

Best-Response Dynamics in Combinatorial Auctions with Item Bidding

Paul Dütting*

Thomas Kesselheim†

Abstract

In a combinatorial auction with item bidding, agents participate in multiple single-item second-price auctions at once. As some items might be substitutes, agents need to strategize in order to maximize their utilities. A number of results indicate that high welfare can be achieved this way, giving bounds on the welfare at equilibrium. Recently, however, criticism has been raised that equilibria are hard to compute and therefore unlikely to be attained.

In this paper, we take a different perspective. We study simple best-response dynamics. That is, agents are activated one after the other and each activated agent updates his strategy myopically to a best response against the other agents' current strategies. Often these dynamics may take exponentially long before they converge or they may not converge at all. However, as we show, convergence is not even necessary for good welfare guarantees. Given that agents' bid updates are aggressive enough but not too aggressive, the game will remain in states of good welfare after each agent has updated his bid at least once.

In more detail, we show that if agents have fractionally subadditive valuations, natural dynamics reach and remain in a state that provides a $1/3$ approximation to the optimal welfare after each agent has updated his bid at least once. For subadditive valuations, we can guarantee an $\Omega(1/\log m)$ approximation in case of m items that applies after each agent has updated his bid at least once and at any point after that. The latter bound is complemented by a negative result, showing that no kind of best-response dynamics can guarantee more than a an $o(\log \log m / \log m)$ fraction of the optimal social welfare.

1 Introduction

In a combinatorial auction, n players compete for the assignment of m items. The players have private preferences over bundles of items as expressed by a valuation function $v_i: 2^{[m]} \rightarrow \mathbb{R}_{\geq 0}$. Our goal in this work is to find a partition of the items into sets S_1, \dots, S_n that maximizes social welfare $\sum_i v_i(S_i)$, based on reported valuations (bids) $b_i: 2^{[m]} \rightarrow \mathbb{R}_{\geq 0}$

with the freedom to impose payments p_1, \dots, p_n on the players.

Even if valuations are known, finding an allocation that maximizes social welfare is typically NP-hard. Furthermore, since valuations are assumed to be private information, some mechanics are needed to extract this information. The traditional approach is to incentivize players to bid truthfully. Insisting on truthfulness has the advantage that for the individual players it is easy to participate as it is not necessary to act strategically. However, truthfulness requires central coordination of the entire allocation and payments.

An alternative approach to this problem that is arguably seen more often in practice is to let players participate in a simpler, non-truthful mechanism and to accept strategic behavior. To derive theoretical performance guarantees, one then seeks to prove bounds on the so-called Price of Anarchy, the worst-case ratio between the optimal social welfare and the welfare at equilibrium. The most prominent example in the context of combinatorial auctions is *item bidding*, where the items are sold through separate single-item auctions.

One can show that for pretty general classes of valuations, such as submodular or the even more general classes fractionally subadditive and subadditive, all equilibria from a broad range of equilibrium concepts obtain a decent fraction of the optimal social welfare. More recently, however, these results have been criticized for ignoring the computational complexity of finding an equilibrium. In fact, by now, there is quite a selection of impossibility results showing that finding exact equilibria is often computationally intractable.

Our approach in this paper is different. We consider simple, best-response dynamics, in which players are activated in a round-robin fashion and players when activated buy their favorite set of items at the current prices, in a myopic way. Christodoulou et al. [7] showed that one instance of such dynamics converges if players' valuation functions are fractionally subadditive. However, they also showed that it takes exponential time. For subadditive valuations, even convergence cannot be guaranteed because any fixed point would be a pure Nash equilibrium, and pure Nash equilibria may not exist (see Appendix A). We show that despite possibly long convergence time or no convergence at all,

*Department of Mathematics, London School of Economics, Houghton Street, London WC2A 2AE, UK. Email: p.d.duetting@lse.ac.uk.

†Max Planck Institute for Informatics and Saarland University, Saarland Informatics Campus, Campus E1 4, 66123 Saarbrücken, Germany. Email: thomas.kesselheim@mpi-inf.mpg.de. Supported in part by the DFG through Cluster of Excellence MMCI.

the social welfare reaches a good level very fast.

1.1 The Setting We study combinatorial auctions with n bidders N and m items M . Each bidder $i \in N$ has a valuation function $v_i: 2^M \rightarrow \mathbb{R}_{\geq 0}$. Our objective is to find a feasible allocation, i.e., a partition of the items, S_1, \dots, S_n , that maximizes social welfare $\sum_{i \in N} v_i(S_i)$. We assume that an allocation of items to bidders is found by distributed strategic behavior of the bidders using item bidding. That is, each bidder $i \in N$ places a bid $b_{i,j}$ on each item $j \in M$. Each item $j \in M$ is assigned to the bidder $i \in N$ with the highest bid $b_{i,j}$ at a price of $p_j = \max_{i' \neq i} b_{i',j}$. Ties are broken in an arbitrary, but fixed manner.

We assume that bidders choose their bids strategically so as to maximize their quasi-linear utilities. Bidder i 's utility u_i as a function of the bids $b = (b_{i'})_{i' \in N}$ is $u_i(b) = v_i(S) - \sum_{j \in S} p_j$, where S is the set of items won by bidder i .

We say that a bid b_i is a best response to the bids b_{-i} if bidder i 's utility is maximized by b_i . That is, $u_i(b_i, b_{-i}) \geq u_i(b'_i, b_{-i})$ for all b'_i . Note that any best response must give bidder i a set of items S that maximizes $u_i(b) = v_i(S) - \sum_{j \in S} p_j$. We call these sets of items demand sets. A (pure) Nash equilibrium in this setting is a profile of bids $b = (b_{i'})_{i' \in N}$ such that for each bidder $i \in N$ his bid b_i is a best response against bids b_{-i} .

We study simple game-playing dynamics in which bidders get activated in turn and myopically choose to play a best response. More formally, starting from an initial bid vector b^0 , in each time step $t \geq 1$, some bidder $i \in N$ is activated and updates his bid b_i^{t-1} from the previous round to a best response to the other players' bids $b_{-i}^{t-1} = b_{-i}^{t-1}$ which do not change from the previous to the current round. The fixed points of such best-response dynamics are Nash equilibria. However, Nash equilibria do not necessarily exist and even if they do best-response dynamics may not converge.

We will evaluate best-response dynamics by the social welfare that they achieve. For bid profile b and corresponding allocation S_1, \dots, S_n we write $SW(b) = \sum_i v_i(S_i)$ for the social welfare at bid profile b . We seek to compare this to the optimal social welfare $OPT(v)$.

1.2 Variants of Best-Response Dynamics Since payments in combinatorial auctions with item bidding are second price, there are typically many ways to choose a best response. Clearly, not all best responses will ensure that good states (in terms of social welfare) will be reached quickly.

EXAMPLE 1.1. (GROSS UNDERBIDDING) Consider a single-item auction with n bidders. Suppose $v_1 = C$

and $v_i = 1$ for $i \geq 2$, where $C \gg 1$. Suppose we start at $b = (0, \dots, 0)$ and the item assigned to bidder 1. A possible best response sequence has bidders update their bids in round-robin fashion, each time increasing the winning bid by ϵ .

EXAMPLE 1.2. (GROSS OVERBIDDING) Consider the same setting as in the previous example. If in the first round of updates the last bidder bids $C + \epsilon$ this will terminate the dynamics.

Note that in both these examples the social welfare after each round of best responses (and on average) is 1, which can be arbitrarily smaller than the optimal social welfare C .

The issue in each of these examples is as follows. Through the bids $b_{i,j}$, the bidders effectively declare additive valuations. The allocation maximizes the declared welfare $DW(b) = \sum_i \sum_{j \in S_i} b_{i,j}$, which usually differs from the actual welfare $SW(b)$. In both examples, there exist update steps in which the declared utility of the bidder, i.e., $u_i^D(b) = \sum_{j \in S_i} b_{i,j} - \sum_{j \in S_i} \max_{k \neq i} b_{k,j}$, is very different from his actual utility. We will prove bounds on the welfare achieved by best-response dynamics that are quantified by the extent to which declared utilities can differ from the actual utilities as captured by the following definitions.

DEFINITION 1.1. Let $\alpha \geq 0$. We call a bid b_i by bidder i against bids b_{-i} α -aggressive if $u_i^D(b) \geq \alpha \cdot \max_{b'_i} u_i(b'_i, b_{-i})$.

DEFINITION 1.2. Let $\beta \geq 1$. A best response dynamic is β -safe if it ensures that $u_i^D(b) \leq \beta \cdot u_i(b)$ for all players i and reachable bid profiles b .

We will usually apply Definition 1.1 when b_i is a best response to b_{-i} . However, it also leaves the freedom to consider approximate best responses. We will see that one way to achieve Definition 1.2 is to require strong no overbidding, but we will also see an example of safe dynamics that allow overbidding. Note that in both cases players will have non-negative actual utilities at all times because $u_i(b^t) \geq \frac{1}{\beta} \cdot u_i^D(b^t) \geq 0$ for every bidder i and time step t .

1.3 Our Results Our first main result is that round-robin best-response dynamics are capable of reaching states with near-optimal social welfare strikingly fast, despite the fact that convergence to equilibrium may take exponentially long or they may not converge at all.

In fact, our result applies to any round-robin bidding dynamics, provided that players choose bids that are aggressive enough but not too aggressive. This, in

particular, includes dynamics in which players choose to play only approximate best responses. Also, their way of making choices does not need to be consistent in any way.

MAIN RESULT 1.1. *In a β -safe round-robin bidding dynamic with α -aggressive bid updates the social welfare at any time step $t \geq n$ satisfies*

$$SW(b^t) \geq \frac{\alpha}{(1 + \alpha + \beta)\beta} \cdot OPT(v).$$

In other words, once every player had the chance to update his bid, the social welfare, at any time step after that, will be within $\alpha/(1 + \alpha + \beta)\beta$ of optimal.

For fractionally subadditive valuations and subadditive valuations there exist round-robin best-response dynamics with $(\alpha, \beta) = (1, 1)$ and $(\alpha, \beta) = (1/\ln m, 1)$. The result for XOS requires access to demand and XOS oracles [11], the result for subadditive valuations requires access to demand oracles and that the greedy algorithm for set cover problems can be executed [16, 2].

Our guarantee on the social welfare achieved by best-response dynamics shows that these dynamics provide a $1/3$ (resp. $\Omega(1/\log m)$) approximation to the optimal social welfare that applies after a single round of bid updates, and at any time step after that.

We also prove a bound on the average social welfare of $1/2(2 + \alpha)\beta$, which improves upon the above bound for large β . In particular, for subadditive valuations it is also possible to achieve $(\alpha, \beta) = (1, \ln m)$. While the point-wise guarantee of this dynamics is only $\Omega(1/\log^2 m)$, its average social welfare is within $\Omega(1/\log m)$ of optimal.

We show that the point-wise welfare guarantee of $1/3$ for fractionally subadditive valuations is tight for the respective mechanism. Our second main result is that the $\Omega(1/\log m)$ bounds are almost best possible in a more general sense.

MAIN RESULT 1.2. *For subadditive valuations no best-response dynamics in which players do not overbid on the grand bundle can guarantee a better than $o(\log \log m / \log m)$ fraction of the optimal social welfare at any time step.*

For round-robin bidding dynamics, this point-wise impossibility result extends to an impossibility for the average social welfare that can be achieved.

The assumption that players do not overbid on the grand bundle is quite natural, and is satisfied by all dynamics that have been proposed in the literature. It obviously applies to strong no-overbidding dynamics, but it also applies to dynamics in which players use weak no-overbidding strategies on the items that they win and bid zero on all other items.

Our proof of the lower bound is based on a non-trivial construction exploiting the algebraic properties of linearly independent vector spaces. It presents an interesting separation from the Price of Anarchy literature, where no such lower bound can be proved.

Finally, we explore to which extent our positive results depend on round-robin activation. We show that our positive results extend to the case where at each step a player is chosen uniformly at random, while the social welfare can be as low as $O(1/n)$ of optimal when the order of activation is chosen adversarially.

1.4 Related Work Best-response dynamics are a central topic in Algorithmic Game Theory. Probably, the best-studied application are congestion games, where best-response dynamics always converge but, except in special cases, take worst-case exponential time before they do so [23, 21, 1]. On the other hand, a number of results show that certain types of best-response dynamics reach states of low social cost quickly [19, 6, 3, 15, 25]. Some of these results extend to weighted congestion games, where equilibria may not exist and best-response sequences may not converge for this reason.

The study of the Price of Anarchy in combinatorial auctions with item bidding was initiated by Christodoulou et al. [7], and subsequently refined and improved upon in [2, 20, 26, 14, 18]. Some of these bounds are based on mechanism smoothness, others are not. They provide welfare guarantees for a broad range of equilibrium concepts ranging from pure Nash equilibria, over (coarse) correlated equilibria, to Bayes-Nash equilibria. For fractionally subadditive valuations there is a smoothness-based proof that shows that the Price of Anarchy with respect to pure Nash equilibria is at most 2 [7, 26]. For subadditive valuations the Price of Anarchy with respect to pure Nash equilibria is also at most 2 [2], but the best smoothness-based proof gives a bound of $O(\log m)$ [2, 26]. In fact, as shown by Roughgarden [24], combinatorial auctions with item bidding achieve (near-)optimal Price of Anarchy among a broad class of “simple” mechanisms.

Also relevant to our analysis in this context is that Christodoulou et al. [7] gave a simple, best-response dynamics for fractionally subadditive valuations, that they called Potential Procedure. They showed that this procedure always converges to a pure Nash equilibrium, but also that it may take exponentially many steps before it converges.

Lately, attempts at proving Price of Anarchy bounds for combinatorial auctions with item bidding have been criticized for not being constructive, in the sense that the computational complexity of finding an

equilibrium remained open. Dobzinski et al. [13], for example, showed that for subadditive valuations computing a pure Nash equilibrium requires exponential communication. Regarding fractionally subadditive valuations they concluded that “if there exists an efficient algorithm that finds an equilibrium, it must use techniques that are very different from our current ones.” Further negative findings were reported by Cai and Papadimitriou [5], who showed that computing a Bayes-Nash equilibrium is PP-hard.

Most recently, Daskalakis and Syrgkanis [8] considered coarse correlated equilibria. They showed that even for unit-demand players (a strict subclass of submodular) there are no polynomial-time no-regret learning algorithms for finding such equilibria, unless $\text{RP} \supseteq \text{NP}$, closing the last gap in the equilibrium landscape. However, they also proposed a novel solution concept to escape the hardness trap, no-envy learning, and gave a polynomial-time no-envy learning algorithm for XOS valuations and complemented this with a proof showing that for this class of valuations every no-envy outcome recovers at least $1/2$ of the optimal social welfare.

Further relevant work comes from Devanur et al. [9], who proposed an alternative to simultaneous second-price auctions, the so-called single-bid auction. This mechanism also admits a polynomial-time no-regret learning algorithm and by a result of [4] achieves optimal Price of Anarchy bounds within a broader class of mechanisms.

A final point of reference are truthful mechanisms for combinatorial auctions. While no mechanism can achieve a better than $1/m^{1/2-\epsilon}$ approximation for submodular valuations with valuation queries alone [12], Dobzinski [10] recently managed to improve a long-standing approximation guarantee of $\Omega(1/\log m)$ for submodular valuations to $\Omega(1/\sqrt{\log m})$ for fractionally subadditive valuations, requiring access to both value and demand oracles.

2 Achieving Aggressive and Safe Bids

As already discussed, best responses are generally not unique in our settings. Our positive results require that updates are *aggressive* and *safe*. In this section we briefly describe how to guarantee these properties for fractionally subadditive (a.k.a. XOS) valuations and subadditive valuations. The missing proofs are provided in Appendix B.

A valuation function is *fractionally subadditive*, or XOS, if there are values $v_{i,j}^\ell \geq 0$ such that $v_i(S) = \max_\ell \sum_{j \in S} v_{i,j}^\ell$. It is *subadditive* if for all $S, T \subseteq M$, $v_i(S \cup T) \leq v_i(S) + v_i(T)$.

The dynamics that we consider approach players in round-robin fashion. When player i is activated he picks

a demand set D at the current prices and updates his bid as described below. Note that here we assume eager updating. This assumption leads to cleaner proofs, but is not necessary. See the full version for details.

2.1 Bid Updates for XOS Valuations For XOS valuations we can update bids as described by [7]. If D is the demand set chosen by player i , let $(v_{i,j}^\ell)_{j \in M}$ be the supporting valuation on this demand set for which $\sum_{j \in D} v_{i,j}^\ell = v_i(D)$, and set $b_{i,j}^t = v_{i,j}^\ell$ for $j \in D$ and $b_{i,j}^t = 0$ otherwise. Note that these update steps can be performed in polynomial time using demand and XOS oracles.

PROPOSITION 2.1. *Starting from an initial bid vector b^0 satisfying strong no-overbidding, the bid updates described above lead to a sequence of bids b^0, b^1, b^2, \dots that is 1-safe and in which each update is a 1-aggressive best response.*

2.2 Bid Updates for Subadditive Valuations

For subadditive functions, it is generally not possible to guarantee $\alpha = 1$ and $\beta = 1$ at the same time. We describe two different, reasonable ways of bid updates.

No-Overbidding Updates Given a bid vector b_{-i} , let $\tilde{u}_i(S, b_{-i}) = v_i(S) - \sum_{j \in S} \max_{k \neq i} b_{k,j}$. That is, $\tilde{u}_i(S, b_{-i})$ is the utility bidder i can derive from buying the set S . Observe that \tilde{u} is subadditive. Let D be an inclusion-wise minimal demand set of bidder i given b_{-i} . We can show that $\tilde{u}_i(S, b_{-i}^t) > 0$ for all $S \subseteq D$ unless $D = \emptyset$. Therefore, by [2] there exists an additive approximation a_i such that (a) $\sum_{j \in D} a_{i,j} \geq 1/\ln m \cdot \tilde{u}_i(D, b_{-i}^t)$ and (b) $\sum_{j \in S} a_{i,j} \leq \tilde{u}_i(S, b_{-i}^t)$ for all $S \subseteq D$ with the property that $a_{i,j} > 0$ for all $j \in D$. We set bids $b_{i,j}^t = a_{i,j} + \max_{k \neq i} b_{k,j}^t$ for $j \in D$ and $b_{i,j}^t = 0$ otherwise. These update steps can be performed in polynomial time with a demand oracle if it is possible to compute the additive approximation, which corresponds to executing the greedy set-cover algorithm on $\tilde{u}_i(\cdot, b_{-i}^t)$.

PROPOSITION 2.2. *Starting from an initial bid vector b^0 that satisfies strong no-overbidding, the bid updates described above lead to a sequence of bids b^0, b^1, b^2, \dots that is 1-safe and in which each update is a $(1/\ln m)$ -aggressive best response.*

Aggressive Updates The basic construction is the same as above except that instead of considering a_i we consider \tilde{a}_i such that $\tilde{a}_{i,j} = \gamma \cdot a_{i,j}$ for all items $j \in D$, where $0 < \gamma \leq \ln m$ is such that $\sum_{j \in D} a_{i,j} = 1/\gamma \cdot \tilde{u}_i(D, b_{-i}^t)$. Note that these bids satisfy: (a) $\sum_{j \in D} \tilde{a}_{i,j} = \tilde{u}_i(D, b_{-i}^t)$ and (b) $\sum_{j \in S} \tilde{a}_{i,j} \leq \gamma \cdot \tilde{u}_i(S, b_{-i}^t)$ for all $S \subseteq D$.

PROPOSITION 2.3. *Starting from an initial bid vector b^0 that satisfies strong no-overbidding, the bid updates described above lead to a sequence of bids that is $\ln m$ -safe and in which each update is a 1-aggressive best response.*

3 Welfare Guarantees

In this section we prove our first main result (Theorem 3.1). The theorem provides a point-wise social welfare guarantee, parametrized in α and β , for round-robin bidding dynamics. It shows that the social welfare is high already after a single round of updates, and remains high at every single step after that.

THEOREM 3.1. *In a β -safe round-robin bidding dynamic with α -aggressive bid updates the social welfare at any time step $t \geq n$ satisfies $SW(b^t) \geq \frac{\alpha}{(1+\alpha+\beta)\beta} \cdot OPT(v)$.*

As we have argued in Proposition 2.1 and Proposition 2.2 there exist round-robin best-response dynamics with $(\alpha, \beta) = (1, 1)$ for fractionally subadditive valuations and $(\alpha, \beta) = (1/\ln m, 1)$ for subadditive valuations. So two corollaries of our theorem are point-wise welfare guarantees of $1/3$ and $\Omega(1/\log m)$ for the respective mechanisms.

We also show a welfare guarantee for the average social welfare, Theorem 3.2 below, that improves upon the pointwise guarantee for large β . Note that the term $(1 - \frac{n}{T})$ is $1 - o(1)$ for $T \in \omega(n)$ and at least $1/2$ for $T \geq 2n$.

THEOREM 3.2. *In a β -safe round-robin bidding dynamic with α -aggressive bid updates the average social welfare in the first T steps satisfies $\frac{1}{T} \sum_{t=1}^T SW(b^t) \geq \frac{\alpha}{(2\alpha+1)\beta} \cdot (1 - \frac{n}{T}) \cdot OPT(v)$.*

This theorem shows that the best-response dynamics described in Proposition 2.3 with $(\alpha, \beta) = (1, \ln m)$, whose point-wise welfare guarantee is only $\Omega(1/\log^2 m)$ by Theorem 3.1, guarantees an average social welfare of $\Omega(1/\log m)$.

In Appendix C we show that the point-wise welfare guarantee of $1/3$ for fractionally subadditive valuations is tight for the respective mechanism. In Section 4 we show that the $\Omega(1/\log m)$ bounds are essentially best possible in a more general sense.

3.1 Proof of Theorem 3.1 The core of our proof of the pointwise welfare guarantee are two lemmata. The first (Lemma 3.2) shows that the declared social welfare after a single round of updates is high when the initial declared welfare is low and the second (Lemma 3.3) shows that the declared welfare after a single round

of updates is high when the initial declared welfare is high. To prove these lemmata we need the following auxiliary lemma.

LEMMA 3.1. *Consider a sequence b^0, \dots, b^n in which bidder i updates his bid in step i . Denote bidder i 's declared utility in step i by $u_i^D(b^i)$. Then, $\sum_{i=1}^n u_i^D(b^i) \leq DW(b^n)$.*

Proof. Consider an arbitrary bidder i . Bidder i updates his bid in step i . Denote the corresponding bid profiles before and after the update by b^{i-1} and b^i . Suppose bidder i 's update buys him the set of items S' . Then

$$u_i^D(b^i) = \sum_{j \in S'} \left(b_{i,j}^i - \max_{k \neq i} b_{k,j}^i \right).$$

For $i > 0$, let $z_j^i = \max_{k \leq i} b_{k,j}^i$ for all j . That is, z_j^i is the maximum bid on item j that is placed by one of the bidders $1, \dots, i$, $z_j^0 = 0$ for all j .

The crucial observation is that $\sum_{j \in S'} (b_{i,j}^i - \max_{k \neq i} b_{k,j}^i) \leq \sum_{j \in M} (z_j^i - z_j^{i-1})$. The reason is as follows. For $j \notin S'$, we have $z_j^i \geq z_j^{i-1}$ by definition. For $j \in S'$, $b_{i,j}^i = z_j^i$ and $\max_{k \neq i} b_{k,j}^i \geq \max_{k < i} b_{k,j}^i = \max_{k < i} b_{k,j}^{i-1} = z_j^{i-1}$.

Summing over all players i we obtain

$$\sum_{i \in N} u_i^D(b^i) \leq \sum_{i \in N} \sum_{j \in M} (z_j^i - z_j^{i-1}).$$

The double sum is telescoping and $z_j^n = \max_k b_{k,j}^n$ and $z_j^0 = 0$ by definition. So,

$$\sum_{i \in N} u_i^D(b^i) \leq \sum_{j \in M} (z_j^n - z_j^0) = \sum_{j \in M} \max_k b_{k,j}^n = DW(b^n),$$

which proves the claim. \square

With the help of this lemma we can now prove our key lemmata.

LEMMA 3.2. *Let S_1^*, \dots, S_n^* be any feasible allocation, in which player i receives items S_i^* . Consider a sequence b^0, \dots, b^n in which bidder i updates his bid in step i using an α -aggressive bid. We have $(\alpha + 1) \cdot DW(b^n) + \alpha \cdot DW(b^0) \geq \alpha \cdot \sum_{i \in N} v_i(S_i^*)$.*

Proof. Consider player i 's action in time step i . Instead of choosing bid b_i^i , he could have bought the set of items S_i^* . As b_i^i is α -aggressive, we get

$$u_i^D(b^i) \geq \alpha \cdot \left(v_i(S_i^*) - \sum_{j \in S_i^*} \max_{k \neq i} b_{k,j}^i \right).$$

Define $p_j^t = \max_i b_{i,j}^t$ for all items j . That is, p_j^t is the maximum bid that is placed on item j in bid profile b^t . We claim that for every $j \in S_i^*$, $\max_{k \neq i} b_{k,j}^i \leq p_j^n + p_j^0$. This is correct because if $b_{k,j}^i$ attains its maximum for $k < i$ then $\max_{k \neq i} b_{k,j}^i \leq p_j^n$ as k 's bid on item j will not change anymore. In the other case, if $k > i$, then $\max_{k \neq i} b_{k,j}^i \leq p_j^0$ because k has not yet changed the bid on item j . Using that both p_j^0 and p_j^n are never negative, the bound follows.

We thus have

$$u_i^D(b^i) + \alpha \cdot \sum_{j \in S_i^*} (p_j^n + p_j^0) \geq \alpha \cdot v_i(S_i^*) .$$

Summing this inequality over all bidders $i \in N$ yields

$$\sum_{i=1}^n u_i^D(b^i) + \alpha \cdot \sum_{i=1}^n \sum_{j \in S_i^*} (p_j^n + p_j^0) \geq \alpha \cdot \sum_{i=1}^n v_i(S_i^*) .$$

We can upper bound the first sum by $DW(b^n)$ using Lemma 3.1. The double sum adds up every $j \in M$ exactly once and we have $\sum_{j \in M} p_j^n = DW(b^n)$ and $\sum_{j \in M} p_j^0 = DW(b^0)$. We obtain

$$(\alpha + 1) \cdot DW(b^n) + \alpha \cdot DW(b^0) \geq \alpha \cdot \sum_{i=1}^n v_i(S_i^*) . \quad \square$$

LEMMA 3.3. *Consider a β -safe bid sequence b^0, \dots, b^n in which player i changes his bid from b^{i-1} to b^i using an α -aggressive bid. Then, $DW(b^n) \geq \frac{\alpha}{\beta} \cdot DW(b^0)$.*

Proof. Consider an arbitrary bidder i and his update from b^{i-1} to b^i . Denote the set of items that bidder i won under bids b^{i-1} by S_i^{i-1} , and the set of items that he wins under bids b^i by S_i^i . So

$$u_i^D(b^{i-1}) = \sum_{j \in S_i^{i-1}} b_{i,j}^{i-1} - \sum_{j \in S_i^{i-1}} \max_{k \neq i} b_{k,j}^{i-1} \quad \text{and,}$$

$$u_i^D(b^i) = \sum_{j \in S_i^i} b_{i,j}^i - \sum_{j \in S_i^i} \max_{k \neq i} b_{k,j}^i .$$

Using that for all $k \neq i$ and all j we have $b_{k,j}^{i-1} = b_{k,j}^i$ we obtain that the difference in declared welfare over all bidders between steps $i-1$ and i is equal to the difference in bidder i 's declared utility at these time steps. Formally,

$$DW(b^i) = \sum_{j \in M \setminus S_i^i} \max_{k \neq i} b_{k,j}^{i-1} + \sum_{j \in S_i^i} b_{i,j}^i$$

$$= \sum_{j \in M} \max_{k \neq i} b_{k,j}^{i-1} + \sum_{j \in S_i^i} b_{i,j}^i - \sum_{j \in S_i^i} \max_{k \neq i} b_{k,j}^i$$

$$= \sum_{j \in M} \max_{k \neq i} b_{k,j}^{i-1} + u_i^D(b_i)$$

$$= \sum_{j \in M \setminus S_i^{i-1}} \max_{k \neq i} b_{k,j}^{i-1}$$

$$+ \sum_{j \in S_i^{i-1}} \max_{k \neq i} b_{k,j}^{i-1} + u_i^D(b_i)$$

$$= \sum_{j \in M \setminus S_i^{i-1}} \max_{k \neq i} b_{k,j}^{i-1} + \sum_{j \in S_i^{i-1}} b_{i,j}^{i-1} + u_i^D(b_i)$$

$$- \sum_{j \in S_i^{i-1}} b_{i,j}^{i-1} + \sum_{j \in S_i^{i-1}} \max_{k \neq i} b_{k,j}^{i-1}$$

$$= DW(b^{i-1}) + u_i^D(b_i) - u_i^D(b^{i-1}) .$$

We now extend this identity to a lower bound on $DW(b^i)$. Since b_i^i is α -aggressive, we have $u_i^D(b^i) \geq \alpha \cdot u_i(b^{i-1})$. Since the bidding sequence is β -safe, $u_i^D(b^t) \leq \beta \cdot u_i(b^t)$ for all t . So,

$$DW(b^i) = DW(b^{i-1}) + u_i^D(b^i) - u_i^D(b^{i-1})$$

$$\geq DW(b^{i-1}) + u_i^D(b^i) - \beta \cdot u_i(b^{i-1})$$

$$\geq DW(b^{i-1}) + u_i^D(b^i) - \frac{\beta}{\alpha} \cdot u_i^D(b^i)$$

$$= DW(b^{i-1}) - \left(\frac{\beta}{\alpha} - 1 \right) \cdot u_i^D(b^i) .$$

Summing this inequality over all bidders $i \in N$ and using the telescoping sum $\sum_{i \in N} (DW(b^i) - DW(b^{i-1})) = DW(b^n) - DW(b^0)$ we obtain

$$DW(b^n) \geq DW(b^0) - \left(\frac{\beta}{\alpha} - 1 \right) \sum_{i \in N} u_i^D(b^i) .$$

Since $\alpha \leq 1$ and $\beta \geq 1$ the factor $(\beta/\alpha - 1) \geq 0$. We can therefore use Lemma 3.1 to conclude that

$$DW(b^n) \geq DW(b^0) - \left(\frac{\beta}{\alpha} - 1 \right) DW(b^n) .$$

This implies the claim. \square

We will use our key lemmata to show a lower bound on the declared welfare. To relate the declared welfare to the social welfare we will use the following lemma.

LEMMA 3.4. *In a β -safe sequence of bid profiles b^0, b^1, \dots for every $t \geq 0$, $DW(b^t) \leq \beta \cdot SW(b^t)$.*

Proof. Consider an arbitrary time step t . Since the bid profile b^t is β -safe we know that for the allocation

T_1, \dots, T_n that corresponds to b^t ,

$$\begin{aligned} \sum_i u_i^D(b^t) &= \sum_i \left(\sum_{j \in T_i} \left(b_{i,j}^t - \max_{k \neq i} b_{k,j}^t \right) \right) \\ &\leq \beta \cdot \sum_i u_i(b) \\ &= \beta \cdot \sum_i \left(v_i(T_i) - \sum_{j \in T_i} \max_{k \neq i} b_{k,j}^t \right). \end{aligned}$$

Rearranging this and using that $\beta \geq 1$ we obtain

$$\begin{aligned} DW(b^t) &= \sum_i \sum_{j \in T_i} b_{i,j}^t \\ &\leq \beta \cdot SW(b^t) - (\beta - 1) \sum_i \sum_{j \in T_i} \max_{k \neq i} b_{k,j}^t \\ &\leq \beta \cdot SW(b^t), \end{aligned}$$

and the claim follows. \square

We are now ready to prove the theorem.

Proof of Theorem 3.1. To prove the guarantee for time step $t \geq n$ consider the bid sequence of length $n + 1$ from b^{t-n} to b^t . At time steps $t - n + 1$ to t each bidder updates his bid exactly once. By the virtue of being a subsequence of a β -safe bidding sequence the sequence b^{t-n}, \dots, b^t is β -safe. Moreover each bid update is α -aggressive.

Applying first Lemma 3.3 and then Lemma 3.2 with b^t taking the role of b^n , b^{t-n} taking the role of b^0 , and setting S_1^*, \dots, S_n^* to the allocation that maximizes welfare we obtain

$$\begin{aligned} (1 + \alpha + \beta) \cdot DW(b^t) &= (\alpha + 1) \cdot DW(b^t) + \alpha \cdot \frac{\beta}{\alpha} DW(b^t) \\ &\geq (\alpha + 1) \cdot DW(b^t) + \alpha \cdot DW(b^{t-n}) \\ &\geq \alpha \cdot OPT(v). \end{aligned}$$

Now, by Lemma 3.4, $DW(b^t) \leq \beta \cdot SW(b^t)$. Combining this with the previous inequality yields

$$(1 + \alpha + \beta) \cdot \beta \cdot SW(b^t) \geq \alpha \cdot OPT(v),$$

as claimed. \square

3.2 Proof of Theorem 3.2 With the proof of the pointwise welfare guarantee at hand we have already done the bulk of the work for proving our guarantee regarding the average welfare. The basic idea is to sum the lower bound on the declared welfare at any given time step as provided by Lemma 3.2 over all time steps

to obtain a lower bound on the average declared welfare, and to turn this into a lower bound on the actual social welfare using Lemma 3.4.

Proof of Theorem 3.2. We first use Lemma 3.2 to relate the declared welfare at time steps t and $t - n$ to the optimal social welfare. Namely, for all $t \geq n$,

$$(\alpha + 1) \cdot DW(b^t) + \alpha \cdot DW(b^{t-n}) \geq \alpha \cdot OPT(v).$$

Next we take the sum over all time steps t and use that $DW(b^t) \geq 0$ to obtain the following lower bound on the average declared welfare

$$\begin{aligned} \frac{1}{T} \cdot \sum_{t=1}^T DW(b^t) &\geq \frac{1}{T} \cdot \sum_{t=n+1}^T DW(b^t) \\ &\geq \frac{\alpha}{\alpha + 1} \cdot \frac{1}{T} \cdot \sum_{t=n+1}^T \left(OPT(v) - DW(b^{t-n}) \right) \\ &\geq \frac{\alpha}{\alpha + 1} \cdot \frac{T - n}{T} \cdot OPT(v) - \frac{\alpha}{\alpha + 1} \cdot \frac{1}{T} \cdot \sum_{t=1}^T DW(b^t). \end{aligned}$$

Solving this inequality for $\frac{1}{T} \cdot \sum_{t=1}^T DW(b^t)$ and using Lemma 3.4 to lower bound $SW(b^t)$ by $1/\beta \cdot DW(b^t)$ we obtain

$$\begin{aligned} \frac{1}{T} \cdot \sum_{t=1}^T SW(b^t) &\geq \frac{1}{\beta} \cdot \frac{1}{T} \cdot \sum_{t=1}^T DW(b^t) \\ &\geq \frac{\alpha}{(2\alpha + 1)\beta} \cdot \frac{T - n}{T} \cdot OPT(v), \end{aligned}$$

which proves the claim. \square

4 Lower Bound for Subadditive CAs

Next we show our second main result (Theorem 4.1), which shows that no best-response dynamics in which bidders do not overbid on the grand bundle can achieve a point-wise welfare guarantee that is significantly better than $1/\log m$. The assumption that bidders do not overbid on the grand bundle seems quite natural, and does allow overbidding on subsets of items. It is satisfied by all dynamics that we have described in Section 2 and more generally by all dynamics that have been proposed in the literature.

THEOREM 4.1. *For every positive integer $k \in \mathbb{N}_{>0}$ there exists an instance with $n = 2$ players, $m = 2^k - 1$ items, and subadditive valuations $v = (v_1, v_2)$ such that in every best-response dynamics in which players do not*

overbid on the grand bundle there exist infinitely many time steps t at which

$$SW(b^t) \leq \frac{1}{\Omega\left(\frac{\log m}{\log \log m}\right)} \cdot OPT(v).$$

To prove this theorem we show that whenever the second player has updated its bid social welfare will be low. This does not imply that the average welfare will be low as well. However, if we restrict attention to round-robin dynamics, then we can extend the construction by adding additional players after the second player that play a low-stakes game on separate items forcing the average welfare to be low as well.

4.1 Proof of Theorem 4.1 Our proof of the lower bound is built around the following family of hard instances, with $n = 2$ players and $m = 2^k - 1$ items. The valuations of the first player are based on an example that demonstrates the worst-case integrality gap for set cover linear programs (see, e.g., [27, Example 13.4]), which has been used in the context of CAs with item bidding before [2]. The crux of our construction is in the design of the second player's valuation function, and its interplay with the valuation function of the first player.

DEFINITION 4.1. *For every positive integer $k \in \mathbb{N}_{>0}$ the hard instance \mathcal{I}_k consists of $n = 2$ bidders and $m = 2^k - 1$ items and the following subadditive valuations:*

1. *First bidder: Number the items from 1 to m and let \mathbf{i} be a k -bit binary vector representing the integer i . Interpret \mathbf{i} as a k -dimensional vector over \mathbb{F}_2 . Write $\mathbf{i} \cdot \mathbf{j}$ as the dot product of the two vectors. Let $S_i = \{j \mid \mathbf{j} \cdot \mathbf{i} = 1\}$. Note that each such set contains $(m+1)/2$ items, and each item is contained in $(m+1)/2$ such sets. For each set of items $T \subseteq M$ let $v_1(T)$ be the minimum number of sets S_i required to cover the items in T .*
2. *Second bidder: Set $\rho = 4\frac{k}{m}$ and $d = k - \log k$. Let \mathcal{D} denote the set of all d -dimensional subspaces of \mathbb{F}_2^k excluding the zero vector. Then for any set of items T let*

$$v_2(T) = \rho \cdot \max_{D \in \mathcal{D}} w_D(T) \quad , \quad \text{where}$$

$$w_D(T) = \begin{cases} |T \cap D| & \text{for } |T \cap D| < \frac{|D|}{2} \\ \frac{|D|}{2} & \text{for } \frac{|D|}{2} < |T \cap D| < |D| \\ |D| & \text{else} \end{cases} .$$

Note that, in the instances just described, the first player has a valuation of $v_1(M) \geq k = \log_2(m+1)$ for the grand bundle, while the second player has a

maximum valuation of $\max_T v_2(T) = \rho \cdot |D| = \rho \cdot (2^d - 1) \leq \rho \cdot 2^d = 4$ for any set of items.

To prove the theorem we first use linear algebra to derive a symmetry property of \mathcal{D} , which together with weak no-overbidding of the first player on the grand bundle implies the existence of a subset of items $D \in \mathcal{D}$ with low prices (Lemma 4.1). Intuitively, this is because the sets of items that the second player is interested in are rather small (of size about $m/\log m$), and there are sufficiently many of these sets. We then show that every demand set of the second player under these prices includes some set of items $D' \in \mathcal{D}$ (Lemma 4.2). In the final step (Lemma 4.3) we show that if the second player buys any such set D' , then the first player's valuation for the remaining items $M \setminus D'$ and hence the overall social welfare is at most $O(\log \log m)$.

LEMMA 4.1. *Let $k \in \mathbb{N}_{>0}$. Consider the hard instance \mathcal{I}_k . For every vector of bids b such that the first player does not overbid on the grand bundle there is a d -dimensional subspace $D \in \mathcal{D}$ such that $\sum_{j \in D} b_{1,j} < \rho \cdot \frac{|D|}{2}$.*

Proof. Since the first player does not overbid on the grand bundle we have $\sum_{j \in M} b_{1,j} \leq v_1(M) = k$, so the average bids are bounded by $\frac{1}{m} \sum_{j \in M} b_{1,j} \leq \frac{k}{m}$.

Observe that the number of d -dimensional subspaces of \mathbb{F}_2^k that contain a vector $0 \neq x \in \mathbb{F}_2^k$ is given as $\binom{k-1}{d-1}_2$, where $\binom{\cdot}{\cdot}_q$ refers to the q -binomial coefficient (see, e.g., [22]). So, in particular, this number is independent of x . Therefore, instead of taking the average over all items M , we can take the average over all sets $D \in \mathcal{D}$ and take the average within such a set, i.e., $\frac{1}{m} \sum_{j \in M} b_{1,j} = \frac{1}{|\mathcal{D}|} \sum_{D \in \mathcal{D}} \frac{1}{|D|} \sum_{j \in D} b_{1,j}$.

In combination, there has to be a D such that $\frac{1}{|D|} \sum_{j \in D} b_{1,j} \leq \frac{1}{m} \sum_{j \in M} b_{1,j} \leq \frac{k}{m}$. Since $\frac{k}{m} < \frac{\rho}{2} = 2\frac{k}{m}$ the claim follows. \square

LEMMA 4.2. *Let $k \in \mathbb{N}_{>0}$. Consider the hard instance \mathcal{I}_k . If the prices p as seen by the second player are such that $\sum_{j \in D} p_j < \rho \cdot |D|/2$ for some $D \in \mathcal{D}$, then each demand set of the second player under these prices includes some $D' \in \mathcal{D}$.*

Proof. By our assumption on the sum of the prices of the items in D , $u(D) = v_2(D) - \sum_{j \in D} p_j = \rho \cdot w_D(D) - \sum_{j \in D} p_j > \rho \cdot \frac{|D|}{2}$. Now, let $S \subseteq M$ be a demand set under v_2 . If $|S \cap D'| < |D'|$ for all $D' \in \mathcal{D}$, then we have $u(S) = v_2(S) - \sum_{j \in S} p_j \leq v_2(S) = \rho \cdot \max_{D' \in \mathcal{D}} w_{D'}(S) \leq \rho \cdot \max_{D' \in \mathcal{D}} \frac{|D'|}{2} < u(D)$. This means, S can only be a demand set if $|S \cap D'| = |D'|$ for some $D' \in \mathcal{D}$. \square

LEMMA 4.3. *Let $k \in \mathbb{N}_{>0}$. Consider the hard instance \mathcal{I}_k . Then for $D' \in \mathcal{D}$ we have $v_1(M \setminus D') \leq k - d$.*

Proof. To show the bound on v_1 , we use that $D' \cup \{0\}$ is a subspace of \mathbb{F}_2^k of dimension d . That is, any basis x_1, \dots, x_d of $D' \cup \{0\}$ can be extended by x_{d+1}, \dots, x_k to a basis of \mathbb{F}_2^k . Let $X = (x_1, \dots, x_k)$. This way, X^{-1} is the matrix that expresses $j \in \mathbb{F}_2^k$ as a linear combination of x_1, \dots, x_k . As x_1, \dots, x_d is a basis of $D' \cup \{0\}$, we know that for every $j \notin D' \cup \{0\}$ the vector $X^{-1}j$ cannot be zero in all components $d + 1, \dots, k$. This implies that the set $M \setminus D'$ can be covered by sets S_i for i being the rows $d + 1, \dots, k$ of X^{-1} . Therefore $v_1(M \setminus D') \leq k - d$. \square

Proof of Theorem 4.1. Any best-response dynamics has to ask every bidder infinitely often. We claim that the social welfare is $O(\log \log m)$ right after each update of the second player. Since the optimal social welfare is $\Omega(\log m)$ this shows the claim.

Let b^t be a bid vector after the second player has made a move. Using Lemma 4.1, we know that there is a set $D \in \mathcal{D}$ with $\sum_{j \in D} b_{1,j}^{t-1} < \rho \cdot \frac{|D|}{2}$. By Lemma 4.2, the second player then buys a superset of some $D' \in \mathcal{D}$. Therefore, right after the second player has updated his bid the first player is allocated a subset of the items $M \setminus D'$. Lemma 4.3 implies that the social welfare for this allocation is no higher than $k - d + \rho 2^d = O(\log \log m)$. \square

5 Beyond Round-Robin Activation

Our positive results make use of the fact that bidders are activated to update their bid in round-robin fashion. That is, between two activations of a bidder, each other bidder is activated exactly once. In this section, we investigate alternative activation protocols. Our first theorem shows that our positive results extend to the case where at each step a random player gets to update his bid.

THEOREM 5.1. *Consider a β -safe sequence of bids that is generated by choosing at each time step a player uniformly at random and letting this player update his bid to an α -aggressive bid. Then for any time step $T \geq n$, $\mathbf{E}[SW(b^T)] \geq \frac{\alpha}{2(1+4\alpha)\beta} \cdot OPT(v)$.*

Our second theorem shows that if an adversary chooses the order in which the players get to update their bids, then it may take exponentially long before states of high welfare are reached.

THEOREM 5.2. *For all n and k , there is an instance of n agents and $(n - 1) \cdot (k + 1)$ items and an activation sequence, such that until each agent has been activated*

$\Omega(2^k)$ times the welfare has never exceeded a $\frac{1}{n-1}$ fraction of the optimum.

We defer the proofs of both these results to the full version, where we also present some additional results on adversarial activation.

6 Concluding Remarks and Outlook

In our analysis we focused on fractionally subadditive and subadditive valuations, which do not exhibit complements. A natural question is whether similar results can be obtained for classes of valuations that exhibit complements. In the full version, we discuss an example with MPH- k valuations that highlights the difficulties that arise. Another interesting follow-up question is whether there is a general result that translates a Price of Anarchy guarantee for a given mechanism that is provable via smoothness into a result that shows that best-response sequences reach states of good social welfare quickly. The example with MPH- k valuations already limits the potential scope of such a result. It would still be interesting to identify natural sufficient conditions. One such condition could be that the mechanism admits some kind of potential function (as the procedure for XOS valuations), but our results already show that this condition is certainly not necessary.

References

- [1] H. Ackermann, H. Röglin, and B. Vöcking. On the impact of combinatorial structure on congestion games. *Journal of the ACM*, 55(6), 2008.
- [2] K. Bhawalkar and T. Roughgarden. Welfare guarantees for combinatorial auctions with item bidding. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 700–709, 2011.
- [3] V. Bilò, A. Fanelli, M. Flammini, and L. Moscardelli. Performance of one-round walks in linear congestion games. *Theory of Computing Systems*, 49(1):24–45, 2011.
- [4] M. Braverman, J. Mao, and S. M. Weinberg. Interpolating between truthful and non-truthful mechanisms for combinatorial auctions. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1444–1457, 2016.
- [5] Y. Cai and C. H. Papadimitriou. Simultaneous bayesian auctions and computational complexity. In *Proceedings of the 15th ACM Conference on Economics and Computation*, pages 895–910, 2014.
- [6] G. Christodoulou, V. S. Mirrokni, and A. Sidiropoulos. Convergence and approximation in potential games. *Theoretical Computer Science*, 438:13–27, 2012.

- [7] G. Christodoulou, A. Kovács, and M. Schapira. Bayesian combinatorial auctions. *Journal of the ACM*, 63(2):11, 2016.
- [8] C. Daskalakis and V. Syrgkanis. Learning in auctions: Regret is hard, envy is easy. *Proceedings of the 57th IEEE Symposium on Foundations of Computer Science*, pages 219–228, 2016.
- [9] N. R. Devanur, J. Morgenstern, V. Syrgkanis, and S. M. Weinberg. Simple auctions with simple strategies. In *Proceedings of the 16th ACM Conference on Economics and Computation*, pages 305–322, 2015.
- [10] S. Dobzinski. Breaking the logarithmic barrier for truthful combinatorial auctions with submodular bidders. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing*, pages 940–948, 2016.
- [11] S. Dobzinski and M. Schapira. An improved approximation algorithm for combinatorial auctions with submodular bidders. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1064–1073, 2006.
- [12] S. Dobzinski and J. Vondrák. Impossibility results for truthful combinatorial auctions with submodular valuations. *Journal of the ACM*, 63(1):5, 2016.
- [13] S. Dobzinski, H. Fu, and R. D. Kleinberg. On the complexity of computing an equilibrium in combinatorial auctions. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 110–122, 2015.
- [14] P. Dütting, M. Henzinger, and M. Starnberger. Valuation compressions in VCG-based combinatorial auctions. In *Proceedings of the 9th Conference on Web and Internet Economics*, pages 146–159, 2013.
- [15] B. Farzad, N. Olver, and A. Vetta. A priority-based model of routing. *Chicago Journal of Theoretical Computer Science*, 2008, 2008.
- [16] U. Feige. On maximizing welfare when utility functions are subadditive. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, pages 41–50, 2006.
- [17] U. Feige, M. Feldman, N. Immorlica, R. Izsak, B. Lucier, and V. Syrgkanis. A unifying hierarchy of valuations with complements and substitutes. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, pages 872–878, 2015.
- [18] M. Feldman, H. Fu, N. Gravin, and B. Lucier. Simultaneous auctions are (almost) efficient. In *Proceedings of the 45th ACM Symposium on Theory of Computing Conference*, pages 201–210, 2013.
- [19] M. X. Goemans, V. S. Mirrokni, and A. Vetta. Sink equilibria and convergence. In *Proceedings of the 46th IEEE Symposium on Foundations of Computer Science*, pages 142–154, 2005.
- [20] A. Hassidim, H. Kaplan, Y. Mansour, and N. Nisan. Non-price equilibria in markets of discrete goods. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, pages 295–296, 2011.
- [21] D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.
- [22] A. Prasad. Counting subspaces of a finite vector space — 1. *Resonance*, 15(11):977–987, 2010.
- [23] R. W. Rosenthal. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.
- [24] T. Roughgarden. Barriers to near-optimal equilibria. In *55th IEEE Annual Symposium on Foundations of Computer Science*, pages 71–80, 2014.
- [25] T. Roughgarden. Intrinsic robustness of the price of anarchy. *Journal of the ACM*, 62(5):32, 2015.
- [26] V. Syrgkanis and É. Tardos. Composable and efficient mechanisms. In *Proceedings of the 45th ACM Symposium on Theory of Computing Conference*, pages 211–220, 2013.
- [27] V. V. Vazirani. *Approximation Algorithms*. Springer-Verlag New York, Inc., New York, NY, USA, 2001. ISBN 3-540-65367-8.

A Non-Existence of Weak No-Overbidding Pure Nash Equilibria

We can also leverage our novel insights regarding hard instances (Definition 4.1) for subadditive combinatorial auctions with item bidding to show that there is no pure Nash equilibrium in which players use weakly no-overbidding strategies.

THEOREM A.1. *Let $k \in \mathbb{N}_{>0}$. Consider the hard instance \mathcal{I}_k with $n = 2$ players and $m = 2^k - 1$ items. There is no pure Nash equilibrium in weakly no-overbidding strategies if $k \geq 8$. This remains true if we define a bid profile to be at equilibrium if no player has a beneficial deviation to a weakly no-overbidding strategy.*

Proof. Assume that b is a weakly no-overbidding pure Nash equilibrium. Suppose the second player wins the set of items $W \subseteq M$ in b , then the first player wins the set of items $M \setminus W$. By weak no-overbidding, we have

$$\sum_{j \in M \setminus W} b_{1,j} \leq v_1(M \setminus W) \quad \text{and} \quad \sum_{j \in W} b_{2,j} \leq v_2(W) .$$

The first player does not win the items in W , which means that $b_{1,j} \leq b_{2,j}$ for all items $j \in W$.

Consequently, we have

$$\begin{aligned}
\sum_{j \in M} b_{1,j} &\leq v_1(M \setminus W) + v_2(W) \\
&\leq v_1(M) + v_2(M) \\
&= k + \rho \cdot 2^d \\
&= k + 4 \cdot \frac{k}{m} \cdot 2^{k - \log k} \\
&= k + 4 .
\end{aligned}$$

By the same argument as in Lemma 4.1, each item $j \in M$ is included in the same number of sets $D \in \mathcal{D}$. Therefore,

$$\frac{1}{|\mathcal{D}|} \sum_{D \in \mathcal{D}} \frac{1}{|D|} \sum_{j \in D} b_{1,j} = \frac{1}{m} \sum_{j \in M} b_{1,j} \leq \frac{k+4}{m} .$$

This implies that there is a set $D \in \mathcal{D}$ such that

$$\frac{1}{|D|} \sum_{j \in D} b_{1,j} \leq \frac{k+4}{m} .$$

Since $k \geq 8$ by assumption, $m > 2k + 8$, and therefore

$$\sum_{j \in D} b_{1,j} \leq \frac{k+4}{m} \cdot |D| < \frac{|D|}{2} .$$

By Lemma 4.2 and because the second player plays a best response, we have $W \supseteq D'$ for some $D' \in \mathcal{D}$.

In the remainder, we will show that this implies that the first player has a beneficial weakly no-overbidding deviation b'_1 .

Let $b'_{1,j} = b_{2,j} + \frac{1}{m}$ for $j \in W$ and $b'_{1,j} = b_{1,j}$ for $j \in M \setminus W$. Observe that in (b'_1, b_2) the first player wins all items M . This bid fulfills the weak no-overbidding property because

$$\begin{aligned}
\sum_{j \in M} b'_{1,j} &= \sum_{j \in W} \left(b_{2,j} + \frac{1}{m} \right) + \sum_{j \in M \setminus W} b_{1,j} \\
&\leq v_2(W) + 1 + v_1(M \setminus W) \\
&\leq v_2(D') + 1 + v_1(M \setminus D') \\
&\leq \rho 2^d + 1 + k - d \\
&= 4 + 1 + \log k \\
&\leq k \\
&= v_1(M) ,
\end{aligned}$$

where the first inequality uses that b is weakly no-overbidding, the second inequality exploits the definition of v_2 , the third inequality holds by Lemma 4.3, and the final inequality holds because we have assumed $k \geq 8$.

The deviation by the first player is beneficial because

$$\begin{aligned}
u_1(b'_1, b_2) &= v_1(M) - \sum_{j \in M} b_{2,j} \\
&= k - d - \sum_{j \in M \setminus W} b_{2,j} + d - \sum_{j \in W} b_{2,j} \\
&\geq u_1(b) + d - v_2(W) \\
&\geq u_1(b) + d - 4 > u_1(b) ,
\end{aligned}$$

where the first inequality uses Lemma 4.3, the second inequality uses that $v_2(W) \leq v_2(D') = 4$, and the final inequality follows from the definition of $d = k - \log k$ and the assumption that $k \geq 8$ and so $d > 4$. \square

B Missing Proofs from Section 2

B.1 Sufficiency of Strong No-Overbidding

We show that in order to have a 1-safe dynamic it suffices that initial bids and the subsequent updates fulfill no-overbidding in the strong sense. A bid vector b is strongly no-overbidding if $\sum_{j \in S} b_{i,j} \leq v_i(S)$ for every bidder i and every set of items S . A best response b_i by bidder i against bids b_{-i} is strongly no overbidding if $\sum_{j \in S} b_{i,j} \leq v_i(S)$.

LEMMA B.1. *If the initial bid vector b^0 is strongly no overbidding and at each time step $t \geq 1$ some bidder i gets to update his bid to a best response, which is strongly no overbidding, then the resulting best-response dynamic is 1-safe.*

Proof. Since the initial bid vector and each update satisfy strong no-overbidding we have $\sum_{j \in S} b_{i,j}^t \leq v_i(S)$ for all bidders i , time steps $t \geq 0$, and sets of items S . Subtracting $\sum_{j \in S} \max_{k \neq i} b_{k,j}^t$ from both sides shows the claim. \square

B.2 Proof of Proposition 2.1

Consider an arbitrary bidder i and his update to bid b_i^t . The bid b_i^t satisfies strong no-overbidding by definition. Hence Lemma B.1 shows that the bid sequence is 1-safe. It remains to show that b_i^t is a 1-aggressive best response.

We first show that the bid b_i^t is a best response to b_{-i}^t . Let S_i denote the set of items that bidder i wins with bid b_i^t against bids b_{-i}^t and let D be the demand set on the basis of which b_i^t is defined. Then,

$$\begin{aligned}
u_i(b^t) &= v_i(S_i) - \sum_{j \in S_i} \max_{k \neq i} b_{k,j}^t \\
&\geq \sum_{j \in S_i} (b_{i,j}^t - \max_{k \neq i} b_{k,j}^t) \\
&\geq \sum_{j \in D} (b_{i,j}^t - \max_{k \neq i} b_{k,j}^t)
\end{aligned}$$

$$\begin{aligned}
&= v_i(D) - \sum_{j \in D} \max_{k \neq i} b_{k,j}^t \\
&\geq \max_S \left(v_i(S) - \sum_{j \in S} \max_{k \neq i} b_{k,j}^t \right),
\end{aligned}$$

where the first inequality uses that v_i is XOS, the second uses that $\max_{k \neq i} b_{k,j}^t = b_{i,j}^t$ for $j \in D \setminus S_i$ and $\max_{k \neq i} b_{k,j}^t \leq b_{i,j}^t$ for $j \in S_i \setminus D$, the following equality exploits the definition of b_i^t , and the final inequality uses that D is a demand set.

To show that b_i^t is 1-aggressive it suffices to show that bidder i 's declared and actual utility at time step t coincide. Since the right-hand side in the preceding chain of inequalities is at least $v_i(S_i) - \sum_{j \in S_i} \max_{k \neq i} b_{k,j}^t$ all inequalities in the chain of inequalities must be equalities. This implies that

$$\begin{aligned}
u_i(b^t) &= v_i(S_i) - \sum_{j \in S_i} \max_{k \neq i} b_{k,j}^t \\
&= \sum_{j \in S_i} (b_{i,j}^t - \max_{k \neq i} b_{k,j}^t) \\
&= u_i^D(b^t),
\end{aligned}$$

as claimed.

B.3 Proof of Proposition 2.2 Consider an arbitrary bidder i and his update to bid b_i^t . We first argue that b_i^t is a best response. We claim that $\tilde{u}_i(S, b_{-i}^t) > 0$ for all $S \subseteq D$ unless $D = \emptyset$. To see this assume by contradiction that there exist a $S \subseteq D$ such that $\tilde{u}_i(S, b_{-i}^t) \leq 0$. Then, by subadditivity of v_i ,

$$\begin{aligned}
\tilde{u}_i(D, b_{-i}^t) &\leq \left(v_i(D \setminus T) - \sum_{j \in D \setminus T} \max_{k \neq i} b_{k,j}^t \right) + \\
&\quad \left(v_i(T) - \sum_{j \in S} \max_{k \neq i} b_{k,j}^t \right) \\
&\leq \tilde{u}_i(D \setminus T, b_{-i}^t),
\end{aligned}$$

which contradicts the definition of D . Because of this the additive approximation a_i has $a_{i,j} > 0$ for all $j \in D$. It follows that $b_{i,j}^t > \max_{k \neq i} b_{k,j}^t$ for all $j \in D$, and so bidder i wins all items $j \in D$, and for the items $j \notin D$ that he wins $\max_{k \neq i} b_{k,j}^t = 0$.

To see that b_i^t is $1/\ln m$ -aggressive observe the following. Let S_i denote the set of items that bidder i wins with bid b_i^t . Then, considering the bid b_i^t defined

on the basis of demand set D , we have

$$\begin{aligned}
u_i^D(b^t) &= \sum_{j \in S_i} \left(b_{i,j}^t - \max_{k \neq i} b_{k,j}^t \right) \\
&\geq \sum_{j \in D} \left(b_{i,j}^t - \max_{k \neq i} b_{k,j}^t \right) \\
&= \sum_{j \in D} a_{i,j} \geq \frac{1}{\ln m} \cdot \tilde{u}_i(D, b_{-i}^t),
\end{aligned}$$

where the first inequality uses that $b_{i,j}^t = \max_{k \neq i} b_{k,j}^t$ for $j \in D \setminus S_i$ and $b_{i,j}^t \geq \max_{k \neq i} b_{k,j}^t$ for $j \in S_i \setminus D$, and the second inequality uses property (a) of bid b_i^t .

That the bid sequence is 1-safe follows from the starting condition and Lemma B.1 by observing that bidder i 's update satisfies strong no-overbidding. Namely, for every $S \subseteq D$,

$$\begin{aligned}
\sum_{j \in S} b_{i,j}^t &= \sum_{j \in S} (a_{i,j} + \max_{k \neq i} b_{k,j}^t) \\
&\leq \tilde{u}_i(S, b_{-i}^t) + \sum_{j \in S} \max_{k \neq i} b_{k,j}^t \\
&= v_i(S),
\end{aligned}$$

where the inequality follows from property (b) of bid b_i^t .

B.4 Proof of Proposition 2.3 The argument that the bid b_i^t chosen by bidder i is a best response and 1-aggressive is identical to the respective argument in the proof of Proposition 2.2, except that this time we collect a factor of 1 instead of $1/\ln m$ when we apply property (a) of bid b_i^t .

To see that the bid sequence is $\ln m$ -safe, consider a point in time $t' \geq t$ after bidder i 's update. In the vector $b_{t'}$, bidder i gets a set $S \subseteq M$ that is possibly different from D . Note that for $j \in S \setminus D$, $b_{i,j}^{t'} = 0$ by our definition. Furthermore, for $j \in S \cap D$, $\max_{k \neq i} b_{k,j}^{t'} \leq \max_{k \neq i} b_{k,j}^t$, because bid updates are only non-zero if an item changes its owner. Therefore, because bidder i wins item j , all new bids have to be zero.

In combination, we have

$$\begin{aligned}
u_i^D(b^{t'}) &= \sum_{j \in S \cap D} \left(\tilde{a}_{i,j} + \max_{k \neq i} b_{k,j}^t - \max_{k \neq i} b_{k,j}^{t'} \right) \\
&\leq \ln m \cdot \left(\tilde{u}_i(S \cap D, b_{-i}^t) + \sum_{j \in S \cap D} (\max_{k \neq i} b_{k,j}^t - \max_{k \neq i} b_{k,j}^{t'}) \right) \\
&= \ln m \cdot u_i(b^{t'}),
\end{aligned}$$

because the sum of $\tilde{a}_{i,j}$ terms is bounded by $\ln m \cdot \tilde{u}_i(S \cap D, b_{-i}^t)$ by definition and the sum of the remaining terms is non-negative.

C Tightness of Point-wise Guarantee for XOS Valuations

The following proposition shows that the point-wise welfare guarantee of $1/3$ for the round-robin best-response dynamics for fractionally subadditive valuations described in Section 2 is tight, even if the valuations are unit demand.

PROPOSITION C.1. *Consider the dynamics described in Section 2.1 There is an input with $n = 3$ players, $m = 3$ items, and unit-demand valuations and an initial bid vector such that when started from this bid vector the social welfare obtained by the dynamics after a single round of bid updates is $1/3 \cdot OPT(v)$.*

Proof. The valuations of all three bidders are unit demand, i.e., for all players i and sets of items S , $v_i(S) = \max_{j \in S} v_{i,j}$. The item valuations $v_{i,j}$ for $1 \leq i, j \leq 3$ are given by the following table:

	item 1	item 2	item 3
player 1	1	0	0
player 2	$1 + \epsilon$	$1 + 2\epsilon$	$1 + 3\epsilon$
player 3	0	0	1

Suppose that the XOS representation of these valuations is that each player has an additive valuation $a^{i,0}$ that is all zero and then one for each item j , $a^{i,j}$, such that $a^{i,j}(k) = v_{i,j}$ for $k = j$ and $a^{i,j}(k) = 0$ otherwise.

Let b^0 be the bid profile in which Player 2 bids $1 + \epsilon$ on item 1, all other bids are 0. That is, $b^0 = (a^{1,0}, a^{2,1}, a^{3,0})$. Suppose that the order of updates is first player 1 gets to update his bid, then player 2, and then player 3.

Player 1 is already playing a best response to b_{-1}^0 , so $b^1 = b^0$. Now, to get b^2 , player 2 updates his bids to a best-response to b_{-2}^1 , which is $a^{2,3}$. That is, he bids zero on the first two items and $1 + 3\epsilon$ on the third. So $b^2 = (a^{1,0}, a^{2,3}, a^{3,0})$. With these bids, however, bidding 0 on all items is a best-response of player 3, therefore $b^3 = b^2$.

Observe that $SW(b^3) = DW(b^3) = 1 + 3\epsilon$, whereas the optimal social welfare is $3 + 2\epsilon$. The claim follows by letting ϵ tend to zero. \square