

On the Complexity of Computing an Equilibrium in Combinatorial Auctions

Shahar Dobzinski*

Hu Fu[†]

Robert Kleinberg[‡]

April 8, 2014

Abstract

We study combinatorial auctions where each item is sold separately but simultaneously via a second price auction. We ask whether it is possible to efficiently compute in this game a pure Nash equilibrium with social welfare close to the optimal one.

We show that when the valuations of the bidders are submodular, in many interesting settings (e.g., constant number of bidders, budget additive bidders) computing an equilibrium with good welfare is essentially as easy as computing, completely ignoring incentives issues, an allocation with good welfare. On the other hand, for subadditive valuations, we show that computing an equilibrium requires exponential communication. Finally, for XOS (a.k.a. fractionally subadditive) valuations, we show that if there exists an efficient algorithm that finds an equilibrium, it must use techniques that are very different from our current ones.

1 Introduction

1.1 Background

Combinatorial auctions have received much attention in recent years. The literature on combinatorial auctions is quite large, but roughly speaking it is fair to say that most papers take an engineering-like approach, by designing algorithms that find allocations with high welfare (e.g., [10, 25]), or designing algorithms that achieve good approximation ratios when bidders play their dominant strategies.

Recently, several papers (e.g., [3, 6, 11, 14, 18, 17, 21, 23, 24]) took a more “existential” approach to analyzing combinatorial auctions. Instead of designing algorithms that give a specific recipe for computing an efficient allocation, a simple game is defined; usually it is assumed that each item is sold separately but simultaneously via some kind of an auction. The implicit assumption is that the players somehow reach an equilibrium, and the quality of this equilibrium is analyzed. The literature considers several variants of such games that differ in the single item auction type (first or second price), and in the solution concept (pure nash, mixed nash, etc.).

The goal of the current paper is to somewhat narrow the gap between these two mindsets. That is, given such a game we want to determine whether it is possible to compute a good equilibrium efficiently.

*Weizmann Institute of Science.

[†]Microsoft Research.

[‡]Cornell University.

1.1.1 The Setting

In a combinatorial auction there is a set M items ($|M| = m$) for sale. There is also a set of N bidders ($|N| = n$). Each bidder i has a valuation function $v_i : 2^M \rightarrow \mathbb{R}$ that denotes the value of bidder i for every possible subset of the items. We assume that the valuation functions are monotone (for all $S \subseteq T$, $v_i(S) \geq v_i(T)$) and normalized ($v_i(\emptyset) = 0$). The goal is to find an allocation of the items (S_1, \dots, S_n) that maximizes the social welfare: $\sum_i v_i(S_i)$. We are interested in algorithms that run in time $\text{poly}(m, n)$.

If each valuation is naively represented by 2^m numbers, reading the input alone will be too slow for our algorithms. Therefore, the standard approach that assumes that the valuations belong to some subclass that can be succinctly represented, or that the valuations are represented by black boxes that can only answer specific type of queries. The standard queries are value query (given a bundle S , what is the value of $v(S)$?) and demand query (given p_1, \dots, p_m , what is a bundle that maximizes $v(S) - \sum_{j \in S} p_j$?). For impossibility results we will also consider the communication complexity model, where the queries are not restricted and we only count the number of bits communicated. See [4] for a thorough description of these models.

Christodoulou, Kovacs, and Schapira [6] were the first to analyze the quality of such equilibria in combinatorial auctions. In particular, the auction format they study is simultaneous second price auctions. In this game the strategy of each bidder is to bid one number $b_i(j)$ for every item j . Player i gets item j if $b_i(j) \geq b_{i'}(j)$ for all players i' (if there are several players with the maximal bid for j , then j is given to some arbitrary one of them). If player i won the set of items S_i then his payment is $\sum_{j \in S_i} \max_{i' \neq i} b_{i'}(j)$. To avoid trivial equilibria, they use the standard no overbidding assumption which states that for every player i and bundle S we have that $v_i(S) \geq \sum_{j \in S} b_i(j)$. See [3, 6] for a discussion.

1.2 Our Goal: The Computational Efficiency of Finding a Good Equilibrium

In this paper we study the basic task of computing a pure equilibrium when the auction format is a simultaneous second price auction. We start with discussing submodular valuations, where for every item j and bundles $S \subseteq T$ we have that $v_i(j|S) \geq v_i(j|T)$.¹ A special class of submodular valuations also of interest for us is the *budget additive valuations*, where the valuation is fully described by a budget b and the value of each single item, with $v(S)$ given by $\min\{b, \sum_{j \in S} v(j)\}$.

In [6] it is proved that if all valuations are submodular then the (pure) price of anarchy is at most 2. Furthermore, they show that this bound can be achieved constructively: the greedy algorithm of [16] — coupled with appropriately chosen prices — actually finds an equilibrium that necessarily has a price of anarchy of 2. As better algorithms are known for the general case (an $\frac{e}{e-1}$ -approximation algorithm [25]) and for some interesting special cases (e.g., a $\frac{4}{3}$ -approximation algorithm for budget additive bidders [5], and an FPTAS for a constant number of budget additive bidders [1]), they ask whether it is possible to find an equilibrium with an approximation ratio better than 2 in polynomial time.

Our first set of results partially answers this question by providing a series of black-box reductions that show that in several interesting cases we can exploit an arbitrary approximation algorithm to efficiently find an equilibrium that provides the same (or better) approximation ratio:

Theorem: Suppose that we are in one of the following settings: (1) a constant number of submodular bidders, or (2) an arbitrary number of budget additive bidders. Then, given any

¹Here and elsewhere in the paper, we use the standard notation $v(A|B)$ to denote $v(A \cup B) - v(B)$, the marginal value of bundle A given bundle B .

allocation (S_1, \dots, S_n) , using only polynomially many value queries it is possible to find a no-overbidding equilibrium with welfare at least $\sum_i v_i(S_i)$.

In fact, we show that the portion of the theorem that deals with a constant number of submodular bidders can be extended to the more general setting of *bounded competition*, where the number of bidders may be arbitrary, but the number of bidders that compete for any single item is at most a constant. (See the definition of t -restricted instances in Section 2.) As an additional consequence of our analysis, we observe that for general submodular valuations one can modify any α -approximation algorithm to compute an equilibrium that provides an α -approximation, at the cost of an additional pseudo-polynomial number of value queries.

Next, we move on to discuss supersets of submodular valuations such as subadditive valuations (those with the property that for every S and T , $v(S) + v(T) \geq v(S \cup T)$). For this class, in [3] it was observed that sometimes there is no pure equilibrium at all. We show that even if equilibrium is known to exist, it is impossible to efficiently find it:

Theorem: It takes exponential communication to find a pure no-overbidding equilibrium in combinatorial auctions with subadditive bidders, even if such equilibrium is known to exist.

This is the first computational impossibility result for finding an equilibrium in combinatorial auctions that we are aware of.²

Of particular interest is the class of XOS valuations that lies between submodular and subadditive valuations in the complement free hierarchy [16]. In [6] it was shown that the price of anarchy is at most 2 also if the valuations are known to be XOS. As for constructively finding the equilibrium, they provide a natural dynamic that always finds an equilibrium, but the communication complexity of this dynamic is exponential.

Whether this dynamic can be altered to converge in polynomial time is not known. In fact, in [6] even the more basic question of whether this dynamic converges in polynomial time if the valuations are known to be submodular (and not just XOS) was mentioned as an open question. We provide a negative answer to this question by showing an instance of combinatorial auctions with two submodular bidders in which the dynamic takes exponential time to converge. This example will lead to a more general impossibility result — in fact our most technically involved result — we therefore take this opportunity to dive into a more technical discussion.

1.3 The Surprising Flexibility of XOS Oracles

Recall the definition of an XOS valuation. A valuation v is called XOS if there exist additive valuations (clauses) a_1, \dots, a_t such that $v(S) = \max_k a_k(S)$. An XOS oracle receives as input a bundle S and returns the *maximizing clause* of S — an additive valuation a_k from the set above such that $v(S) = a_k(S)$.

Let us now describe the exponential time algorithm of [6] for finding an equilibrium with XOS valuations. For simplicity, assume that there are only two players. Start with some arbitrary allocation (S_1, S_2) . The algorithm queries the XOS oracle of player 1 for the maximizing clause a of S_1 . Player 1 bids $b_1(j) = a(j)$ for each $j \in S_1$ and $b_1(j) = 0$ otherwise. Player 2 now calculates his demand S'_2 at price $b_1(j)$ for each item j and is allocated S'_2 . The algorithm continues similarly: we query the XOS oracle of player 2 for the maximizing clause a' of S'_2 . Player 2 bids $b_2(j) = a'(j)$ for each $j \in S_2$ and $b_2(j) = 0$ otherwise. Player 1 now calculates his demand S'_1 at price $b_2(j)$ for each item j and is allocated S'_1 .

²The single related exception we are aware of is [7] which shows that a Bayesian equilibrium is hard to find if the distributions are correlated.

What we have here is essentially two players that are both following best reply strategies. Will this process ever end? First, it is not hard to prove that at every step the sum of prices (i.e., winning bids) goes up. Since the sum of prices of an allocation is at most the welfare of an allocation (by the definition of an XOS valuation) and since the number of possible bundles is finite, the process ends after finitely many steps. The proof that the algorithm ends with an allocation that is a two-approximation essentially follows from the earlier algorithm of [8] that is identical to the exponential-time algorithm above except that the algorithm of [8] ends after each player responds exactly once.

As noted above, the dynamic may take exponentially many steps to end. However, we observe that the dynamic is actually not completely well defined as there are many possible different XOS oracles for a single valuation v , each may lead to a different path; some of the paths may potentially end after polynomially many steps.

To see that there are many possible different XOS oracles for a single valuation v , consider a submodular valuation v (recall that every submodular valuation is also XOS). It is known [8] that the following algorithm finds a maximizing clause a of a bundle S : arbitrarily order the items, and rename them for convenience to $1, 2, \dots, |S|$. Now for every $j \in S$, let $a(j) = v(j|\{1, \dots, j-1\})$. Observe that different orderings of the items result in different maximizing clauses.

This observation may seem useless, but in fact our positive results for submodular bidders are all based on it. More generally, our algorithms can be seen as a variant of the best-reply algorithm of [6], where instead of best-reply strategies we use better-reply strategies.³ We show that specific implementations of the XOS oracle can guarantee a fast termination of the better-reply algorithm. (The implementation of the XOS oracle also depends on the valuations of the other players.)

We are unable to extend our algorithms for submodular bidders to XOS bidders, and we also do not know how to prove that finding an equilibrium with XOS bidders requires exponential communication. However, we do make the observation that all existing algorithms for general XOS valuations that make use of XOS oracles (including pure approximation algorithms that do not take incentives issues into account, e.g., [9]) work with *any* XOS oracle and do not assume a specific implementation. For this kind of algorithms we can prove an impossibility result.

Definition: A no-overbidding equilibrium (S_1, S_2) is called *traditional* with respect to some XOS oracles \mathcal{O}_{v_1} and \mathcal{O}_{v_2} if for each bidder i and item j , if $j \in S_i$ then $b_i(j)$ equals the price of j in $\mathcal{O}_{v_i}(S_i)$, and if $j \notin S_i$ then $b_i(j) = 0$.

In other words, an equilibrium is called traditional if the prices are consistent with the XOS oracles of the players. Notice that the equilibrium obtained by the algorithm of [6] is indeed traditional.

Theorem: Let A be a deterministic algorithm that always produces a traditional equilibrium with respect to some XOS oracles \mathcal{O}_{v_i} . Suppose that A is allowed to make only XOS queries to the oracles \mathcal{O}_{v_i} 's in addition to demand and value queries. Then, A makes an exponential number of queries.

The proof is inspired by a proof of [20] for the hardness of finding an equilibrium in games using only queries that are analogous to value queries. Our proof differs in several technical aspects: first, it holds also for the stronger and more complicated demand queries and not just value queries (for that we use techniques from [2]). Second, the additional structure of our setting implies a more subtle construction, e.g., we have to prove a new isoperimetric inequality for odd graphs.

³A strategy is a *better reply* for a player if it increases his utility over the current strategy that he is playing (but does not necessarily maximize his profit).

Organization

In Section 2, we present algorithms that, given (approximately) efficient allocations, return equilibria in simultaneous second-price auctions with weakly better social welfare (than that of the input), for submodular valuations. We bound the algorithms’ running time for several scenarios. In Section 3, we present a communication complexity lower bound showing that finding or deciding the existence of equilibrium cannot be done with polynomial communication for subadditive valuations. In Section 4, we present our impossibility result for XOS valuations, showing that existing types of algorithms for finding equilibria for these valuations can be forced to run for exponential time. Finally, since our algorithms are essentially local search ones, in Appendix A we discuss the connection between the PLS completeness of finding an allocation that locally maximizing the welfare and finding an equilibrium.

2 Algorithms for Bidders with Submodular Valuations

In this section we provide algorithms that find good equilibria when the valuations are submodular. Our goal is to take an allocation and convert it with “little” computational overhead to an equilibrium that has at least the same welfare as the initial allocation. We first present a generic process that will be our main workhorse, after which we show that different implementations of this process allow us to compute an equilibrium with little computational overhead for several different settings. In particular, this will allow us to prove the main result of this section: one can take any allocation of combinatorial auctions with a constant number of submodular bidders and find an equilibrium with at least the welfare of the given allocation, using only polynomially many value queries (Subsection 2.4). We note that a similar auction was considered in [12] in a slightly different context.

The Iterative Stealing Procedure

1. Start with an arbitrary allocation (S_1, S_2, \dots, S_n) .
2. Each bidder i arbitrarily orders the items in S_i and sets, for every $j \in S_i$, $b_i(j) = v_i(j|\{1, \dots, j-1\} \cap S_i)$. For $j \notin S_i$ let $b_i(j) = 0$.
3. If there exist some players i, i' and item $j \in S_{i'}$ such that $v_i(j|S_i) > b_{i'}(j)$ then let $S_{i'} = S_{i'} - \{j\}$ and $S_i = S_i + \{j\}$. We say that i *steals* the item from i' . Return to Step (2). If there are no such players i, i' and no item j , the process ends.

Remark 2.1. *It is interesting to note the relationship of the iterative stealing procedure to the exponential-time algorithm of [6] (which was described in the introduction). For simplicity assume that there are only two bidders. Step 2 can be essentially seen as a call to an XOS oracle [8]. In the XOS algorithm each bidder “steals” from the other bidders the set of items that maximizes his profit, i.e., best-responds given the bids of the others. In contrast, in our procedure whenever a player steals an item he essentially plays a “better reply”, i.e. plays a strategy better than his current one. To see that, observe that since the valuations are submodular, if there exists a set of items that a player can steal to maximize the profit, then there exists a single item that the player can steal to improve his profit.*

To give a full specification of the process, one has to specify how each player orders his items (Step 2) and which player steals which item if there are multiple possible steals in Step 3. The

implementation of Step 3 will turn out to be less important to us, and we will focus on considering different implementations of Step 2. Before this, we show that stealing can only improve the welfare:

Claim 2.2. Every steal increases the welfare.

Proof. Consider bidder i stealing item j from bidder i' . Let (S_1, \dots, S_n) be the allocation before the steal and (S'_1, \dots, S'_n) be the allocation after the steal. We show that $v_i(S_i) + v_i(S_{i'}) > v_i(S'_i) + v_i(S'_{i'})$ which suffices to prove the claim since the allocation of bidders other than i, i' did not change. To see this, observe that

$$v_i(S'_i) + v_i(S'_{i'}) = v_i(S_i) + v_i(j|S_i) + v_i(S_{i'}) - v_i(j|S_{i'} - \{j\})$$

Observe that by submodularity $v_{i'}(j|S_{i'} - \{j\}) \leq b_{i'}(j)$ (since $b_{i'}(j)$ is the marginal value of j given some subset of $S_{i'} - \{j\}$). On the other hand, since bidder i steals item j , $v_i(j|S_i) > b_{i'}(j)$, and therefore $v_i(j|S_i) > v_{i'}(j|S_{i'} - \{j\})$. This finishes the proof. \square

2.1 Specific Implementations

We first show that the iterative stealing procedure ends in pseudo-polynomial time for any implementation of Step (2). We then show two specific implementations of Step (2): one showing that if all valuations are budget additive then regardless of the initial allocation the procedure terminates in polynomial time, and the other showing that the number of steals is at most m^n . We get that:

Theorem 2.3. *There exist implementations of the iterative stealing procedure such that:*

- *If there exists an α -approximation algorithm for combinatorial auctions with submodular bidders then a no-overbidding equilibrium with the same approximation guarantee can be found by running the approximation algorithm and then an additional pseudo-polynomial number of value queries.*
- *If there exists an α -approximation algorithm for some subclass of combinatorial auctions with budget additive bidders then a no-overbidding equilibrium with the same approximation guarantee can be found by running the approximation algorithm and then an additional polynomial number of value queries.*
- *If there exists an α -approximation algorithm for combinatorial auctions with submodular bidders then a no-overbidding equilibrium with the same approximation guarantee can be found by running the approximation algorithm and then an additional $\text{poly}(m, n) \cdot m^n$ value queries.*

We start with proving the first bullet, and the next two bullets will be proven in the next two subsections. Notice that it is clear that the number of value queries that the iterative stealing procedure makes is $\text{poly}(m, n) \cdot (\text{number of steals})$.

2.2 Pseudo-Polynomial Convergence Time

In this subsection we assume that all valuations are rational numbers.

Claim 2.4. Let $v_{max} = \max_i v_i(M)$. Let Δ be such that $v_i(j|S)$ is a multiple of Δ , for every bidder i , item j , and subset S . The number of steals that the iterative stealing procedure makes is at most $n \cdot \frac{v_{max}}{\Delta}$.

Proof. By Claim 2.2 after every step the welfare increases by at least Δ . An obvious upper bound on the welfare is $n \cdot \frac{v_{max}}{\Delta}$. Thus the procedure terminates after that number of steals in every implementation. \square

Next we show a specific implementation that guarantees a slightly better bound than this last result.

Claim 2.5. For each bidder i and item j , let $\Delta_j^i = |\{v_i(j|S)\}_{S \in 2^M}|$. There exists an implementation of the iterative stealing procedure that terminates after at most $\sum_{i \in N} \sum_{j \in M} \Delta_j^i$.

Proof. We start with some allocation (S_1, \dots, S_n) . Each bidder i will maintain some order $>_i$ on the items according to which Step 2 is implemented. After each steal we possibly update $>_i$. We begin with some order $>_i$ such that for each bidder i it holds that each item $j \in S_i$ precedes every item $j' \notin S_i$. The internal order of items within these two groups is arbitrary. When bidder i steals item j from bidder i' we change $>_i$ and $>_{i'}$ so that j is the last item in both $>_i$ (within the new S_i) and in $>_{i'}$ (among all items). The order of the rest of the items remains the same.

For every item j let the price of item j be $p_j = \max_i b_i(j)$. An easy observation is that using this implementation the price p_j of an item can only go up. Consider bidder i that holds item j . First, right after bidder i steals item j from some other bidder p_j can only go up since the price is simply bidder i 's marginal value for j . If another item j' was stolen from i then the price does not change if $j >_i j'$ and cannot decrease by submodularity if $j >_i j'$. If bidder i steals another item j' then j' is the smallest item in $>_i$ and the price of j did not change as well.

Using the fact that the price of an item can only increase it is easy to get a bound on the number of steals. Recall that after bidder i steals item j the price p_j increases to $v_i(j|S_i)$ and any other price $p_{j'}$ cannot decrease. Thus bidder i may steal each item j at most Δ_j^i times. Consequently, the total number of steals is at most $\sum_{i \in N} \sum_{j \in M} \Delta_j^i$. \square

2.3 Budget Additive Bidders

Recall that a valuation v is budget additive if there exists some b such that for every S we have that $v(S) = \min(b, \sum_{j \in S} v(\{j\}))$. This b is known as the *budget* for the valuation. In the sequel we will denote by b_i the budget of bidder i .

The procedure for budget additive bidders is similar to the iterative stealing procedure presented earlier. We start with some allocation (S_1, \dots, S_n) . Each bidder i will maintain some order $>_i$ on the items according to which Step 2 is implemented. After each steal we will possibly update $>_i$. We begin with some order $>_i$ such that for each bidder i it holds that each item $j \in S_i$ precedes every item $j' \notin S_i$. The internal order of items within these two groups is arbitrary. When bidder i steals item j from bidder i' we change $>_i$ and $>_{i'}$ so that j is the last item in both $>_i$ (within the new S_i) and $>_{i'}$ (among all items). The order of the rest of the items remains the same. For every item j let the price of item j be $p_j = \max_i b_i(j)$. Similarly to the proof of Claim 2.5, using this implementation the price of an item can only go up.

Given some allocation (S_1, \dots, S_n) we call an item $j \in S_i$ *loose* for bidder i if $p_j < v_i(\{j\})$; otherwise the item is called *tight*. A loose item j is *strongly loose* if $p_j = 0$, otherwise it is *weakly loose*. We need a few easy claims:

Claim 2.6. For any allocation, every bidder i has at most one weakly loose item.

Proof. Let j be a loose item and T be the set of tight items held by bidder i . We know that j is weakly loose and therefore $v_i(\{j\}) + T = b_i$. Thus the marginal value of any additional item is 0, and therefore any other item $j' \notin \{j\} + T$ is strongly loose. \square

Claim 2.7. If item j that is tight for bidder i (given the current allocation) is stolen then i will not steal item j later.

Proof. After item j is stolen by i' , the new price of j is greater than $v_i(\{j\})$. Since by submodularity for every bundle S , $v_i(\{j\}) \geq v_i(j|S)$, and since the price of j only goes up, i never steals item j from any other bidder later in the algorithm. \square

Claim 2.8. If item j that is weakly loose for i is stolen then i will not steal item j again during the algorithm before some tight item $j' \in S_i$ was stolen from him.

Proof. The marginal value of j for i does not increase unless some item in $j' \in S_i$, $j' >_i j$, is stolen from bidder i . Observe that since j is weakly loose, every other item j' preceding j by $>_i$ must be tight. \square

Claim 2.9. If item j that is strongly loose for i is stolen then i will not steal item j again during the algorithm before some tight or weakly loose item $j' \in S_i$ was stolen from bidder i .

Proof. The marginal value of j for i does not increase unless some item in $j' \in S_i$, $j' >_i j$, is stolen from him. Observe that since j is strongly loose, for the marginal value to change, item j' cannot be strongly loose as well. \square

We are now ready to give an upper bound on the number of steals. An easy corollary of Claim 2.7 is that there are at most $n \cdot m$ steals of tight items. The next claim shows that at least one of every $O(nm)$ consecutive steals must be a steal of a tight item, and this gives us an upper bound of $(n^2 \cdot m^2)$ on the total number of steals.

Claim 2.10. Every sequence of $O(nm)$ steals must contain at least one steal of an item that is tight.

Proof. Since the price of each item only increases, a strongly loose item, once stolen, will never be a strongly loose item again. Therefore altogether there may be at most m steals of strongly loose items, and we will consider from now on only steals of weakly loose and tight items. Among every $nm + 1$ such steals, there is at least one bidder who lost the same item twice; if in both occasions the item was weakly loose, then by Claim 2.8, the bidder has stolen this item back only because she has lost some tight item. This shows that at least one of these $nm + 1$ steals was one of a tight item. \square

2.4 Constant Number of Submodular Bidders

We now turn to the main result of this section, which concerns combinatorial auctions with a constant number of bidders — or, more generally, those in which the number of bidders competing for any single item is bounded above by a constant. To get some intuition, it will be instructive to consider first the case of two bidders. Let (S_1, S_2) be some allocation. Set bids $b_i(j)$ as in the process using some arbitrary order. If we happen to arrive at equilibrium, then we are already done. Else, call bidder i a top competitor of item j if $v_i(\{j\}) \geq v_{i'}(\{j\})$, where i' is the other bidder. If there is some bidder i and item $j \in S_i$ such that i is a top competitor for j , then if we use an order where j is first, player i' does not want to steal item j from player i . Hence we can imagine that player i' 's valuation is $v_{i'}(\cdot - \{j\})$ (which is still submodular), and look for an equilibrium using the set of items $M - \{j\}$.

Suppose now that all items are not held by their top competitors. Since we are not at equilibrium, some player i can steal some item j from the other player i' . Observe that i is a top competitor for item j (since either bidder i' or bidder i is the top competitor of j and we assumed that i' is not the top competitor). Hence, after stealing the item we are back to the previous case: bidder i puts the item first in the order, and we recurse again. We will show the process ends after at most m steals with an equilibrium.

We start the formal description of the implementation with a few definitions. Fix valuations v_1, \dots, v_n . For every item j let its *competitors* be $C_j = \{i | v_i(\{j\}) > 0\}$. Bidder i is a *top competitor* for item j if $v_i(\{j\}) \geq v_{i'}(\{j\})$, for every other bidder i' . We say that an instance is *t-restricted* if for every item j we have that $|C_j| \leq t$. Notice that every instance is n -restricted. We will use $v_{i,S}$ to denote the valuation $v_i(\cdot | S)$, and will use the terms “prices” and “bids” interchangeably when there is no confusion.

We now give a recursive implementation of the process for t -restricted instances with n bidders and m items. Denote the maximum number of steals that this procedure makes by $f_m(t)$. We will show that $f_m(t) \leq 1 + f_m(t-1) + f_{m-1}(t)$, $f_m(2) \leq m$ for every m , and that $f_1(t) \leq 1$ for every t . Hence, $f_m(n) \leq m^{n-1}$, as needed.⁴ We now describe the process, TOP-STEAL($(v_1, \dots, v_n), (S_1, \dots, S_n), M, t$). The procedure takes in a profile of valuations (v_1, \dots, v_n) , an initial allocation (S_1, \dots, S_n) on a set of items M where each item is t -restricted, and returns an allocation (S'_1, \dots, S'_n) and bids $\{b_i(j)\}$ for each bidder i and item j . We will show that the procedure produces a no-overbidding equilibrium whose social welfare is no worse than that of the initial allocation. The implementation itself consists of the definition of TOP-STEAL in three disjoint cases.

In each of the procedures, we assume that no item belongs to a non-competitor in the initial allocation. This is without loss of generality because at the beginning of each procedure we can check this condition for all items and move items from a non-competitor to a competitor, which only increases social welfare. Items that no one competes for have no effect on equilibria or social welfare and are ignored.

Procedure TOP-STEAL($(v_1, \dots, v_n), (S_1, \dots, S_n), M, t$) (**for** $|M| = 1$)

If the item does not belong to a top competitor, then one top competitor steals the item.

Procedure TOP-STEAL($(v_1, \dots, v_n), (S_1, \dots, S_n), M, t$) (**for** $|M| > 1$, $t = 2$)

1. *If there exist an item j and a bidder i such that $j \in S_i$ and i is a top competitor for j :*
 - (a) Let (S'_1, \dots, S'_n) and $\{b_{i'}(j')\}_{i' \in N, j' \in M}$ be the allocation and prices returned by TOP-STEAL($(v_{i, \{j\}}, v_{-i}), (S_i \setminus \{j\}, S_{-i}), M \setminus \{j\}, t$).
 - (b) Let $b_{i'}(j') = b'_{i'}(j')$ for every $j' \neq j$. Let $b_i(j) = v_i(\{j\})$ and $b_{i'}(j) = 0$ for each $i' \neq i$. Return $(S'_i \cup \{j\}, S'_{-i})$ and the prices $\{b_i(j)\}_{i \in N, j \in M}$.
2. *Else, there is no item j and bidder i such that $j \in S_i$ and i is a top competitor for j :*
 - (a) Each bidder i arbitrarily orders his items and sets, for every $j \in S_i$, $b_i(j) = v_i(j | \{1, \dots, j-1\} \cap S_i)$. For $j \notin S_i$, set $b_i(j) = 0$.
 - (b) If (S_1, \dots, S_n) and the prices $\{b_i(j)\}_{i \in N, j \in M}$ are an equilibrium then return this allocation and prices.
 - (c) Otherwise, find some bidder i that can steal some item j from bidder i' , such that i is a top competitor for j .
 - (d) Define (S'_1, \dots, S'_n) to be the allocation after i steals j from i' ; thus $S'_k = S_k$ for $k \neq i, i'$, $S'_i = S_i \cup \{j\}$, $S'_{i'} = S_{i'} \setminus \{j\}$.

⁴Formally, one can prove by induction that $f_m(t) \leq \binom{m+t-1}{t-1} - 1$.

(e) Return TOP-STEAL($(v_1, \dots, v_n), (S'_1, \dots, S'_n), M, t$).

Procedure TOP-STEAL($(v_1, \dots, v_n), (S_1, \dots, S_n), M, t$) (**for** $|M| > 1$, **arbitrary** $t > 2$)

1. *If there exist an item j and a bidder i such that $j \in S_i$ and i is a top competitor for j :*
 - (a) Let (S'_1, \dots, S'_n) and $\{b_{i'}(j')\}_{i' \in N, j' \in M}$ be the allocation and prices returned by TOP-STEAL($(v_{i, \{j\}}, v_{-i}), (S_i \setminus \{j\}, S_{-i}), M \setminus \{j\}, t$).
 - (b) Let $b_{i'}(j') = b'_{i'}(j')$ for every $j' \neq j$. Let $b_i(j) = v_i(\{j\})$ and $b_{i'}(j) = 0$ for each $i' \neq i$. Return $(S'_i \cup \{j\}, S'_{-i})$ and the prices $\{b_i(j)\}_{i \in N, j \in M}$.
2. *Else, there is no item j and bidder i such that $j \in S_i$ and i is a top competitor for j :*
 - (a) For each bidder i , let T_i be the set of all items for which i is a top competitor.
 - (b) Let (S'_1, \dots, S'_n) and $\{b_i(j)\}_{i \in N, j \in M}$ be the allocation and prices returned by TOP-STEAL($(v_{1, |T_1|}, \dots, v_{n, |T_n|}), (S_1, \dots, S_n), M, t - 1$).
 - (c) If (S'_1, \dots, S'_n) and the prices $\{b_i(j)\}_{i \in N, j \in M}$ are an equilibrium with respect to v_1, \dots, v_n then return this allocation and prices.
 - (d) Otherwise, find some bidder i that can steal some item j from bidder i' , such that i is a top competitor for j .
 - (e) Define (S'_1, \dots, S'_n) to be the allocation after i steals j from i' ; thus $S'_k = S_k$ for $k \neq i, i'$, $S'_i = S_i \cup \{j\}$, $S'_{i'} = S_{i'} \setminus \{j\}$.
 - (f) Return TOP-STEAL($(v_1, \dots, v_n), (S'_1, \dots, S'_n), M, t$)

We now analyze the running time of TOP-STEAL. We start with the base case:

Claim 2.11. When $|M| = 1$, TOP-STEAL($(v_1, \dots, v_n), (S_1, \dots, S_n), M, t$) reaches a no-overbidding equilibrium after one steal (i.e., $f_1(t) \leq 1$).

Proof. When the item j reaches a top competitor i no other bidder i' wants to steal it since $v_i(j) \geq v_{i'}(j)$ by the definition of a top competitor. \square

The following is the key claim in this proof:

Claim 2.12. Let j be an item and let i be a top competitor for this item. Let $v'_i(S) = v_i(S|j)$. Let $S' = (S'_1, \dots, S'_n)$ be a no-overbidding equilibrium with respect to the valuations (v'_i, v_{-i}) where bidder i' bid for item j' is $b'_{i'}(j')$. Then $S = (S'_i \cup \{j\}, S'_{-i})$ is an equilibrium with no overbidding with respect to v_1, \dots, v_n when for item $j' \neq j$ each bidder i' bids $b_{i'}(j') = b'_{i'}(j')$, and for item j bidder i bids $b_i(j) = v_i(j)$ and every other bidder i' bids $b_{i'}(j) = 0$.

Proof. We first show that there is no overbidding: this is certainly true for every bidder $i' \neq i$ since the allocation and bidding of i' did not change from the first equilibrium S' in which there was no overbidding. In addition, bidder i does not overbid either. This is certainly true for every bundle S where $j \notin S$ (since $v'_i(S) = v_i(S)$). For $j \in S$ we have that:

$$v_i(S) = v_i(\{j\}) + v_i(S \setminus \{j\}|j) = v_i(\{j\}) + v'_i(S \setminus \{j\}) \geq b_i(j) + \sum_{j' \neq j, j' \in S} b'_{i'}(j') = \sum_{j' \in S} b_{i'}(j')$$

To see that S is an equilibrium, observe that for every item $j' \notin S'_i \cup \{j\}$ it holds that $v_i(j'|S'_i \cup \{j\}) = v'_i(j'|S'_i)$. Therefore bidder i will not steal any item from any other bidder i' . Similarly, since the valuation of every other bidder $i' \neq i$ did not change and the prices of all items but item j are the same, the only possible steal may involve some bidder i' stealing item j from bidder i . However, since i is a top competitor for j we have $b_i(j) = v_i(\{j\}) \geq v'_i(\{j\}) \geq v'_i(j|S'_i)$ and therefore this steal is not profitable for bidder i' . \square

Lemma 2.13. ⁵ $\text{TOP-STEAL}((v_1, \dots, v_n), (S_1, \dots, S_n), M, 2)$ returns a no-overbidding equilibrium after $f_m(2) \leq m$ steals.

Proof. We prove this by induction on the number of items m , where the case of $m = 1$ is proved in Claim 2.11. We assume correctness for $m - 1$ and prove for m . We divide into two cases:

Case 1: *There exist an item j and a bidder i such that $j \in S_i$ and i is a top competitor for j :* the fact that this is a no-overbidding equilibrium follows since by the induction hypothesis $\text{TOP-STEAL}((v_{i, \{j\}}, v_{-i}), (S_i \setminus \{j\}, S_{-i}), M \setminus \{j\}, 2)$ returns a no-overbidding equilibrium with respect to $(v_{i, \{j\}}, v_{-i})$ and then we can apply Claim 2.12. Notice that by the induction hypothesis, $\text{TOP-STEAL}((v_{i, \{j\}}, v_{-i}), (S_i \setminus \{j\}, S_{-i}), M \setminus \{j\}, 2)$ makes at most $m - 1 < m$ steals.

Case 2: *There is no item j and bidder i such that $j \in S_i$ and i is a top competitor for j :* if we reached an equilibrium in the first step we are obviously done. Otherwise, some bidder i steals item j from another bidder i' . Notice that since $|C_j| = 2$ and since $i' \in C_j$ by assumption we have that i is a top competitor of j . The recursive call to TOP-STEAL in Step 2(e) will satisfy Case 1, so by the first part of this proof the recursive call will reach a no-overbidding equilibrium after at most $m - 1$ steals. Combining this with the one steal in Step 2(d) we have at most m steals in total, as claimed. \square

Lemma 2.14. $\text{TOP-STEAL}((v_1, \dots, v_n), (S_1, \dots, S_n), M, t)$ returns a no-overbidding equilibrium after $f_m(t) \leq 1 + f_m(t - 1) + f_{m-1}(t)$ steals.

Proof. We again prove using induction on the number of items and on the size of t where the base case $m = 1$ was proven in Claim 2.11 and the case of arbitrary m and $t = 2$ was proven in Lemma 2.13. We divide again into two cases:

Case 1. *There exist an item j and a bidder i such that $j \in S_i$ and i is a top competitor for j :* the fact that this is a no-overbidding equilibrium follows since by the induction hypothesis, $\text{TOP-STEAL}((v_{i, \{j\}}, v_{-i}), (S_i \setminus \{j\}, S_{-i}), M \setminus \{j\}, t)$ returns a no-overbidding equilibrium with respect to $(v_{i, \{j\}}, v_{-i})$ and then we can apply Claim 2.12. Notice that by the induction hypothesis, $\text{TOP-STEAL}((v_{i, \{j\}}, v_{-i}), (S_i \setminus \{j\}, S_{-i}), M \setminus \{j\}, t)$ makes at most $f_{m-1}(t)$ steals.

Case 2. *There is no item j and bidder i such that $j \in S_i$ and i is a top competitor for j :* first, observe that if v_1, \dots, v_n is a t -restricted instance, then $v_{1, |T_1}, \dots, v_{n, |T_n}$ is a $(t - 1)$ -restricted instance. If running $\text{TOP-STEAL}()$ on this $(t - 1)$ -restricted instance results in an equilibrium S' , the lemma is proved and we made only at most $f_m(t - 1)$ steals.

Otherwise, we claim that S' is not an equilibrium since some bidder i can steal some item j from bidder i' , and i is a top competitor for j . This is true since for every bidder i and item j we have $v_i(j|S'_i) = v'_i(j|S'_i)$ unless i is a top competitor for j . The recursive call to TOP-STEAL in Step 2(f) will satisfy Case 1, so by the first part of this proof the recursive call will reach a

⁵One could in principle incorporate the next two lemmas into one by starting the induction from $t = 1$. We prefer to keep this lemma and explicitly prove the case of $t = 2$ since it provides an explicit simple proof for the important setting of two bidders.

no-overbidding equilibrium after at most $f_{m-1}(t)$ steals. Combining this with the $f_m(t-1)$ steals in Step 2(e) and the one steal in Step 2(b) we have at most $1 + f_m(t-1) + f_{m-1}(t)$ steals in total, as claimed. \square

3 An Impossibility Result for Subadditive Valuations

Unlike combinatorial auctions with submodular valuations, with subadditive valuations we have no hope to efficiently find an equilibrium.

Theorem 3.1. *It requires exponential communication to determine whether there is a no-overbidding equilibrium in an instance of a combinatorial auction with subadditive bidders. Furthermore, even in instances where an equilibrium is guaranteed to exist, exponential communication is needed to find it. The results hold even if there are only two players and even if the value of every bundle is in $\{0, 1, 2\}$.*

We prove only the first statement, and the second one is an easy corollary of the first. We prove the theorem for two players; extending the theorem for more than two players is straightforward by adding bidders with valuations that are identically zero. For the proof we use a set-pair system:

Definition 3.2. A set-pair system $\mathcal{S} = \{(S_1^r, S_2^r)\}_r$ is *good* if the following holds:

1. Every $(S_1^r, S_2^r) \in \mathcal{S}$ it holds that $|S_1^r| = |S_2^r| = \frac{m}{4}$ and $S_1^r \cap S_2^r = \emptyset$.
2. For every $r \neq l$ we have that $0 < |S_1^r \cap S_2^l| \leq \frac{m}{8}$.

It is well known that good set-pair systems that contain exponentially many allocations exist (e.g., [19]). Given a large good set-pair system, we reduce the disjointness problem to the problem of finding a no-overbidding equilibrium in the simultaneous second-price auction involving valuations to be defined below. The disjointness problem involves two players, each player i having a string $A^i = \{0, 1\}^t$, and the problem is to determine whether there is some k such that $A_k^1 = A_k^2 = 1$. It is known that deciding disjointness (or finding k if we are given that such k exists) requires $\Omega(t)$ bits of communication [15]. In our reduction, the size t will be equal to the size of the good set system, which is exponential in m .

We construct the following valuation functions ($i = 1, 2$) given bit vector $A^i \in \{0, 1\}^t$. As assumed by normalization, $v_i(\emptyset) = 0$; for $S \neq \emptyset$,

$$v_i(S) = \begin{cases} 2, & \text{if } |S| \geq \frac{3m}{4} + 1; \\ 2, & \text{if } \exists r \text{ such that } S_r^i \subseteq S \text{ and } A_r^i = 1; \\ 1, & \text{otherwise.} \end{cases}$$

It is straightforward to verify that these value functions are subadditive. The theorem is a corollary of the following two claims:

Claim 3.3. If there exists k such that $A_k^1 = A_k^2 = 1$ then there is a no-overbidding equilibrium.

Proof. Consider the following strategies: each bidder i bids some small $\epsilon > 0$ for every $j \in S_i^k$ and 0 otherwise. Denote the outcome allocation in the simultaneous second-price auction of these two strategies by (S_1, S_2) . We claim that these strategies form an equilibrium (observe that there is no overbidding in these strategies). First, observe that no item gets positive bids from both bidders simultaneously (by the properties of a good set-pair system) and that each positively bids on all items of some bundle that has a value of 2 for him. Thus, for every bidder i , $v_i(S_i) = 2$. Since $v_i(M) = 2$ as well and since the total payment of each bidder is 0, there is no strategy that gives bidder i a profit larger than 2 (regardless of the strategy of the other bidder). \square

Claim 3.4. If there is no k such that $A_k^1 = A_k^2 = 1$ then there is no equilibrium with no overbidding.

Proof. Suppose that there is some no-overbidding equilibrium, and denote by $b_i(j)$ the bid of bidder i for item j in that equilibrium. Let the allocation in the equilibrium be (S_1, S_2) . Observe that by the properties of a good set-pair system and the definition of the valuations it either holds that $v_1(S_1) \neq 2$ or that $v_2(S_2) \neq 2$. Thus, we may assume without loss of generality that $v_1(S_1) \leq 1$.

Let us call a set of items U *unprotected* if $v_1(U) = 2$ and $\sum_{j \in U} b_2(j) < 1$. Note that if U is unprotected and bidder 1 bids $b_2(j) + \epsilon$ for every $j \in U$ and 0 for every $j \notin U$, then for sufficiently small positive ϵ this satisfies no-overbidding⁶ (as the sum of all bids is less than 1) and strictly increases bidder 1's profit. Thus, existence of an unprotected set contradicts our hypothesis that (b_1, b_2) is not an equilibrium.

Let T , $|T| = m/4$, be such that $v_1(T) = 2$. Observe that by the properties of a good set-pair system it holds that $v_2(T) = 1$. If $\sum_{j \in T} b_2(j) < 1$ then T is an unprotected set, so suppose henceforth that $\sum_{j \in T} b_2(j) = 1$. For every $j \notin T$ we have $v_2(T + \{j\}) = 1$ by the properties of a good set system, since $v_1(T) = 2$ and $T + \{j\}$ has only one element that does not belong to T . Now the assumption that bidder 2 does not overbid implies $b_2(j) = 0$ for all $j \notin T$. Consequently $\sum_{j \in M} b_2(j) = \sum_{j \in T} b_2(j) = 1$. Now choose any j' such that $b_2(j') > 0$ and observe that $M - \{j'\}$ is an unprotected set. \square

4 XOS Valuations: An Impossibility Result for Traditional Algorithms

This section proves an impossibility result for XOS algorithms: we essentially show that every efficient algorithm for XOS valuations must use techniques very different from traditional algorithms for XOS valuations. Section 4.2 proves the impossibility result itself, but we start by getting a feel of it in Section 4.1, by negatively answering a related open question from [6]: does the XOS algorithm of [6] end in polynomial time if the valuations are submodular?

4.1 Exponential Convergence for the XOS Algorithm with Submodular Bidders

We now show that if the bidders have submodular valuations, then the algorithm of [6] may terminate after exponentially many steps, even when there are only two bidders. The key for this example is the XOS oracle: we choose the prices in a legitimate way according to the algorithm, but one that will lead to a long convergence path. The dependence of this lower bound on the implementation of the XOS oracle is unavoidable; as discussed earlier, the stealing procedure presented in Section 2 is in fact a special case of the algorithm of [6] that essentially differs only in the implementation of the XOS oracle.

4.1.1 The Example

Let m be odd and let $m' = \lfloor m/2 \rfloor$. We now present a family of submodular valuations that is parametrized by a family of positive constants k_S^i , for every bidder i and S such that $|S| = m' + 1$

⁶Strictly speaking, no-overbidding only requires that the strategies employed in the equilibria be no-overbidding, and allow deviations to be overbidding. In our argument, disallowing overbidding deviations only possibly enlarges the set of equilibria, and makes our lower bounds even stronger.

and some $\epsilon > 0$. Each member of this family is defined as follows:

$$v_i(S) = \begin{cases} |S| & \text{if } |S| \leq m'; \\ m' + \frac{1}{2} + k_S^i \cdot \epsilon & \text{if } |S| = m' + 1; \\ m' + 1 & \text{if } |S| \geq m' + 2. \end{cases}$$

Notice that v_i is indeed a submodular function as long as $k_S^i \cdot \epsilon < \frac{1}{2}$, which is easy to guarantee. Obviously, the key to getting an exponential path is choosing the k_S^i 's carefully. To simplify the presentation, we will choose the k_S^i 's as we proceed.

We start with an allocation (S_1, S_2) such that $|S_1| = m'$ and $|S_2| = m - m' = m' + 1$. Set $k_{S_1}^i = 0$ for $i = 1, 2$. Bidder 1 now bids $b_1(j) = 1$ for every item $j \in S_1$, except a single item $j' \in S_1$ where $b_1(j') = 1/2$. The bid on the rest of the items is 0. Notice that this is a valid XOS clause of S_1 , and hence a legitimate step in the algorithm. Now set $k_{S_2+j'}^2 = 1$. Notice that the best response of bidder 2 is to take $S_2 + j'$ (regardless of the choice of the other k_S^2 's).

The allocation now is $(S'_1, S'_2) = (S_1 - j', S_2 + j')$. Choose some $j'' \neq j'$ in S'_2 . Set $k_{S'_2}^2 = 2$. Let $b_2(j) = 1$ for every item $j \in S_2$ other than $j'' \in S'_2$ and $b_2(j'') = 1/2 + 2\epsilon$ for j'' . The bid on the rest of the items is 0 again. Observe that, as long as we set $k_{S'_1+j''}^1 = 3$, the best response of player 1 is to take $S'_1 + j''$. It should be clear that we can iterate this construction to produce an exponential path such that the two players keep exchanging a single item in each best response. To be specific, we can use a monotone Gray code [22] to explicitly construct such a path (we only use the numbers that have hamming weight between m' and $m' + 1$). The zeros in every codeword represent items allocated to player 1, and the ones represent items allocated to player 2.

4.2 The Impossibility Result

Recall that an XOS oracle for an XOS valuation v returns for each bundle S a maximizing clause for S in v . We refer to the coefficients of this maximizing clause as “prices.” In this section we denote an XOS oracle for a valuation v by \mathcal{O}_v . We now prove a lower bound on deterministic algorithms that find a *traditional* equilibrium. For simplicity, we restrict ourselves to the case where there are only two players (the impossibility result trivially extends to any number of players by adding players with valuations that are identically zero).

Definition 4.1. A no overbidding equilibrium (S_1, S_2) is called *traditional* with respect to some XOS oracles \mathcal{O}_{v_1} and \mathcal{O}_{v_2} if for each bidder i and item j , if $j \in S_i$ then $b_i(j)$ equals the price of j in $\mathcal{O}_{v_i}(S_i)$, and if $j \notin S_i$ then $b_i(j) = 0$.

Notice that the equilibrium obtained by the algorithm of [6] is indeed traditional.

Theorem 4.2. *Let A be a deterministic algorithm that always produces a traditional equilibrium with respect to some XOS oracles \mathcal{O}_{v_i} . Suppose that A is allowed to make only XOS queries to the oracles \mathcal{O}_{v_i} in addition to demand and value queries. Then, A makes an exponential number of queries in the worst case. The theorem holds even if there are only two bidders with identical valuations.*

The proof consists of three parts. First, we will present a family of XOS valuations (sensitive valuations) and give conditions for equilibrium if all valuations belong to this family. In the second part, we will show that demand queries to sensitive valuations are almost useless: specifically, every demand query can be simulated by polynomially many value queries. Finally, we prove that any algorithm for finding equilibrium when bidders have sensitive valuations that uses value queries makes exponentially many value queries in the worst case. We get our bound by combining it all together.

4.2.1 The proof part I: conditions for equilibrium

Let $m > 20$ be an odd integer and let $m' = \lfloor m/2 \rfloor$. Given non-negative k_S^i 's such that for every i and bundle S , $|S| = m' + 1$, we have $\frac{1}{4} > k_S^i > 0$, consider the valuation v_i that is the maximum of the following additive valuations (and hence XOS by definition):

1. There is an additive valuation C_j for every item j that gives a value of $m' - 20$ for j and value of 0 for every other item.
2. For every bundle S , $|S| = m'$, there is an additive valuation A_S that gives a value of 1 to every item $j \in S$ and value 0 for any other item.
3. For every bundle S , $|S| = m' + 1$, there is an additive valuation $M_{S,j}$, for one item $j \in S$. $M_{S,j}$ gives a value of $\frac{1}{4} + k_S^i$ for item j , a value of 1 for every item $j' \neq j$, $j' \in S$, and 0 for every item that is not in S .
4. For every bundle S , $|S| = m' + 10$, there is an additive valuation B_S that gives a value of $\frac{m'+1}{m'+10}$ for every item in S and 0 otherwise.

We call such valuations *sensitive* valuations. An XOS oracle for a sensitive valuation is called *standard* if for every bundle S it returns an additive valuation that is specified in the above definition. It will sometimes be easier to work with an explicit description of a sensitive valuation:

Proposition 4.3. *Let v'_i be a sensitive valuation defined by k_S^i 's. Let v_i be the following valuation:*

$$v_i(S) = \begin{cases} m' - 20 & \text{if } |S| \leq m' - 20, \\ |S| & \text{if } m' - 20 < |S| \leq m', \\ m' + \frac{1}{4} + k_S^i & \text{if } m' + 1 \leq |S| < m' + 10, \\ m' + 1 & \text{if } |S| \geq m' + 10. \end{cases}$$

For every S we have that $v_i(S) = v'_i(S)$.

Proof. We divide the analysis into cases, according to the size of S :

1. $|S| \leq m' - 20$: $C_j(S) = m' - 20$, while for every bundle S' and item j , $M_{S',j}(S) = A_{S'}(S) \leq |S| \leq m' - 20$ and $B_{S'}(S) = \frac{m'+1}{m'+10} \cdot |S| < m' - 20$.
2. $m' - 20 < |S| \leq m'$: $A_S(S) = |S|$. On the other hand, for every other bundle S' and item j' , $C_{j'}(S) = m' - 20$, $B_{S'}(S) = \frac{m'+1}{m'+10} \cdot |S| < |S|$, and $M_{S',j'}(S) \leq A_S(S)$.
3. $m' + 1 \leq |S| \leq m' + 9$: let $T \subseteq S$ be a maximum value bundle of size $m' + 1$. Observe that $v_i(T) = v'_i(T)$. There exists some $j \in T$ such that $M_{T,j}(S) = m' + \frac{1}{4} + k_S^i$. However, for every other bundle S' and item j' , $A_{S'}(S) \leq M_{S',j'}(S) \leq v_i(T)$, $C_{j'}(S) = m' - 20$ and $B_{S'}(S) \leq \frac{m'+1}{m'+10} \cdot (m' + 9) = m' + 1 - \frac{m'+1}{m'+10} \leq m' + \frac{1}{4} \leq v_i(S)$.
4. $|S| \geq m' + 10$: for some bundle $S' \subseteq S$, $|S'| = m' + 10$ we have that $B_{S'}(S) = m' + 1$, and this value is larger than the maximum value of every other additive valuation in the support of v'_i .

□

Definition 4.4. A bundle S of size $m' + 1$ and a corresponding XOS clause $M_{S,j}$ are called a *j -local maximum* of a valuation v with respect to an item $j \in S$ if $v(S) \geq v(M - S + \{j\})$.

The definition of j -local maximum and of $M_{S,j}$ implies that the XOS clause of S puts a weight strictly smaller than 1 on j , therefore, if $v(S) < v(M - S + \{j\})$, it may not be true that $S' = M - S + \{j\}$ is j -local maximum, because even though $v(S') > v(M - S' + \{j\}) = v(S)$, the XOS clause for S' may be $M_{S',j'}$ for an item $j' \neq j$. The reference to the XOS clause is important for traditional equilibria, in which a bidder allocated a j -local maximum necessarily bids strictly less than 1 on j .

Proposition 4.5. *Let v be a sensitive valuation and let (S_1, S_2) be a traditional equilibrium of two bidders with the same valuation v with respect to the standard XOS representation. Then, either $|S_1| = m'$, $|S_2| = m - m' = m' + 1$, and S_2 is a j -local maximum of v (for some item $j \in S_2$), or $|S_2| = m'$, $|S_1| = m - m'$, and S_1 is a j -local maximum of v (for some item $j \in S_2$).*

Proof. We divide again into several cases. We only consider cases where $|S_1| \leq m'$, otherwise we have that $|S_2| \leq m'$ and the proof is symmetric.

1. $|S_1| \leq m' - 20$: in this case, $v_1(S_1) = m' - 20$ and $|S_2| = m - |S_1| \geq m - m' + 20 > m' + 10$. Hence, $v(S_2) = m' + 1$. In particular, bidder 2 bids 0 on $m - m' - 10$ items (since we are at a *traditional* equilibrium and in the standard XOS representation at most $m' + 10$ items get a positive value). In this case bidder 1 is better off bidding $\epsilon > 0$ on each of these items and 0 on all other items, since his value will increase to at least $m - m' - 10$ and his payment will be 0.
2. $m' - 20 < |S_1| \leq m'$: here bidder 1 bids 1 on every item $j \in S_1$. Bidder 2 is using either a B_S clause or an $M_{S,j}$ clause. If bidder 2 uses a B_S clause then the price of every item in that clause is less than 1. Furthermore, $|S_2| \geq m' + 10$ and therefore $|S_1| \leq m - m' - 10$. This implies that if bidder 1 increases his bid to 1 for any single item in S_2 : his value will increase by 1 while his payment will increase by a smaller amount. Notice that this is still a no-overbidding strategy for bidder 1.

Hence, the only possibility is that bidder 2 uses an $M_{S,j}$ clause. Without loss of generality $|S_2| = m - m'$ (otherwise there are some items in S_2 that bidder 2 bids 0 on). Assume that S_2 is not a j -local maximum. Denote $S'_1 = M - S_2 + \{j\}$ and observe that $v(S'_1) > v(S_2)$. Notice that Bidder 1 can increase his profit by taking S'_1 : his profit in this case is $v(S'_1) - p_j = v(S'_1) - (v(S_2) - m') > m'$, and it is easy to see that $v(S_1) \leq m'$ since $|S_1| \leq m'$.

□

4.2.2 The proof part II: simulating demand queries by value queries

We now bound the number of value queries needed to simulate a demand query on a sensitive valuation by adapting techniques from [2]. Recall that the input for a demand query is a price for each item, and the output is a bundle that maximizes the profit. Since we show an impossibility result and since the values of all bundles of sizes other than $m' + 1$ are known, we may consider only demand queries that may return a bundle of size $m' + 1$.

Definition 4.6. Let dq be a demand query. A bundle S , $|S| = m' + 1$, is *covered* by dq if there exists some sensitive valuation v such that when querying v for dq a profit-maximizing bundle is S .

The main idea of this section is to show that for every dq the number of covered bundles is bounded:

Lemma 4.7. Consider some demand query dq . Let $\mathcal{S} = \{S \mid S \text{ is covered by } \text{dq}\}$. Then, $|\mathcal{S}| \leq \text{poly}(m)$.

Before proving Lemma 4.7, notice that it immediately implies that:

Corollary 4.8. Any demand query dq for a sensitive valuation v can be simulated by polynomially many value queries.

Proof. It suffices to only use value queries on all sets covered by dq , of which there are only polynomially many. Among these sets and all the others (whose values we already know), we return a profit-maximizing one. \square

Proof of Lemma 4.7

Let S be a bundle covered by dq , and let v be a sensitive valuation such that $\text{dq}(v) = S$. Denote the price of every item j in dq by p_j . Let $t \in \text{argmin}_j p_j$. Notice that $p_t \leq \frac{1}{2}$, otherwise any bundle of size m' is more profitable than S . Since the profit from $\{t\}$ is $m' - 20 - p_t$ and yet the demand query returns S it holds that:

$$\begin{aligned} m' + \frac{1}{4} + k_S - \sum_{j \in S} p_j &\geq m' - 20 - p_t \\ m' + \frac{1}{4} + k_S - \sum_{j \in S} p_j &\geq m' - 20 - \frac{1}{2} \\ 21 &> \sum_{j \in S} p_j \end{aligned} \tag{1}$$

Next we observe that there is no set $S' \subseteq M - S$, $|S'| = 9$, such that $\sum_{j \in S'} p_j < \frac{1}{2}$. This is true since if such S' exists then $S \cup S'$ is more profitable than S (observe that $v(S \cup S') = m' + 1$ since $|S \cup S'| = m' + 10$):

$$v(S \cup S') - \sum_{j \in S \cup S'} p_j > m' + 1 - \sum_{j \in S} p_j - \frac{1}{2} \geq v(S) - \sum_{j \in S} p_j$$

Therefore, for all items $j \in M - S$ but at most nine of them, it holds that $p_j \geq \frac{1}{18}$. Denote these expensive items by E and let $C = M - S - E$ (i.e., the set of cheap items).

The crux of the proof is that if some bundle G is covered by dq , then it cannot contain more than $18 \cdot 21 = 378$ items from E , otherwise $21 < \sum_{j \in G} p_j$, a contradiction to (1). Thus, each G covered by dq is composed of up to 378 items from E (of which there are at most $(|E| + 1)^{378}$ possible choices, where we use $|E| + 1$ and not just $|E|$ since less than 378 items might be chosen from E), up to nine items from C (of which the number of possible choices is 2^9), and additional items from S (we use the bound $(|S| + 1)^{387}$ since we have to choose up to $387 = 378 + 9$ items to exclude from S). A bound on the number of bundles that can be constructed this way — and hence also a bound on the number of bundles covered by dq — is therefore:

$$(|E| + 1)^{378} \cdot 2^9 \cdot (|S| + 1)^{387} \leq m^{378} \cdot 2^9 \cdot m^{387} < 1000 \cdot m^{765}$$

4.2.3 The proof part III: the power of value queries

We now show that exponentially many value queries are needed in order to find a j -local maximum of a sensitive valuation v . The adversary will construct the hard valuation v by following the

algorithm: whenever the algorithm queries some bundle S , $|S| = m' + 1$ the adversary will make sure that S is not a local maximum. Obviously, since v is sensitive there is no use for the algorithm to query bundles of size different from $m' + 1$, since the values of those bundles are known in advance. We will assume that whenever the algorithm makes a value query to some bundle S , it also makes an XOS query to the same bundle (this only makes our bound stronger). Specifically, we prove that:

Lemma 4.9. *Any deterministic algorithm for finding a j -local maximum in sensitive valuations makes at least $\frac{2^{0.75m'-1}}{m'}$ queries in the worst case.*

In the proof we use a graph G that is composed of $\binom{m}{m'+1}$ vertices, each of which is associated with a different bundle of size $m' + 1$. Two vertices S and S' of G are connected by an edge if and only if $S = M - S' + \{j\}$, for some $j \in S'$. This graph is known as the *odd graph* $O_{m'+1}$. We will refer to the subsets of M (of size $m' + 1$) and vertices of $O_{m'+1}$ interchangeably, and the meaning should be clear, for example, when we say a set S is a neighbor of another set S' . We will need the following isoperimetric inequality on odd graphs (this inductive proof is inspired by [13]):

Proposition 4.10 (Isoperimetric inequalities for odd graphs). *Let $O_n = (V, E)$ be an odd graph. Then, for any $S \subseteq V$, $k = |S|$, at least $(n - \frac{4}{3} \log k)k$ edges have exactly one end in S . As a corollary, since O_n is n -regular, the number of neighbors of S is at least $(n - \frac{4}{3} \log k) \frac{k}{n}$.*

Proof. We use the following binary representation of the vertices of O_n : the t -th index of $v \in V$ is 1 if and only if v corresponds to a set that contains the t -th element.

Let $\mathcal{E}(k)$ denote the maximum possible number of edges inside a set $S \subseteq V$ of size k . By the regularity of O_n , the number of edges having exactly one end in S is at least $nk - 2\mathcal{E}(k)$. It therefore suffices to show $\mathcal{E}(k) \leq \frac{2k \log_2 k}{3}$, and we prove this by induction on k . For the base case of $k = 2$ observe that $\mathcal{E}(2) = 1$.

We now assume correctness for $k - 1$ and prove for k . Consider a set S of size k . We claim that there exist two distinct indices x_1, x_2 such that for at most two-thirds of the elements $v \in S$ it holds that $v_{x_1} \neq v_{x_2}$. To see this, consider some v and two randomly chosen distinct indices x'_1 and x'_2 . The probability that $x'_1 = 1$ is exactly $\frac{n}{2n-1}$ since the Hamming weight of the representation of every vertex is exactly n . The conditional probability that $x'_2 = 0$ given that $x'_1 = 1$ is $\frac{n-1}{2n-2} = \frac{1}{2}$. The probability of the event $x'_1 = 0, x'_2 = 1$ is the same, by symmetry, and therefore the probability that $v_{x'_1} \neq v_{x'_2}$ is $2 \cdot \frac{n}{2n-1} \cdot \frac{1}{2} \leq \frac{2}{3}$. By linearity of expectation, for randomly chosen x'_1 and x'_2 , the expected number of vertices $v \in S$ such that $v_{x'_1} \neq v_{x'_2}$ is at most $\frac{2|S|}{3}$. Hence, there exist two indices x_1, x_2 such that at most two-thirds of the elements in S have exactly one 1 at the two coordinates. Call this set of elements S_0 and let $S_1 = S - S_0$. Let $\ell = |S_0| \leq \frac{2k}{3}$.

The edges in S are those inside S_0 , those inside S_1 , and the edges between S_0 and S_1 . Crucially, there are at most ℓ edges between S_0 and S_1 : consider some $v \in S_0$ and assume without loss of generality that $v_{x_1} = 0$ and $v_{x_2} = 1$. The number of neighbors of v in S is at most 1 since v has no neighbors whose coordinates at x_1 and x_2 are both 0, and at most one neighbor whose coordinates at x_1 and x_2 are both 1. We have:

$$\begin{aligned} \mathcal{E}(k) &\leq \max_{1 \leq \ell \leq \frac{2k}{3}} (\ell + \mathcal{E}(\ell) + \mathcal{E}(k - \ell)) \leq \max_{1 \leq \ell \leq \frac{2k}{3}} \left[\ell + \frac{2\ell \log \ell}{3} + \frac{2(k - \ell) \log(k - \ell)}{3} \right] \\ &\leq \frac{2}{3} \left(k + \max_{1 \leq \ell \leq \frac{2k}{3}} [\ell \log \ell + (k - \ell) \log(k - \ell)] \right) \leq \frac{2}{3} \left(k + k \log \left(\frac{k}{2} \right) \right) = \frac{2k \log k}{3}. \end{aligned}$$

□

Denote by Q_i the i -th query that the algorithm makes (to a bundle of size $m' + 1$). For the first query Q_1 we will return a value of $m' + \frac{1}{4} + \epsilon$ and fix the appropriate clause $M_{Q_1, j}$, for some j . At this point the algorithm cannot be sure if Q_1 is a local maximum, since the value of $M - S + \{j\}$ is still undetermined. In general, when the algorithm makes the l -th query Q_l , tentatively we would like to set $v(Q_l) = m' + \frac{1}{4} + l \cdot \epsilon$, and set the XOS clause of Q_l to $M_{Q_l, j}$, where j is such that Q_l was not queried yet.

The problem with this is that all neighbors of Q_l may have been queried, so a local maximum may quickly be found by the algorithm. To fix this, whenever the algorithm queries Q_l , before determining $v(Q_l)$ and the maximizing clause, we first consider “small” patches of the graph that will be cut off by the removal of Q_l ; we will first fix the values for vertices in these patches: in each small connected component that is cut off, the values assigned is increasing from vertices furthest from Q_l to those closest. At least, we give the largest value to Q_l . We will mark these vertices as “colored”, and remove them from the graph, together with Q_l , for future rounds. Importantly, the XOS clause $M_{Q_l, j}$ for Q_l is chosen such that j is in an uncolored neighbor of Q_l , and therefore the algorithm cannot know whether it has found a local maximum or not.

More formally, the process is as follows. Let $\epsilon > 0$ be such that $2^m \cdot \epsilon < \frac{1}{4}$.

1. Set $x = 1$. Let $\mathcal{C} = \mathcal{Q} = \emptyset$.
2. For each query Q_i :
 - (a) If $Q_i \in \mathcal{C}$, return the already fixed value and clause, and go for the next query. Otherwise let $\mathcal{Q} = \mathcal{Q} + \{Q_i\}$ and proceed.
 - (b) For each connected component CC of $G - \mathcal{C} - \mathcal{Q}$ of size less than $c = 2^{0.75m'-1}$:
 - i. Let D be the diameter of CC , let $x = x + D + 1$. For each $S \in CC$ that is of distance d from Q_i let $v(S) = m' + \frac{1}{4} + (x - d) \cdot \epsilon$. Let (S, S', \dots, Q_i) be a shortest path from S to Q_i in CC . Let $S \cap S'$ be $\{j\}$. Let the maximizing clause of S be $M_{S, j}$.
 - ii. Add CC to \mathcal{C} .
 - (c) Update $x = x + 1$. Let $v(Q_i) = m' + \frac{1}{4} + x \cdot \epsilon$. If there is an uncolored neighbor S' of Q_i , let $S' \cap Q_i$ be $\{j\}$ and let the maximizing clause of Q_i be $M_{Q_i, j}$. Otherwise the process terminates.

Since the distance d in step 2(b)i is also bounded by the diameter of the connected component, we are guaranteed that the value we assign during the process is always increasing. Therefore until we cannot find an uncolored neighbor in step (2c), the algorithm does not find a j -local maximum and therefore by Proposition 4.5 it does not find a traditional equilibrium. At the moment we cannot do this, the connected components newly colored in that step together with Q_i forms one connected component CC^* that was not colored in previous steps, and therefore is of size at least c . The crucial point is that \mathcal{Q} contains the set of neighbors of CC^* , and by the isoperimetric inequality, $|\mathcal{Q}| \geq |CC^*| \cdot (1 - \frac{4 \log |CC^*|}{3m'}) = \frac{2^{0.75m'-1}}{m'}$, as needed.

Acknowledgments

We thank Noam Nisan for pointing our attention to [13].

References

- [1] Nir Andelman and Yishay Mansour. Auctions with budget constraints. In *SWAT*. 2004.

- [2] Ashwinkumar Badanidiyuru, Shahar Dobzinski, and Sigal Oren. Optimization with demand oracles. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. ACM, 2012.
- [3] Kshipra Bhawalkar and Tim Roughgarden. Welfare guarantees for combinatorial auctions with item bidding. In *SODA*, 2011.
- [4] L. Blumrosen and N. Nisan. Combinatorial auctions (a survey). In N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [5] Deeparnab Chakrabarty and Gagan Goel. On the approximability of budgeted allocations and improved lower bounds for submodular welfare maximization and gap. *SIAM Journal on Computing*, 39(6):2189–2211, 2010.
- [6] G. Christodoulou, A. Kovács, and M. Schapira. Bayesian combinatorial auctions. *Automata, Languages and Programming*, pages 820–832, 2008.
- [7] Shahar Dobzinski, Noam Nisan, and Sigal Oren. Economic efficiency requires interaction. In *ACM STOC*. 2014.
- [8] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. In *ACM STOC*, 2005.
- [9] Shahar Dobzinski and Michael Schapira. An improved approximation algorithm for combinatorial auctions with submodular bidders, 2006.
- [10] U. Feige. On maximizing welfare when utility functions are subadditive. *SIAM J. Comput.*, 39(1):122–142, 2009.
- [11] Michal Feldman, Hu Fu, Nick Gravin, and Brendan Lucier. Simultaneous auctions are (almost) efficient. In *Proceedings of the 45th annual ACM symposium on Symposium on theory of computing*, pages 201–210. ACM, 2013.
- [12] Hu Fu, Robert Kleinberg, and Ron Lavi. Conditional equilibrium outcomes via ascending price processes with applications to combinatorial auctions with item bidding. In *ACM Conference on Electronic Commerce*, 2012.
- [13] Sergiu Hart. A note on the edges of the n-cube. *Discrete Mathematics*, 14(2):157–163, 1976.
- [14] Avinatan Hassidim, Haim Kaplan, Yishay Mansour, and Noam Nisan. Non-price equilibria in markets of discrete goods. In *ACM Conference on Electronic Commerce*, pages 295–296, 2011.
- [15] Eyal Kushilevitz and Noam Nisan. *Communication Complexity*. Cambridge University Press, New York, NY, USA, 1997.
- [16] B. Lehmann, D. Lehmann, and N. Nisan. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior*, 55(2):270–296, 2006.
- [17] Renato Paes Leme, Vasilis Syrgkanis, and Eva Tardos. Sequential auctions and externalities. In *SODA*, pages 869–886, 2012.
- [18] Brendan Lucier and Allan Borodin. Price of anarchy for greedy auctions. In *SODA*, 2010.

- [19] Noam Nisan. The communication complexity of approximate set packing and covering. In *ICALP*, pages 868–875, 2002.
- [20] Noam Nisan. The computational complexity of pure nash. 2009. Blog post in Turing’s Invisible Hand. November 19.
- [21] Tim Roughgarden. The price of anarchy in games of incomplete information. In *ACM Conference on Electronic Commerce*, pages 862–879, 2012.
- [22] Carla D. Savage and Peter Winkler. Monotone Gray codes and the middle levels problem. *Journal of Combinatorial Theory, Series A*, 70(2):230–248, 1995.
- [23] Vasilis Syrgkanis and Eva Tardos. Bayesian sequential auctions. In *ACM EC*, 2012.
- [24] Vasilis Syrgkanis and Eva Tardos. Composable and efficient mechanisms. In *STOC*, 2013.
- [25] Jan Vondrák. Optimal approximation for the submodular welfare problem in the value oracle model. In *STOC*, 2008.

A Local Maximum vs. Nash Equilibrium

Consider combinatorial auction with submodular bidders. A *local maximum* is an allocation of the items (S_1, \dots, S_n) such that for every $j \in S_i$ and player i' we have that $\sum_k v_k(S_k) \geq v_i(S_i - \{j\}) + v_{i'}(S_{i'} + \{j\}) + \sum_{k \neq i, i'} v_k(S_k)$. Already in [6] it was observed that if the valuations are submodular then for every allocation (S_1, \dots, S_n) that is a local maximum there are bids that together constitute a no overbidding equilibrium (this is the reasons why the price of stability in these games is 1: the welfare maximizing allocation is obviously a local maximum).

One could have hoped that to find an equilibrium with a good approximation one could just find a local maximum with a good approximation. However, finding a local maximum turns out to be PLS complete. To see this, we start with the PLS complete problem of finding a maximal cut and convert it to a combinatorial auction with two submodular bidders with identical valuations. Given a weighted graph, the set of items will be the set of items and the value of a set of vertices S is the weight of all edges that have at least one end point at S . Observe that a local maximum in the combinatorial auction corresponds to a maximal cut, and vice versa.

Interestingly, in our reduction it is easy to find a nash equilibrium in polynomial time: this is a direct consequence of our algorithm for a constant number of submodular bidders (the number of bidders is two in our reduction). This in particular proves that the set of equilibria strictly contains the set of local maxima (hence it is easier to find an equilibrium than a local maximum).