Best-Response Dynamics in Combinatorial Auctions with Item Bidding☆

Paul Dütting[a], Thomas Kesselheim[b]

[a]Department of Mathematics, London School of Economics, Houghton Street, London WC2A 2AE, UK, p.d.duetting@lse.ac.uk
[b]University of Bonn, Institute of Computer Science, Endenicher Allee 19a, 53115 Bonn, Germany, thomas.kesselheim@uni-bonn.de

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 social welfare can be achieved this way, giving bounds on the welfare at equilibrium. Recently, however, criticism has been raised that equilibria of this game are hard to compute and therefore unlikely to be attained.

In this paper, we take a different perspective by studying simple best-response dynamics. 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 reach and remain in states of high welfare after each agent has updated his bid at least once.

1. Introduction

In a combinatorial auction, n agents compete for the assignment of m items. The agents have private preferences over bundles of items as expressed by a valuation function $v_i\colon 2^{[m]} \to \mathbb{R}_{\geq 0}$. Our goal in this work is to find a partition of the items into sets $S_1, \ldots, S_n$ that maximizes social welfare $\sum_i v_i(S_i)$, based on reported valuations (bids) $b_i\colon 2^{[m]} \to \mathbb{R}_{\geq 0}$ with the freedom to impose payments $p_1, \ldots, p_n$ on the agents.

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 agents to bid truthfully. Insisting on truthfulness has the advantage that for the individual agent 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 agents 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 of this approach 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 agents are activated in a round-robin fashion and agents when activated buy their favorite set of items at the current prices, in a myopic way. Christodoulou et al. [1] showed that one instance of such dynamics converges if agents' valuation functions are fractionally subadditive. However, they also showed that it takes exponential

☆An extended abstract appeared in the Proceedings of the 28th ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain. Part of the work of P. Dütting was done while he was at Google Research in Zürich, Switzerland.

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, the social welfare reaches a good level very fast.

1.1. The Setting

We study combinatorial auctions with n agents N and m items M. Each agent $i \in N$ has a valuation function $v_i \colon 2^M \to \mathbb{R}_{\geq 0}$. Our objective is to find a feasible allocation, i.e., a partition of the items, $S_1, \ldots, S_n$, that maximizes social welfare $\sum_{i \in N} v_i(S_i)$. We assume that an allocation of items to agents is found by distributed strategic behavior of the agents using item bidding, and focus on the original proposal where the price of an item equals the second highest bid on that item. That is, each agent $i \in N$ places a bid $b_{i,j}$ on each item $j \in M$. Each item $j \in M$ is assigned to the agent $i \in N$ 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 agents choose their bids strategically so as to maximize their quasi-linear utilities. agent 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 agent i.

We say that a bid b_i is a best response to the bids b_{-i} if agent 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 agent 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 agent $i \in N$ his bid b_i is a best response against bids b_{-i}.

We study simple game-playing dynamics in which agents 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 agent $i \in N$ is activated and updates his bid b_i^{t-1} from the previous round to a best response to the other agents' bids $b_{-i}^t = b_{-i}^{t-1}$ 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 $S_1, \ldots, S_n$ 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, second-price auction with n agents. Suppose $v_1 = C$ and $v_i = 1$ for $i \geq 2$, where $C \gg 1$. Suppose we start at $b = (0, \ldots, 0)$ and the item assigned to agent 1. A possible best response sequence has agents update their bids in round-robin fashion, each time increasing the winning bid by ϵ.*

For the first $\Omega(1/n\epsilon)$ rounds, 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. After $O(1/n\epsilon)$ rounds all agents but the first will have dropped out, and a social-welfare maximizing state will be reached.

Example 1.2 (Gross Overbidding)**.** *Consider the same setting as in the previous example. If in the first round of updates the last agent bids $C + \epsilon$ this will terminate the dynamics.*

Here, the dynamics converge after a single round of bid updates, but at this point the dynamics are stuck in a highly inefficient state.

These two examples illustrate two extremes of unnatural behavior: In the first example, the first agent grossly understates his value and therefore utility. It would have been absolutely fine and, in fact, more conducive if the first agent had bid more aggressively early on. In the second example, the last agent grossly overstates his value and utility for the item. He turns out to be lucky in this case, but his strategy is rather risky. He might end up paying much more than his value for the item.

We will see that natural dynamics avoid these pitfalls. In these dynamics, like for example the one by Christodoulou et al. [1], the declared utility closely matches the actual utility. We will argue that this enables these dynamics to reach states of high social welfare surprisingly fast. Our welfare bounds will be parameterized by the extent to which the declared utilities can differ from the actual utilities, which also means that they capture a broad range of dynamics, including dynamics in which agents only use approximate best responses.

To formally state our conditions, we need the following definitions. In a combinatorial auction with item bidding, the bids $b_{i,j}$ effectively express additive valuations. The allocation $S_1, \dots, S_n$ maximizes the declared welfare $DW(b) = \sum_i \sum_{j \in S_i} b_{i,j}$, which usually differs from the actual welfare $SW(b) = \sum_i v_i(S)$. The declared utility is given by $u_i^D(b) = \sum_{j \in S_i} b_{i,j} - \sum_{j \in S_i} \max_{k \neq i} b_{k,j}$, whereas the actual utility is given by $u_i(b) = v_i(S_i) - \sum_{j \in S_i} \max_{k \neq i} b_{k,j}$.

Our conditions are:

Definition 1.3 (α-aggressive)**.** Let $\alpha \geq 0$. We call a bid b_i by agent 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 1.4 (β-safe)**.** Let $\beta \geq 1$. A bidding dynamic is *β-safe* if it ensures that $u_i^D(b) \leq \beta \cdot u_i(b)$ for all agents i and reachable bid profiles b.

Definition 1.3 requires a lower bound on the declared utilities. It prevents effects like that in Example 1.1. We will usually apply it when b_i is a best response to b_{-i}. In this case, it means that the declared utility has to be at least an α fraction of the actual utility. However, it also leaves the freedom to consider approximate best responses. Definition 1.4 states an upper bound on the declared utilities. It rules out situations as that in Example 1.2. One way to achieve it is to require strong no overbidding (i.e., that agents do not overbid on any bundle), but we will also see an example of a safe dynamic that allows overbidding. Note that in both cases it is guaranteed that agents 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 agent i and time step t.

1.3. Our Results

In this work we consider combinatorial auctions with item bidding, and show welfare guarantees for bidding dynamics under natural assumptions on the bidding behavior. Our work identifies factors that allow the dynamics to reach states of high social welfare quickly, and factors that prevent the dynamics from reaching these states.

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 dynamic provided that agents choose bids that are aggressive enough but not too aggressive. This, in particular, includes dynamics in which agents 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. *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 agent 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.

It is rather straightforward to verify that the best-response dynamic of Christodoulou et al. [1] for fractionally subadditive valuations has $(\alpha, \beta) = (1, 1)$. So a first implication of our first main result is that this dynamic, which may take exponentially long to reach an equilibrium, at which point it guarantees at least $1/2$ of the optimal social welfare, actually reaches $1/3$ of the optimal social welfare after a single round of bid updates, and at any point after that.

A second implication of our first main result concerns subadditive valuations, where, as already mentioned, pure Nash equilibria may not exist and hence best-response dynamics may not converge. For this class of valuations, we show that there is a best-response dynamic with $(\alpha, \beta) = (1/\ln m, 1)$. In fact, this is a more or less immediate consequence of arguments previously used in the Price of Anarchy literature. What's new is that by our first main result this dynamic achieves a $\Omega(1/\log m)$ approximation to the optimal social welfare that starts to apply after a single round of bid updates.

We note that both these dynamics are obtained under standard computational assumptions. The result for fractionally subadditive valuations requires access to demand and XOS oracles [2]. XOS oracles—which are standard in computer science—are given a valuation function v_i and a set of items S and return an additive function a_i such that $a_i(S) = v_i(S)$ and $a_i(T) \leq v_i(T)$ for all $T \neq S$. The result for subadditive valuations requires access to demand oracles and the ability to compute an approximate supporting additive valuation for a fixed set [3, 4, 5].

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 for subadditive valuations are essentially best possible in a very general sense.

Main Result 2. *For agents with subadditive valuations no best-response dynamics in which agents do not overbid on the grand bundle can guarantee a $\omega(\log\log m/\log m)$ 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 agents 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 weak no-overbidding dynamics, in which agents do not overbid on the set of 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 an agent 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 [6, 7, 8]. On the other hand, a number of results show that certain types of best-response dynamics reach states of low social cost quickly [9, 10, 11, 12, 13]. 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. [1], and subsequently refined and improved upon in [4, 14, 15, 16, 17]. These works provide welfare guarantees for a broad range of equilibrium concepts ranging from pure Nash equilibria to (coarse) correlated equilibria and Bayes-Nash equilibria. Some of these bounds are based on mechanism smoothness, others are not. For fractionally subadditive valuations there is a smoothness-based proof that shows that the Price of Anarchy for pure Nash equilibria is at most 2 [1, 15]. For subadditive valuations the Price of Anarchy for pure Nash equilibria is also at most 2 [4], but the best smoothness-based proof gives a bound of $O(\log m)$ [4, 15]. In fact, as shown by Roughgarden [18], combinatorial auctions with item bidding achieve (near-)optimal Price of Anarchy among a broad class of "simple" mechanisms.

Christodoulou et al. [1] gave a polynomial-time algorithm for computing a pure Nash equilibrium for submodular valuations. They furthermore 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. [19], 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 [20], who showed that computing a Bayes-Nash equilibrium is PP-hard.

Most recently, Daskalakis and Syrgkanis [21] considered coarse correlated equilibria. They showed that even for unit-demand agents (a strict subclass of submodular) there are no polynomial-time no-regret learning algorithms for finding such equilibria, unless $\mathsf{RP} \supseteq \mathsf{NP}$, 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 fractionally subadditive 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 in this context comes from Devanur et al. [22], 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 Braverman et al. [23], achieves optimal Price of Anarchy bounds within a broader class of mechanisms.

Our work on best-response dynamics in combinatorial auctions is also closely related to iterative combinatorial auction formats. These include the simultaneous multi-round auction (SMRA) (see [24]), which is also based on item bidding, and the combinatorial clock auction (CCA) [25], which allows combinatorial bids. It is interesting to note that practical implementation of these mechanisms are usually complemented with a variety of rules that restrict the allowed bids—such as activity rules, minimum bid increments, and/or monotonicity rules—that are designed to achieve fast progress and near-optimal social welfare at termination. Welfare guarantees for the SMRA can be found in [26, 27, 28] and welfare guarantees for the CCA can be found in [29, 30]. However, to the best of our knowledge, only pseudo-polynomial running time guarantees are known for these auctions, and the understanding of the tradeoff between running time and performance at termination is rather limited.

A final point of reference are approximation algorithms for combinatorial auctions. It is known that no mechanism can achieve a better than $1/m^{1/2-\epsilon}$ approximation for submodular valuations with valuation queries alone [31]. Relying on demand queries, Dobzinski et al. [32] gave a 1/2-approximation using techniques, which in the meantime became standard in Price-of-Anarchy analyses. Feige [3] used more sophisticated techniques to improve the guarantees to $1 - 1/e$ for fractionally subadditive and 1/2 for subadditive valuations. These results only give approximation guarantees but no truthful mechanism. For truthfulness, Dobzinski [33] recently managed to improve on a long-standing approximation guarantee of $\Omega(1/\log m)$ for submodular valuations to $\Omega(1/\sqrt{\log m})$ for XOS valuations. In the meantime, Assadi and Singla [34] further improved this bound to $\Omega(1/(\log \log m)^3)$.

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 $S, T \subseteq M$, $v_i(S \cup T) \leq v_i(S) + v_i(T)$.

In the description of our update procedures, we reference two types of oracles. A *demand oracle* takes as input a vector of item prices $(p_j)_{j \in M}$ and returns a demand set, that is, a set that maximizes $v_i(S) - \sum_{j \in S}$. An *XOS oracle* takes as input a set S and returns $v'_{i,j} \geq 0$ such that $v_i(S) = \sum_{j \in S} v'_{i,j}$ and $v_i(S') \geq \sum_{j \in S'} v'_{i,j}$ for every set S'. Note that this is possible if and only if v_i is XOS. The function v'_i is called additive supporting function or just supporting valuation.

The dynamics that we consider approach agents in round-robin fashion. When agent 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.

We say that bids $b_{i,j}$ for all i and all j satisfy *strong no-overbidding* if for every agent i and all sets of items S it holds that $\sum_{j \in S} b_{i,j} \leq v_i(S)$.

2.1. *Bid Updates for XOS Valuations*

For XOS valuations we can update bids as described by Christodoulou et al. [1]. In this dynamic, agents are activated one by one and when its their turn they choose an arbitrary demand set D and bid an additive supporting function on that set and zero on all other items. That is, if D is the demand set chosen by agent i, let $(v_{i,j}^{\ell})_{j \in M}$ be a 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 $b^0, b^1, b^2, \ldots$ 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. For our first dynamic for subadditive valuations we proceed as follows. We again consider agents one by one. When it is agent i's turn, we let this agent choose an inclusion-minimal demand set D. That is, a demand set D such that no strict subset $D' \subset D$ is also a demand set.

Now, to determine agent i's bid on D when facing bids b_{-i}^t by the agents other than i, we look at $\tilde{u}_i(S, b_{-i}^t) = v_i(S) - \sum_{j \in S} \max_{k \neq i} b_{k,j}^t$ for all S, i.e., the utility that agent i can derive from buying the set S. Observe that $\tilde{u}_i(\,\cdot\,, b_{-i}^t)$ is subadditive for every b_{-i}^t. We can show that $\tilde{u}_i(S, b_{-i}^t) > 0$ for all $S \subseteq D$ unless $D = \emptyset$. Therefore, by [4, 5] 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.

Note that if we have access to demand oracles, we can find an inclusion-minimal demand set D with polynomially many queries. We first query for an arbitrary demand set D. Then we issue a demand query on all subsets of D of size $|D| - 1$. If none of these queries yields a set D' with the same utility, then we know that D is inclusion minimal. Otherwise, one of the queries will return a set D' with $D' \subset D$ that yields the same utility as D, and we continue the process with D'. Note that since $|D'| < |D|$, this process will terminate after at most m iterations.

Assuming access to demand oracles, the update steps can therefore be performed in polynomial time 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 $b^0, b^1, b^2, \ldots$ that is 1-safe and in which each update is a $(1/\ln m)$-aggressive best response.*

Aggressive Updates. For our second dynamic for subadditive valuations, 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. Our theorem does not require agents to play exact best responses, and it also does not require that all agents use the same strategy for updating their bids.

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 remark at this point that, in case of fractionally subadditive valuations and initial bids being $b^0 = 0$, the argument for $(\alpha, \beta) = (1, 1)$ can be simplified and improved to show a guarantee of $1/2$.[1] However, as we show in Appendix C, the point-wise welfare guarantee of $1/3$ from any starting bids is tight for the respective mechanism.

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 $\frac{1}{T} \sum_{t=1}^{T} SW(b^t) \geq \frac{\alpha}{(2\alpha+1)\beta} \cdot \left(1 - \frac{n}{T}\right) \cdot OPT(v)$.*

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 Section 4 we show that the $\Omega(1/\log m)$ bounds are essentially best possible for best-response dynamics in a very 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 3.4) 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 3.5) 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 3.3. *Consider a sequence $b^0, \ldots, b^n$ in which agent i updates his bid in step i. Denote agent 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 agent i. agent i updates his bid in step i. Suppose agent 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 agents $1, \ldots, i$, $z_j^0 = 0$ 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 agents 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. $\square$

[1] For the Potential Procedure of [1], this already follows from the arguments in [32] in combination with monotonicity of declared welfare.

Our first key lemma shows that if the initial declared welfare is low, we reach a state of high declared welfare after a single round of bid updates. We use that bid updates are aggressive, which causes bids to be high enough.

Lemma 3.4. *Let $S_1^*, \dots, S_n^*$ be any feasible allocation, in which agent i receives items S_i^*. Consider a sequence $b^0, \dots, b^n$ in which agent 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 agent 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 agents $i \in N$ 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 3.3. 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

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

as claimed. □

In our second key lemma, we show that the declared welfare never drops drastically. So, in particular, if we start from a state of high declared welfare, we will still be in a state of high declared welfare after each bidder updated his bid although the declared welfare is not necessarily monotone. To prove the lemma, we use the fact that previous bids were safe—so they are not too high—and new bids are aggressive—so they are high enough.

Lemma 3.5. *Consider a β-safe bid sequence $b^0, \dots, b^n$ in which agent 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 agent i and his update from b^{i-1} to b^i. Denote the set of items that agent 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

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

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 agents between steps $i-1$ and i is equal to the difference in agent i's declared utility at these time steps. Formally,

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

$$
\begin{aligned}
&= \sum_{j \in M} \max_{k \neq i} b_{k,j}^{i-1} + \sum_{j \in S_i^i} b_{i,j}^i - \sum_{j \in S_i^i} \max_{k \neq i} b_{k,j}^i \\
&= \sum_{j \in M} \max_{k \neq i} b_{k,j}^{i-1} + u_i^D(b_i) \\
&= \sum_{j \in M \setminus S_i^{i-1}} \max_{k \neq i} b_{k,j}^{i-1} + \sum_{j \in S_i^{i-1}} \max_{k \neq i} b_{k,j}^{i-1} + u_i^D(b_i) \\
&= \sum_{j \in M \setminus S_i^{i-1}} \max_{k \neq i} b_{k,j}^{i-1} + \sum_{j \in S_i^{i-1}} b_{i,j}^{i-1} + u_i^D(b_i) - \sum_{j \in S_i^{i-1}} b_{i,j}^{i-1} + \sum_{j \in S_i^{i-1}} \max_{k \neq i} b_{k,j}^{i-1} \\
&= DW(b^{i-1}) + u_i^D(b_i) - u_i^D(b^{i-1}) \ .
\end{aligned}
$$

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,

$$
\begin{aligned}
DW(b^i) &= DW(b^{i-1}) + u_i^D(b^i) - u_i^D(b^{i-1}) \\
&\geq DW(b^{i-1}) + u_i^D(b^i) - \beta \cdot u_i(b^{i-1}) \\
&\geq DW(b^{i-1}) + u_i^D(b^i) - \frac{\beta}{\alpha} \cdot u_i^D(b^i) \\
&= DW(b^{i-1}) - \left(\frac{\beta}{\alpha} - 1 \right) \cdot u_i^D(b^i) \ .
\end{aligned}
$$

Summing this inequality over all agents $i \in N$ 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 3.3 to conclude that

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

which concludes the proof. ☐

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. Note that this lemma captures the intuition that bidders never overbid drastically when using safe bids.

Lemma 3.6. *In a β-safe sequence of bid profiles $b^0, b^1, b^2, \ldots$ 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 $T_1, \ldots, T_n$ that corresponds to b^t,

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

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

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

and the claim follows. ☐

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 agent updates his bid exactly once. By the virtue of being a subsequence of a β-safe bidding sequence the sequence $b^{t-n}, \ldots, b^t$ is β-safe. Moreover each bid update is α-aggressive.

Applying first Lemma 3.5 and then Lemma 3.4 with b^t taking the role of b^n, b^{t-n} taking the role of b^0, and setting $S_1^*, \ldots, S_n^*$ 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 3.6, $DW(b^t) \leq \beta \cdot SW(b^t)$. Combining this with the previous inequality yields

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

as claimed. □

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 3.4 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 3.6.

Proof of Theorem 3.2. We first use Lemma 3.4 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

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

Solving this inequality for $\frac{1}{T} \cdot \sum_{t=1}^{T} DW(b^t)$ and using Lemma 3.6 to lower bound $SW(b^t)$ by $1/\beta \cdot DW(b^t)$ we obtain

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

which proves the claim. □

4. Impossibility for Subadditive Valuations

Next we show our second main result (Theorem 4.1), which shows that no best-response dynamics in which agents 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 agents 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 $n = 2$ agents, $m = 2^k - 1$ items, and subadditive valuations $v = (v_1, v_2)$ such that in every best-response dynamics in which agents 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 agent has updated his 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 agents after the second agent 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 $n = 2$ agents and $m = 2^k - 1$ items. The valuations of the first agent are based on an example that demonstrates the worst-case integrality gap for set cover linear programs (see, e.g, [35, Example 13.4]), and has been used in the context of combinatorial auctions with item bidding before [4]. The crux of our construction is in the design of the second agent's valuation function, and its interplay with the valuation function of the first agent.

Definition 4.2. For every positive integer $k \in \mathbb{N}_{>0}$ the hard instance $\mathcal{I}_k$ consists of $n = 2$ agents and $m = 2^k - 1$ items and the following subadditive valuations:

1. First agent: 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 agent: Set $\rho = 4\frac{k}{m}$ and $d = k - \log_2 k$. 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

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

$$w_D(T) = \begin{cases} 0 & \text{for } |T| = 0 \\ \frac{|D|}{2} & \text{for } 0 < |T \cap D| < |D| \\ |D| & \text{else} \end{cases} .$$

Note that, in the instances just described, the first agent has a valuation of $v_1(M) \geq k = \log_2(m+1)$ for the grand bundle, while the second agent 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 the fact that the first agent does not overbid on the grand bundle implies the existence of a subset of items $D \in \mathcal{D}$ with low prices (Lemma 4.3). Intuitively, this is because the sets of items that the second agent is interested in are rather small (of size about $m/\log_2 m$), and there are sufficiently many of these sets. We then show that every demand set of the second agent under these prices includes some set of items $D' \in \mathcal{D}$ (Lemma 4.4). In the final step, we show that if the second agent buys any such set D', then the first agent's valuation for the remaining items $M \setminus D'$ and hence the overall social welfare is at most $O(\log\log m)$ (Lemma 4.5).

Lemma 4.3. *Let $k \in \mathbb{N}_{>0}$. Consider the hard instance $\mathcal{I}_k$. For every vector of bids b such that the first agent 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 agent 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 independent of x. Namely, it is given by $\binom{k-1}{d-1}_2$, where $\binom{\cdot}{\cdot}_q$ refers to the q-binomial coefficient (see, e.g., [36]). 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 conclusion, 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. □

Lemma 4.4. *Let $k \in \mathbb{N}_{>0}$. Consider the hard instance $\mathcal{I}_k$. If the prices p as seen by the second agent 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 agent 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 $D' \in \mathcal{D}$. □

Lemma 4.5. *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 $x_1, \ldots, x_d$ of $D' \cup \{0\}$ can be extended by $x_{d+1}, \ldots, x_k$ to a basis of $\mathbb{F}_2^k$. Let $X = (x_1, \ldots, x_k)$. This way, X^{-1} is the matrix that expresses $j \in \mathbb{F}_2^k$ as a linear combination of $x_1, \ldots, x_k$. As $x_1, \ldots, x_d$ 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 $d+1, \ldots, k$. This implies that the set $M \setminus D'$ can be covered by sets S_i for i being the rows $d+1, \ldots, k$ of X^{-1}. Therefore $v_1(M \setminus D') \leq k - d$. □

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

Let b^t be a bid vector after the second agent has made a move. Using Lemma 4.3, 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 4.4, the second agent then buys a superset of some $D' \in \mathcal{D}$. Therefore, right after the second agent has updated his bid the first agent is allocated a subset of the items $M \setminus D'$. Lemma 4.5 implies that the social welfare for this allocation is no higher than $k - d + \rho 2^d = O(\log \log m)$. □

5. Beyond Round-Robin Activation

Our positive results make use of the fact that agents are activated to update their bid in round-robin fashion. That is, between two activations of an agent, each other agent is activated exactly once. In this section, we investigate alternative activation protocols.

5.1. Randomized Activation

We first show that our positive results extend to the case where at each step a random agent gets to update his bid.

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

The key difference to the previous positive results is as follows. In the case of round-robin activation, we could bound the price that an agent has to pay for an item j at any time by the sum of the maximum bid before the first and after the n-th step. As now, in the case of random activation, an agent can potentially be activated multiple times during the first n steps, this is not true anymore. Instead, we can show the following lemma.

Lemma 5.2. *Consider a sequence of bids that is generated by choosing at each time step an agent uniformly at random and letting this agent update his bid. Then, for all items $j \in M$ and all lengths of the sequence $T \geq 0$, we have*

$$\mathbf{E}\left[\max_{t \leq T} \max_i b_{i,j}^t\right] \leq \left(1 - \frac{1}{n}\right)^{-T} \mathbf{E}\left[\max_i b_{i,j}^T\right] .$$

The proof can be found in Appendix D. The overall idea is to bound the probability that an agent who causes a high bid is activated again. Using this lemma, we can follow a similar pattern as when proving Theorem 3.1.

Proof of Theorem 5.1. Since all of our arguments apply starting from any vector of bids, we can without loss of generality assume that T is the final of a sequence of n bid updates, and so $T = n$. Let N' be the set of agents that are selected to bid at least once during this sequence of bid updates. Denote by $S_1^*, \ldots, S_n^*$ the allocation that maximizes social welfare. By a variant of Lemma 3.4, which does not make use of round-robin activation and is given as Lemma D.1 in Appendix D, we have

$$DW(b^T) + \alpha \sum_{j \in M} \max_{t \leq T} \max_i b_{i,j}^t \geq \alpha \sum_{i \in N'} v_i(S_i^*) \ .$$

Note that $DW(b^T)$, $\max_{t \leq T} \max_i b_{i,j}^t$, and N' are now random variables. Taking expectations of both sides, we get

$$\mathbf{E}\left[DW(b^T) + \alpha \sum_{j \in M} \max_{t \leq T} \max_i b_{i,j}^t\right] \geq \mathbf{E}\left[\alpha \sum_{i \in N'} v_i(S_i^*)\right] \ .$$

By linearity of expectation, this implies

$$\mathbf{E}\left[DW(b^T)\right] + \alpha \sum_{j \in M} \mathbf{E}\left[\max_{t \leq T} \max_i b_{i,j}^t\right] \geq \alpha \sum_{i \in N} \mathbf{Pr}\left[i \in N'\right] v_i(S_i^*) \ .$$

The probability of each agent to be selected at least once is $\mathbf{Pr}\left[i \in N'\right] = 1 - \left(1 - \frac{1}{n}\right)^T$. Lemma 5.2 shows that $\mathbf{E}\left[\sum_{j \in M} \max_{t \leq T} \max_i b_{i,j}^t\right] \leq \left(1 - \frac{1}{n}\right)^{-T} \mathbf{E}\left[DW(b^T)\right]$.

We obtain

$$\left(1 + \alpha \left(1 - \frac{1}{n}\right)^{-T}\right) \mathbf{E}\left[DW(b^T)\right] \geq \alpha \left(1 - \left(1 - \frac{1}{n}\right)^T\right) \cdot \sum_{i \in N} v_i(S_i^*) \ ,$$

and therefore

$$\mathbf{E}\left[DW(b^T)\right] \geq \alpha \cdot \frac{1 - \left(1 - \frac{1}{n}\right)^T}{1 + \alpha \left(1 - \frac{1}{n}\right)^{-T}} \cdot \sum_{i \in N} v_i(S_i^*) \ .$$

Finally, we use Lemma 3.6 to relate the declared social welfare to the actual social welfare and the fact that $T = n \geq 2$ to lower bound $1 - (1 - 1/n)^n \geq 1/2$ and upper bound $(1 - 1/n)^{-n} \leq 4$. This yields,

$$\mathbf{E}\left[SW(b^T)\right] \geq \frac{\alpha}{2(1 + 4\alpha)\beta} \cdot OPT(v) \ ,$$

as claimed. □

5.2. Adversarial Activation

We conclude by showing that our positive results that show quick convergence to states of high welfare no longer apply if an adversary chooses the order in which agents get to update their bids. Our result concerns XOS valuations, and 1-safe bidding sequences in which each bid update is to a 1-aggressive best response. It applies even if agents update their bids as in the Potential Procedure of [1]. That is, unless the activated agent already plays a best response, he chooses an arbitrary demand set and bids his supporting additive valuation on the respective set and zero on all other items.

Theorem 5.3. *For every $\epsilon > 0$, n, and k, there is an instance with n agents with XOS valuations and $(n-1)\cdot(k+1)$ items, an initial bid vector b^0, and an activation sequence such that, even if each activated agent updates his bid as in the Potential Procedure, until each agent has been activated $\Omega(2^k)$ times the welfare has never exceeded a $\frac{1+\epsilon}{n-1}$ fraction of the optimum.*

At the core of our proof (in Appendix E) is the following proposition that applies even if agents have unit-demand valuations, i.e., an agent's valuation for a set of items is the maximum value for any item in the set. It shows the existence of a cyclic activation pattern in which each agent gets to update his bid, but the dynamic remains in states of low welfare. The construction assumes that agents also update their bid if this does not strictly improve their utility, and that ties among multiple best responses are broken in our favor.

Proposition 5.4. *For every $\epsilon > 0$ and n, there is an instance of n agents with unit-demand valuations for $n-1$ items, an initial bid vector b^0, and a cyclic activation pattern in which every agent is activated at least once and bid updates are as in the Potential Procedure except that updates need not be strict improvements and ties among multiple best responses are broken in our favor, but the social welfare is always at most a $\frac{1+\epsilon}{n-1}$ fraction of the optimal welfare.*

Proof. There are n agents and $n-1$ items. agent i's valuation for a set $S \subseteq M$ is given as $v_i(S) = \max_{j \in S} v_{i,j}$. For agent 1, we let $v_{1,1} = \ldots, v_{1,n-1} = 1+\epsilon$. For agent $i > 1$, define $v_{i,i-1} = 1$ and $v_{i,j} = 0$ for $j \neq i-1$. The social optimum assigns item j to agent $j+1$ and has welfare $n-1$.

In the initial bid vector b^0 all agents bid zero. The activation scheme is as follows: In every odd step agent 1 makes a move, while in even steps agents $i > 1$ are activated in a round-robin way. That is, the activation works repeatedly as $1, 2, 1, 3, 1, 4, \ldots, 1, n-1, 1, n$.

With this activation order, it's possible that agent 1 bids $1+\epsilon$ on item t the t-th time he is activated, while agents $i > 1$, when activated, see a bid of $1+\epsilon$ on the item they are interested in, and therefore bid 0 on all items. This way the social welfare at any time step $t \geq 1$ is $1+\epsilon$. □

Our proof in the appendix combines this construction with several copies of the exponential lower-bound construction of Theorem 3.4 in [1], and thus ensures that each update is a strict improvement and unique.

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 Appendix F, we discuss an example with MPH-k valuations [37] 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 in Appendix F 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] G. Christodoulou, A. Kovács, M. Schapira, Bayesian combinatorial auctions, Journal of the ACM 63 (2) (2016) 11.

[2] S. Dobzinski, M. Schapira, An improved approximation algorithm for combinatorial auctions with submodular bidders, in: Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA'06, 2006, pp. 1064–1073.

[3] U. Feige, On maximizing welfare when utility functions are subadditive, in: Proceedings of the 38th Annual ACM Symposium on Theory of Computing, STOC'06, 2006, pp. 41–50.

[4] K. Bhawalkar, T. Roughgarden, Welfare guarantees for combinatorial auctions with item bidding, in: Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms, SODA'11, 2011, pp. 700–709.

[5] S. Dobzinski, Two randomized mechanisms for combinatorial auctions, in: Proceedings of APPROX-RANDOM, Vol. 4627, 2007, pp. 89–103.

[6] R. W. Rosenthal, A class of games possessing pure-strategy nash equilibria, International Journal of Game Theory 2 (1) (1973) 65–67.

[7] D. Monderer, L. S. Shapley, Potential games, Games and Economic Behavior 14 (1) (1996) 124–143.

[8] H. Ackermann, H. Röglin, B. Vöcking, On the impact of combinatorial structure on congestion games, Journal of the ACM 55 (6).

[9] M. X. Goemans, V. S. Mirrokni, A. Vetta, Sink equilibria and convergence, in: Proceedings of the 46th IEEE Symposium on Foundations of Computer Science, FOCS'05, 2005, pp. 142–154.

[10] G. Christodoulou, V. S. Mirrokni, A. Sidiropoulos, Convergence and approximation in potential games, Theoretical Computer Science 438 (2012) 13–27.

[11] V. Bilò, A. Fanelli, M. Flammini, L. Moscardelli, Performance of one-round walks in linear congestion games, Theory of Computing Systems 49 (1) (2011) 24–45.

[12] B. Farzad, N. Olver, A. Vetta, A priority-based model of routing, Chicago Journal of Theoretical Computer Science 2008.

[13] T. Roughgarden, Intrinsic robustness of the price of anarchy, Journal of the ACM 62 (5) (2015) 32.

[14] A. Hassidim, H. Kaplan, Y. Mansour, N. Nisan, Non-price equilibria in markets of discrete goods, in: Proceedings of the 12th ACM Conference on Electronic Commerce, EC'11, 2011, pp. 295–296.

[15] V. Syrgkanis, É. Tardos, Composable and efficient mechanisms, in: Proceedings of the 45th ACM Symposium on Theory of Computing Conference, STOC'13, 2013, pp. 211–220.

[16] P. Dütting, M. Henzinger, M. Starnberger, Valuation compressions in VCG-based combinatorial auctions, in: Proceedings of the 9th Conference on Web and Internet Economics, WINE'13, 2013, pp. 146–159.

[17] M. Feldman, H. Fu, N. Gravin, B. Lucier, Simultaneous auctions are (almost) efficient, in: Proceedings of the 45th ACM Symposium on Theory of Computing Conference, STOC'13, 2013, pp. 201–210.

[18] T. Roughgarden, Barriers to near-optimal equilibria, in: 55th IEEE Annual Symposium on Foundations of Computer Science, FOCS'14, 2014, pp. 71–80.

[19] S. Dobzinski, H. Fu, 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, SODA'15, 2015, pp. 110–122.

[20] Y. Cai, C. H. Papadimitriou, Simultaneous bayesian auctions and computational complexity, in: Proceedings of the 15th ACM Conference on Economics and Computation, EC'14, 2014, pp. 895–910.

[21] C. Daskalakis, V. Syrgkanis, Learning in auctions: Regret is hard, envy is easy, in: Proceedings of the 57th IEEE Symposium on Foundations of Computer Science, 2016, pp. 219–228.

[22] N. R. Devanur, J. Morgenstern, V. Syrgkanis, S. M. Weinberg, Simple auctions with simple strategies, in: Proceedings of the 16th ACM Conference on Economics and Computation, EC'16, 2015, pp. 305–322.

[23] M. Braverman, J. Mao, 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, SODA'16, 2016, pp. 1444–1457.

[24] P. Milgrom, Putting Auction Theory to Work, Cambridge University Press, Cambridge, UK, 2004.

[25] D. Porter, S. Rassenti, A. Roopnarine, V. Smith, Combinatorial auction design, Proceedings of the National Academy of Sciences of the United States of America 100 (2003) 11153–11157.

[26] P. Milgrom, Putting auction theory to work: The simultaneous ascending auction, Journal of Political Economy 108 (2000) 245–272.

[27] H. Fu, R. Kleinberg, R. Lavi, Conditional equilibrium outcomes via ascending price processes with applications to combinatorial auctions with item bidding, in: Proceedings of the 13th ACM Conference on Electronic Commerce, EC'12, 2012, p. 586.

[28] N. Bousquet, Y. Cai, A. Vetta, Welfare and rationality guarantees for the simultaneous multiple-round ascending auction, in: Proceedings of the 11th Conference on Web and Internet Economics, WINE'15, 2015, pp. 216–229.

[29] N. Bousquet, Y. Cai, C. Hunkenschröder, A. Vetta, On the economic efficiency of the combinatorial clock auction, in: Proceedings of 27th ACM-SIAM Symposium on Discrete Algorithms, SODA'16, 2016, pp. 1407–1423.

[30] M. Dupré la Tour, A. Vetta, The combinatorial clock auction: the effects of strategic behaviour and the price increment rule on social welfare, in: Proceedings of the 2018 ACM Conference on Economics and Computation, EC'18, 2018, pp. 91–108.

[31] S. Dobzinski, J. Vondrák, Impossibility results for truthful combinatorial auctions with submodular valuations, Journal of the ACM 63 (1) (2016) 5.

[32] S. Dobzinski, N. Nisan, M. Schapira, Approximation algorithms for combinatorial auctions with complement-free bidders, in: Proceedings of the 37th Annual ACM Symposium on Theory of Computing, 2005, pp. 610–618.

[33] 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, STOC'16, 2016, pp. 940–948.

[34] S. Assadi, S. Singla, Improved truthful mechanisms for combinatorial auctions with submodular bidders, in: Proceedings of the 60th IEEE Symposium on Foundations of Computer Science, FOCS'19, 2019, pp. 233–248.

[35] V. V. Vazirani, Approximation Algorithms, Springer Verlag, New York, NY, USA, 2001.

[36] A. Prasad, Counting subspaces of a finite vector space — 1, Resonance 15 (11) (2010) 977–987.

[37] U. Feige, M. Feldman, N. Immorlica, R. Izsak, B. Lucier, V. Syrgkanis, A unifying hierarchy of valuations with complements and substitutes, in: Proceedings of the 29th AAAI Conference on Artificial Intelligence, AAAI'15, 2015, pp. 872–878.

A. Non-Existence of Weak No-Overbidding Pure Nash Equilibria for Subadditive Valuations

We can also leverage our novel insights regarding hard instances (Definition 4.2) for subadditive valuations to show that there need not be a pure Nash equilibrium in weakly no-overbidding strategies, even if agents only consider deviations to weakly no-overbidding strategies.

A bid b_i of agent i against bids b_{-i} of the agents other than i that wins him the set of items S satisfies *weak no-overbidding* if $\sum_{j \in S} b_{i,j} \leq v_i(S)$.

Theorem A.1. *Let $k \in \mathbb{N}_{>0}$. Consider the hard instance $\mathcal{I}_k$ for subadditive valuations with $n = 2$ agents 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 agent 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 agent wins the set of items $W \subseteq M$ in b, then the first agent wins the set of items $M \setminus W$. By weak no-overbidding, we have

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

The first agent 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

$$\begin{aligned}\sum_{j \in M} b_{1,j} &\leq v_1(M \setminus W) + v_2(W) \\ &\leq v_1(M) + v_2(M) \\ &= k + \rho \cdot 2^d \\ &= k + 4 \cdot \frac{k}{m} \cdot 2^{k - \log_2 k} \\ &= k + 4 \ .\end{aligned}$$

By the same argument as in Lemma 4.3, 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 4.4 and because the second agent 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 agent 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 agent wins all items M. This bid fulfills the weak no-overbidding property because

$$\begin{aligned}\sum_{j\in M} b'_{1,j} &= \sum_{j\in W}\left(b_{2,j}+\frac{1}{m}\right)+\sum_{j\in M\setminus W} b_{1,j}\\ &\le v_2(W)+1+v_1(M\setminus W)\\ &\le v_2(D')+1+v_1(M\setminus D')\\ &\le \rho 2^d+1+k-d\\ &= 4+1+\log_2 k\\ &\le k\\ &= v_1(M)\ ,\end{aligned}$$

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 4.5, and the final inequality holds because we have assumed $k \geq 8$.

The deviation by the first agent is beneficial because

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

where the first inequality uses Lemma 4.5, the second inequality uses that $v_2(W) \leq v_2(D') = 4$, and the final inequality follows from the definition of $d = k - \log_2 k$ and the assumption that $k \geq 8$ and so $d > 4$. □

B. Proofs Omitted from Section 2

In this appendix we prove the propositions that establish the existence of aggressive and safe bidding dynamics for XOS and subadditive valations.

B.1. Sufficiency of Strong No-Overbidding

We first 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.

Lemma B.1. *If the initial bid vector b^0 satisfies strong no-overbidding and at each time step $t \geq 1$ some agent i gets to update his bid to a best response, which satisfies strong 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 $\sum_{j\in S} b^t_{i,j} \leq v_i(S)$ for all agents i, time steps $t \geq 0$, and sets of items S. Subtracting $\sum_{j\in S} \max_{k\neq i} b^t_{k,j}$ from both sides shows the claim. □

B.2. Proof of Proposition 2.1

Consider an arbitrary agent i and his update to bid b^t_i. The bid b^t_i satisfies strong no-overbidding by definition. Hence Lemma B.1 shows that the bid sequence is 1-safe. It remains to show that b^t_i 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 agent 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,

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

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 agent 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

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

B.3. Proof of Proposition 2.2

Consider an arbitrary agent 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 $\tilde{u}_i(T, b_{-i}^t) \leq 0$. Then, by subadditivity of v_i,

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

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 agent 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 agent 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 agent i's update satisfies strong no-overbidding. Namely, for every $S \subseteq D$,

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

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 agent 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 agent i's update. In the vector $b_{t'}$, agent 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 agent i wins item j, all new bids have to be zero.

In combination, we have

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

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 the Point-Wise Welfare 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 $n = 3$ agents, $m = 3$ 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 agents are unit demand, i.e., for all agents 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
agent 1	1	0	0
agent 2	$1+\epsilon$	$1+2\epsilon$	$1+3\epsilon$
agent 3	0	0	1

Suppose that the XOS representation of these valuations is that each agent 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 agent 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 agent 1 gets to update his bid, then agent 2, and then agent 3.

agent 1 is already playing a best response to b_{-1}^0, so $b^1 = b^0$. Now, to get b^2, agent 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 agent 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. □

D. Proof of Theorem 5.1

In this appendix we provide additional details for the proof of Theorem 5.1. We first prove Lemma 5.2. Afterwards, we state and prove Lemma D.1.

D.1. Proof of Lemma 5.2

For a fixed T, let $y_j = \max_{t \leq T} \max_i b_{i,j}^t$ and $p_j^t = \max_i b_{i,j}^t$ for $t \leq T$. We first show that for all $x > 0$

$$\mathbf{Pr}\left[y_j \geq x\right] \leq \left(1 - \frac{1}{n}\right)^{-T} \mathbf{Pr}\left[p_j^T \geq x\right] \tag{D.1}$$

To show (D.1), we use that y_j is defined to be $\max_{t' \leq T} p_j^{t'}$. That is, if $y_j \geq x$, there has to be a $t' \in \{0, 1, \ldots, T\}$ for which $p_j^1 < x, \ldots, p_j^{t'-1} < x, p_j^{t'} \geq x$. Note that for different t' these are disjoint events, so

$$\mathbf{Pr}\left[y_j \geq x\right] = \sum_{t'=0}^{T} \mathbf{Pr}\left[p_j^1 < x, \ldots, p_j^{t'-1} < x, p_j^{t'} \geq x\right] \quad .$$

Let us fix t' and consider the event that $p_j^1 < x, \ldots, p_j^{t'-1} < x, p_j^{t'} \geq x$. If $t' > 0$, in step t' an agent i has been selected that whose bid has set $p_j^{t'} \geq x$; if $t' = 0$, the initial bid of some agent i on item j is at least x. We have have $p_j^T < x$ only if this agent i is selected to update his bid in steps $t'+1, \ldots, T$. This happens with probability $1 - \left(1 - \frac{1}{n}\right)^{T-t'} \leq 1 - \left(1 - \frac{1}{n}\right)^T$. Formally, we have

$$\mathbf{Pr}\left[p_j^T < x \;\middle|\; p_j^1 < x, \ldots, p_j^{t'-1} < x, p_j^{t'} \geq x\right] \leq 1 - \left(1 - \frac{1}{n}\right)^T \quad .$$

This implies

$$\mathbf{Pr}\left[p_j^T \geq x, p_j^1 < x, \ldots, p_j^{t'-1} < x, p_j^{t'} \geq x\right] \geq \left(1 - \frac{1}{n}\right)^T \mathbf{Pr}\left[p_j^1 < x, \ldots, p_j^{t'-1} < x, p_j^{t'} \geq x\right] \quad .$$

We thus obtain

$$\begin{aligned}
\mathbf{Pr}\left[y_j \geq x\right] &= \sum_{t'=0}^{T} \mathbf{Pr}\left[p_j^1 < x, \ldots, p_j^{t'-1} < x, p_j^{t'} \geq x\right] \\
&\leq \left(1 - \frac{1}{n}\right)^{-T} \sum_{t'=0}^{T} \mathbf{Pr}\left[p_j^T \geq x, p_j^1 < x, \ldots, p_j^{t'-1} < x, p_j^{t'} \geq x\right] \\
&= \left(1 - \frac{1}{n}\right)^{-T} \mathbf{Pr}\left[p_j^T \geq x\right] \quad .
\end{aligned}$$

This concludes the proof of (D.1).

To show the lemma, let $\epsilon > 0$. We use that the expectation of a non-negative random variable X can be approximated by $\sum_{k=0}^{\infty} \epsilon \cdot \mathbf{Pr}\left[X \geq k \cdot \epsilon\right] \leq \mathbf{E}\left[X\right] \leq \sum_{k=1}^{\infty} \epsilon \cdot \mathbf{Pr}\left[X \geq k \cdot \epsilon\right]$. Applying this approximation and using (D.1), we get

$$\begin{aligned}
\mathbf{E}\left[p_j^T\right] &\geq \sum_{k=1}^{\infty} \epsilon \mathbf{Pr}\left[p_j^T \geq k\epsilon\right] \\
&\geq \sum_{k=0}^{\infty} \left(1 - \frac{1}{n}\right)^T \epsilon \mathbf{Pr}\left[y_j \geq k\epsilon\right] - \epsilon \\
&\geq \left(1 - \frac{1}{n}\right)^T \mathbf{E}\left[y_j\right] - \epsilon \quad .
\end{aligned}$$

As this holds for all $\epsilon > 0$, we also have

$$\mathbf{E}\left[p_j^T\right] \geq \left(1 - \frac{1}{n}\right)^T \mathbf{E}\left[y_j\right] \quad .$$

D.2. Lemma D.1 and its Proof

Next we state and prove Lemma D.1, which we used in the proof of Theorem 5.1.

Lemma D.1. *Let $S_1^*, \ldots, S_n^*$ be any feasible allocation, in which agent i receives items S_i^*. Consider a sequence $b^0, \ldots, b^T$ in which each agent from N' updates his bid at least once using an α-aggressive bid. We have $(\alpha + 1) \cdot DW(b^T) + \alpha \cdot \sum_{j \in M} \max_{t \leq T} \max_i b_{i,j}^t \geq \alpha \cdot \sum_{i \in N'} v_i(S_i^*)$.*

To prove this lemma we need the following auxiliary lemma.

Lemma D.2. *Consider a sequence $b^0, \ldots, b^T$ in which agents from N' update their bid at least once. For $i \in N'$, let t_i denote the time of the last update for agent i. Then, $\sum_{i \in N'} u_i^D(b^{t_i}) \leq DW(b^T)$.*

Proof. Without loss of generality, let $N' = \{1, \ldots, n'\}$ and $t_1 < t_2 < \ldots < t_{n'}$. Consider any $i \in N'$ and let agent i's update buy him the set of items S'. Then

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

For $i \in N'$, let $z_j^i = \max_{k < i} b_{k,j}^{t_i}$ for all j, $z_j^0 = 0$. That is, z_j^i is the highest "final" bid on item j.

We observe that

$$\sum_{j \in S'} (b_{i,j}^{t_i} - \max_{k \neq i} b_{k,j}^{t_i}) \leq \sum_{j \in M} (z_j^i - z_j^{i-1}) \ .$$

This is for the following fact. For $j \notin S'$, we have $z_j^i \geq z_j^{i-1}$ by definition. For $j \in S'$, $b_{i,j}^{t_i} = z_j^i$ and $\max_{k \neq i} b_{k,j}^{t_i} \geq \max_{k < i} b_{k,j}^{t_i} = \max_{k < i} b_{k,j}^{t_{i-1}} = z_j^{i-1}$.

By summing over all agents $i \in N'$, we obtain

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

The double sum is telescoping and $z_j^T = z_j^{t_{n'}} = \max_{k \leq n'} b_{k,j}^T \leq \max_k b_{k,j}^T$ and $z_j^0 = 0$ by definition. So,

$$\sum_{i \in N'} u_i^D(b^{t_i}) \leq \sum_{j \in M} (z_j^T - z_j^0) = \sum_{j \in M} \max_k b_{k,j}^T = DW(b^T) \ ,$$

as claimed. □

We are now ready to prove the lemma.

Proof of Lemma D.1. For $i \in N'$, let t_i denote the last time agent i updates his bid. Instead of choosing bid $b_i^{t_i}$, he could have bought the set of items S_i^*. As $b_i^{t_i}$ is α-aggressive, we get

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

Let $y_j = \max_t \max_k b_{k,j}^t$.

We thus have

$$u_i^D(b^{t_i}) + \alpha \cdot \sum_{j \in S_i^*} y_j \geq \alpha \cdot v_i(S_i^*) \ .$$

Summing this inequality over all agents $i \in N'$ yields

$$\sum_{i \in N'} u_i^D(b^{t_i}) + \alpha \cdot \sum_{i \in N'} \sum_{j \in S_i^*} y_j \geq \alpha \cdot \sum_{i \in N'} v_i(S_i^*) \ .$$

The first sum is at most $DW(b^T)$ by Lemma D.2. The double sum covers each $j \in M$ at most once, therefore it is bounded by $\sum_{j \in M} y_j$. Consequently,

$$DW(b^T) + \alpha \cdot \sum_{j \in M} y_j \geq \alpha \cdot \sum_{i \in N'} v_i(S_i^*) \ ,$$

as claimed. □

E. Proof of Theorem 5.3

Our proof of Theorem 5.3 combines the construction that we used to prove Proposition 5.4 with the following exponential lower-bound construction.

Lemma E.1 (Theorem 3.4 of [1])**.** *For every k there is an instance with two agents, A and B, and k items, with fractionally subadditive valuations v_A and v_B defined by additive functions $(a_A^t)_{t\in\mathbb{N}}$ and $(a_B^t)_{t\in\mathbb{N}}$ such that in the Potential Procedure, started from initial bid vector b^0 in which both agents bid zero and with agent A making the first move, agent $z \in \{A, B\}$ plays a_z^t the t-th time he gets to update his bid and it takes at least $\Omega(2^k)$ steps before the procedure converges.*

Proof of Theorem 5.3. As in the proof of Proposition 5.4 we use n agents, we start with the initial bid vector b^0 in which all agents bid zero, and we consider agent 1 being activated in every odd step and the remaining agents being activated in round-robin fashion in even steps.

We use $m = (n-1)\cdot(k+1)$ items. Items $1,\ldots,n-1$ are used to mimic the sequence of Proposition 5.4. The remaining items are grouped into $n-1$ sets of size k, namely $C_i := \{n-1+(i-2)k+1,\ldots,n-1+(i-1)k\}$ for $i > 2$, and on each of these sets agent 1 follows the steps of the exponential-length sequence of Lemma E.1 with one of the other $n-1$ agents, with agent 1 taking the role of agent A and agent $i > 1$ taking the role of agent B.

To define the valuations, for $z \in \{A, B\}$, $i = 2,\ldots,n$, and $t \geq 1$, let $a_{z,i}^t$ be the additive valuation functions defined in Lemma E.1 that are used by agent $z \in \{A, B\}$ after the t-th update, using the items C_i.

We first define the valuation function v_i for agent $i > 1$. Namely, given some $\epsilon > 0$, let the valuation function v_i of agent $i > 1$ be defined as

$$v_i(S) = \max\{\mathbf{1}_{i-1\in S}, \epsilon\cdot\max_t a_{B,i}^t(S)\} \ .$$

That is, agent i has a high value to buy item $i-1$. He also has a very small value for items C_i according to the valuations of agent B in the exponential lower-bound construction using the items C_i.

For agent 1, we define the valuation function by setting $v_1(S) = \max_t v_1^t(S)$, where v_1^t is the additive valuation function that is used when agent 1 updates his bid for the t-th time. It is designed in such a way that the t-th update is a best response in the game on C_i with agent $i = (t-1) \mod (n-1) + 1$, who has just updated his bid, and makes the bid of agent 1 move from item $i-1$ to i, which agent $i+1$ is interested in, who will be activated next.

To define v_1^t formally, observe that when agent 1 makes his t-th update, some of the other agents have performed $\lceil\frac{t}{n-1}\rceil$ updates so far, the others only $\lfloor\frac{t}{n-1}\rfloor$. Let the respective sets of agents be denoted by $N'(t)$ and $N''(t)$. Based on this, define

$$v_1^t(S) = (1+\epsilon)\cdot\mathbf{1}_{(t-1) \bmod (n-1)+1\in S} + \epsilon\cdot\sum_{i\in N'(t)} a_{A,i}^{\lceil\frac{t}{n-1}\rceil}(S) + \epsilon\cdot\sum_{i\in N''(t)} a_{A,i}^{\lfloor\frac{t}{n-1}\rfloor}(S) \ .$$

By these definitions, the bids on items $1,\ldots,n-1$ change exactly the way as in the proof of Proposition 5.4 as long as there are still changes on items C_i for $i > 1$. By Lemma E.1 it takes at least $\Omega(2^k)$ updates until such a set C_i reaches a stable state. Therefore, our constructed best-response sequence has low welfare at least until every agent $2,\ldots,n$ has updated his bid at least $\Omega(2^k)$ times. Moreover, every update is the unique best response. □

F. Negative Result for MPH-k Valuations

The maximum over positive hypergraph-k or MPH-k hierarchy [37] comprises valuation functions with different degrees of complementarity, as parametrized by k. A valuation function v_i belongs to MPH-k if there are values $v_{i,T}^\ell \geq 0$ such that $v_i(S) = \max_\ell \sum_{T\subseteq S,|T|\leq k} v_{i,T}^\ell$. Any (monotone) valuation function can be captured with $k = m$. Fractionally subadditive valuations are precisely the case $k = 1$.

Observe that for a usual valuation function even in MPH-2, the only bids that fulfill strong no-overbidding are zero on every item. Therefore, it is not possible that agents bid α-aggressively for $\alpha > 0$ and satisfy no-overbidding in the strong sense at the same time. However, as our dynamics in Section 2.2 demonstrates, strong no-overbidding is not a necessary requirement for good welfare guarantees. Unfortunately, the case is different for MPH-k. Below we show a negative result for the valuation class MPH-3. It relies on ties regarding identical bids and multiple best responses being broken to the disadvantage of the dynamics.

Proposition F.1. *There are valuation functions for n agents on $O(n)$ items that belong to MPH-3 such that round-robin best-response dynamics only reach states that achieve a $O(\frac{1}{n})$-fraction of the optimal social welfare.*

Proof. For a given k, we define an instance with $k+4$ items and $2k+4$ agents as follows. agent $i \in [k-1]$ has a valuation of 3 for the bundles $\{i, k+1, k+2\}$ and $\{i, k+3, k+4\}$, with no value for the subsets. agent k has a valuation of 3 for the bundles $\{k, k+1, k+3\}$ and $\{k, k+2, k+4\}$, with no value for the subsets. Furthermore, there are $k+4$ agents $k+1, \ldots, 2k+4$, each of which has a valuation of 1 for exactly one (distinct) item $j \in [k+4]$. Note that due to agents $k+1, \ldots, 2k+4$, the optimal social welfare is $k+4$. Our best-response sequence will never reach a state with social welfare higher than 3.

We assume that ties are broken as follows. agent $k+1, \ldots, 2k+4$ never get an item if there is an equal bid from an agent $i \in [k]$. Among the agents $i \in [k]$, on items $k+1$ and $k+3$, agent k is preferred to $k-1$, agent $k-1$ to $k-2$, and so on. On items $k+2$ and $k+4$, agents $i \in [k-1]$ are preferred to agent k, agent $k-1$ is preferred to $k-2$, agent $k-2$ to $k-3$, and so on.

Now consider the round-robin best-response dynamics in which agents get activated in the order they are indexed. Throughout the bidding dynamics agents $k+1, \ldots, 2k+4$ will bid truthfully on their respective items. The other agents bid as follows. In odd rounds agents $i = 1, \ldots, k-1$ buy items $\{i, k+1, k+2\}$, bidding 1 on each of them. Afterwards, agent k buys items $\{k, k+1, k+3\}$, again bidding 1 on each of them. In even rounds, agents $i = 1, \ldots, k-1$ buy items $\{i, k+3, k+4\}$, bidding 1 each, making agent k buy items $\{k, k+2, k+4\}$.

Note that at every point in this sequence, only the agent that has just updated his bid gets a bundle of items of any positive value. This value is 3. □