 3.1. We first prove a crucial lemma that concerns the best response of a single player against "posted prices" on the goods. More precisely, for a γ-conservative player with an α-subadditive valuation v and a vector p of nonnegative prices on the goods, we use $u(a, p)$ to denote the utility the player gets with the bid vector a:

$$u(a,p) = v(S) - \sum_{j \in S} p(j), \tag{4}$$

where $S = \{j \in M : a(j) \geq p(j)\}$ is the set of goods that the player wins. In our proof of Theorem 3.1, we apply this lemma to each player, where the prices p are set by the other players' bids.

Lemma 3.3 (How To Bid Against Fixed Posted Prices) *Let O be a set of goods and $p(j)$ a price for the good $j \in O$. Let $v : 2^O \to \mathbb{R}$ be an α-subadditive valuation function on this set. There is a 1-conservative bid vector a such that*

$$u(a,p) \geq \frac{v(O)}{\alpha} - \sum_{j \in O} p(j).$$

Proof: For convenience, we define the "translated utility" by $\tilde{u}(a,p) = u(a,p) + \sum_{j \in O} p(j)$. This quantity was called the "player payoff" in the summary in Section 1.3. We show that there is a 1-conservative bid vector a with $\tilde{u}(a,p) \geq v(O)/\alpha$.

The proof is by induction on the number of goods. Our inductive hypothesis is: for every set O of ℓ goods, prices p on O, and α-subadditive valuation v, there exists a 1-conservative bid vector a and a partition $S_1, S_2, \ldots, S_k$ of O such that $\tilde{u}(a,p) \geq v(S_1) + v(S_2) + \cdots + v(S_k)$.

For the base case (when $\ell = 1$), let p denote the price of the good. If $v(O) \geq p$, then for every (1-conservative) a between p and $v(O)$, $\tilde{u}(a,p) = v(O)$. If $p > v(O)$, then for every 1-conservative bid $a \leq v(O)$, $\tilde{u}(a,p) = p \geq v(O)$.

For the inductive step, let O be a set of $\ell > 1$ goods. There are now two cases, depending on whether or not there is a set A for which

$$\sum_{j \in A} p(j) > v(A). \tag{5}$$

If not, then setting $a = p$ yields a 1-conservative bid that wins the entire set O and hence achieves $\tilde{u}(a,p) = v(O)$.

For the second case, choose an arbitrary (non-empty) subset A that satisfies (5). By the inductive hypothesis, there is a 1-conservative bid vector a on $O \setminus A$ and a partition $S_1, S_2, \ldots, S_k$ of the goods $O \setminus A$ such that

$$\begin{aligned} \tilde{u}(a,p) &= v(U) + \sum_{j \in O \setminus (A \cup U)} p(j) \\ &\geq v(S_1) + v(S_2) + \cdots + v(S_k), \end{aligned} \tag{6}$$

where U is the goods of $O \setminus A$ won with the bid vector a.

Extend a to a bid vector on all of O by padding it with zeros on A; it remains 1-conservative. Applying inequalities (5) and (6) completes the inductive step:

$$\begin{aligned} \tilde{u}(a,p) &\geq v(U) + \sum_{j \in O \setminus (A \cup U)} p(j) + \sum_{j \in A} p(j) \\ &\geq v(S_1) + v(S_2) + \cdots + v(S_k) + v(A). \end{aligned}$$

Finally, using the definition of translated utility and α-subadditivity (3), we conclude that our bid vector a satisfies $u(a,p) \geq v(O)/\alpha - \sum_{j \in O} p(j)$, as desired. ■

We next establish another lemma that is used in all of our price of anarchy bounds. The lemma shows how to use the conservative bidding assumption to charge the sum of the maximum bids on items to the social welfare of the corresponding allocation.

Lemma 3.4 *For every γ-conservative bid profile $\mathbf{b} = \{b_1, b_2, \ldots, b_n\}$,*

$$\sum_{j \in M} \max_k b_k(j) \leq \gamma \cdot SW(\mathbf{b}).$$

Proof: The maximum bid on an item is the bid by the player who won that item in the corresponding allocation $\mathbf{X}(\mathbf{b})$. The γ-conservative assumption implies that the sum of a player's bids on the goods allocated to it is at most γ times the player's value for these goods. Hence,

$$\begin{aligned}
\sum_{j \in M} \max_k b_k(j) &= \sum_{i=1}^{n} \sum_{j \in X_i(\mathbf{b})} b_i(j) \\
&\leq \sum_{i=1}^{n} \gamma \cdot v_i(X_i(\mathbf{b})) \\
&= \gamma \cdot SW(\mathbf{b}).
\end{aligned}$$

■

We now prove our upper bound on the pure POA of CAs with item bidding (Theorem 3.1).

Proof of Theorem 3.1: Fix a valuation profile $\mathbf{v} = \{v_1, v_2, \ldots, v_n\}$ and let $OPT = \{O_1, O_2, \ldots, O_n\}$ denote an optimal allocation for it. Suppose the bid profile $\mathbf{b} = \{b_1, b_2 \ldots, b_n\}$ is a pure Nash equilibrium, with the corresponding allocation $\mathbf{X}(\mathbf{b}) = \{X_1(\mathbf{b}), X_2(\mathbf{b}), \ldots, X_n(\mathbf{b})\}$.

Fix a player i and define $p_i(j) = \max_{k \neq i} b_k(j)$ for each good j. Invoking Lemma 3.3, there is a feasible bid vector a_i such that

$$u_i(a_i, p_i) \geq \frac{v_i(O_i)}{\alpha} - \sum_{j \in O_i} p_i(j),$$

where u_i is defined as in (4). Combining this inequality with the Nash equilibrium condition we derive

$$\begin{aligned}
u_i(\mathbf{b}) &\geq u_i(a_i, \mathbf{b}_{-i}) \\
&= u_i(a_i, p_i) \\
&\geq \frac{v_i(O_i)}{\alpha} - \sum_{j \in O_i} p_i(j) \\
&\geq \frac{v_i(O_i)}{\alpha} - \sum_{j \in O_i} \max_k b_k(j), \qquad (7)
\end{aligned}$$

where in the last step we replace $\max_{k\neq i} b_k(j)$ with the only larger quantity $\max_k b_k(j)$. Using the fact that $v_i(X_i(\mathbf{b})) \geq u_i(\mathbf{b})$ for every player i, inequality (7), and Lemma 3.4, we have

$$\begin{aligned} SW(\mathbf{b}) &= \sum_{i=1}^{n} v_i(X_i(\mathbf{b})) \\ &\geq \sum_{i=1}^{n} u_i(\mathbf{b}) \\ &\geq \sum_{i=1}^{n} \left[\frac{v_i(O_i)}{\alpha} - \sum_{j\in O_i} \max_k b_k(j) \right] \\ &\geq \frac{SW(OPT)}{\alpha} - \gamma \cdot SW(\mathbf{b}). \end{aligned}$$

Rearranging terms completes the proof. ■

4 The Bayes-Nash POA

This section treats the incomplete information model. Section 4.1 shows by example that the guarantee of Theorem 3.1 does not carry over to this more general model. Section 4.2 covers two important lemmas. Section 4.3 proves the main result of this section (Theorem 4.6): the Bayes-Nash POA in every CA with item bidding and γ-conservative bidders with α-subadditive valuations is at most $\alpha(\gamma+1)\mathcal{H}_m$, where $\mathcal{H}_m = 1 + \frac{1}{2} + \cdots + \frac{1}{m}$ is the mth Harmonic number. Section 4.4 shows that, when bidders' valuations are not independent, the Bayes-Nash POA can be polynomial in the number of goods.

4.1 Separation of the Pure POA and Bayes-Nash POA

The following example shows that the Bayes-Nash POA of CAs with item bidding and 1-conservative bidders with subadditive valuations is larger than 2, the pure POA (Theorem 3.1). This separation stands in contrast to many previous works, where the worst-case pure POA coincides with that of several more general equilibrium concepts (including the Bayes-Nash POA); see also the discussion before Lemma 4.3.

Example 4.1 (The Bayes-Nash POA Is Larger Than 2) Consider 2 players and 8 goods, divided into two sets X_0 and X_1 with 4 goods each. Let $a = 0.24$, $b = 0.521$, and $p = 0.97$.

Each player $i = 0, 1$ has a deterministic valuation on the set X_{i+1} and a distribution over additive valuations on the set X_i. (All additions to indices are interpreted modulo 2.) Precisely, for each player i and good $j \in X_i$, $v_i(\{j\})$ is a with probability p and is b with probability $1-p$. The full valuation is then

$$v_i(S) = \begin{cases} 1 & \text{for } \emptyset \neq S \subset X_{i+1} \\ 2 & S = X_{i+1} \\ \sum_{j\in S} v_i(\{j\}) & \text{for } S \subseteq X_i \\ \max\{v_i(S\cap X_0), v_i(S\cap X_1)\} & \text{otherwise.} \end{cases}$$

Every valuation in the support is subadditive. Since one possible allocation is to award set X_0 to player 1 and set X_1 to player 0, $\mathbb{E}[SW(OPT)] \geq 4$.

Now consider the strategy profile in which player i bids $v_i(j)$ for every $j \in X_i$ and 0 for every $j \in X_{i+1}$. These bids are 1-conservative. The expected welfare of this strategy profile is

$$\begin{aligned}\mathbb{E}[v_0(X_0)] + \mathbb{E}[v_1(X_1)] &= \sum_i \sum_{j \in X_i} \mathbb{E}[v_i(j)] \\ &= 2 \times 4 \times (ap + b(1-p)) \\ &= 1.98744.\end{aligned}$$

We next verify that this strategy profile is a Bayes-Nash equilibrium, which proves that the Bayes-Nash POA is larger than 2.

By symmetry, we only need to check the equilibrium conditions for player 0. Player 1's bids act as posted prices for player 0. We show that for each fixed valuation of player 0, the maximum expected utility the player can achieve, given the distribution over posted prices induced by player 1's strategy, is achieved by the suggested strategy profile.

Fix a valuation v_0 for player 0. In the proposed strategy file, this player's expected utility is $v_0(X_0) \geq 4a = 0.96$. Since $v_0(S) = \max\{v_0(S \cap X_0), v_0(S \cap X_1)\}$ for each set S and player 0 currently wins all of X_0 without paying anything, we can focus on deviating bids that are non-zero only on goods of X_1.

First, suppose that player 0 bids at least b on some good of X_1. Since deviating bids must be 1-conservative, $2b > 1$, and $2a + b > 1$, player 0 bids at least b on only one good and at least a on at most one other good of X_1. If the player bids at least b on one good and at least a (but less than b) on another, then its expected utility is at most $(1-p)^2(1-b) + (1-p)p(1-a) + p(1-p)(1-a-b) + p^2(1-2a) = 1 - 2pa - b(1-p)$. If it bids at least b on one good and less than a on the rest, then its expected utility is at most $(1-p)(1-b) + p(1-a) = 1 - ap - b(1-p)$. In either case, its expected utility is at most $1 - ap - b(1-p) = 0.75157 < 0.96$. Second, if player 0 bids at least a but less than b on a strict subset of the goods in X_1, and less than a on the rest, then its expected utility is at most $1 - a < 0.96$. Finally, if player 0 bids at least a but less than b on all four goods in X_1, then its expected utility is at most $2p^4 + (1 - (1-p)^4 - p^4) - 4ap = 0.954 < 0.96$.

4.2 Two Lemmas on Fractionally Subadditive Valuations

This section provides the two main ingredients for our upper bound on the Bayes-Nash POA of CAs with item bidding. Both are stated in terms of fractionally subadditive valuations, defined next.

Definition 4.2 (Fractionally Subadditive Valuation) *A valuation v on a set M of goods is β-*fractionally subadditive *if for each subset $T \subseteq M$ there is a bid vector a that satisfies:*

(i) $\sum_{j \in S} a(j) \leq v(S)$ for every subset $S \subseteq T$;

(ii) $\sum_{j \in T} a(j) \geq v(T)/\beta$.

Put differently, Definition 4.2 asserts that for each target set T of goods, there is a 1-conservative bid vector with sum of bids at least $v(T)/\beta$ on the target set.[8]

The first lemma in this section plays a role analogous to that of Lemma 3.3. Because the Bayes-Nash POA is larger than 2 (Example 4.1), we need a more general statement that applies

[8]By strong duality, a valuation is fractionally subadditive if and only if for every "fractional cover" of a subset T of goods — meaning nonnegative coefficients $\lambda_1, \ldots, \lambda_k$ for sets $S_1, \ldots, S_k$ such that $\sum_{i\,:\, j \in S_i} \lambda_i \geq 1$ for every good $j \in T$ — the value of $\sum_{i=1}^{k} \lambda_i v(S_i)$ is at least $v(T)/\beta$. When $\beta = 1$, these are precisely the XOS ("XOR-of-OR-of-singletons") valuations [8].

beyond pure Nash equilibria. Essentially, our upper bound on the pure POA (Theorem 3.1) does not extend to the incomplete information model because it is not — and, by Example 4.1, cannot be — a "smoothness proof" in the sense of [17]. Unlike in [17], the Nash equilibrium hypothesis is invoked for player i in the proof of Theorem 3.1 with a candidate deviation a_i that depends on the bid vectors $\mathbf{b}_{-i}$ of the other players. The next lemma differs from Lemma 3.3 in that we exhibit a suitable bid vector for a player that is independent of the goods' prices — albeit only for β-fractionally subadditive valuations, rather than α-subadditive ones.

Lemma 4.3 (How To Bid Against All Posted Prices) *Let O be a set of goods. Let $v : 2^O \to \mathbb{R}^+$ be a β-fractionally subadditive valuation function on this set. There is a 1-conservative bid vector a such that, for every set of prices p on O,*

$$u(a,p) \geq \frac{v(O)}{\beta} - \sum_{j \in O} p(j). \tag{8}$$

Proof: Let a denote the bid vector guaranteed by Definition 4.2 for the valuation v and target set O. We can assume that a is zero outside of the set O. Let p be an arbitrary price vector. Let X be the goods of O that the player wins when it bids a. Using the two conditions of Definition 4.2 and the fact that $a(j) \leq p(j)$ for every good j that the player does not win, we have

$$\begin{aligned} u(a,p) &\geq v(X) - \sum_{j \in X} p(j) \\ &\geq \sum_{j \in X} a(j) - \sum_{j \in X} p(j) \\ &\geq \sum_{j \in X} a(j) - \sum_{j \in X} p(j) + \sum_{j \in O/X} (a(j) - p(j)) \\ &= \sum_{j \in O} a(j) - \sum_{j \in O} p(j) \\ &\geq \frac{v(O)}{\beta} - \sum_{j \in O} p(j). \end{aligned}$$

■

Lemma 4.3 applies to fractionally subadditive valuations, but how do these relate to subadditive valuations? The next lemma answers this question: every valuation that is α-subadditive in the sense of (3) is $\alpha\mathcal{H}_m$-fractionally subadditive, where $\mathcal{H}_m \approx \ln m$ is the mth Harmonic number. This result is tight up to a constant factor (Example 4.5).

Lemma 4.4 (α-Subadditive Implies $\alpha\mathcal{H}_m$-Fractionally Subadditive) *For every $\alpha \geq 1$ and set M of m goods, every α-subadditive valuation v over M is also $\alpha\mathcal{H}_m$-fractionally subadditive, where $\mathcal{H}_m = 1 + \frac{1}{2} + \cdots + \frac{1}{m}$.*

Proof: Given an α-subadditive valuation v and a target set $T \subseteq M$ of goods, we exhibit a bid vector a satisfying conditions (i) and (ii) of Definition 4.2, with $\beta = \alpha\mathcal{H}_m$. We use a variation of the greedy Set Cover algorithm (Algorithm 1) for this purpose.

To verify that the output a of Algorithm 1 satisfies condition (i) of Definition 4.2, fix $S \subseteq T$. Consider the iteration of Algorithm 1 in which the ℓth good of S is added to C. The algorithm

Algorithm 1 Certifying $(\alpha\mathcal{H}_m)$-fractional subadditivity

Input: α-subadditive valuation v on a set M of goods. Target set $T \subseteq M$.
Initialization: $C = \emptyset$, $a(j) = 0$ for all j in M.
While $C \neq T$,

- $A = \arg\min_{A' \subseteq T} \frac{v(A')}{|A' \setminus C|}$;
- for $j \in A \setminus C$, $a(j) = \frac{v(A)}{|A \setminus C| \cdot \mathcal{H}_m}$;
- $C = C \cup A$;

Output: the bid vector a.

could have chosen the set S in this iteration; by its greedy criterion, it chose a set A satisfying

$$\frac{v(A)}{|A \setminus C|} \leq \frac{v(S)}{|S \setminus C|} \leq \frac{v(S)}{(|S| - \ell + 1)}.$$

Thus,

$$\sum_{j \in S} a(j) \leq \frac{v(S)}{\mathcal{H}_m} \left(1 + \frac{1}{2} + \cdots + \frac{1}{|S|}\right) \leq v(S),$$

as required.

To show that condition (ii) of Definition 4.2 holds with $\beta = \alpha\mathcal{H}_m$, let $A_1, A_2, \ldots, A_t$ denote the sets chosen by Algorithm 1. By the definition (3) of α-subadditivity, we have

$$\sum_{j \in T} a(j) = \sum_{k=1}^{t} \sum_{j \in A_k} a(j) = \sum_{k=1}^{t} \frac{v(A_k)}{\mathcal{H}_m} \geq \frac{v(T)}{\alpha\mathcal{H}_m},$$

which completes the proof. ■

Example 4.5 (Lemma 4.4 Is Tight) Adapting a known lower bound for the worst-case integrality gap of Set Cover linear programs (see e.g. [18, Example 13.4]) shows that Lemma 4.4 is tight, in the worst case, up to a constant factor. In more detail, for a positive integer k we construct a valuation function on a set of $m = 2^k - 1$ goods. Number the goods from 1 to m and let $\mathbf{i}$ denote a k-bit binary vector representing the integer i. We interpret $\mathbf{i}$ as a k-dimensional vector over $GF[2]$ and use $\mathbf{i} \cdot \mathbf{j}$ to denote the dot product of the two vectors (modulo 2). Define $S_i = \{j : \mathbf{i} \cdot \mathbf{j} = 1\}$. Each such set contains $(m+1)/2$ goods and each good is contained in $(m+1)/2$ sets.

Define a valuation v by setting $v(S)$ to be the smallest number of sets of the form S_i whose union contains S. This valuation is subadditive. We show that every 1-conservative bid vector a satisfies $\sum_{j \in M} a(j) \leq 2v(M)/\log_2 m$.

The problem of maximizing the sum of bids subject to no overbidding can be formulated as a linear program

$$\begin{aligned} &\max \sum_{j \in M} a(j) \\ &\text{subject to } \sum_{j \in S} a(j) \leq v(S) && \forall S \subseteq M; \\ &a(j) \geq 0 && \forall j, \end{aligned}$$

with dual linear progam

$$\begin{aligned}
&\min \sum_{S\subseteq M} \lambda_S v(S) && \\
&\text{subject to} \sum_{S\,:\,j\in S} \lambda_S \geq 1 && \forall j; \\
&\lambda_S \geq 0 && \forall S \subseteq M.
\end{aligned}$$

For the valuation v defined above, setting $\lambda_{S_i} = 2/(m+1)$ for every $i = 1, 2, \ldots, m$ and $\lambda_S = 0$ otherwise yields a feasible dual solution with objective value $2m/(m+1) \leq 2$. Linear programming duality implies that $\sum_{j\in M} a(j) \leq 2$ for every 1-conservative bid vector a.

On the other hand, we claim that $v(M)$ is at least $k = \log_2(m+1)$. Consider $p < k$ sets $S_{i_1}, \ldots, S_{i_p}$. Let A be the matrix formed by the vectors $\mathbf{i_1}, \mathbf{i_2}, \ldots, \mathbf{i_p}$. This matrix has rank less than k, so its null space has rank at least one. Thus, there is a vector $\mathbf{j}$ such that $\mathbf{i_t} \cdot \mathbf{j} = 0$ for each t. By definition, j does not belong to any of the sets $S_{i_1}, \ldots, S_{i_p}$, so these sets do not cover M. As k sets are required to cover M, $v(M) \geq k$ and the subadditive valuation v is not β-subadditive for any $\beta < (\log_2 m)/2$.

4.3 Upper Bound on the Bayes-Nash POA

We now prove our main result for the incomplete information model.

Theorem 4.6 *For every CA with item bidding that has γ-conservative players with α-subadditive valuations, the Bayes-Nash POA is at most $\alpha(\gamma+1)\mathcal{H}_m$.*

Proof: Let $\mathcal{D} = D_1 \times D_2 \cdots \times D_n$ denote the common prior over players' valuations. The Bayes-Nash equilibrium strategy for each player i is a mixed strategy B^{v_i} for each $v_i \in V_i$. Let $B_{-i}^{v_{-i}}$ denote the joint distribution of mixed strategies of all players other than i for the valuation profile v_{-i}.

We fix a valuation profile for all players $\mathbf{v} = \{v_1, v_2, \ldots, v_n\}$. Let $\mathbf{O^v} = \{O_1^{\mathbf{v}}, O_2^{\mathbf{v}} \ldots, O_n^{\mathbf{v}}\}$ be a corresponding optimal allocation. Set $\beta = \alpha\mathcal{H}_m$. By Lemmas 4.3 and 4.4, for each player i there is a bid vector a_i such that, for every price vector p on $O_i^{\mathbf{v}}$,

$$u_i(a_i, p) \geq \frac{v_i(O_i^{\mathbf{v}})}{\beta} - \sum_{j\in O_i^{\mathbf{v}}} p(j).$$

In particular, if the prices are set equal to the maximum bids by other players in any bid profile then the above guarantee holds. So, if player i considers deviation a_i in the Bayes-Nash equilibrium, for each bid b_i in the mixed Bayes-Nash strategy B^{v_i} of player i, the Bayes-Nash condition implies:

$$\begin{aligned}
E_{w_{-i}|v_i, b_{-i}\sim B_{-i}^{w_{-i}}}\left[u_i(b_i, b_{-i})\right] &\geq E_{w_{-i}|v_i, b_{-i}\sim B_{-i}^{w_{-i}}}\left[u_i(a_i, b_{-i})\right] \\
&\geq E_{w_{-i}|v_i, b_{-i}\sim B_{-i}^{w_{-i}}}\left[\frac{v_i(O_i^{\mathbf{v}})}{\beta} - \sum_{j\in O_i^{\mathbf{v}}} \max_{k\neq i} b_k(j)\right] \\
&= \frac{v_i(O_i^{\mathbf{v}})}{\beta} - E_{w_{-i}|v_i, b_{-i}\sim B_{-i}^{w_{-i}}}\left[\sum_{j\in O_i^{\mathbf{v}}} \max_{k\neq i} b_k(j)\right]
\end{aligned}$$

The subscript $w_{-i}|v_i$ indicates that w_{-i} is drawn from the distribution $\mathcal{D}$ conditioned on v_i. Finally combined with the fact that for any bid profile, $v_i(X_i(\mathbf{b})) \geq u_i(\mathbf{b})$,

$$E_{w_{-i}|v_i, b_{-i}\sim B_{-i}^{w_{-i}}}\left[v_i(X_i(\mathbf{b}))\right] \geq \frac{v_i(O_i^{\mathbf{v}})}{\beta} - E_{w_{-i}|v_i, b_{-i}\sim B_{-i}^{w_{-i}}}\left[\sum_{j\in O_i^{\mathbf{v}}} \max_{k\neq i} b_k(j)\right]$$

The next expression is obtained in two steps. First consider all possible bids b_i for player i and take expectation according to B^{v_i}. The right hand side does not depend on b_i hence remains unchanged. On the left hand side, we now consider expected value to player i as b_i varies. We can obtain similar expressions for each player, sum these over all players and take expectation as valuation profile $\mathbf{v}$ is drawn from $\mathcal{D}$.

$$\begin{aligned}
&\mathop{\mathbb{E}}_{\mathbf{v}}\left[\sum_{i=1}^{n} E_{w_{-i}|v_i, b\sim B^{v_i,w_{-i}}}\left[v_i(X_i(b))\right]\right] \\
&\geq \mathop{\mathbb{E}}_{\mathbf{v}}\left[\sum_{i=1}^{n} v_i(O_i^{\mathbf{v}})/\beta\right] - \mathop{\mathbb{E}}_{\mathbf{v}}\left[\sum_{i=1}^{n} E_{w_{-i}|v_i, b_{-i}\sim B_{-i}^{w_{-i}}}\left[\sum_{j\in O_i^{\mathbf{v}}} \max_{k\neq i} b_k(j)\right]\right]
\end{aligned} \tag{9}$$

To keep things simple, we will process each term in the above expression independently.

Let's look at the first term.

$$\begin{aligned}
&\mathop{\mathbb{E}}_{\mathbf{v}}\left[\sum_{i=1}^{n} \mathop{\mathbb{E}}_{w_{-i}|v_i, b\sim B^{v_i,w_{-i}}}\left[v_i(X_i(b))\right]\right] \\
&= \sum_{i=1}^{n} \sum_{v_i, w_{-i}|v_i} \sum_{v_{-i}|v_i} D(w_{-i}|v_i) D(v_i) D(v_{-i}|v_i) \mathop{\mathbb{E}}_{b\sim B^{v_i,w_{-i}}}\left[v_i(X_i(b))\right] \\
&= \sum_{i=1}^{n} \sum_{v_i, w_{-i}|v_i} D(w_{-i}, v_i) \mathop{\mathbb{E}}_{b\sim B^{v_i,w_{-i}}}\left[v_i(X_i(b))\right] = E_{w,B}[SW(b)]
\end{aligned}$$

The simplification follows as none of the summands depend on v_{-i}. The second term in (9) is just $\mathbb{E}[SW(OPT)]$. Finally we simplify the last expression in (9). Here, in the second step we use the fact that the distributions D_i are independent.

$$\begin{aligned}
&\mathop{\mathbb{E}}_{\mathbf{v}}\left[\sum_{i=1}^{n} \mathop{\mathbb{E}}_{w_{-i}|v_i, b_{-i}\sim B_{-i}^{w_{-i}}}\left[\sum_{j\in O_i^{\mathbf{v}}} \max_{k\neq i} b_k(j)\right]\right] \\
&= \mathop{\mathbb{E}}_{\mathbf{v}}\left[\sum_{i=1}^{n} \sum_{w_{-i}|v_i} D(w_{-i}|v_i) \mathop{\mathbb{E}}_{b_{-i}\sim B_{-i}^{w_{-i}}}\left[\sum_{j\in O_i^{\mathbf{v}}} \max_{k\neq i} b_k(j)\right]\right] \\
&= \mathop{\mathbb{E}}_{\mathbf{v}}\left[\sum_{i=1}^{n} \sum_{w_{-i}} D(w_{-i}) \mathop{\mathbb{E}}_{b_{-i}\sim B_{-i}^{w_{-i}}}\left[\sum_{j\in O_i^{\mathbf{v}}} \max_{k\neq i} b_k(j)\right]\right] \\
&= \mathop{\mathbb{E}}_{\mathbf{v}}\left[\sum_{i=1}^{n} \sum_{w} D(w_{-i}) D(w_i) \mathop{\mathbb{E}}_{b\sim B^{w}}\left[\sum_{j\in O_i^{\mathbf{v}}} \max_{k\neq i} b_k(j)\right]\right]
\end{aligned}$$

In the last step we introduce a new copy for player i's valuation along with the corresponding mixed strategy bids. Currently nothing depends on these hence is without loss. Continuing with the simplification,

$$
\begin{aligned}
&\mathop{\mathbb{E}}_{\mathbf{v}}\left[\sum_{i=1}^{n} \mathop{\mathbb{E}}_{w_{-i}|v_i, b_{-i}\sim B_{-i}^{w_{-i}}}\left[\sum_{j\in O_i^{\mathbf{v}}} \max_{k\neq i} b_k(j)\right]\right] \\
&\geq \mathop{\mathbb{E}}_{\mathbf{v}}\left[\sum_{i=1}^{n}\sum_{w} D(w) \mathop{\mathbb{E}}_{b\sim B^w}\left[\sum_{j\in O_i^{\mathbf{v}}} \max_{k\neq i} b_k(j)\right]\right] \\
&\geq \mathop{\mathbb{E}}_{\mathbf{v}}\left[\sum_{i=1}^{n}\sum_{w} D(w) \mathop{\mathbb{E}}_{b\sim B^w}\left[\sum_{j\in O_i^{\mathbf{v}}} \max_{k} b_k(j)\right]\right] \\
&= \mathop{\mathbb{E}}_{\mathbf{v}}\left[\sum_{w} D(w) \mathop{\mathbb{E}}_{b\sim B^w}\left[\sum_{i=1}^{n}\sum_{j\in O_i^{\mathbf{v}}} \max_{k} b_k(j)\right]\right] \\
&= \mathop{\mathbb{E}}_{\mathbf{v}}\left[\sum_{w} D(w) \mathop{\mathbb{E}}_{b\sim B^w}\left[\gamma SW(b)\right]\right] \\
&= \mathop{\mathbb{E}}_{\mathbf{v}}\left[\gamma \mathop{\mathbb{E}}_{w,b\sim B^w}[SW(b)]\right] = \gamma \mathop{\mathbb{E}}_{w,B}[SW(b)]
\end{aligned}
$$

Note that the expression $\max_{k\neq i} b_k(j)$ only grows when replaced by $\max_k b_k(j)$. The second-last step follows from lemma 3.4. Expression (9) now becomes,

$$
\mathop{\mathbb{E}}_{w,B}[SW(b)] \geq \frac{\mathbb{E}[SW(OPT)]}{\beta} - \gamma \mathop{\mathbb{E}}_{w,B}[SW(b)]
$$

Reorganizing we obtain the bound of $\beta(\gamma+1)$. ■

4.4 The Bayes-Nash POA with Correlated Valuations

An assumption in the proof of theorem 4.6 was the player's valuations were independent. This assumption is crucial. In the example below, where player valuations are correlated, any Bayes-Nash equilibrium of the one-conservative CA with Item Bidding has POA $\Omega(n^{1/4})$.

Example 4.7 There are $n+\sqrt{n}$ total items and $2n$ players. Players occur in pairs. Each pair contains one player each of types I and II. A valuation from the correlated distribution $\mathcal{D}$ is drawn as follows. First $\sqrt{n}$ items are picked at random, this is the "common pool". The remaining n items are randomly matched with the n pairs of players. For each pair, this is the "reserve item". A player of type I values any non-empty subset of the $\sqrt{n}+1$, common pool plus one reserve item, iitems at one. Player of type II values only the corresponding reserve item at $n^{-1/4}$.

First note that we can assume in a Bayes-Nash equilibrium a player of type II always bids $n^{-1/4}$. If the player is bidding lower than $n^{-1/4}$, then bidding higher can only tilt the outcome in its favor. But since this is an equilibrium raising its bid should not change its expected utility; and since it has a single minded valuation also not change the outcome. Hence raising his bid to $n^{-1/4}$ is without loss.

A player of type I when bidding in an equilibrium sees a set of $\sqrt{n}+1$ items where each item is equally likely to be the reserve item. Note that it has to bid more than $n^{-1/4}$ to win the reserve item. Moreover, due to the 1-conservative condition, it can bid more than $n^{-1/4}$ on at most $n^{1/4}$ items. Thus only with probability $n^{1/4}/(\sqrt{n}+1) \approx n^{-1/4}$, it will win the reserve item. In all the other cases the reserve item is picked up by the type II player. Hence we conclude:

$$E_{\mathbf{v}\sim\mathcal{D}}[SW(\mathbf{X}(\mathbf{b}))] \leq \sqrt{n} + n \cdot n^{-1/4} + n \cdot n^{-1/4}$$

The social welfare in the Bayes-Nash is $O(n^{3/4})$. While in the optimal, all players of type I will snag their reserve items yielding a social welfare of at least n. Hence the Price of Anarchy is $\Omega(n^{1/4})$.

5 PoA Bound for Coarse-Correlated Equilibria

A proof similar to that of Theorem 4.6 establishes the same upper bound of $\beta(\gamma+1)$ on the worst-case POA of coarse correlated equilibria in the full information model. Due to the inclusion properties of the equilibrium concepts the same bound applies to the mixed POA and correlated POA.

Theorem 5.1 *If all players have valuation functions v_i that are β-fractionally subadditive and all players are γ-conservative, then the coarse correlated POA of a combinatorial auction with item bidding is $\beta(\gamma+1)$.*

Proof: Let $\mathcal{B}$ be the distribution over different bid profiles in the coarse correlated equilibrium. Let $\mathbf{O} = \{O_1, O_2, \ldots, O_n\}$ be the optimal allocation for the players' valuation profiles. The coarse correlated equilibrium condition for a player i is

$$\mathop{\mathbb{E}}_{\mathcal{B}}[u_i(\mathbf{b})] \geq \mathop{\mathbb{E}}_{\mathcal{B}}[u_i(a_i', \mathbf{b}_{-i})].$$

Fix a player i. From Lemma 4.3 there exist a 1-conservative deviation a_i such that for every bid profile $\mathbf{b}$, setting highest bids by players other than i as prices on the items we have,

$$u_i(a_i, b_{-i}) \geq \frac{v_i(O_i)}{\beta} - \sum_{j\in O_i} \max_{k\neq i} b_k(j)$$

Applying the coarse-correlated equilibrium condition for player i, with deviation a_i,

$$\mathop{\mathbb{E}}_{\mathcal{B}}[u_i(b)] \geq \mathop{\mathbb{E}}_{\mathcal{B}}[u_i(a_i, b_{-i})] \geq \frac{v_i(O_i)}{\beta} - \mathop{\mathbb{E}}_{\mathcal{B}}\left[\sum_{j\in O_i} \max_{k\neq i} b_k(j)\right]$$

Summing this over all players and using the fact that for each bid profile $\mathbf{b}$ and player i, $v_i(X_i(\mathbf{b})) \geq u_i(\mathbf{b})$ we obtain,

$$\begin{aligned}
\mathop{\mathbb{E}}_{\mathcal{B}}[SW(\mathbf{b})] &\geq \sum_{i=1}^{n} \mathop{\mathbb{E}}_{\mathcal{B}}[u_i(\mathbf{b})] \\
&\geq \sum_{i=1}^{n} \left[\frac{v_i(O_i)}{\beta} - \mathop{\mathbb{E}}_{\mathcal{B}}\left[\sum_{j\in O_i} \max_{k\neq i} b_k(j)\right]\right] \\
&\geq \frac{SW(\mathbf{O})}{\beta} - \mathop{\mathbb{E}}_{\mathcal{B}}\left[\sum_{i=1}^{n}\sum_{j\in O_i} \max_{k} b_k(j)\right] \\
&\geq \frac{SW(\mathbf{O})}{\beta} - \gamma \cdot \mathop{\mathbb{E}}_{\mathcal{B}}[SW(\mathbf{X}(\mathbf{b}))]
\end{aligned}$$

In the second step we replace $\max_{k \neq i} b_k(j)$ with $\max_k b_k(j)$ which is clearly larger. The last step follows from lemma 3.4. Reorganizing the above inequality establishes that the POA is at most $\beta(\gamma + 1)$. ■

6 Conclusion

Our work suggests a number of interesting open questions.

- What is the POA of Bayes-Nash equilibria in combinatorial auctions with item bidding and subadditive valuations? There is a significant gap between our upper and lower bounds.
- Identify necessary and sufficient conditions for the existence of a pure Nash equilibrium in a combinatorial auction with item bidding and subadditive valuations. (Cf., Example 2.1.)
- As outlined in the Introduction, our results quantify the equilibrium efficiency loss that results from forcing bidders with subadditive valuations to submit additive bids. Is there a more general theory about the equilibrium efficiency loss caused by this type of "valuation compression"?

References

[1] L. M. Ausubel and P. Milgrom. The lovely but lonely Vickrey auction. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*, chapter 1, pages 17–40. MIT Press, 2006.

[2] J. Bae, E. Beigman, R. Berry, M. Honig, and R. Vohra. Sequential bandwidth and power auctions for distributed spectrum sharing. *Selected Areas in Communications, IEEE Journal on*, 26(7):1193–1203, 2008.

[3] A. Blum, M. T. Hajiaghayi, K. Ligett, and A. Roth. Regret minimization and the price of total anarchy. In *40th Annual ACM Symposium on Theory of Computing (STOC)*, pages 373–382, 2008.

[4] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, and M. Kyropoulou. On the efficiency of equilibria in generalized second price auctions. In *12th ACM Conference on Electronic Commerce (EC)*, pages 81–90, 2011.

[5] G. Christodoulou, A. Kovács, and M. Schapira. Bayesian combinatorial auctions. In *35th International Colloqium on Automata, Languages and Programming (ICALP)*, pages 820–832, 2008.

[6] S. Dobzinski. Two randomized mechanisms for combinatorial auctions. In *10th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems*, pages 89–103, 2007.

[7] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007.

[8] U. Feige. On maximizing welfare when utility functions are subadditive. *SIAM Journal on Computing*, 39(1):122–142, 2009.

[9] I. Giotis and A. Karlin. On the equilibria and efficiency of the GSP mechanism in keyword auctions with externalities. In *The 4th International Workshop on Internet and Network Economies (WINE)*, pages 629–638, 2008.

[10] J. C. Harsanyi. Games with incomplete information played by 'bayesian' players, parts I, II, and III. *Management Science*, 14, 1967-68.

[11] A. Hassidim, H. Kaplan, M. Mansour, and N. Nisan. Non-price equilibria in markets of discrete goods. In *12th ACM Conference on Electronic Commerce*, pages 295–296, 2011.

[12] B. Lehmann, D. J. Lehmann, and N. Nisan. Combinatorial auctions with decreasing marginal utilities. In *3rd ACM Conference on Electronic Commerce (EC)*, pages 18–28, 2001.

[13] R. Paes Leme and É. Tardos. Pure and Bayes-Nash price of anarchy for generalized second price auction. In *51st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 735–744, 2010.

[14] B. Lucier and A. Borodin. Price of anarchy for greedy auctions. In *21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 537–553, 2010.

[15] B. Lucier and R. Paes Leme. Gsp auctions with correlated types. In *12th ACM Conference on Electronic Commerce (EC)*, pages 71–80, 2011.

[16] P. Milgrom. *Putting Auction Theory to Work*. Cambridge University Press, 2004.

[17] T. Roughgarden. Intrinsic robustness of the price of anarchy. In *41st ACM Symposium on Theory of Computing (STOC)*, pages 513–522, 2009.

[18] V. V. Vazirani. *Approximation Algorithms*. Springer, 2003.

[19] H. P. Young. *Strategic Learning and Its Limits*. Oxford University Press, 2004.